

Hitachi Ops Center Administrator

10.8.2

User Guide

Learn how to create, edit, and manage parity groups, pools, and volumes, protect data, monitor block storage, and create and manage file storage resources with Ops Center Administrator. You can also manage virtual storage machines and migrate volumes. For instructions on configuring storage for high availability, please refer to the *Hitachi Ops Center Administrator High Availability User Guide*.

© 2019, 2022 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 Vantara LLC (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 Vantara LLC at <https://support.hitachivantara.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 Vantara LLC.

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; and
2. Verifying that your 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 and Lumada are trademarks or registered trademarks of Hitachi, Ltd., in the United States and other countries.

AI/AS/400e, DB2, Domino, DS6000, DS8000, Enterprise Storage Server, eServer, FICON, FlashCopy, GDPS, HyperSwap, IBM, Lotus, MVS, OS/390, PowerHA, PowerPC, RS/6000, S/390, System z9, System z10, Tivoli, z/OS, z9, z10, z13, z14, z/VM, and z/VSE are registered trademarks or trademarks of International Business Machines Corporation.

Active Directory, ActiveX, Bing, Edge, Excel, Hyper-V, Internet Explorer, the Internet Explorer logo, Microsoft, the Microsoft corporate logo, the Microsoft Edge 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.

Copyright and license information for third-party and open source software used in Hitachi Vantara products can be found at <https://www.hitachivantara.com/en-us/company/legal.html>.

Contents

Preface.....	8
Intended Audience.....	8
Product version.....	8
Release notes.....	8
Document conventions.....	8
Conventions for storage capacity values.....	10
Accessing product documentation.....	11
Getting help.....	11
Comments.....	11
Chapter 1: Getting started.....	12
Overview of Ops Center Administrator.....	12
Unified management of block storage and file storage.....	13
Dashboard.....	14
Analyzing data in the dashboard.....	17
Inventory and resource information.....	19
Onboarding a storage system workflow.....	20
Configuring a storage system workflow.....	20
Creating and attaching volumes to servers workflow.....	21
Chapter 2: Adding a storage system.....	22
Chapter 3: Configure block storage.....	25
Managing port security and settings.....	25
Login WWN and iSCSI Name details.....	26
Editing iSCSI port settings.....	27
Adding a fabric switch.....	28
Supported fabric switches.....	29
Creating parity groups	30
Creating parity groups, basic method.....	31
Creating parity groups or hot spares, advanced method	33
Managing free and spare disks.....	35
Parity groups inventory.....	35
External parity groups inventory.....	38
Discovering external storage.....	40

Initializing external parity groups.....	41
Creating external parity groups.....	41
Deleting external parity groups.....	42
Enabling accelerated compression in Ops Center Administrator	43
Creating a pool.....	43
Pools inventory.....	44
Creating a pool, basic method.....	45
How Ops Center Administrator calculates pool sizes for the basic method of pool creation.....	49
Creating a pool, advanced method.....	49
Pool details.....	52
Update pool tier settings.....	58
Create, attach, and protect volumes with local replication.....	59
Creating, attaching, and protecting volumes with local replication.....	59
Managing storage systems.....	64
Storage system inventory.....	64
Storage system details	65
Using Storage Navigator for advanced storage configuration.....	71
Enabling the Storage Navigator Launch Settings option.....	72
Accessing Storage Navigator.....	73
Using Ops Center Administrator with Storage Advisor Embedded	73
Using NAS Manager for advanced file storage configuration	74
Accessing NAS Manager.....	74
Administrator types and responsibilities.....	74
Updating a storage system.....	75
Managing servers.....	75
Adding servers	75
Scanning host groups.....	79
Scanning Host Groups.....	79
Using Scan Host Groups to Migrate Information.....	80
Exporting servers.....	81
Server and server groups inventory.....	81
Server details.....	83
Updating a server.....	85
Deleting servers.....	88
Editing LUN paths.....	88
Reusing existing LUN paths	91
Managing server groups.....	93
Adding a server group.....	93
Updating a server group	93
Server group details.....	94
Managing volumes	94

Volumes inventory	95
About capacity saving	98
Creating volumes.....	98
Attaching existing volumes to a server.....	100
Attaching like volumes.....	104
Updating a volume.....	105
Volume details.....	107
Updating the tiering policy for a volume	108
Shredding volumes.....	109
Managing host groups.....	110
Host groups inventory	110
Creating host groups.....	110
Editing host groups.....	111
Host group details.....	112
Adding and deleting WWNs from a host group	113
Adding volumes to host groups as logical units.....	114
Chapter 4: Configure file storage	116
Managing file pools.....	116
File pool inventory.....	116
Creating file pools.....	117
File pool details.....	119
Expanding a file pool.....	120
Managing virtual file servers.....	120
Virtual file server inventory.....	120
Creating a virtual file server	121
Virtual file server details.....	121
Managing file systems.....	121
File system inventory	121
Creating file systems	122
File system details	122
Updating a file system.....	123
Managing shares and exports	123
Shares and exports inventory.....	123
Creating shares and exports.....	124
Export details.....	124
Updating an export	124
Share details	125
Adding existing account domain groups to a share.....	125
Updating a share.....	126

Chapter 5: Protecting data.....	128
Data protection concepts.....	128
Selecting volumes to protect.....	129
Protecting volumes with local replication	129
Protecting volumes by cloning	130
Protecting volumes by creating snapshots	130
Protecting existing volumes with high availability.....	131
Viewing replication groups.....	136
Replication groups inventory.....	137
Editing replication groups.....	138
Importing HA replication pairs from Ops Center Protector.....	140
Chapter 6: Manage virtual storage machines.....	142
Virtual storage machine inventory	142
Creating virtual storage machines	143
Removing undefined resources from a virtual storage machine	144
Moving defined volumes to a virtual storage machine.....	145
Adding resources to a virtual storage machine.....	146
Virtual storage machine details	147
Adding storage systems to a virtual storage machine.....	147
Adding undefined volumes to a virtual storage machine.....	148
Adding undefined host groups to a virtual storage machine.....	149
Removing storage systems from a virtual storage machine.....	150
Removing defined volumes from a virtual storage machine.....	150
Removing undefined volumes from a virtual storage machine.....	151
Removing host groups from a virtual storage machine.....	152
Chapter 7: Migrate volumes.....	154
Volume migration workflow.....	154
External migration of volumes over 4TiB.....	156
Attaching volumes to storage when the source is managed in Ops Center Administrator (for external migration).....	156
Discovering external storage.....	157
Discovering and virtualizing volumes.....	159
Migrating volumes	162
Managing external volumes	164
Managing migration tasks	165
Editing a volume migration task	166
Interrupting an in-progress volume migration	167
Workflow for cleanup after migration.....	168
Deleting a volume migration task.....	170
Unvirtualizing external volumes.....	170

Detaching volumes from storage.....	171
Chapter 8: Monitor block storage.....	172
Monitoring capacity.....	172
Monitoring hardware	173
Monitoring data protection	174
Monitoring jobs	175
Setting up to receive SNMP alerts.....	177
Chapter 9: Administer Ops Center Administrator.....	180
Administering security.....	180
Privileges and roles.....	180
Registering an AD domain using a DNS server	181
Adding an account domain	182
Assigning user roles.....	182
Registering an AD domain using the virtual appliance manager	184
Updating an AD domain registered using the virtual appliance manager	185
Adding a login warning banner.....	185
Backing up and restoring system settings	186
System elements preserved during backup.....	187
Downloading audit logs.....	187
Downloading maintenance information	187
Managing the Ops Center Administrator service with Docker.....	188
Managing the Ops Center Administrator service with Podman.....	189
Managing the JDK in Ops Center Administrator environment.....	190
Viewing the JDK version used by Ops Center Administrator.....	190
Replacing the JDK used by Ops Center Administrator.....	190
Troubleshooting the JDK replacement.....	191
Index.....	192

Preface

Hitachi Ops Center Administrator is an infrastructure management solution that unifies storage management solutions such as storage provisioning, data protection, and storage management; simplifies the management of large-scale data centers by providing smarter software services; and is extensible to provide better programmability and control.

Intended Audience

This document is intended for system administrators, Hitachi Vantara representatives, and authorized service providers who configure and operate Virtual Storage Platform storage systems with Hitachi Ops Center Administrator.

Readers of this document should be familiar with the following:

- RAID storage systems and their basic functions.
- Volume creation and management.
- Pool creation and management.
- Parity group creation and management.

Product version

This document revision applies to Hitachi Ops Center Administrator version 10.8.2 or later.

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 located on Support Connect at <https://knowledge.hitachivantara.com/Documents>.






Document conventions

This document uses the following typographic conventions:

Convention	Description
Bold	<ul style="list-style-type: none"> Indicates text in a window, including window titles, menus, menu options, buttons, fields, and labels. Example: Click OK. Indicates emphasized words in list items.
<i>Italic</i>	<ul style="list-style-type: none"> 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: <code>pairedisplay -g group</code> <p>(For exceptions to this convention for variables, see the entry for angle brackets.)</p>
Monospace	Indicates text that is displayed on screen or entered by the user. Example: <code>pairedisplay -g oradb</code>
< > angle brackets	<p>Indicates variables in the following scenarios:</p> <ul style="list-style-type: none"> Variables are not clearly separated from the surrounding text or from other variables. Example: <code>Status-<report-name><file-version>.csv</code> 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	<p>Indicates that you have a choice between two or more options or arguments. Examples:</p> <p>[a b] indicates that you can choose a, b, or nothing.</p> <p>{ a b } indicates that you must choose either a or b.</p>

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

Icon	Label	Description
	Note	Calls attention to additional information.

Icon	Label	Description
	Tip	Provides helpful information, guidelines, or suggestions for performing tasks more effectively.
	Important	Highlights information that is essential to the completion of a task.
	Caution	Warns the user of adverse conditions and/or consequences (for example, disruptive operations, data loss, or a system crash).
	CAUTION	Warns the user of a hazardous situation that, if not avoided, could result in major or minor injury.
	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 capacity values (for example, logical device capacity, cache memory capacity) are calculated based on the following values:

Logical capacity unit	Value
1 block	512 bytes
1 cylinder	Mainframe: 870 KiB

Logical capacity unit	Value
	Open-systems <ul style="list-style-type: none"> ▪ OPEN-V: 960 KiB ▪ Others: 720 KiB
1 KiB	1,024 (2^{10}) bytes
1 MiB	1,024 KiB or $1,024^2$ bytes
1 GiB	1,024 MiB or $1,024^3$ bytes
1 TiB	1,024 GiB or $1,024^4$ bytes
1 PiB	1,024 TiB or $1,024^5$ bytes
1 EiB	1,024 PiB or $1,024^6$ bytes

Accessing product documentation

Product user documentation is available on the Hitachi Vantara Support Website: <https://knowledge.hitachivantara.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

The [Hitachi Vantara Support Website](https://support.hitachivantara.com/en_us/contact-us.html) is the destination for technical support of products and solutions sold by Hitachi Vantara. To contact technical support, log on to the Hitachi Vantara Support Website for contact information: https://support.hitachivantara.com/en_us/contact-us.html.

[Hitachi Vantara Community](https://community.hitachivantara.com) is a global online community for Hitachi Vantara 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.hitachivantara.com, register, and complete your profile.

Comments

Please send us your comments on this document to doc.comments@hitachivantara.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 Vantara LLC.

Thank you!

Chapter 1: Getting started

Hitachi Ops Center Administrator is a unified software management tool that reduces the complexity of managing storage systems by simplifying the setup, management, and maintenance of storage resources.

The following sections describe Ops Center Administrator concepts.

Overview of Ops Center Administrator

Ops Center Administrator reduces infrastructure management complexities and enables a simplified approach to managing storage infrastructures. It has intuitive graphical user interfaces and recommended configuration practices to streamline system configurations and storage management processes.

You can leverage Ops Center Administrator to easily provision new storage capacity for business applications without requiring in-depth knowledge of the underlying infrastructure resource details. It provides centralized management while reducing the number of steps to configure, optimize, and deploy new infrastructure resources.



Some of the key Ops Center Administrator capabilities include:

- Simplified user experience for managing infrastructure resources. Visual aids enable easy viewing and interpretation of key management information, such as used and available capacity and guide features to help quickly determine the appropriate next steps for a given management task.
- Recommended system configurations to speed initial storage system setup and accelerate new infrastructure resource deployments.

- Integrated configuration workflows with Hitachi recommended practices to streamline storage provisioning and data protection tasks.
- Common, centralized management for supported storage systems.
- A REST-based API to provide full management programmability and control in addition to unified file-based management support.
- Automated SAN zoning during volume attach and detach. Optional auto-zoning eliminates the need for repetitive zoning tasks to be performed on the switch.

Unified management of block storage and file storage

Ops Center Administrator can be used to onboard and configure both block storage and file storage if NAS modules are included in the chassis of a supported storage system.

Understand block and file storage

- Block storage:

In block storage, volumes of storage are created. A server-based operating system can connect to each block of storage and control it as an individual hard drive. Each storage block can be individually formatted with the required file system, such as NTFS or VMFS. Block storage systems are typically deployed in a Storage Area Network (SAN) environment.

From the dashboard of Ops Center Administrator, you can discover, register, and onboard a block storage system.

- File storage:

Ops Center Administrator supports unified onboarding and configuration of file storage in the form of NAS modules.

If a supported storage system includes NAS modules, the file storage is automatically added with the block storage. File pools and other file resources can then be created in the Ops Center Administrator interface or by using the API.

Adding block and file storage together

Ops Center Administrator enables you to add block and file storage in a single step. The only requirements are the service processor (SVP) IP address, user name, and password. When the file storage is added, the cluster is automatically registered in Ops Center Administrator.

Unified configuration

After a storage system is onboarded, all block and file resources can be configured and managed from a single Storage System window. File pool creation workflow incorporates best practices that simplify workflow and enhance usability. The file pools are used to easily create virtual file servers, file systems, and shares and exports. File system creation automatically mounts and formats the new file system.

Unified reporting

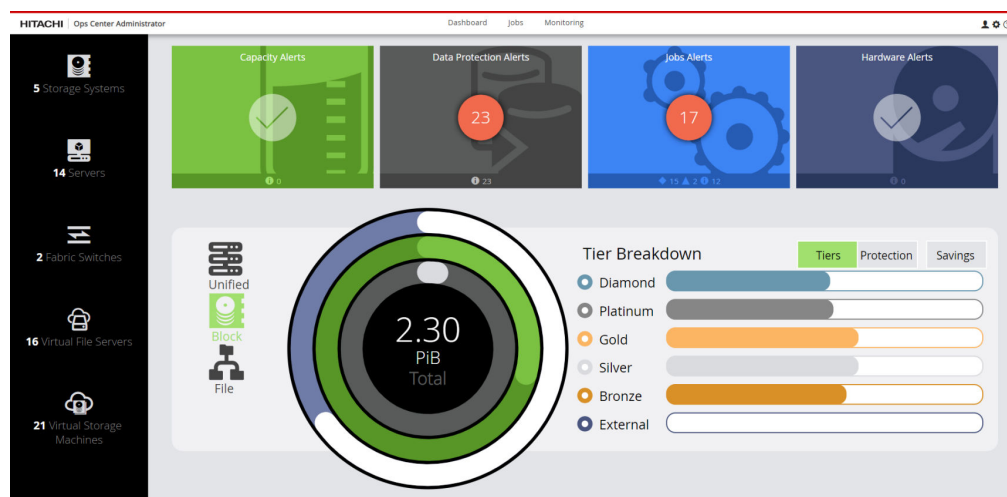
Capacity is reported for all aggregated storage systems in the dashboard.

Capacity is also displayed for each storage system in the Storage System detail window.

Three views of capacity are available: file only, block only, or a unified view of block and file.

Dashboard

After a storage system has been onboarded to Ops Center Administrator, the dashboard displays as soon as you log in. The Ops Center Administrator dashboard gives the tools to easily configure, manage, and monitor storage systems.



From the Ops Center Administrator dashboard, you can access managed resources and provision storage in the context of a given storage system or server. The provided templates and configurations make it possible to quickly and easily provision a storage system without knowing the details of the underlying hardware and software.

The top navigation menu gives access to Jobs and Monitoring windows. Links to the following settings are available, based on the user role:

- Tier Management
- Security
- SNMP
- Change Local Password
- Ops Center Protector

The dashboard has three distinct sections:

Resource side window

Gives quick access to review the configuration of your storage systems, servers, and fabric switches. If the storage system includes NAS modules, virtual file servers can also be accessed.

Alert tiles

Four alert tiles represent various aspects of the health of the storage system. When Ops Center Administrator detects a problem with a storage system environment, a number appears in the tile. The number indicates the number of alerts for that aspect of the storage system. Click the alert tile to go directly to a summary of the problems.

Resource summary

The middle area, with the information gauge, gives a summary of the capacity allocated from the registered storage systems.

Resource side window

The resource side window enables quick access to storage systems and to servers.

- Click Storage Systems to view and add storage systems.
- Click Servers to view and add servers.
- Click Fabric Switches to view and add fabric switches.
- Click Virtual File Servers to view and add virtual file servers. Displays if the storage system includes NAS modules.
- Click Virtual Storage Machines to view virtual storage machines and move volumes to a VSM or add and remove undefined resources.

Alert tiles

Across the top of the dashboard are tiles that display alerts for storage capacity, data protection, jobs, and hardware.

If a tile includes a circled check mark, there are no alerts for that part of the storage system, and everything is functioning normally. A number in a red circle within a tile indicates one or more problems with that part of the storage system.

You can click a tile for Capacity Alerts, Data Protection Alerts, or Hardware Alerts to view the summary for the category in the Monitoring tab.

The Jobs Alert tile displays the number of jobs in the last 24 hours with a status of Failed or Success with Errors.

Resource summary

The circular information gauge displays capacity metrics for the available storage.

- If the storage systems include file storage, you can click Block or File next to the information gauge to view a legend and capacity values for either type of storage. Click Unified to view a legend and capacity values for both block and file.
- For block-only storage systems, the numerical data for each capacity parameter in the ring is displayed to the left of the information gauge.
- The number in the center of the rings shows the total usable capacity of all storage systems. The total usable capacity is the capacity available from all the parity groups across all storage systems.

If you do not have any parity groups configured on the storage system, this number is zero and all other data points in the capacity visualization are zero.



Note: To understand uninitialized raw capacities, review the available unused disks on the detail window for each storage system.

- The light gray ring indicates the sum of all pool capacity available across all storage systems. The dark gray indicates the parity group capacity that is not yet allocated to pools.

If you do not have any pools created, the light gray ring indicates zero. As you create pools, this number increases to eventually become equal to the total usable capacity when you have consumed all parity groups for pool creation.



Note: Allocated to Pools plus Unallocated to Pools equals the Total Usable Capacity in the center of information gauge.

- The light green ring (Thin Used) indicates the storage utilization. As you create volumes on the pools and start consuming capacity, the utilization of thin pools increases, and the value in green increases.

If Thin Used starts to increase and get closer to your total pool capacity, the pools might be starting to fill up.

- Physical capacity allocated to file pools is indicated by medium blue in the File view and by light blue in the Unified view.
- File pool utilization is indicated by light blue in the File view and by medium blue in the Unified view.
- File over-commit capacity is represented by darkest gray in the outer ring of the File view.
- The subscribed capacity of all volumes, as a percentage, is represented by white in the outermost ring in the Block and Unified views. If the white ring extends outside the circle, this indicates oversubscription. Capacity subscription beyond the total available capacity is not an issue if your capacity utilization is well within the total capacity.
- Physical capacity, or total usable capacity across all parity groups, is represented by dark blue in the outermost ring in the Block and Unified views.

If you notice the total pool capacity (light gray) and Thin used (light green) values getting closer to total capacity, you might be running out of storage on one or more storage systems and might need to add disks to increase storage capacity. Review the information gauge for each storage system to identify which storage system needs additional capacity. In addition, inspect disks for each storage system to determine if there is unused capacity available for parity group creation.

The right side of the resource summary offers alternate views:

- *Protection*: is the breakdown of data protection metrics including a representation of types of protected, unprotected, and secondary capacity and gauge of the total percentage of capacity protected.
- *Tier Breakdown*: is a visualization of the amount of each tier that is allocated to pools.
- *Savings*:
 - *Data Reduction*: The ratio of logical used capacity to the physical used capacity, for all compression and deduplication technologies.



Note: The value displayed as the capacity of data after reduction includes the size of meta-data and garbage data generated by the storage system in addition to user data. The value might temporarily be greater than the capacity of data before reduction.

- *Capacity Efficiency*: The ratio of Thin Free plus Thin Used to the physical used capacity. Capacity efficiency is only calculated for volumes on Thin, Tiered, and Snap pools.
 - If disk-based compression is in use alone or in combination with controller-based compression, or with compression accelerator, the physical used capacity is the result from disk-based compression alone.
 - If only controller-based compression and/or compression accelerator are in use, the physical used capacity is the result from controller-based compression and/or compression accelerator.
 - If no compression is in use, then physical used capacity is the used capacity of the pool).

Analyzing data in the dashboard

The dashboard is a visual display of the important information required to analyze the overall capacity utilization and health of your storage system. It has visual indicators such as total usable capacity, current utilization, a data protection summary, and monitoring alerts.


Analyzing data shown in alert tiles

The alert tiles collectively present the health of the storage system environment. In one glance you can verify that the overall health is sound if you see no alerts on the alert tiles. Not having alerts indicates that there are no capacity or hardware issues in the environment, no failed jobs in the last 24 hours, and that the data protection is working without any issues.

If there are any alerts, you can drill down to the relevant alert window to investigate the cause. Ops Center Administrator has alerts for capacity utilization, hardware, data protection, and jobs status.

Analyzing data in the information gauge

The information gauge gives a visual indication of the total capacity of all storage systems managed by Ops Center Administrator.



The capacity indicated in the center of the ring is the total usable capacity available via the configured parity groups. After you add a storage system and configure parity groups, the total capacity indicator will show the capacity from the newly added storage system. The Thin Used capacity (light green ring) indicates the total capacity that is currently being used. If the usage is around 70-80% of the total capacity, you might receive capacity alerts based on the thresholds set by your storage administrator. The default thresholds are 70% and 80%, and can be changed during pool creation.

The light gray ring that gives a sum of capacities of all pools in the systems should be closer to 100% capacity. This means that you are using your entire parity group capacity by allocating it to pools. If the Thin Used capacity ring (light green) nears the total capacity (light gray ring) then you might run out of pool capacity soon. In such a case, expand the pool to consume more capacity.

If you notice that the total pool capacity (light gray ring) and Thin used (light green ring) values are getting closer to total capacity, you might be running out of disk capacity on one or more storage systems and would need to add disk space to increase storage capacity. Before adding disk space:

- Review the information gauge for each storage system to identify which storage system needs additional capacity.
- Review unused disks for each storage system to determine if any raw unused capacity is available for parity group creation.

Capacity subscription beyond the total available capacity should not be an issue if your thin capacity utilization is well within the total capacity.

Analyzing tier metrics

As parity groups are created, the various disk types become categorized into tiers. The tiers and corresponding disk types are as follows.

Table 1 Tier definitions

Tier	Disk type
Diamond	SCM NVMe
Platinum	FMD, FMD DC2, SSD, SSD(RI) and FMD HDE, SSD NVMe
Gold	SAS 15 k
Silver	SAS 10 k

Tier	Disk type
Bronze	SAS 7.2 k




Note: Adding all tier capacities together equals the Total Usable Capacity in the center of the information gauge.

Analyzing data protection metrics

The balance of your protected primary volumes and secondary volumes depends on the number of copies you have chosen to maintain, and on the type of the data protection technology being used. If you set aside more volumes for data protection, the overall usable capacity might be affected. However, if you have a large amount of unprotected data, you should consider data protection options.



Note: These data protection capacity numbers are based on oversubscribed allocations and, as a result, correlate with the overall oversubscription percentage, not the usable capacity numbers represented in the rest of the information gauge.



Inventory and resource information


The inventory windows display details about the storage system resources in Ops Center Administrator. These resources include storage systems, servers, ports, and pools, volumes, parity groups, external parity groups (if the storage system has external storage), and replication groups. If the storage system has NAS modules, resources will also include file pools, virtual file servers, file systems, and shares and exports. Common actions can be performed on the inventory windows, such as the following:

- You can select one or more resources and delete them.
When you delete a storage system, you disassociate it from Ops Center Administrator.
When you delete a pool or volume, the resource is de-provisioned and removed from the storage system.
- When a parity group is deleted, it is removed from the storage system and the disks used to create the parity group are no longer in use. You can delete the parity group to reconfigure the storage system with some other RAID configuration or simply decommission the array. If the parity group is in use by a pool, the parity group deletion will fail.

- You can select one or more of the same type of resources and update their properties. The properties that can be updated depend on the type of resource.
- You can click a particular resource to see more details about it on its resource detail window.
 - When you delete a block pool, the parity groups used by the pool will no longer be in In Use status. The pool volumes on these parity groups will be formatted and the parity group will eventually be in Available status.
 - When you delete a volume, the pool subscription will go down. Volume deletion will fail if the volume participates in data protection or is attached to a server.
 - When you delete a file pool, the underlying related block pool is deleted.
 - When you delete a server, the server is disassociated from Ops Center Administrator. You will no longer be able to provision volumes to the server (or its WWNs). Server deletion will fail if it has volumes attached to it.

Onboarding a storage system workflow

When you onboard a storage system, you associate it with Ops Center Administrator so that you can manage it from the dashboard. Then you can add servers to which you can attach volumes. Optionally, you can add fabric switches to auto-create zones during volume provisioning.



Note: If you use resource groups, make sure that the SVP username used to onboard storage systems in Ops Center Administrator has access to all custom resource groups and meta resource groups.

Configuring a storage system workflow

Configure a storage system to manage storage using Ops Center Administrator.



Note: Ops Center Administrator retrieves and displays any resources that already exist in the storage system.

Complete the following to configure a storage system:

- Enable port security: Port security must be enabled for a host storage domain (HSD) to be created on the port. If port security is disabled, Ops Center Administrator does not select the port for HSD creation. If the user manually selects the port with security disabled, the host group is created.
- Create parity groups so you can create pools.
- Create pools so you can create volumes.



Creating and attaching volumes to servers workflow

Ops Center Administrator shows a workflow in a single window to create volumes and immediately attach them to servers. Optionally, you can protect volumes as part of the same workflow.

Begin by choosing a server on which you want to create and attach volumes. As part of the workflow, Ops Center Administrator automatically attaches volumes to servers after you select settings. Finally, you are provided the option of immediately applying data protection.



Chapter 2: Adding a storage system

Onboarding a storage system is the process of associating it with Ops Center Administrator. After the storage system is onboarded, manage it from the Ops Center Administrator dashboard.



Note: Do not onboard the same storage system to multiple Ops Center Administrator instances. This may exhaust SVP resources, which causes the SVP to be slow or unresponsive to the Administrator server.

You can add multiple storage systems at the same time by specifying an SVP IP address that has multiple storage systems. You cannot add storage systems with different credentials at the same time.

To add a storage system without an SVP, you can specify either of the GUM IP addresses.

If a supported storage system includes NAS modules, the file storage is added automatically along with the block storage.

Ops Center Administrator also supports onboarding of 4-node clusters, which is a requirement for GAD Enhanced for NAS. GAD Enhanced for NAS takes advantage of global-active device to cluster two VSP G200, G/F400, G/F600, G/F800 systems with NAS modules across two sites.



Note: Onboarding a storage system also migrates server information from Hitachi Storage Advisor Embedded. For data integrity purposes, the Hitachi Storage Advisor Embedded provisioning function will be turned off.

See "Appendix A: Migrating to Ops Center Administrator" in the *Hitachi Ops Center Administrator Getting Started Guide* for details.

Before you begin

Ops Center Administrator needs access to all resources groups on the storage system so that the workflows function correctly. Verify that the service processor (SVP) user name used to onboard a storage system in Ops Center Administrator has access to all custom resource groups and meta resource groups.

The user must be a member of the Administration Users Group.

If the storage system includes NAS modules, the user must also be a member of the Support Personnel Group.

**Note:**

To receive storage systems alerts from Ops Center Administrator, the storage system must be set to SNMPv3. If storage systems have SNMPv1 or SNMPv2c settings configured, you can still manage onboarded storage systems, but alert information for these storage systems is not shown in Ops Center Administrator. For Ops Center Administrator to receive alerts from onboarded storage systems, change the SNMP to SNMPv3, then add the Ops Center Administrator server as an SNMP trap destination.

To launch Storage Navigator under one of the following conditions, the account used for onboarding must be authenticated locally by Storage Navigator:

- The storage system is VSP G1x00, F1500.
- The Storage Navigator Launch Setting option is disabled.



Note: To get storage system alerts from two or more VSP G1x00/VSP F1500 storage systems, enable Use a unique SNMP engine ID for all VSP G1x00 and VSP F1500 storage systems. To do this, see "Edit Advanced System Settings wizard" in the *Hitachi Virtual Storage Platform G1x00 and F1500 System Administrator Guide*.



Note: Some storage systems support upgrading the controllers by Controller Board replacement. After Controller Board replacement, update the storage system data in Ops Center Administrator by manually refreshing or waiting for the automatic refresh to complete. You can confirm that the data in Ops Center Administrator is successfully updated by verifying the following information:

- The MODEL in the storage system inventory and storage system details is updated.
- The report showing that the storage system model is updated is reported in the Jobs window.

Procedure

1. On the Ops Center Administrator dashboard, click **Storage Systems**.
2. Click the plus sign (+) to add a storage system.
3. Enter values for the following parameters in the **Onboard Storage System** window.

IP Address:

For a storage system with an SVP, enter the IP address (IPv4) of the external service processor for the storage system you want to discover.

User name and password:

Log in as a user who has administrator privileges on this storage system. For example, you can log in with the username `maintenance`.



Note: If the storage system has a virtual SVP, you can specify the SVP access port number following the IP address in the IP address field. The syntax is **IP-address:Port-number**.

4. Click **Submit**.

Result

When you successfully add a 4-node cluster, a message displays requesting that you ensure both storage systems in the global-active device pair are added before performing procedures.

The Jobs tab is updated with a job called `Create Storage System`. If you are adding multiple storage systems, there a job for each one.

Wait a while for Administrator to add the storage system. Refresh the Jobs tab to verify that the storage system has been onboarded.

The dashboard shows that the displayed number of storage systems has been incremented by one. Additionally, when you click Storage Systems, you are redirected to the storage system inventory window where you can see the newly added storage system.

When a storage system is onboarded, Ops Center Administrator goes through an initialization process where it gathers the information about the current configuration of the storage system. During this time, you see that the ports, volumes, pools, and parity groups in the storage system are "Not accessible". After the initialization is complete, you can see the port, pool, volume, and parity group information in the **Storage System detail** window.

Next steps

1. From the parity groups inventory window for the storage system, create parity groups to convert the raw disk capacity into usable capacity.
2. From the settings menu, access the tier definitions before creating pools.

Chapter 3: Configure block storage

Use Ops Center Administrator to configure the block storage.

Managing port security and settings

Before a host storage domain (HSD) can be created on a port, port security must be enabled. By default, security is disabled on supported storage systems. If port security is disabled, Ops Center Administrator does not select the port for host storage domain (HSD) creation.

If you select a port that has security disabled, you are limited to the default host group. If you manually select the port with security disabled, the host group is created.


Procedure

1. On the Ops Center Administrator dashboard, click **Storage Systems** to see the inventory of storage systems and capacity information.
2. Click a storage system to see its configuration of pools, ports, volumes, and parity groups.
3. Click **Ports** to see the configured storage ports for the storage system. The Ports window appears, shown as follows:



Note: Ports with NVMe mode are not displayed.


ALERTS	ID	WWN	SPEED	FABRIC	CONNECTION TYPE	SECURITY	VSM PORT	ATTRIBUTE
<input type="radio"/>	CL1-A	50:06:0E:80:07:27:46:00	AUTO	On	P-to-P	Enabled	No	Target
<input type="radio"/>	CL1-B	50:06:0E:80:07:27:46:01	AUTO	On	P-to-P	Enabled	No	Target
<input type="radio"/>	CL1-C	50:06:0E:80:07:27:46:02	AUTO	On	P-to-P	Enabled	No	Target
<input type="radio"/>	CL3-A	50:06:0E:80:07:27:46:20	AUTO	On	P-to-P	Enabled	No	Target
<input type="radio"/>	CL3-B	50:06:0E:80:07:27:46:21	AUTO	On	P-to-P	Enabled	No	Initiator
<input type="radio"/>	CL3-C	50:06:0E:80:07:27:46:22	AUTO	On	P-to-P	Enabled	No	Initiator
<input type="radio"/>	CL5-A	50:06:0E:80:07:27:46:40	AUTO	On	P-to-P	Enabled	No	RCU Target

4. To modify ports, do one of the following:
 - Select one or more Fibre Channel ports, then click **Edit Fibre Port** to open the **Edit Fibre Port** window, where you can change the security setting.
 - Select an iSCSI port, then click **Edit iSCSI Port** to open the **Edit iSCSI Port** window, where you can change security settings and IPv4 or IPv6 settings.
 5. For VSP G1x00, F1500 storage systems, you can update port attributes by selecting one of the following:
 - **Target:** Fibre target port
 - **Initiator:** MCU initiator port
 - **RCU Target:** RCU target port
 - **External:** External initiator port
 6. The VSP 5000 series can have either the **Target** attribute or the **Bidirectional** attribute. You can select one or the other in the **Edit Fibre Port** window.
-  **Note:** The bidirectional port attribute includes all the port attributes (such as Target and External). Ports with this attribute are available on some of storage system models.
7. Click **Enable Security** or **Disable Security**, then click **OK**.
A job is started to update the port security.


Modifying port attributes in Storage Navigator

You can open the **Ports inventory** window in Storage Navigator by clicking Open the Ports/ Hosts Groups/iSCSI Targets window in Storage Navigator to view and modify port attributes.


For more information, see the Storage Navigator online help.

 **Note:** Any changes you make in Storage Navigator may not be reflected in Ops Center Administrator for a few minutes.

Next steps

To view the Login WWN and iSCSI Name details, click  in the upper right-hand corner of the window.

Login WWN and iSCSI Name details

To view the Login WWN and iSCSI Name details, click  in the upper right-hand corner of the Ports window.

WWN	PORT ID	HOST GROUP NAME	STATUS
13:57:08:16:96:00:01:00	CL1-A	HostGroup-1A001	Logged-in
13:57:08:16:96:00:01:01	CL1-A	HostGroup-1A001	Logged-in
13:57:08:16:96:00:01:02	CL1-A	HostGroup-1A001	Logged-in
13:57:08:16:96:00:01:03	CL1-A	HostGroup-1A001	Logged-in
13:57:08:16:96:00:02:00	CL1-A	HostGroup-1A002	Logged-in
13:57:08:16:96:00:02:01	CL1-A	HostGroup-1A002	Logged-in
13:57:08:16:96:00:02:02	CL1-A	HostGroup-1A002	Logged-in
13:57:08:16:96:00:02:03	CL1-A	HostGroup-1A002	Logged-in
13:57:08:16:96:02:01:00	CL1-C	HostGroup-1C001	Logged-in
13:57:08:16:96:02:01:01	CL1-C	HostGroup-1C001	Logged-in
13:57:08:16:96:02:01:02	CL1-C	HostGroup-1C001	Logged-in
13:57:08:16:96:02:01:03	CL1-C	HostGroup-1C001	Logged-in
13:57:08:16:96:02:02:00	CL1-C	HostGroup-1C002	Logged-in
13:57:08:16:96:02:02:01	CL1-C	HostGroup-1C002	Logged-in
13:57:08:16:96:02:02:02	CL1-C	HostGroup-1C002	Logged-in
13:57:08:16:96:02:02:03	CL1-C	HostGroup-1C002	Logged-in
13:57:08:16:96:04:01:00	CL1-E	HostGroup-1E001	Logged-in
13:57:08:16:96:04:01:01	CL1-E	HostGroup-1E001	Logged-in
13:57:08:16:96:04:01:02	CL1-E	HostGroup-1E001	Logged-in
13:57:08:16:96:04:01:03	CL1-E	HostGroup-1E001	Logged-in

To see the iSCSI name details, click the iSCSI tab.

iSCSI NAME	PORT ID	iSCSI TARGET ALIAS	iSCSI TARGET NAME	STATUS
iqn.1994-04.jp.co.hitachi.rsd.950.1.30070...	CL1-G	10-G00	iqn.1994-04.jp.co.hitachi.rsd.950.1.30070...	Logged-in
iqn.1994-04.jp.co.hitachi.rsd.950.1.30070...	CL1-G	10-G00	iqn.1994-04.jp.co.hitachi.rsd.950.1.30070...	Logged-in
iqn.1994-04.jp.co.hitachi.rsd.950.1.30070...	CL1-G	10-G00	iqn.1994-04.jp.co.hitachi.rsd.950.1.30070...	Logged-in
iqn.1994-04.jp.co.hitachi.rsd.950.1.30070...	CL1-G	10-G00	iqn.1994-04.jp.co.hitachi.rsd.950.1.30070...	Logged-in
iqn.1994-04.jp.co.hitachi.rsd.950.1.30070...	CL1-G	HostGroup-1G001	iqn.1994-04.jp.co.hitachi.rsd.950.1.30070...	Logged-in
iqn.1994-04.jp.co.hitachi.rsd.950.1.30070...	CL1-G	HostGroup-1G001	iqn.1994-04.jp.co.hitachi.rsd.950.1.30070...	Logged-in
iqn.1994-04.jp.co.hitachi.rsd.950.1.30070...	CL1-G	HostGroup-1G001	iqn.1994-04.jp.co.hitachi.rsd.950.1.30070...	Logged-in
iqn.1994-04.jp.co.hitachi.rsd.950.1.30070...	CL1-G	HostGroup-1G001	iqn.1994-04.jp.co.hitachi.rsd.950.1.30070...	Logged-in
iqn.1994-04.jp.co.hitachi.rsd.950.1.30070...	CL1-G	HostGroup-1G002	iqn.1994-04.jp.co.hitachi.rsd.950.1.30070...	Logged-in
iqn.1994-04.jp.co.hitachi.rsd.950.1.30070...	CL1-G	HostGroup-1G002	iqn.1994-04.jp.co.hitachi.rsd.950.1.30070...	Logged-in
iqn.1994-04.jp.co.hitachi.rsd.950.1.30070...	CL1-G	HostGroup-1G002	iqn.1994-04.jp.co.hitachi.rsd.950.1.30070...	Logged-in
iqn.1994-04.jp.co.hitachi.rsd.950.1.30070...	CL1-G	HostGroup-1G002	iqn.1994-04.jp.co.hitachi.rsd.950.1.30070...	Logged-in
iqn.1994-04.jp.co.hitachi.rsd.950.1.30070...	CL1-G	HostGroup-1G003	iqn.1994-04.jp.co.hitachi.rsd.950.1.30070...	Logged-in
iqn.1994-04.jp.co.hitachi.rsd.950.1.30070...	CL1-G	HostGroup-1G003	iqn.1994-04.jp.co.hitachi.rsd.950.1.30070...	Logged-in
iqn.1994-04.jp.co.hitachi.rsd.950.1.30070...	CL1-G	HostGroup-1G003	iqn.1994-04.jp.co.hitachi.rsd.950.1.30070...	Logged-in
iqn.1994-04.jp.co.hitachi.rsd.950.1.30070...	CL1-G	HostGroup-1G003	iqn.1994-04.jp.co.hitachi.rsd.950.1.30070...	Logged-in
iqn.1994-04.jp.co.hitachi.rsd.950.1.30070...	CL1-G	HostGroup-1G004	iqn.1994-04.jp.co.hitachi.rsd.950.1.30070...	Logged-in
iqn.1994-04.jp.co.hitachi.rsd.950.1.30070...	CL1-G	HostGroup-1G004	iqn.1994-04.jp.co.hitachi.rsd.950.1.30070...	Logged-in
iqn.1994-04.jp.co.hitachi.rsd.950.1.30070...	CL1-G	HostGroup-1G004	iqn.1994-04.jp.co.hitachi.rsd.950.1.30070...	Logged-in
iqn.1994-04.jp.co.hitachi.rsd.950.1.30070...	CL1-G	HostGroup-1G004	iqn.1994-04.jp.co.hitachi.rsd.950.1.30070...	Logged-in

You can filter the display by:

- Port ID
- Status (Logged-In or Logged-Out)
- Host Group Name (Fibre tab only)
- iSCSI Target Alias (iSCSI tab only)
- iSCSI Target Name (iSCSI tab only)

Editing iSCSI port settings

Change the settings for port security and for IPv4 and IPv6.

Procedure

1. Navigate to the **Ports** window and select the **iSCSI** tab.
2. Select the port, then click **Edit iSCSI** settings.

3. Update the port security or the settings for IPv4 and IPv6, then click **Submit**.

Adding a fabric switch

You can add, update, or delete a fabric switch after onboarding a storage system in the Ops Center Administrator interface.

After you add a fabric switch, you can choose to auto-create zones during volume provisioning. A fabric switch is required for any process that uses auto-select, such as host group creation and auto-selection of ports while attaching volumes to servers.

Before you begin

Verify the following:

- Servers and ports are connected according to the manufacturer's instructions.
- There is an active zone set with at least one dummy zone available.
- The Ops Center Administrator server is connected to the same IP network and has access to the SNMP broadcast of Fibre Channel switches.
- You have the required information about the fabric switch:
 - Virtual Fabric ID (required only for Cisco switches)
 - Fabric Switch Type
 - Fabric Switch IP Address
 - Port Number
 - Username
 - Password
- You have the Admin role for the fabric switch.

Procedure

1. On the Ops Center Administrator dashboard, select **Fabric Switches** to open the **Fabric Switches** window.
2. Click the plus sign (+) to open the **Add Fabric Switches** window.

VIRTUAL FABRIC ID	FABRIC SWITCH TYPE	FABRIC SWITCH IP ADDRESS	PORT NUMBER	USERNAME	PASSWORD
Virtual Fabric Id	Fabric Switch Type	Fabric Switch IP Address	22	Username	*****

3. Enter the following information from the configuration of the switch you are adding:
 - **Virtual Fabric ID:** For Cisco switches, the VSAN ID. Not applicable to Brocade switches.
 - **Fabric Switch Type:** Select **Brocade** or **Cisco**.
 - **Fabric Switch IP Address**
To add or update a core switch, use the Management IP address of the switch or the Active CP IP address.
 - **Port Number**
 - **Username**
 - **Password**
4. Click **Submit**.

Result

A job is created to add the fabric switch.

Supported fabric switches

Ops Center Administrator supports the following Brocade® and Cisco® fabric switches.

- Brocade: Fabric OS 7.x through 8.2.2a
- Cisco: MDS NX-OS Release 6.2(9) or later

Creating parity groups

Parity groups are the basic units of storage capacity. Creating parity groups converts the raw disk capacity in your storage system into usable capacity.

Ops Center Administrator parity group concepts

Ops Center Administrator has a simple one-click method of parity group creation that is based on best practices applicable to the disks in the storage system.

An advanced method of parity group creation is also available. You can choose to use the advanced method if there is not a need to rely on best practices.



Note: You cannot use Ops Center Administrator to create parity groups on the following storage systems.

- VSP 5000 series
- VSP G1x00, F1500

These parity groups must be created by an authorized service representative, but you can initialize them in Ops Center Administrator. You can enable encryption in Storage Navigator.

Creating parity groups in Ops Center Administrator also creates LDEVs that can be consumed for pool creation.

Encryption can be enabled during parity group creation if prerequisites are met, including a storage system with an Encryption Disk Board.

Normal practice is to use all available disk capacity when creating parity groups to ensure that all the storage system capacity is usable. There can be exceptions to this practice, for example:

- If the entire capacity of the storage system is not needed.
- If there is a need to create fewer parity groups to reserve more disks as spares.

Viewing parity groups

Access the **Parity Groups** window by clicking Storage Systems on the dashboard and then clicking a storage system tile. Click the Parity Groups tile.

A summary of parity groups includes disk type, number of parity groups, capacity, and available spares, all sorted by disk type.

- You can click Manage Spare Disks to open the Disk Management window and set free disks as spare disks, or spares as free disks.



Note: This feature is not available for the following storage systems:

- VSP 5000 series
- VSP G1x00, F1500

- Click any parity group tile to view its details.
- Click the plus sign (+) to open the Create Parity Groups window.

When you open the Create Parity Groups window, the basic method is selected by default. Encryption can be enabled using either method.

- Basic option: Creating a parity group using the basic option needs no input, but you can change the RAID type or the number of parity groups. Ops Center Administrator applies best practices for creating these parity groups.
- Advanced option: This option allows the user to fully configure the RAID layout of the parity group by selecting the specific disks to assign for parity group creation.

Enabling parity group encryption in Storage Navigator

Click Open the **Parity Groups** window in Storage Navigator to enable parity group encryption.

For more information, see the Storage Navigator online help.



Note: Any changes you make in Storage Navigator might not be reflected in Ops Center Administrator for a few minutes.

Creating parity groups, basic method

The following procedure describes the basic option for creating a parity group and enabling encryption.



Note: You cannot use Ops Center Administrator to create parity groups on the following storage systems.

- VSP 5000 series
- VSP G1x00, F1500

These parity groups must be created outside of Ops Center Administrator by an authorized service representative.

Before you begin

- Register the storage system.
- Identify the target storage system name.
- Identify the total capacity that you expect to use. Plan to use all of the available disks in the system when you create parity groups.

Procedure

1. On the Ops Center Administrator dashboard, select **Storage Systems** in the resource side window to see the inventory of registered storage systems.
2. Click a storage system to create and configure the parity groups for it.
3. Click **Parity Groups** to see the inventory of all parity groups in the storage system.
4. Click the plus sign (+). In the **Create Parity Groups** window, review the list of unused disk types in the storage system. This information is grouped by disk type, disk speed, and disk capacity, and includes the following details:
 - Number of available disks.
 - Available spares detected for each disk type, disk speed, and capacity.
 - Number of new or additional spares to reserve. This calculation is based on the total spares needed based on best practices, and the number of existing spares in the system.
 - The correct RAID configuration for the disk type.
 - The number of parity groups that can be created.
 - The total usable capacity that is available based on the number of parity groups and the RAID configuration.

Create Parity Groups

Storage System 410031

DISK TYPE	TOTAL DISKS	AVAILABLE DISKS	AVAILABLE SPARES	ADDITIONAL SPAR...	RAID TYPE	PARITY GROUPS	TOTAL CAPACITY	ENCRYPTION
FMD DC2 1.60TB	32	9	2	0	RAID6 (6D+2P)	1	8.73 TiB	ON

ENCRYPTION: ON OFF

Basic Advanced Cancel Submit

5. Decide whether the RAID configuration for each disk type is acceptable. Choose one of the following options:
 - Accept the RAID configuration, which uses the full capacity of the installed drives.
 - Change the RAID configuration or create fewer parity groups. Ops Center Administrator shows the number of parity groups that can be created for the new RAID configuration and the corresponding usable capacity.
6. (Optional) Click Encryption **ON** to use encryption.

Prerequisites for enabling encryption:

- The storage system must have an Encryption Disk Board.
- Encryption License Key must be installed.
- The Key Management Server must be configured on the SVP.



Note: Encryption cannot be disabled in Ops Center Administrator.

7. Click **Submit**.

Result

A job is started to create the parity group for the storage system. This job includes the following tasks:

- Identifies the appropriate number and position for the spare disk.
- Assigns a spare disk.
- Creates the required number of parity groups for the requested RAID layout.
- Creates and quick formats the necessary volumes on the parity group so that it is ready for pool creation.
- The job might create sub-jobs when multiple parity groups are being created. Each sub-job shows the status of the parity groups being created.

Next steps

- Verify the status of the parity group creation job by clicking Jobs.



Note: Parity group creation might take a long time.

Creating parity groups or hot spares, advanced method

The advanced option enables users to configure parity groups that are not based on the provided best practices.

You can fully configure the RAID layout of the parity group by selecting the specific disks and hot spares to assign for parity group creation.



Note: You cannot use Ops Center Administrator to create parity groups on the following storage systems.

- VSP 5000 series
- VSP G1x00, F1500

These parity groups must be created outside of Ops Center Administrator by an authorized service representative.

Procedure

1. On the dashboard, click **Storage Systems**, then click a storage system tile.

- Click **Parity Groups** to open the **Parity Groups** window, then click the plus sign to open the **Create Parity Groups** window.
- Click **Advanced**.

- Select a **RAID Type** and a **RAID Layout**.

RAID Type	RAID Layout
RAID5	3D+1P, 4D+1P, 6D+1P, 7D+1P
RAID6	6D+2P, 12D+2P, 14D+2P
RAID1+0	2D+2D

- (Optional) Click Encryption **ON** to use encryption.

Prerequisites for enabling encryption:

- The storage system must have an Encryption Disk Board.
- Encryption License Key must be installed.
- The Key Management Server must be configured on the SVP.



Note: Encryption cannot be disabled in Ops Center Administrator.

- Choose available disks. Select enough disks to match the RAID Layout.



Note: FMD HDE drives with encryption disabled can be added to parity groups but the capacity cannot be used in pools.

- Click **Submit**.

Result

A job is started to create the parity group for the storage system.

Next steps

- Verify the status of the job by clicking Jobs.



Note: Parity group creation might take time.

- Create a pool.

Managing free and spare disks

You can set free disks to spare and spare to free.



Note: This feature is not available for the following storage systems:

- VSP 5000 series
- VSP G1x00, F1500

If you choose to use the basic method to create parity groups, Ops Center Administrator will automatically review the available spare disks and allocate more spare disks if needed.

If you choose to create parity groups using the advanced method, you should review the number of spare disks in the parity groups inventory summary. To assign more or fewer spare disks, use disk management.

Procedure

1. On the dashboard, click **Storage Systems**.
2. Click a storage system tile, then click **Parity Groups**.
3. Click **Manage Spare Disk** to open the **Disk Management** window and view disks not being used in any parity group.
4. Click **Disks to Set to Free** or **Disks to Set to Spare** and select one or more disks. Selecting Disks to Set to Spare reserves the disks as a spare disks and they cannot be used for parity group creation. Selecting Disks to Set to Free unallocates the spare disks and return then to the pool of available disks.
5. Click **Submit**.

Parity groups inventory

After parity groups are created, they are displayed for each storage system.

To view parity groups for a storage system, click Storage Systems on the dashboard, then click a storage system tile to view the associated parity groups, pools, volumes, and ports. Click Parity Groups to open the Parity Groups window.



Note: If a parity group is created outside of Ops Center Administrator and consists of multiple disk types, only one disk type is displayed in the Disk Type column.

You can perform the following actions on this window:

- If there are hardware alerts related to the parity groups, you can click the Hardware Alerts tile to view details for different types of components on the Monitoring window for the storage system.
- To assign free and spare disks, click Manage Spare Disks to open the Disk Management window.



Note: This feature is not available for the following storage systems:

- VSP 5000 series
- VSP G1x00, F1500

- Select one or more parity groups and do one of the following:

- To initialize parity groups, click Initialize Parity Groups.

You can create and initialize pool volumes on a parity group if pool volumes have not been created. You can also format any pool volumes that are in blocked state. You can initialize parity groups under the following conditions:

- When the parity group either does not have any pool volumes or at least one of the pool volumes is in Blocked state. In this case, you will notice that the parity group status is UNINITIALIZED.
- When the pool volumes on the parity group do not account for the entire parity group capacity and there are unused partitions on the parity group. In this case you, will notice that the parity group status is Available but the available capacity is much less than the total capacity. This is supported only when the parity group status is AVAILABLE_PHYSICAL, IN_USE, or UNINITIALIZED.
- To delete the parity group, click Delete. When you delete a parity group, it is removed from the storage system and the disks used to create the parity group are no longer in use. You can delete the parity group to reconfigure the storage system with some other RAID configuration or simply decommission the array.



Note: You cannot delete the following parity groups:

- Parity groups on VSP 5000 series and VSP G1x00, F1500 storage systems.
- Parity groups that contain NAS startup volumes.

- To enable accelerated compression on parity groups, click Enable Compression on Parity Groups.

Compression can be enabled for unencrypted parity groups that were created using FMD DC2 or FMD HDE disks. The parity groups must also be in IN_USE status.



Note: The IN_USE status means that the parity group is in a pool. You can create a pool on an existing parity group that was created using FMD DC2 or FMD HDE drives. Then identify the parity group in the detail window for the pool and enable compression in the Parity Groups window.

Parity groups that contain NAS startup volumes cannot be compressed.

After compression is enabled, the status of the parity group is IN_USE.

- Click the plus sign (+) to create a new parity group.



Note: Parity groups on the following storage systems are created outside Ops Center Administrator by an authorized service representative:

- VSP 5000 series
- VSP G1x00, F1500

- Use the filters to view the status of parity groups. Click a filter to apply it, then click it again to remove it.
 - **AVAILABLE:** Parity group is not being used for any storage pools. It is available for pool creation.
 - **AVAILABLE_PHYSICAL:** Parity group is not being used for any storage pools. Compression is enabled on the parity group and physical capacity of the parity group is available for pool creation. Not applicable to parity groups that are not enabled for compression.
 - **FORMATTED:** At least one of the pool volumes in the parity group is being formatted. The full format takes longer than a quick format because the service processor fully scans the hard drive.
 - **QUICK_FORMATTING:** At least one of the pool volumes in the parity group is being formatted using the quick format method. To make the formatting process quick, the drive is not fully checked, the files are still there, and the volume could be re-built to gain access to the files again.
 - **IN_USE:** Parity group is being used by a storage pool.
 - **UNINITIALIZED:** The parity group either has no volumes or at least one of the pool volumes is in Blocked status, or one or more partitions is uninitialized and has a size greater than 16,787,456 blocks.
 - **UNSUPPORTED_ATTACHED:** At least one of the pool volumes has a path to a storage port.
 - **UNSUPPORTED_INACCESSIBLE_RESOURCE_GROUP:** Parity group and at least one of its pool volumes are in different resource groups and user does not have access to one of the resource groups.

External parity groups inventory

External parity groups are parity groups in a storage system that are connected to a storage system that has been onboarded in Ops Center Administrator. You can display a list of external parity groups in a storage system.

Procedure

1. Click the **External Parity Groups** tile on the detail window for a storage system. If the storage system has no external parity groups, the tile does not display.

HITACHI Ops Center Administrator Dashboard jobs Monitoring

Dashboard > Storage Systems > 39303 > External Parity Groups

External Parity Groups

Search Filter

0 selected

Loaded all 320 items

ID	EXTERNAL STORAGE SYSTEM	EXTERNAL VENDOR	EXTERNAL MODEL	STATUS	TOTAL	EXTERNAL PATHS
1-1	93010505	HITACHI	HUS	UNSUPPORTED_ATTACHED	15.00 GiB	CL4-8_50060E801013A75A_0
1-10	410438	HITACHI	VSP Gx00	UNINITIALIZED	2.75 GiB	CL4-8_50060E801228C636_0
1-100	410438	HITACHI	VSP Gx00	UNSUPPORTED_ATTACHED	1.00 GiB	CL4-8_50060E801228C636_101
1-101	410438	HITACHI	VSP Gx00	UNSUPPORTED_ATTACHED	1.00 GiB	CL4-8_50060E801228C636_102
1-102	410438	HITACHI	VSP Gx00	UNSUPPORTED_ATTACHED	1.00 GiB	CL4-8_50060E801228C636_105
1-103	410438	HITACHI	VSP Gx00	UNSUPPORTED_ATTACHED	100.00 GiB	CL4-8_50060E801228C636_106
1-104	410438	HITACHI	VSP Gx00	UNSUPPORTED_ATTACHED	100.00 GiB	CL4-8_50060E801228C636_107
1-105	410438	HITACHI	VSP Gx00	UNSUPPORTED_ATTACHED	100.00 GiB	CL4-8_50060E801228C636_108
1-106	410438	HITACHI	VSP Gx00	UNSUPPORTED_ATTACHED	1.00 GiB	CL4-8_50060E801228C636_109
1-107	410438	HITACHI	VSP Gx00	UNSUPPORTED_ATTACHED	1.00 GiB	CL4-8_50060E801228C636_111
1-108	410438	HITACHI	VSP Gx00	UNSUPPORTED_ATTACHED	1.00 GiB	CL4-8_50060E801228C636_112
1-109	410438	HITACHI	VSP Gx00	UNSUPPORTED_ATTACHED	1.00 GiB	CL4-8_50060E801228C636_114
1-11	210944	HITACHI	HUS VM	AVAILABLE	4.00 GiB	CL4-8_50060E80132AC900_6
1-110	410438	HITACHI	VSP Gx00	UNSUPPORTED_ATTACHED	100.00 GiB	CL4-8_50060E801228C636_115
1-111	410438	HITACHI	VSP Gx00	UNSUPPORTED_ATTACHED	7.00 GiB	CL4-8_50060E801228C636_117
1-112	410438	HITACHI	VSP Gx00	UNSUPPORTED_ATTACHED	7.00 GiB	CL4-8_50060E801228C636_118
1-113	410438	HITACHI	VSP Gx00	UNSUPPORTED_ATTACHED	7.00 GiB	CL4-8_50060E801228C636_119
1-114	410438	HITACHI	VSP Gx00	UNSUPPORTED_ATTACHED	7.00 GiB	CL4-8_50060E801228C636_120
1-115	410438	HITACHI	VSP Gx00	UNSUPPORTED_ATTACHED	7.00 GiB	CL4-8_50060E801228C636_121
1-116	410438	HITACHI	VSP Gx00	UNSUPPORTED_ATTACHED	7.00 GiB	CL4-8_50060E801228C636_122
1-117	410438	HITACHI	VSP Gx00	UNSUPPORTED_ATTACHED	500.00 GiB	CL4-8_50060E801228C636_124

show 200 items 2 Next



Note: External parity groups are not available for compression. If external parity groups have been used to create a pool, you can use their capacity to create volumes in Ops Center Administrator.

Use the filters to view the status of the external parity groups. Click a filter to apply it, then click it again to remove it.

EXTERNAL STORAGE SYSTEM

The ID of the external storage system.

EXTERNAL VENDOR

The external storage system vendor.

EXTERNAL MODEL

The external storage system model.

STATUS

The parity group's status can be one of the following:

- **IN_USE:** The external parity group is being used by a storage pool.
- **AVAILABLE:** The external parity group is not being used for any storage pools. It is available for pool creation and expansion.
- **UNINITIALIZED:** The external parity group either has no volumes, at least one of the pool volumes is in Blocked status, or one or more partitions is uninitialized and has a size greater than 16,777,216 blocks.
- **FORMATTING:** The volumes in the external parity groups are being formatted.
- **UNSUPPORTED_ATTACHED:** At least one of volumes has a path to a storage port or volumes are used for multiple pools.
- **UNSUPPORTED_INACCESSIBLE_RESOURCE_GROUP:** The external parity group and at least one of its pool volumes are in different resource groups and the user does not have access to one of the resource groups when using another configuration tool.

TOTAL

The total capacity of the storage system in GiB.

EXTERNAL PATHS

The path of the volume in the external storage system.

Discovering external storage

Before you begin

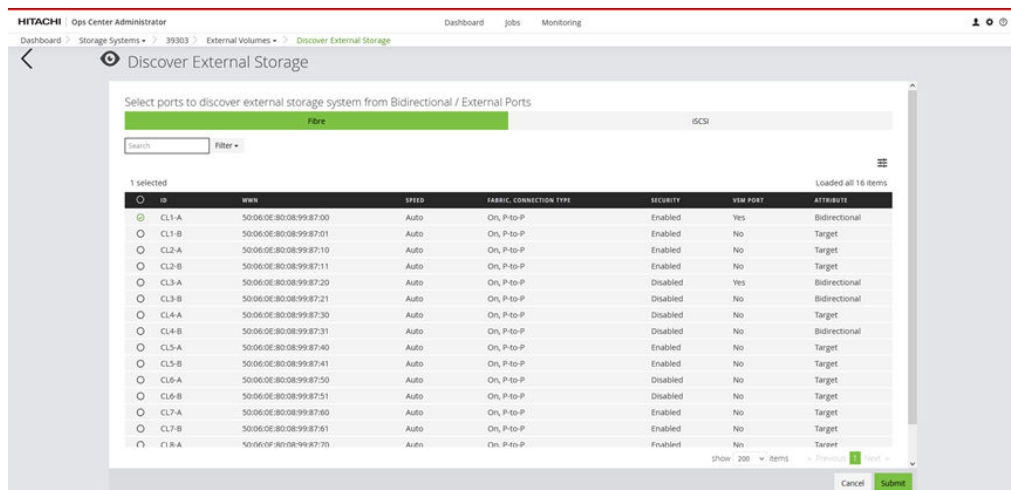
You must first attach volumes to the target storage system.

Procedure

1. Click the **External Parity Groups** tile on the detail window for a storage system. If the storage system has no external parity groups, the tile does not display.
2. Click the **Discover External Storage** (🔍) icon and select the ports that are attached to the storage system.



Note: The bidirectional port attribute includes all port attributes (such as Target and External). Ports with this attribute are available on some of storage system models.



3. Click **Submit**.

Initializing external parity groups

You can create and initialize pool volumes on an external parity group if pool volumes have not been created. You can also format any pool volumes that are in blocked state. You can initialize external parity groups under the following conditions:

- When the external parity group either does not have any pool volumes or at least one of the pool volumes is in Blocked state. In this case, the external parity group status is UNINITIALIZED.
- When the pool volumes on the external parity group do not account for the entire external parity group capacity and there are unused partitions on the external parity group. In this case, the external parity group status is AVAILABLE_PHYSICAL, IN_USE, or UNINITIALIZED.



Note: If you try to initialize an external parity group in a Blocked state, the job may take a long time to complete, or fail with a timeout error because the quick format option is not available on the storage system. If the job fails, check the parity group state in the External Parity Groups Inventory.

Procedure

1. Navigate to **External Parity Groups** on the storage system.
2. To initialize external parity groups, click **Initialize External Parity Groups**.

Creating external parity groups

You can create external parity groups with virtualizing volumes in external storage systems.

Before you begin

You must attach volumes to storage and discover external storage.

Procedure

1. Navigate to **External Parity Groups** on the storage system.
2. Click the plus (+) sign to open the **Create External Parity Groups** window. A list of discovered LUNs appears.

Select Discovered Volumes. To update discovered volumes information, discover external storage systems from the external volumes inventory.

☐ Select All

DEVICE ID	PRODUCT ID	S/N	CAPACITY	MAPPED	PORT	LUN ID	WWN/ISCSI INFO
HITACHI 504099990000	VSP Gx00	445566	10.00GiB	YES	CL1-A	0	50:06:0E:80:06:CC:7E:00
HITACHI 504099990001	VSP Gx00	445566	10.00GiB	YES	CL1-A	1	50:06:0E:80:06:CC:7E:00
HITACHI 504099990002	VSP Gx00	445566	10.00GiB	YES	CL1-A	2	50:06:0E:80:06:CC:7E:00
HITACHI 504099990003	VSP Gx00	445566	10.00GiB	YES	CL1-A	3	50:06:0E:80:06:CC:7E:00
HITACHI 504099990004	VSP Gx00	445566	10.00GiB	YES	CL1-A	4	50:06:0E:80:06:CC:7E:00
HITACHI 504099990005	VSP Gx00	445566	10.00GiB	YES	CL1-A	5	50:06:0E:80:06:CC:7E:00
HITACHI 504099990006	VSP Gx00	445566	10.00GiB	YES	CL1-A	6	50:06:0E:80:06:CC:7E:00

Showing 100 of 100 Discovered Volumes

Select Mapped Status

Select Vendor

Select Serial Number

Cancel Submit

You can filter the discovered volumes by mapped status, vendor, or serial number.

3. Select the discovered LUNs and click **Submit**.


Next steps

After creating external parity groups completes successfully, you can create or update pools with the external parity groups.

Deleting external parity groups

You can delete external parity groups.

Procedure

1. Navigate to **External Parity Groups** on the storage system.
2. Select an external parity group or groups to delete and click the **Delete External Parity Groups** () icon.



Note: You cannot delete the following external parity groups:

- External parity groups on VSP 5000 series and VSP G1x00, F1500 storage systems.
- External parity groups that contain NAS startup volumes.

Enabling accelerated compression in Ops Center Administrator

Accelerated compression can be enabled for parity groups on FMD DC2 or FMD HDE drives that are part of the following storage systems:

- VSP 5000 series
- VSP G1x00, F1500
- VSP G200, G/F400, G/F600, G/F800
- VSP G/F350, G/F370, G/F700, G/F900

This enables you to realize more virtual capacity than the actual usable capacity.



Note: Accelerated compression is not compatible with encryption. If a parity group is encrypted when created, the data cannot be compressed.

Ops Center Administrator only uses the physical capacity of the FMD DC2 or FMD HDE parity groups for which write is assured.

Accelerated compression can be enabled in the in the Parity Groups inventory window. The parity groups must be in IN_USE status.

If a pool includes parity groups that have accelerated compression enabled, you can view the following in the detail window for the pool:

- **Expansion Ratio:** The ratio of the total capacity of FMD DC2 or FMD HDE pool volumes with respect to the total capacity of FMD DC2 or FMD HDE pool volumes assured for writing.
- **Compression Ratio:** The ratio of data size compression.
- **Savings %:** Displays the percentage of the capacity reduced by data compression with respect to the used capacity before the data compression. This will not display a value until data is compressed.

If the savings percentage for a pool is not high enough, you can provision additional pool volumes to the pool from the parity groups using FMD DC2 or FMD HDE drives.

Creating a pool

Use the basic option to take advantage of tiers that are based on best practices.

To get more flexibility and to not need to take advantage of best practices, you can use the advanced option to select specific parity groups.

The types of pools are as follows:

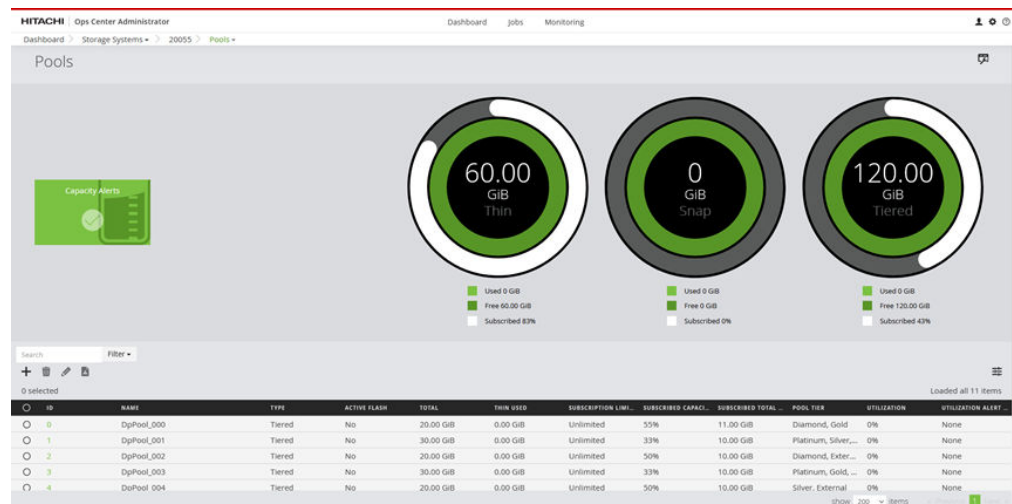
- DP (Dynamic Provisioning), which is functionality that allocates virtual volumes to a host and uses the physical capacity that is necessary according to the data write request.
- Tiered, which is used with Dynamic Provisioning and places data in a hardware tier according to the I/O load. For example, 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.
- HTI (Thin Image), which stores snapshot data in pools. A pool consists of multiple pool-VOLs. The pool-VOLs contain the snapshot data. A pool can contain up to 1,024 pool-VOLs.

Pools inventory

Access the **Pools** window to add, update, and delete pools.

From the dashboard, click Storage Systems and then click a storage system tile to view parity groups, pools, volumes, and ports.

Click Pools to view the inventory of pools.



View pool details in either a tile view or a list view. Capacity utilization is represented by a color bar:

- Green: under 70%
- Orange: from 70% to 80%
- Red: over 80%

A key in a tile or row indicates that the pool is encrypted.

You can filter the volumes as follows:

- Free Space: Filter the pools by the amount of free space they have.
- Size of Pool: Filter by pool size.
- Pool Type: Filter by the pool type: Thin, Tiered, or Snap.
- Pool Tier: Filter by the pool tier: Diamond, Platinum, Gold, Silver, Bronze, or External.

- Active Flash Enabled: Filter by whether Active Flash is enabled.
- DDM Enabled: If True, the pool can be used to migrate volumes greater than 4TiB.
- Encryption: Filter by whether encryption is used. Select YES, NO, PARTIAL, or UNKNOWN.
- FMD Compression: Filter by whether a pool is using a parity group that is enabled for accelerated compression. Select YES, NO, PARTIAL, or UNKNOWN.
- Deduplication Enabled: Select YES or NO.

The following actions are available in this window:

- Click Open the Pools window in Storage Navigator to open the Pools window of Storage Navigator in a separate browser window.
- Select one or more pools, then click Delete to delete the pools. Deleted pools are de-provisioned and removed from the storage system.
- Select a pool, then click Edit to update it in the Update Storage Pool window.
- Select a Tiered Pool, then click Update Tiering Settings to update it in the Update Tiering Settings window.
- Click the plus sign (+) to add a pool in the Create Pool window.



Note: Any changes you make in Storage Navigator might not be reflected in Ops Center Administrator for a few minutes.

Creating a pool, basic method

Use the basic method of pool creation to create pools based on best practices.



Note: You cannot create a pool with external parity groups using the basic method. You can only use the advanced method. See [Creating a pool, advanced method \(on page 49\)](#) for more information.

Before you begin

Create and configure parity groups on the storage system.

You need a minimum of four parity groups of the Bronze, Silver, or Gold tiers, or one parity group of the Platinum or Diamond tier to create a pool using the basic method. Otherwise, you can use the advanced method of pool creation.

License requirements:

- For a Dynamic Tiering pool: Dynamic Tiering
- For a Thin pool: Dynamic Provisioning
- For a Thin Image pool: Thin Image
- For active flash: active flash

Procedure

1. On the Ops Center Administrator dashboard, click **Storage Systems** to see the inventory of registered storage systems.

2. Click a storage system to create a pool for it.
3. Click **Pools**.
4. Click the plus sign (+) to open the **Create Pool** window.

By default, the **Basic** option is selected.

Create Pool

Pool Label:

Storage System:

Intend to use for Snap? ☐ Yes ☒ No

0 GiB

☐ Diamond
☐ Platinum
☐ Gold

Select capacity from Tiers to allocate to Pool

Tier	Capacity	Available
Diamond	0 GiB	20.51 TiB Available
Platinum	0 GiB	87.50 TiB Available
Gold	0 GiB	0 GiB Available
Silver	0 GiB	0 GiB Available
Bronze	0 GiB	0 GiB Available

Utilization Threshold (Low) %:

Utilization Threshold (High) %:

Subscription Limit: ☒ Unlimited

If this threshold is exceeded, snapshots will be disabled to ensure there is enough capacity for user data.

Basic Advanced Cancel Submit

5. Enter a **Pool Label**.

Pool labels can contain only alphanumeric characters, hyphens, and underscores. Initial hyphens are not allowed.

6. In the **Select capacity from Tiers to Allocate to Pool** pane, you can choose storage from 1, 2, or a maximum of 3 tiers (**Diamond**, **Platinum**, **Gold**, **Silver**, or **Bronze**).

- If you select only one tier, you can use the **Intend to use for ?** toggle to decide whether to use Snap pools.
 - Click **Yes** to create Snap pools.



Note: Use this option to use data protection to create snapshots. You can also create a pool for snapshots later by returning to the Create Pool window.

- Click **No** to create Dynamic Provisioning (HDP) pools. These pools can also be used to create snapshots.
- If you select two or three tiers, the system creates Tiered pools. The following options are available for a Tiered pool:
 - **Tiering:** Select **Automatic** or **Manual** monitoring and tier relocation.
 - **Cycle Time:** Select the cycle of performance monitoring and tier relocation. For automatic tiering, select one of the monitoring cycles from drop down. Cycle time can be 0.5, 1, 2, 4, 8 or 24 hours.
 - **Monitoring Period:** When 24 hours is selected in the cycle time list, specify the time from 00:00 to 23:59 (default value), in which performance monitoring is to be performed. Set one or more hours between the starting time and the ending time. The starting time must not be later than the ending time.
 - **Monitoring:** Specifies the monitoring mode. If you perform the tier relocation weighted to the past period monitoring result, select **Continuous**. If you perform the tier relocation on the specified cycle, select **Periodical**.
 - **Relocation Speed:** Specifies the window relocation speed. You can set the speed to: 1 (Slowest), 2 (Slower), 3 (Standard), 4 (Faster), or 5 (Fastest). The default is 3 (Standard). If the speed specified is slower than 3 (Standard), the data drive load is low when tier relocation is performed.
 - **Buffer Space for Relocation:** Specifies the storage area that is necessary for processing tier relocation. Set this as percentage of the tier size. For each tier enter an integer value from 2 to 40 as the percentage to set for that tier. The default value for each tier is 2%.
 - **Buffer Space for New window Assignment:** Specifies the storage area that is necessary for processing new window assignments. Set this as a percentage of the tier size. For each tier enter an integer value from 0 to 50 as the percentage to set for that tier. For tier 1, this value must be between 0 and 50. The default value is 8% for all drive types.



Note: **Tiering**, **Cycle Time**, **Monitoring Period** and **Relocation Speed** can only be set in environments with an SVP.

Table 2 Tier definitions

Tier	Disk type
Diamond	SCM NVMe
Platinum	FMD, FMD DC2, SSD, SSD(RI), FMD HDE, SSD NVMe
Gold	SAS 15 k
Silver	SAS 10 k
Bronze	SAS 7.2 k

7. Click a selected tier to view the available storage capacity and select a capacity size.
8. Click **Tier Management** to see the disk type, capacity, and speed of each pool category.
9. Review the high and low pool utilization thresholds. The thresholds serve as the Warning and Critical thresholds for monitoring capacity. Adjust the thresholds if needed.
If you are creating a Thin (HDP) pool you can also choose whether to permanently suspend snapshots when usage exceeds the Critical threshold to reserve capacity for user data. If the threshold is exceeded, the pairs become suspended (PSUE) and the S-VOLs can never accept read-write processes. You can still write to P-VOLs. If this option is selected, a message displays in the Utilization graph in the detail window for the pool.



Note: This option is available for the following storage systems:

- VSP 5000 series
- VSP E series
- VSP G1x00, F1500 with microcode version 80-05-4x and later.
- VSP G200, G/F400, G/F600, G/F800 with firmware version 83-04-4x and later.
- VSP G/F350, G/F370, G/F700, G/F900
- VSP N400, N600, N800

10. Select the **Subscription Limit %** check box to set the limit to Unlimited.



Note: Subscription Limit % is not applicable to Snap pools or to the following storage systems:

- VSP 5000 series with firmware versions 90-01-6x
- VSP E990 with firmware version 93-02-02-60/8x or earlier
- VSP G/F350, G/F370, G/F700, G/F900

11. Click **Submit**.

Result

A job is started to allocate the storage capacity and create the pool.

Next steps

- Verify the status of the pool creation job by clicking Jobs.
- Create a volume.

How Ops Center Administrator calculates pool sizes for the basic method of pool creation

Ops Center Administrator incorporates best practices to calculate the best pool sizes based on the media available in the storage system. When the basic pool creation option is used, pool size is determined as follows:

- Ops Center Administrator identifies all the parity groups of the same disk type, disk capacity, disk speed, RAID type and layout and determines the available capacity of their usable LDEVs. Only parity groups in the Available and Quick Formatting states are eligible for pool creating using the basic method.
- Ops Center Administrator then uses combinations of these parity groups where four or more parity groups of the Bronze, Silver, or Gold tiers can be added together, or there is a parity group of the Platinum or Diamond tier, to compute various possible pool sizes that can be created for the combination of disk type, disk capacity, disk speed, RAID type and layout.



Note: If your parity groups do not meet these requirements, you can only use the advanced method of pool creation.

- Based on the Storage System model, Ops Center Administrator determines the maximum pool size.
- Ops Center Administrator then displays the pools in increasing order of pool size.

Creating a pool, advanced method

You can use the advanced method of pool creation to select parity groups filtered by disk type, speed, layout, and RAID level. This method does not use best practices.



Note: Best practice is not to mix different disk types in a Thin pool. If you do choose to create such a pool, Ops Center Administrator identifies the tier in this pool as "mixed".

Before you begin

- Create and configure parity groups on the storage system.
- License requirements:
 - For a Dynamic Tiering pool: Dynamic Tiering
 - For a Thin pool: Dynamic Provisioning
 - For a Thin Image pool: Thin Image
 - For active flash: active flash

Procedure

1. On the Ops Center Administrator dashboard, click **Storage Systems** to see the inventory of registered storage systems.
2. Click a storage system to create a pool for it.
3. Click **Pools**.
4. Click the plus sign (+) to open the **Create Pool** window, then click **Advanced**.

5. Enter a **Pool Label**.
Pool labels can contain only alphanumeric characters, hyphens, and underscores. Initial hyphens are not allowed.
6. Select **Yes** under **Intend to enable DDM** to use the pool to migrate volumes greater than 4 TiB.
7. Click the **Type** list and choose one of the following: **HDP**, **Tiered**, or **Snap**. You can create snapshots on DP pools and on Snap pools.

The following options are available for a tiered pool :

- **Tiering:** Select **Automatic** or **Manual** monitoring and tier relocation.
- **Cycle Time:** Select the cycle of performance monitoring and tier relocation. For automatic tiering, select one of the monitoring cycles from drop down. Cycle time can be 0.5, 1, 2, 4, 8 or 24 hours.
- **Monitoring Period:** When 24 hours is selected in the cycle time list, specify the time from 00:00 to 23:59 (default value), in which performance monitoring is to be performed. Set one or more hours between the starting time and the ending time. The starting time must not be later than the ending time.
- **Monitoring:** Specifies the monitoring mode. If you perform the tier relocation weighted to the past period monitoring result, select **Continuous**. If you perform the tier relocation on the specified cycle, select **Periodical**.
- **Relocation Speed:** Specifies the window relocation speed. You can set the speed to: 1 (Slowest), 2 (Slower), 3 (Standard), 4 (Faster), or 5 (Fastest). The default is 3 (Standard). If the speed specified is slower than 3 (Standard), the data drive load is low when tier relocation is performed.

- **Buffer Space for Relocation:** Specifies the storage area that is necessary for processing tier relocation. Set this as percentage of the tier size. For each tier enter an integer value from 2 to 40 as the percentage to set for that tier. The default value for each tier is 2%.
- **Buffer Space for New window Assignment:** Specifies the storage area that is necessary for processing new window assignments. Set this as a percentage of the tier size. For each tier enter an integer value from 0 to 50 as the percentage to set for that tier. For tier 1, this value must be between 0 and 50. The default value is 8% for all drive types.



Note: Tiering, Cycle Time, Monitoring Period and Relocation Speed can only be set in environments with an SVP.

8. Choose whether to enable active flash, which needs capacity from the Platinum or Diamond tier.
9. Click **Tier Management** to see the definitions of disk type, capacity, and speed of each pool category.
10. Select one or more parity groups to use to create the pool. External parity groups can also be used. You can scroll through the parity groups to use the search and filter functions.
 - For a Snap pool, select one or more parity groups with identical disk types.
 - For an HDP pool, select one or more parity groups with identical disk types.
 - For a Tiered pool, select one or more parity groups with two or three disk types.

Ops Center Administrator shows the total size of the pool using the selected parity groups.

11. Set the utilization thresholds or use the default threshold settings.
The thresholds are used as the Warning and Critical thresholds for Capacity Monitoring. If you are creating a (HDP) pool you can also choose whether to permanently suspend snapshots when usage exceeds the Critical threshold to reserve capacity for user data. If the threshold is exceeded, the pairs become suspended (PSUE) and the S-VOLs can never accept read-write processes. You can still write to P-VOLs. If this option is selected, a message displays in the Utilization graph in the detail window for the pool.



Note: This option is available for the following storage systems:

- VSP 5000 series
- VSP E series
- VSP G1x00, F1500 with microcode version 80-05-4x and later.
- VSP G200, G/F400, G/F600, G/F800 with firmware version 83-04-4x and later.
- VSP G/F350, G/F370, G/F700, G/F900
- VSP N400, N600, N800

12. The subscription limit can be set higher than 100%. Select the **Subscription Limit %** check box to set the limit to Unlimited.



Note: Subscription Limit % is not applicable to Snap pools or to the following storage systems:

- VSP 5000 series with firmware versions 90-01-6x
- VSP E990 with firmware version 93-02-02-60/8x or earlier
- VSP G/F350, G/F370, G/F700, G/F900

13. Click **Submit**.

Result

A job is started to allocate the storage capacity and create the pool.

Next steps

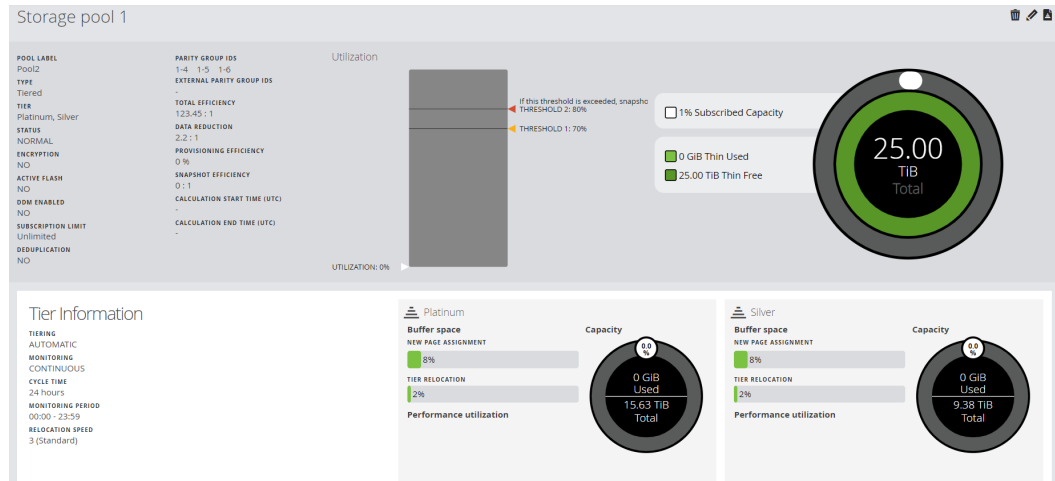
- Verify the status of the pool-creation job by clicking Jobs.
- Create volumes.

Pool details

To access pool details, click Storage Systems. Then click a storage system to view all resources. Click Pools and then click a pool tile.



Note: An external pool cannot be edited or deleted.



Note: Some values cannot be displayed in an environment without an SVP. In that case, a hyphen displays.

In the pool details window, you can:

- View the attributes of the pool.
- View a graph of utilization and thresholds.

- If the pool is based on a parity group using accelerated compression ("FMD DC2 Compression"), you can also view compression details, including expansion ratio and compression ratio.
- For a tiered pool, view the following tier information:

Pool tier usage for the pools supported on the storage system. Ops Center Administrator supports Bronze, Silver, Gold, and Platinum or Diamond tiers. For each tier, you can view statistics for buffer space, tier relocation and performance utilization. Click Edit to update the storage from tier that has been allocated to the pool.

- View volumes carved from the pool and perform procedures on them. Available procedures include deletion, updating, attaching and detaching to and from servers, protecting, unprotecting, and restoring. Click a volume in tile or list view to open the detail window for the volume.

The Provisioning Status can be any of the following:

- Attached: the volume is attached to a server.
- Unattached: the volume is not attached to a server.
- Unmanaged: the volume has only a LUN path associated. For example, a volume might be attached to a server that is not known to Ops Center Administrator.

The Data Protection can be any of the following:

- Protected: the volume is protected.
 - Unprotected: the volume is not protected.
 - Secondary: the volume is a snapshot or clone of another volume.
- If the volume uses compression or deduplication and compression, you can view the compression ratio and deduplication ratio.
 - Delete the pool.
 - Edit the pool configuration by clicking the pencil icon to open the Update Storage Pool window.
 - Edit a tiered pool configuration by clicking Update Tiering Settings to open the Update Tiering Settings window.
 - After a pool is created, you can edit the pool label or expand the pool. Basic and advanced options, like those for pool creation, are available for pool expansion. The basic option includes choices for a new/expanded pool size based on the current set of parity groups in the pool. You can choose to expand an existing tier in the pool or add a new tier to the pool. If you add a tier to a thin pool, it will be expanded and converted into a tiered pool. Using the advanced option enables you to add more parity groups to the pool and increase capacity.



Note: To expand a HDP pool, you must add a parity group of the same RAID layout, the same disk type, the same RAID level, and the same disk speed.

Click Submit to update the window and start a job that will update the pool.

If the pool contains capacity that has had compression technology applied, you can view the following:

Total Efficiency:

The ratio of the total saving effect achieved by accelerated compression, capacity saving (compression and deduplication), snapshot, and Hitachi Dynamic Provisioning.

Shows for the following storage systems:

- VSP E series
- VSP G/F350, G/F370, G/F700, G/F900 with firmware version 88-03-0x or later
- Compression and Deduplication metrics:
 - Software Saving: The capacity reduction ratio before and after capacity saving.
 - Software Compression: The capacity compression ratio before and after capacity saving.
 - Software Deduplication: The capacity deduplication ratio before and after capacity saving.
 - Software Pattern Matching: The capacity reduction ratio before and after capacity saving pattern matching.
- FMD Compression:
 - FMD Saving: The capacity reduction ratio before and after capacity saving.
 - FMD Compression: The capacity compression ratio before and after accelerated compression.
 - FMD Pattern Matching: The capacity reduction ratio before and after accelerated compression pattern matching.
- Calculation End Time: The end date and time for the calculation. The date and time are displayed in UTC.
- Calculation Start Time: The start date and time for the calculation. The date and time are displayed in UTC.
- Provisioning Efficiency: The efficiency ratio achieved by Dynamic Provisioning.

Capacity Efficiency:

Capacity efficiency is the ratio of Thin Free plus Thin Used to the physical used capacity. Capacity efficiency is only calculated for volumes on Thin, Tiered, and Snap pools.

Shows for the following storage systems:

- VSP E series
- VSP G/F350, G/F370, G/F700, G/F900 with firmware version 88-03-0x or later
- Compression and Deduplication metrics:
 - Software Saving: The capacity reduction ratio before and after capacity saving.
 - Software Compression: The capacity compression ratio before and after capacity saving.
 - Software Deduplication: The capacity deduplication ratio before and after capacity saving.
 - Software Pattern Matching: The capacity reduction ratio before and after capacity saving pattern matching.
- FMD Compression:
 - FMD Saving: The capacity reduction ratio before and after capacity saving.
 - FMD Compression: The capacity compression ratio before and after accelerated compression.
 - FMD Pattern Matching: The capacity reduction ratio before and after accelerated compression pattern matching.
- Calculation End Time: The end date and time for the calculation. The date and time are displayed in UTC.
- Calculation Start Time: The start date and time for the calculation. The date and time are displayed in UTC.
- Provisioning Efficiency: The efficiency ratio achieved by Dynamic Provisioning.

Data Reduction:

The ratio of logical used capacity to the physical used capacity, for all compression and deduplication technologies.




Note: The value displayed as the capacity of data after reduction includes the size of meta-data and garbage data generated by the storage system in addition to user data. For this reason, the value might be greater temporarily than the capacity of data before reduction.




The following actions on volumes are available on this window:


- Click the plus sign (+) to add volumes on the Create Volumes window.
- Click a volume to view details and attach, update or delete the volume.
- Select a volume, then click Edit to update the volume by editing the name or updating the Capacity Saving setting and command device settings. You can also expand the capacity.



- You can also edit the names of multiple volumes and update command device settings at the same time if you select multiple volumes.

- Select one or more volumes to perform one of the following actions:
 - To delete volumes, click Delete. When you delete a volume, it is de-provisioned and removed from the storage system.

 **Note:** If deduplication or compression has been enabled on the volume and the volume has data, it cannot be deleted in Ops Center Administrator. Access Storage Navigator to delete the volume.

 - To attach volumes to one or more servers on the Attach Volumes window, click Attach Volumes.
 - To detach volumes from servers, click Detach Volumes.
-  **Note:** The command device configuration will be disabled if all the attached LUN paths are removed.
- To apply data protection on the Protect Volumes window, click Protect Volumes.
-  **Note:** Command devices cannot be protected. Do not perform this action for command devices.
-  **Note:** Ops Center Administrator can protect volumes with high availability in existing replication group with following conditions:

 - P-VOL STATUS and S-VOL STATUS of all paired volumes in the replication group is PAIR.
 - Status of the corresponding replication in Ops Center Protector is OK.
- To unprotect a volume, select it, then click Unprotect.
-  **Note:** Ops Center Administrator can unprotect volumes with high availability in the replication group with following conditions:

 - P-VOL STATUS and S-VOL STATUS of all paired volumes in the replication group is PAIR.
 - Status of the corresponding replication in Ops Center Protector is OK.
-  **Note:** A volume in an "External" replication group cannot be unprotected.
- To restore a secondary volume to the primary volume, click Restore to open the Restore Volume window.
- To attach the volume to a storage system to prepare for volume migration, click Attach to Storage.
- To migrate volumes, click Migrate Volumes.
-  **Note:** Command devices cannot be migrated. Do not perform this action for command devices.
- To detach volumes from storage, click Detach from Target Storage.



Note: The command device configuration will be disabled if all the attached LUN paths are removed.

- To shred volumes that have been migrated, click Shred Volumes.
- Select a tiered volume, then click Update Tiering Policy to update the tiering policy.
- To attach like volumes, click Attach Like Volumes.

Update pool tier settings

You can edit the tiering for a single pool.

Procedure

1. In the detail window for a pool, click **Update Tiering Settings** to update the following tiering settings:
 - **Tiering:** Select **Automatic** or **Manual** monitoring and tier relocation.
 - **Cycle Time:** Select the cycle of performance monitoring and tier relocation. For automatic tiering, select one of the monitoring cycles from drop down. Cycle time can be 0.5, 1, 2, 4, 8 or 24 hours.
 - **Monitoring Period:** When 24 hours is selected in the cycle time list, specify the time from 00:00 to 23:59 (default value), in which performance monitoring is to be performed. Set one or more hours between the starting time and the ending time. The starting time cannot be set later than the ending time.
 - **Monitoring:** Specify the monitoring mode. If you perform the tier relocation weighted to the past period monitoring result, select **Continuous**. If you perform the tier relocation on the specified cycle, select **Periodical**.
 - **Relocation Speed:** Specify the window relocation speed. You can set the speed to: 1 (Slowest), 2 (Slower), 3 (Standard), 4 (Faster), or 5 (Fastest). The default is 3 (Standard). If the speed specified is slower than 3 (Standard), the data drive load is low when tier relocation is performed.
 - **Buffer Space for Relocation:** Specify the storage area that is necessary for processing tier relocation. Set this as percentage of the tier size. For each tier enter an integer value from 2 to 40 as the percentage to set for that tier. The default value for each tier is 2%.
 - **Buffer Space for New window Assignment:** Specify the storage area that is necessary for processing new window assignments. Set this as a percentage of the tier size. For each tier enter an integer value from 0 to 50 as the percentage to set for that tier. For tier 1, this value must be between 0 and 50. The default value is 8% for all drive types.



Note: Tiering, Cycle Time, Monitoring Period and Relocation Speed can only be set in environments with an SVP.

2. Click **Submit**.

Create, attach, and protect volumes with local replication

You can create volumes and attach them to servers and then apply data protection in a single workflow by first selecting a server.

Creating, attaching, and protecting volumes with local replication

You can start the create-and-attach workflow by selecting a server and creating volumes.

When you create volumes in Ops Center Administrator you can:

- Create multiple volumes of the same size or different sizes at the same time.
- Select the specific pool for volume creation or let Ops Center Administrator automatically select the best pool based on utilization.
- Specify a common label and starting label suffix for identical volumes that are the same size and have the same pool requirement.

Before you begin

- Create parity groups.
- Create pools.
- Add servers.

Procedure

1. On the Ops Center Administrator dashboard, click **Servers** to see the inventory of servers.
2. Do one of the following to open the **Create, Attach and Protect Volumes** window:
 - Select a server, click **Attach Volumes**, and select **Create, Attach, and Protect Volumes with Local Replication**.
 - Click a server tile to open the **Server <ID>** details window. Click the cylinder icon (🔧) in the upper right area and select **Create, Attach and Protect Volumes with Local Replication**.

3. Configure volumes for the specified storage system.

You can switch to another storage system by using the **Storage System** list. To add the volume to a virtual storage machine, use the **Virtual Storage Machine** list.

- Select the number of volumes.
- Enter the volume label and select a suffix for it.
- Select the volume size.
- Select the volume unit: **GiB**, **TiB**, or **Blocks**.
- (Optional) Select a **Volume ID Range** to specify one for the new volume for granular management of volumes. For instance, for ease of management, you can assign a certain range of IDs to certain departments in your organization. ID ranges can be specified in Decimal or Hexadecimal. The default selection is **Auto**, which means that Ops Center Administrator automatically selects the ID for the volume. You can also specify a **Virtual ID / Range**.
- Select the pool type: **Thin** or **Tiered**.
- For a Thin pool, select the pool tier: **Diamond**, **Platinum**, **Gold**, **Silver**, or **Bronze**. If the storage system has available capacity from external storage, you can also select the **External** tier.
- (Optional) Select the pool from the list of available pools. The default selection is **Auto Selected**, which means that Ops Center Administrator selects the best pool for provisioning the volume based on utilization and tier requirements.



Note: DDM pools are not available and cannot be used to create volumes.

- Optionally, select a type of **Capacity Saving**: **Compression** or **Deduplication and Compression**.



Note: You can set capacity saving for volumes based on tiered pools for the following storage systems:

- VSP 5000 series
- VSP G1x00, F1500 with microcode version 80-05-4x or later.



Note: If you choose **Deduplication and Compression** and later want to update the volume to **Compression** you must first disable **Capacity Saving**.

- j. (Optional) For a tiered pool, select the **Tiering Policy** from the list. Tiering policy choices available in the list depend on the choice of pool that was made in the previous step. Tiering policy choice is not available for auto-selected pools.
4. When you have made your choices, click the plus sign (+) to add a volume row to the list of volumes being created. Add more rows as needed.
5. Click **Next** to attach volumes to the selected servers.

# OF VOL	LABEL	LABEL SUFFIX	SIZE	POOL TYPE	TIER	POOL
1	2525		1 GiB	Tiered		Auto Selected

6. The **Host Mode** is set by default to the server operating system. You can make a selection if needed.
The server OS Type is provided when the server is added to Ops Center Administrator.
7. The prepopulated **Host Mode Option** depends on the **Host Mode** selection. The default Host Mode Option can be changed manually.
Default values are set only for **VMWARE EX** and **WIN EX** host modes. The default for all other Host Modes is none.
Ops Center Administrator identifies all host groups containing any of the server WWNs. If all of those host groups have the same host mode and host mode options, those settings are prepopulated with the same settings in the host groups.
8. Select **Mandate LUN Alignment**.
This option specifies whether to assign the same LUN number to multiple servers for a volume. If **Yes** is specified, the same LUN number is always assigned, and if **No** is specified, the same value is not necessarily used. If attachment is to only one server, this setting has no effect.
9. The **Auto Create Zone** is set to **No** by default, which means that no auto-zone is created. You can set it to **Yes** to automatically create zones.
Physical connections must be available on the switch. If these exist, Ops Center Administrator creates an auto-zone between the ports you selected in the **Attach Volumes** window.

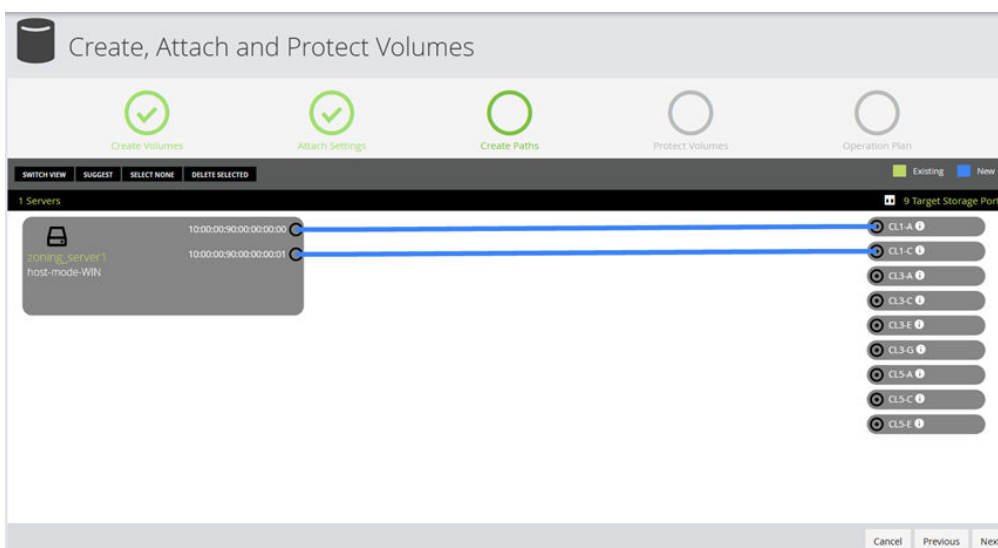
Note: If there was already a zone connected between those two ports, an existing zone is used. (A new one will not be created.)

10. You can confirm which Volume IDs will be used by clicking **View Proposed Volume ID Selection**.

To mandate using volume IDs shown in pop-up dialog, select **Yes** under **Mandate Using Displayed Volume ID**.

Note: The **Mandate Using Displayed Volume ID** selection will be reset to **No** if you move back to the previous step.

11. Click **Next** to view options for creating and editing LUN paths. The displayed path configuration of the **Create Paths** window depends on the selected **Host Mode** and **Host Mode Option** set in the **Attach Settings** window.



12. In the **Create Paths** window, you can view servers and their WWNs, along with ports on the storage system. Ops Center Administrator scans for existing host groups on the storage array and try to reuse them by default.

For VSP 5000 series storage systems, you can click a port to display a list of ports with their DKC locations, CTL locations, and redundancy levels.

Port Redundancy Level				
NO.	PORT ID	DKC LOCATI...	CTL LO...	REDUNDANC...
1	CL4-A	DKC-1	CTL11	DKC
2	CL2-A	DKC-1	CTL11	DKC
3	CL6-A	DKC-1	CTL11	DKC
4	CL8-A	DKC-1	CTL11	DKC
5	CL1-J	DKC-2	CTL21	DKC
6	CL3-J	DKC-2	CTL21	DKC
7	CL5-J	DKC-2	CTL21	DKC
8	CL7-J	DKC-2	CTL21	DKC
9	CL2-J	DKC-3	CTL31	DKC
10	CL4-J	DKC-3	CTL31	DKC

Total Ports: 15

OK

The following options are available for managing LUN paths:

- If you connect more than one server to the same port, the **Share a host group with all servers** box displays. Select the box to add the servers to a single host group.



Note: If host groups already exist on that port and a single host group cannot be created, the check box will not appear.

To prevent the volume from being added to an existing path, click the path to highlight it, then click **Delete Selected**.

Existing paths are populated as follows: all existing host groups with one or more server WWN and the exact same host mode and host mode options selected in the Attach Settings window are populated as paths.

- Click **Suggest** to populate automatically selected paths. By default, the least-used ports are selected. Suggest paths need both server and storage ports to be logged into the fabric switches in the Ops Center Administrator inventory.



Note: For VSP 5000 series storage systems, the first port suggested has the fewest LUN definitions and the second port suggested has the fewest number of LUNs in the highest redundancy level.

- To manually create a path, click a WWN and a port to connect them with a blue connector line. Click the connector again and then click **Delete Selected** to delete the connection.
- Optionally, you can click **Switch View** to use tables to create paths by selecting Server Ports to attach to Storage Ports.

13. When you are satisfied with the paths, click **Next** to view options for protecting volumes.

The screenshot shows the 'Protect Volumes' configuration window. On the left, a sidebar lists configuration steps: 'Create Volumes' (checked), 'Attach Settings' (checked), 'Create Paths' (checked), 'Protect Volumes' (active), and 'Operation Plan' (disabled). The main area is titled 'Protect Volumes' and includes a 'Schedule' section with tabs for 'Hourly', 'Daily', 'Weekly', and 'Monthly'. The 'Hourly' tab is selected, showing a 'Start Minute' of 0 and a frequency of 'Every 1 Hours'. A note states: 'Snap will start every 1 hours of the UTC time. Current time in UTC is 03/09/22 12:51 AM.' Below this is a 'Replication Group Name' field with a 'Use New' dropdown. A 'Pool for Snapshot' dropdown is set to 'Auto Selected'. At the bottom, a 'Volumes' table is displayed:

# OF VOL.	LABEL	LABEL SUFFIX	SIZE	POOL TYPE	TIER	POOL
1	2525		1 GB	Tiered		Auto Selected

Navigation buttons at the bottom include 'Cancel', 'Previous', 'Next', and 'Submit'.

14. Click **Next** to continue to the optional Operation Plan step, you can confirm settings specified in each window and can confirm the resources to be used in the job, including LUN candidates.



Note: Internet Explorer 11 does not support the Operation Plan page. Use another supported browser to access this feature.

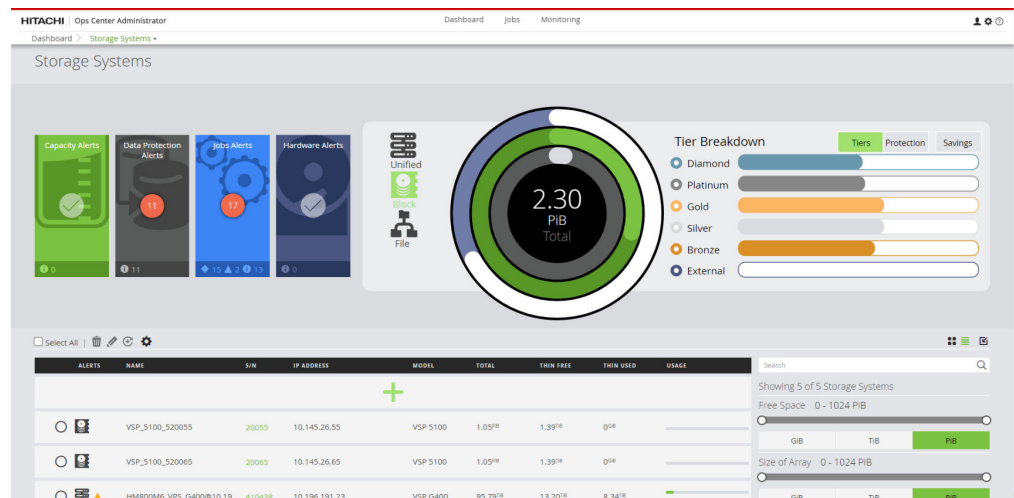
15. You can specify the LUN range for volumes in the **LUN Settings** pop up window which is launched by clicking LUN Settings button above the Planned Path Configuration table.
16. You can make it mandatory to use the LUNs displayed in the table by specifying Yes to **Mandate Using Displayed LUN**.
17. When you are satisfied with all the settings, click **Submit**.

Managing storage systems

Access the storage system details window for an overview of registered storage systems. If you need to configure the storage system, you will need to access the service processor with Storage Navigator.

Storage system inventory

From the dashboard, click Storage Systems to access the Storage Systems window for an overview of registered storage systems.



The Storage Systems window includes a summary section identical to the summary on the dashboard. This includes monitoring tiles, capacity summary and data protection summary.

You can search for a specific storage system by searching for its serial number in the search box. You can also use various filters to find a specific storage system.

From the storage systems window you can view and manage individual storage systems.

View storage system details in either a tile view or a list view. Capacity utilization is represented by a color bar:

- Green: under 70%.
- Orange: from 70% to 80%.
- Red: over 80%.

The following actions are available:

- Click a storage system to view details in the Storage System *<serial number>* window, where you can also access and manage the associated parity groups, pools, volumes, and ports.
- To delete a storage system, select it, then click Delete.
- To edit the user name and password for a storage system, select it, then click Edit to open the Update Storage System *<serial number>* window.
- To update storage systems manually, select them and then click Refresh.

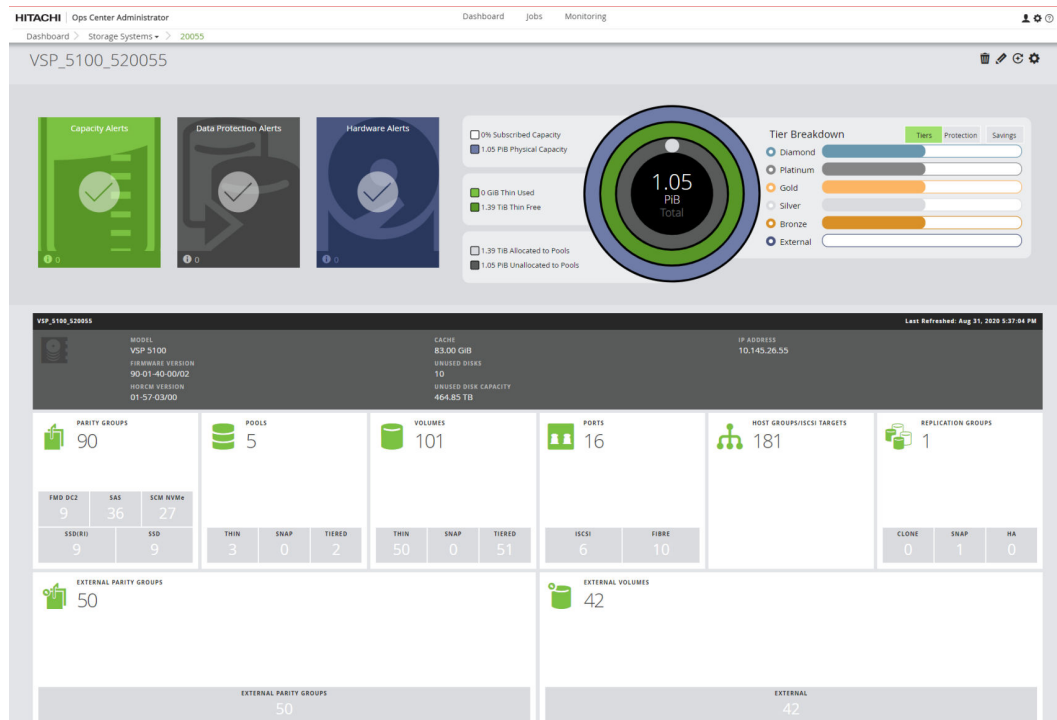
You can view the status of the updates in the Jobs window. The date and time of the last update is displayed in the detail window for each storage system.

- To set up high availability, click High Availability Setup.

Storage system details

You can view details about a storage system and access resources.

On the dashboard, click Storage Systems and then click a Storage System *serial_number* tile.



If the storage system has no SVP, you can click Settings to access the following:

- Click Launch Storage Advisor Embedded to open Storage Advisor Embedded in a new browser.
- Click Switch CTL to connect to switch the IP address that Ops Center Administrator is using to communicate with the storage system. Use this option if Ops Center Administrator cannot to communicate with the storage system.

Click Savings to view the Savings Breakdown for the following storage systems:

- VSP E series
- VSP G/F350, G/F370, G/F700, G/F900 with firmware version 88-03-0x or later.

For all other storage systems, Savings includes Data Reduction and Capacity Efficiency.

- **Total Efficiency:** The ratio of the total saving effect achieved by accelerated compression, capacity saving (compression and deduplication), snapshot, and Dynamic Provisioning.
- **Data Reduction:** The ratio of logical used capacity to the physical used capacity, for all compression and deduplication technologies.



Note: The value displayed as the capacity of data after reduction includes the size of meta-data and garbage data generated by the storage system in addition to user data. The value might temporarily be greater than the capacity of data before reduction.

- If disk-based compression is in use, either alone or in combination with controller-based compression, or with compression accelerator, the physical used capacity is the result from disk-based compression alone.
- If only controller-based compression and/or compression accelerator are in use, the physical used capacity is the result from controller-based compression and/or compression accelerator.
- If no compression is in use, then physical used capacity is the used capacity of the pools.
- **Capacity Efficiency:** The ratio of Thin Free plus Thin Used to the physical used capacity. Capacity efficiency is only calculated for volumes on Thin, Tiered, and Snap pools.
- **Provisioning Efficiency:** The efficiency ratio achieved by Dynamic Provisioning.
- **Snapshot Efficiency:** The efficiency ratio achieved by snapshot.

Click Details to view the foregoing metrics, plus the following components of Total Efficiency:

- Software Compression and Deduplication:
 - Software Saving: The capacity reduction ratio before and after capacity saving.
 - Software Compression: The capacity compression ratio before and after capacity saving.
 - Software Deduplication: The capacity deduplication ratio before and after capacity saving.
 - Software Pattern Matching: The capacity reduction ratio before and after capacity saving pattern matching.
- FMD Compression:
 - FMD Saving: The capacity reduction ratio before and after accelerated compression.
 - FMD Compression: The capacity compression ratio before and after accelerated compression.
 - FMD Pattern Matching: The capacity reduction ratio before and after accelerated compression pattern matching.
 - Calculation Period (UTC): The start date and time for the calculation and the end date and time for the calculation.

The following actions are available:

- Click Delete to delete the storage system. This action is available only to users with the System Administrator role.
- Click Edit to update the storage system user name and password. This action is available only to users with the System Administrator role.
- Click Refresh to manually update the storage system.

You can view the status of the update in the Jobs window. The date and time of the last update displays in the storage system detail window.

- Click Settings in the upper-right corner to start the following:
 - Storage Navigator to perform any advanced storage management functions that are not available in Ops Center Administrator.

If you have System Administrator, Storage Administrator, and Security Administrator roles, you can launch Storage Navigator directly in the dashboard or directly in the Actions > Other Functions mode, which includes the Data Retention Utility.
 - Configure storage system encryption using Storage Navigator.
 - Manage license keys using Storage Navigator.
 - If you onboarded any of the following storage systems, you can access Storage Advisor Embedded:
 - VSP E series
 - VSP G/F350, G/F370, G/F700, G/F900Click Launch Storage Advisor Embedded to log in to Storage Advisor Embedded.
 - If the storage system includes optional NAS module, you can start NAS Manager to perform functions that are not available in Ops Center Administrator.
- Review the information gauge regarding available capacity in storage system.
- Use the Alerts tiles to access alerts related to the storage system.

- View attributes about the storage system, including information that Ops Center Administrator discovered about the system, such as model, serial number, and SVP IP address.

Field	Description	
Model	Model name of the storage system.	
S/N	Storage system serial number.	
SVP Version	SVP firmware version.	
HORCM Version	A data replication component residing on the server.	
Cache Capacity	Total cache capacity configured on the storage system.	
Unused Disks	Disks that are not assigned as hot spares or used to create parity groups.	
Unused Disk Capacity	Total unused capacity of all disks.	
SVP IP Address	Address of the service processor on the storage system. The RMI port number will also be displayed if the port number is specified at the time of onboarding. The RMI port number can only be specified for virtual SVP.	
Maintenance Utility IP Addresses 1 and 2	Configured automatically when the storage system is onboarded.	
Health Status (for a cluster, if any)	Robust	Operating normally, with no failures in the cluster interconnect, the management network, or quorum device communication.
	Degraded	Operating, but one or more nodes has failed or there has been a failure in connectivity.
	Online	The cluster node is a part of the cluster and is exchanging heartbeats (periodic signals) with the other cluster members.
	Offline	Offline: The cluster node is no longer exchanging heartbeats with the other cluster members. This might be caused by a restart or a fault condition. Services cannot be migrated to a cluster node in this state.
	Dead	No heartbeats from this cluster node for a significant period. Services

Field	Description	
		cannot be migrated to a cluster node in this state.
	Unknown	Unknown: The node has not been online since the cluster was started.

- If you notice that the number of unused disks and unused disk capacity is too high, it implies that you have additional raw capacity in the system that can be converted into usable capacity by creating parity groups. Click Create Parity Groups to configure unused disks into parity groups.
- View the total and unused disk capacity for the different pool tiers.
- Click a tile in the lower portion of the window to display the details for parity groups, pools, volumes, or ports for the storage system.

Using Storage Navigator for advanced storage configuration

Storage Navigator is the element manager for a supported storage system block module. It is a factory-installed application running on the SVP, which is directly connected to the storage system. You can access Storage Navigator from the storage system detail window for advanced configuration options while performing management functions such as remote replication, volume migration, and resource group management with Ops Center Administrator. You can access Storage Navigator for the following storage systems:

- VSP 5000 series
- VSP E series
- VSP G200, F/G400, F/G600, F/G800
- VSP G130, G/F350, G/F370, G/F700, G/F900

When you launch Storage Navigator, the name of the user who onboarded the storage system shows in the Storage Navigator UI by default. However, you can also modify the Ops Center Administrator settings so that the name of the logged-in user is shown instead. For detailed information, see [Enabling the Storage Navigator Launch Settings option \(on page 72\)](#).



Note: If the name contains a character that cannot be used in Storage Navigator, it is replaced with an asterisk (*).

Storage systems with an SVP must be managed using Storage Navigator. Storage systems without an SVP must be managed with Storage Advisor Embedded.



Note: If you use Storage Navigator to manage resource groups, ensure that the user who adds storage systems in Ops Center Administrator has access to all custom resource groups and meta-resource groups.

Use the Storage Navigator online help for information on completing advanced storage configuration tasks.

To use Storage Navigator running on HARMAN AIR, ensure that Storage Device Launcher is installed on your client machine. For information on how to install Storage Device Launcher, see the SVOS Administration documentation for your storage system.

Prerequisites

To launch Storage Navigator from Ops Center Administrator, you must complete the following tasks:

- Register a firewall exception. Verify the port to use, which is described in "Ports requirements" in the *Hitachi Ops Center Administrator Getting Started Guide*.
- If the host name is used for the certificate, resolve the name of the Ops Center Administrator server from the storage system (SVP).
- For VSP 5000 series, register the Ops Center Administrator server certificate in the storage system.

Enabling the Storage Navigator Launch Settings option

If you enable the Enable Launch by Login User option, the settings of selected storage systems and newly-onboarded storage systems are updated. This includes overwriting external user account authentication settings and storage launch settings that might be configured by other management servers. for example:

- The Hitachi Command Suite server is used as an external authentication server.
- Other Ops Center Administrator servers are using the Enable Launch by Login User setting.

To enable the Storage Navigator Storage Launch Settings option:

Procedure

1. From the **Settings** menu, click **Storage Launch Settings**.

Storage Launch Settings

Device Manager-Storage Navigator Launch Settings

Enable Launch by Login User (Selected) | Disable Launch by Login User

Enabling/Disabling the setting will change the in-context launch behavior for onboarded storage systems with certain models. The setting on the storage system side will be done at the time of onboarding.
For already-onboarded storage systems, select from the table below to reflect the setting (Optional).

☐ Select All

	NAME	S/N	IP ADDRESS	MODEL
<input type="radio"/>	VSP_5100_520055	20055	10.145.26.55	VSP 5100
<input type="radio"/>	VSP_5100_520065	20065	10.145.26.65	VSP 5100
<input type="radio"/>	HM800M6_VPS_G400@10.196.191.23	410438	10.196.191.23	VSP G400
<input type="radio"/>	HM850H3_VSP_G900@10.196.191.32	415248	10.196.191.32	VSP G900
<input type="radio"/>	HM800S2_VSP_G200@10.196.191.20	430002	10.196.191.20	VSP G200

Showing 5 of 5 Storage Systems

Free Space 0 - 1024 PiB

Size of Array 0 - 1024 PiB

Migration Tasks Exist ALL

False True

Cancel Submit

2. Select **Enable Launch by Login User**.
3. Select the storage systems on which to implement the setting for already-onboarded storage systems.

4. Click **Submit**.

Accessing Storage Navigator

Start Storage Navigator from the Settings menu in the Storage System details window.

You can launch Storage Navigator in one of three ways:

Launch Device Manager - Storage Navigator Login Page

Launch the Storage Navigator login page.

In the login dialog box, use the user account assigned to you by your security administrator.

If you have System Administrator, Storage Administrator, and Security Administrator roles, the following two ways of launching Storage Navigator are available to you:

Launch Device Manager - Storage Navigator Dashboard

Launch Storage Navigator dashboard directly.

Launch Device Manager - Storage Navigator Dashboard (Data Retention Utility)

Launch the Storage Navigator dashboard directly in the Actions > Other Functions mode, which includes the Data Retention Utility.

Using Ops Center Administrator with Storage Advisor Embedded

Storage Advisor Embedded enables you to manage a single storage system with a simple interface and intuitive navigation and is included with the following storage systems:

- VSP 5000 series
- VSP E series
- VSP G/F350, G/F370, G/F700, G/F900

Accessing Storage Advisor Embedded

Access Storage Advisor Embedded from the detail window of the storage system. Click Settings and select Launch Storage Advisor Embedded.

Using Storage Advisor Embedded

You can access Storage Advisor Embedded to view storage system performance in real time and review capacity usage and data reduction.



Note: To review historical performance, you can access the performance monitor in Storage Navigator.

Storage Advisor Embedded allows you to quickly and easily create one-time snapshots and Snap Clones.

You can also use Storage Advisor Embedded to create pools and volumes and update ports.



Note: Resources created outside Ops Center Administrator will take a few minutes to update in Ops Center Administrator.

Tier management

The Tier Management window displays the tier definitions. You can edit the tier names (Diamond, Platinum, Gold, Silver, Bronze, and External).

Access the Tier Management window by clicking Settings and selecting Tier Management.

Using NAS Manager for advanced file storage configuration

NAS Manager is the element manager for NAS modules. It is a factory-installed application running on the NAS modules. You can access NAS Manager from the Storage System details window for advanced configuration options while performing management functions such as data protection, server settings and network configuration.

Use the NAS Manager online help to obtain procedure information for advanced file storage configuration tasks.

Accessing NAS Manager

Start NAS Manager from the Settings menu available in the **Storage System** details window. In the login dialog box, use the user account assigned to you by your administrator. The tasks that you can do on the system depend on your user role.

See the Administrator types and responsibilities list for roles that are available for use in NAS Manager and the permissions that each role gives to the users.

Administrator types and responsibilities

This section describes the types of storage system administrators and defines their expected roles in managing the system and the associated storage subsystems.

Global Administrators

Global Administrators can manage everything in the system: file systems, file services, or file system related features and functions, storage devices and their components. Also, the Global Administrator creates and manages SMU user profiles (Server Administrators, Storage Administrators, Server+Storage Administrators, and other Global Administrators). Global Administrators also control what servers and storage devices each administrator can access.

Storage Administrators

Storage Administrators manage storage devices, as specified in the administrator profile created by the Global Administrator.

Storage Administrators can manage only storage devices and their components (racks, physical disks, SDs, and storage pools). Storage Administrators cannot manage file systems, file services, or file system related features and functions, and they cannot manage users.

Server Administrators

Server Administrators manage servers and clusters, as specified in the administrator profile created by the Global Administrator. Server Administrators cannot manage storage devices.

Server Administrators can manage file systems and file services such as CIFS Shares, NFS Exports, and they can manage file system related features and functions such as snapshots, quotas, and migration policies and schedules.

Server+Storage Administrators

Server+Storage Administrators manage servers, clusters, and storage devices, as specified in the administrator profile created by the Global Administrator.

Server+Storage administrators can manage everything Server Administrators and Storage Administrators can manage: file systems, file services, or file system related features and functions, and they can also manage storage devices and their components.

Updating a storage system

Update the user name and password for a storage system. For this action, the System Administrator role must be assigned to the user.

Procedure

1. Log into Storage Navigator and create a new user.
2. Access the **Update Storage System** window one of these ways:
 - In the **Storage Systems** window, select a storage system, then click **Edit**.
 - In the storage system detail window, click **Edit**.
3. Edit the **Username** or **Password**, then click **Submit**.

Managing servers

Ops Center Administrator supports provisioning storage to a group of server WWNs by allowing logical servers to be managed by Ops Center Administrator. After servers are onboarded in Ops Center Administrator, storage volumes can be provisioned to the servers by creating Host Storage domains and optionally creating zones to provide a path between the storage volume and server.

Adding servers

Add servers so you can attach volumes. You can add multiple server parameters from a file, or add one server at a time.

You can add servers using one of the following methods:

- Manually add information for one server at a time.
- Import a CSV (comma-separated values) file with information for one server in each row.

The CSV file must have the following headings:

- For Fibre:
 - Name, OSType, WWNs (comma-separated list of WWNs).
 - Description, IPAddress and WWNsUserDefinedNames (comma-separated list of user-defined names for WWNs) are optional headings.
- For iSCSI:
 - Name, OSType, IscsiName (comma-separated list of names).
 - Description, IPAddress, ChapUser, ChapSecret and IscsiNamesUserDefinedNames (comma-separated list of user-defined names for iSCSI Names) are optional headings.

Valid OSType values are as follows:


- AIX
- HP_UX
- LINUX
- NETWARE
- OVMS
- SOLARIS
- TRU64
- VMWARE
- VMWARE_EX
- WIN
- WIN_EX



Note: When you use Internet Explorer, limit the number of servers in a CSV file to 500 servers. If the browser does not work, reduce the number of servers and try again.

Procedure

1. On the Ops Center Administrator dashboard, click **Servers**. Then click the plus sign (+) to open the **Add Server** window.



Add Servers

CSV Import

Fibre Servers

SERVER NAME	DESCRIPTION	IP ADDRESS	OS TYPE	
Fibre_Server_Example	Example Fibre server	1.2.3.4	HP_UX	✕
<div>WWN LIST</div> <div>10:00:00:00:C9:71:C6:F4, 10:00:00:00:C9:71:C6:F4</div>				
<div>WWN USER-DEFINED NAMES</div> <div>Fibre_server1, Fibre_server2</div>				

iSCSI Servers

SERVER NAME	DESCRIPTION	IP ADDRESS	OS TYPE	CHAP USER	
iSCSI_Server_Example	Example iSCSI server	5.6.7.8	HP_UX	sysadmin	✕
<div>CHAP SECRET</div> <div>*****</div>		<div>ISCSI LIST</div> <div>iqn.1991-05.com:microsoft:ap01vm01.ap01vm01.dom.local, iqn.1991-05.com:microsoft:ap01vm01.ap01vm01.dom.local</div>			
<div>ISCSI NAMES USER-DEFINED NAMES</div> <div>iSCSI_server1, iSCSI_server2</div>					

Cancel

Reset

Submit

2. In the **Add Server** window, do one of the following:

- Click the upper plus sign (+) to browse for the CSV file or drag the file to the plus sign. The values from the file populate the window. Example:

```
Name,Description,IPAddress,OSType,WWNS,WWNsUserDefinedNames
Esxi,ESXI
HOST,10.30.90.200,VMWARE_EX,10:00:00:05:33:26:f7:21,Esxi_HBA_1
Win,WINDOWS
HOST,10.30.91.80,WIN_EX,"10:00:00:05:33:26:f7:37,10:00:00:05:3
3:26:f7:36","HOST_HBA_1,HOST_HBA_2"
ESXi_Cisco_1,ESXi_HOST connected to Cisco
Fabric,,VMWARE_EX,"10:00:00:05:33:26:e0:fc,10:00:00:05:33:26:e
0:fd","Fabric_HBA_1"
ESXi_Cisco_2,ESXi_HOST connected to Cisco
Fabric,,VMWARE_EX,"100000053326df1a,100000053326df1b","",Fabric
_HBA_2"
```

- To add both Fibre and iSCSI servers, use the following format:

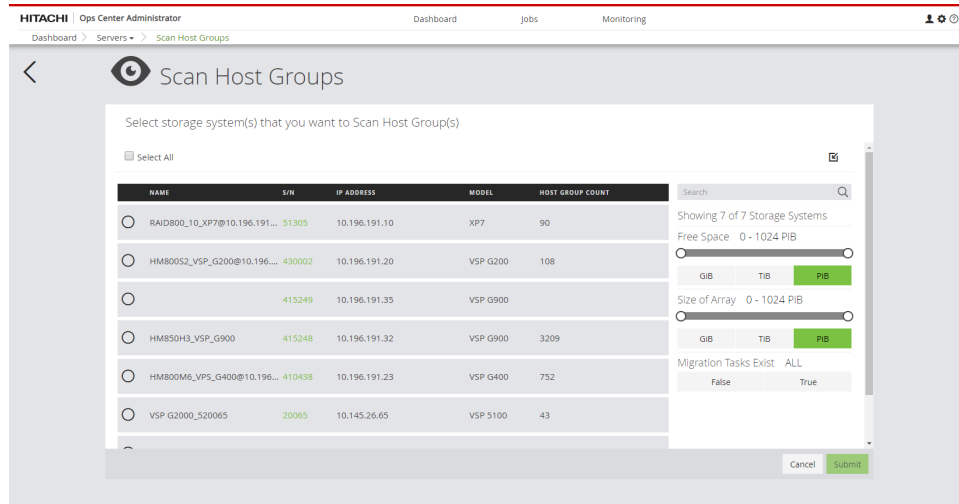
```
Name,Description,IPAddress,OSType,WWNS,WWNsUserDefinedNames,Is
csiNames,IscsiNamesUserDefinedNames,CHAPUserName,CHAPUserSecre
t
linux-iscsi,test dummy host,20.10.10.10,Linux,,,"iqn.
linuxiscsi-1,iqn. linux-iscsi-2,eui.1234567890abCDef","linux-
iscsi-HBA-1,linux-iscsi-HBA-2",,
-windows-iscsi-uni-chap,test dummy
host,,20.10.10.20,Win,,,"iqn.-windows-iscsi-unichap,""host-
HBA-1""",chapUserName,chapUserSecret
-windows-iscsi-bi-chap,test dummy host,20.10.10.30,Win,,,"iqn.-
windows-iscsi-bichap,""windows-iscsi-bi-chap-
HBA-1""",chapUserName,chapUserSecret
-vmware-iscsi-longest,test dummy
host,,20.10.10.40,VMWARE,,,"iqn.123456789012345678901234567890
12345678901234567890123456789012345678901234567890123456789012
34567890123456789012345678901234567890123456789012345678901234
5678901234567890123456789012345678901234567890123456789012345
6789012345678901234567890123456789012345678901234567890123456
789,eui.3234567890abCDef"
ed801h,Windows,10.197.73.57,WIN,10:00:00:90:fa:b4:a8:71,"ed801
h-HBA-1"
ed800n,ESX
Host,10.197.73.7,VMWARE,10:00:00:90:fa:55:85:5d,"ed800n-HBA-1"
-linux,test dummy
host,,10.10.10.10,Linux,10:10:10:10:10:10:10:10,"ed801h-HBA-1"
-windows,test dummy
host,,10.10.10.20,Win,10:10:10:10:10:10:10:20,"host-HBA-1"
-vmware,test dummy
host,,10.10.10.30,VMWARE,10:10:10:10:10:10:10:30
```

- Click the plus sign (+) in the table to add a row and enter the required information for Fibre Channel or iSCSI. You can add more servers by clicking the plus sign again.
- (Optionally) You can use the **WWN List** or **iSCSI List** to add/edit a server.
- (Optionally) You can add comma-separated user-defined names for WWNs or iSCSI names in the order they are specified.

3. Click **Submit** to add the servers.

Next steps

Create volumes and attach them to the server.



4. Click **Submit**.

Using Scan Host Groups to Migrate Information

When migrating host information, it's important to understand how the Scan Host Groups feature functions depending on the environment parameters:

- If you select multiple storage systems, server creation runs in ascending order of the storage system serial number. Ops Center Administrator ignores host groups without WWNs because it only sees host group configuration and cannot determine whether a host group contains WWNs of clustered servers.
- The Ops Center Administrator server cannot have both FC and iSCSI protocols. Different protocols in host groups result in different servers. If the host mode conflicts between host groups, Ops Center Administrator creates only one server based on the host group found first.
- When Ops Center Administrator finds a host group with an external array connection and both storage systems are already onboarded, the host groups are ignored. If only the external storage system is onboarded, Ops Center Administrator creates a server. (Ops Center Administrator cannot check if the host group is used by an external array or a host.) For host groups with the same WWNs set, Ops Center Administrator creates a single server. Even if the name of each host group is different, they are treated as a single server.
- Ops Center Administrator determines the port to use when attaching the volume to the newly created server. If the same WWN is found in multiple host groups during a single scan, the result depends on the order of storage systems in the scan.

In addition, be aware of the following WWN-related information:


- You may need to edit the mapping between servers and WWNs after scanning completes.
- If a server already exists with the same WWNs as scanned host groups, Ops Center Administrator does nothing.
- Ops Center Administrator adds newly found WWNs to existing servers if the server name and host group name are the same.

- Ops Center Administrator creates a new server with the scanned WWNs if there is no matching server name and the host group contains WWNs that no other existing servers use.
- Existing servers are never deleted (new WWNs may be added, but there is no WWN removal or server deletion).
- If host groups over multiple storage systems are scanned and the host groups have volumes provisioned from each storage system:
 - After scanning both storage systems, Ops Center Administrator creates a server from the host groups over the multiple storage arrays.
 - A scan reconciles servers and host groups based on a set of WWNs in each host group.

If a server is created following the scan of one storage system and if it has volumes provisioned from a second (or additional) storage system:

1. Create a server object after scanning the first storage system.
2. During the scan of the second storage system, reconcile its host groups with the existing server.

Exporting servers

You can export servers added to Ops Center Administrator by clicking the download icon () at the top right corner of the server window. You can then download a CSV file with information about all added servers.

The CSV file will have the following information:

- Name
- Description
- IPAddress
- OSType
- WWNs (for FC servers)
- IscsiNames (for iSCSI servers)
- CHAPUserName (for iSCSI servers)
- WWNsUserDefinedNames (for FC servers)
- IscsiNamesUserDefinedNames (for iSCSI servers)



Note: The export function does not save the "CHAPUserSecret" because that information is confidential. You will need to enter the Secret to import this CSV file. Ops Center Administrator warns you of this fact before you export.

Server and server groups inventory

The Servers window displays all servers and includes a graphic summary based on operating systems.

HITACHI | Ops Center Administrator
Dashboard > Servers

Dashboard
Jobs
Monitoring

HITACHI

Servers

OS TYPE	COUNT
Windows	8
Linux	5
VMware	1

Server

Server Group

Search

Filter

+

🔍

📄

🔗

🔗

0 selected

Loaded all 14 items

ID	ALERTS	SERVER NAME	SERVER IP ADDRESS	PROTOCOL	WINRM/SSH INITIATOR NAME	OS TYPE	VOLUME COUNT	REPLICATION TYPE
1	—	dummyGAD4portA	19.90.5.1	Fibre	20:20:20:20:20:20:40, 20:20:20:20:20:20:41, 20:20:20:20:20:20:42	WIN	5	Clone
2	—	dummyGAD4portB	19.90.5.2	Fibre	30:20:20:20:20:20:40, 30:20:20:20:20:20:41, 30:20:20:20:20:20:42	VMWARE	7	Clone
3	—	dummyGAD2portC	19.90.5.3	Fibre	40:20:20:20:20:20:40, 40:20:20:20:20:20:41	WIN	12	Clone, Snap on S...
4	—	dummyGAD2portD	19.90.5.4	Fibre	50:20:20:20:20:20:40, 50:20:20:20:20:20:41	WIN	12	—
5	—	dummyGAD2portE	19.90.5.5	Fibre	60:20:20:20:20:20:40, 60:20:20:20:20:20:41	LINUX	4	—
6	—	dummyGAD2portF	19.90.5.6	Fibre	70:20:20:20:20:20:40, 70:20:20:20:20:20:41	LINUX	3	Snap on Snap
7	—	zoning_server1	15.15.15.15	Fibre	10:00:00:90:00:00:00:00, 10:00:00:90:00:00:00:01	WIN	2	—
8	—	zoning_server2	15.15.15.16	Fibre	10:00:00:90:00:00:00:02, 10:00:00:90:00:00:00:03, 10:00:00:90:00:00:00:04	WIN	1	—
9	—	zoning_server3	15.15.15.17	Fibre	10:00:00:90:00:00:00:06, 10:00:00:90:00:00:00:07	WIN	0	—
10	—	zoning_server4	15.15.15.18	Fibre	10:00:00:90:00:00:00:08, 10:00:00:90:00:00:00:09	WIN	0	—
11	—	zoning_server5	15.15.15.20	Fibre	10:00:00:90:00:00:00:10, 10:00:00:90:00:00:00:11	LINUX	0	—
12	—	zoning_server6	15.15.15.21	Fibre	10:00:00:90:00:00:00:12, 10:00:00:90:00:00:00:13, 10:00:00:90:00:00:00:14	LINUX	0	—
13	—	zoning_server7	15.15.15.22	Fibre	10:00:00:90:00:00:00:16, 10:00:00:90:00:00:00:17	LINUX	0	—

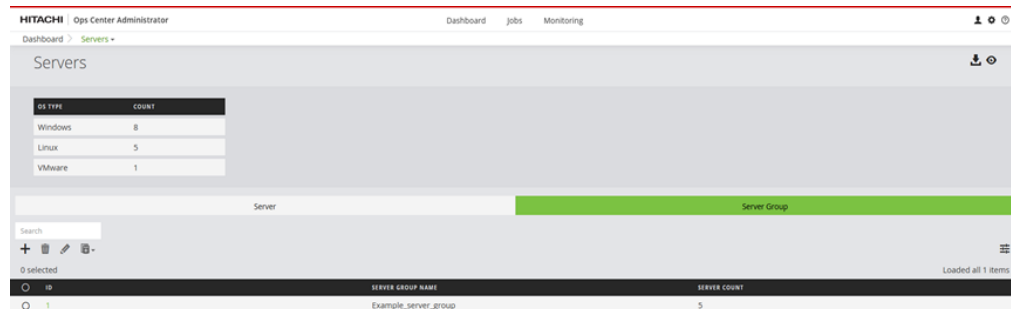
show 200 items

To find the server you are looking for, you can either search by server ID or narrow down the list of servers using filters available.

You can click server tiles on this window to access details of individual servers and manage their volumes.

The following actions are available on this window:

- To add a server, click the plus sign (+) to open the Add Servers window.
- You can select one or more servers and perform the following actions:
 - Click Delete to delete the server.
 - Click Edit to update server parameters.
 - Click Create and Attach Volumes to use the provided workflow to create new volumes, attach them to the server, and optionally, protect the volumes.
 - Click Attach Existing Volumes to attach existing volumes to the server.
 - Click Use Existing LUN Paths to access the Use Existing LUN Paths window.
- Click the Server Group tab to add, edit and delete server groups.



- Click the plus sign (+) to open the Add Server Group window.
- Select a server group, then click Edit to open the Update Server Group window.
- Click a Server Group ID to access the server group detail window.
- To manage volumes for a server group, select it, then click Attach. Then select Attach Existing Volumes, Create, Attach and Protect Volumes with Local Replication, or Create, Attach and Protect Volumes with High Availability.

Server details

The **Server details** window shows all the details about the server and a list of volumes, if any, attached to the server.

This window also shows path details for every volume attached to the server, including the storage ports and server WWNs used to create HSDs and the host mode options set up on the HSD.

You can perform server actions such as removing the server, editing the server or provisioning volumes to server by attaching existing volumes or creating and attaching new volumes.

You can perform volume procedures such protecting or detaching volumes.

On the dashboard, click Servers to open the **Servers inventory** window. Then click a server tile to open the Server <ID> detail window for the selected server.

HITACHI | Ops Center Administrator

Dashboard > Servers > 7

Server 7

SERVER NAME	IP ADDRESS	CAPACITY (THIN USED / TOTAL)	THIN FREE
zoning_server1	15.15.15.15 DESCRIPTION Windows	4.02 GiB / 4.10 TiB	4.09 TiB WIN


Search Filter ▾

0 selected

Loaded all 2 items

>	VOLUME ID	VIRTUAL ID	S/N	VIRTUAL S/N	VOLUME NAME	DATA PROTECT.	POOL ID	ATTRIBUTE	STATUS	PROTECTION CT.	DATA PROTECT.	CAPACITY SAVIN.	TOTAL	THIN USED	USAGE
>	○ 216 (00:00:D8)	216 (00:00:D8)	39303	39303	hsai092	Unprotected	0	Thin	Normal	—	No	No	100.00 GiB	0.00 GiB	0%
>	○ 2776 (00:0E:E...)	2776 (00:0E:E...)	39303	39303	hsa-external...	Unprotected	15	Thin	Normal	—	No	No	4.00 TiB	4.02 GiB	0%

You can perform the following actions at the server level:

- Click Delete to remove the server.
- Click Edit to update server parameters.
- Click the cylinder icon () and select one of the following:
 - Attach Existing Volumes to attach volumes to the server.
 - Create and Attach Volumes to create volumes and attach them to the server.
- Click the arrow next to one of the volumes to display the volume's Port, WWN, LUN, Host Mode, Host Mode Options, Host Group Name, and ALUA mode.
- A server using the iSCSI protocol might have volumes with a mutual CHAP user. Edit the mutual CHAP user by clicking the arrow next to one of the volumes, then clicking Edit Mutual CHAP User. You can edit the mutual CHAP user name and the mutual CHAP secret.
- Edit the host group settings of a server using the Fibre protocol by clicking the details of an attached volume, then clicking Edit. In the Edit Host Group window you can set the path to preferred or non-preferred and change the host group name.

The following actions are available for volumes:

- If there is a **Data Protection Failed** alert, you can click it to see failed data protection jobs on the **Data Protection Monitoring** window.
- Click **Edit** to rename a volume on the **Update Volumes** window. You can also update the size and command device settings.
- You can also rename multiple volumes and enable or disable command device at the same time if you select multiple volumes.
- Click **Edit LUN Path** to open the **Edit Paths** window where you can update LUN paths.

Note: The command device configuration will be disabled if all the attached LUN paths are removed.

- Click **Detach Volumes** to detach volumes from the server.

Before confirming, you can choose whether to remove SAN zones.

Note: The command device configuration will be disabled if all the attached LUN paths are removed.

- Click Protect Volumes to protect volumes using Ops Center Administrator data protection technologies.



Note: Command devices cannot be protected. Do not perform this action for command devices.



Note: Ops Center Administrator can protect volumes with high availability in existing replication group with following conditions:

- P-VOL STATUS and S-VOL STATUS of all paired volumes in the replication group is PAIR.
- Status of the corresponding replication in Ops Center Protector is OK.

- To unprotect a volume, select it, then click Unprotect.



Note: A volume in an "External" replication group cannot be unprotected.

- Click Unprotect Volumes to remove data protection applied to volumes.



Note: Ops Center Administrator can unprotect volumes with high availability in the replication group with following conditions:

- P-VOL STATUS and S-VOL STATUS of all paired volumes in the replication group is PAIR.
- Status of the corresponding replication in Ops Center Protector is OK.

- Click Restore Volumes to restore from a backup.
- Click Attach to Storage to attach volumes to target storage to prepare for volume migration where source and target storage systems are both managed in Ops Center Administrator.
- Click Migrate Volumes to migrate volumes.



Note: Command devices cannot be migrated. Do not perform this action for command devices.

- For a tiered volume, click Update Tiering Policy to update the tiering policy for the volume.
- Click Attach Like Volumes to attach like volumes for local replication or for high availability.

Updating a server

You can change the parameters for an existing server.

Procedure

1. Access the **Update Server** window in one of these ways:
 - From the dashboard, click **Servers** to open the inventory of servers. Select a server, then click **Edit**.
 - On the detail window for a server, click the pencil icon.

2. Edit any of the following for a server with the Fibre Channel protocol:

- **Server Name**
- **Description**
- **IP Address**
- **OS Type**
- **Modify Existing WWNs**
- **Add New WWNs**

- a. If needed, add new WWNs, then click the plus sign to add them to the list.
- b. (Optionally) Click **WWN List** to select and update existing WWNs.



Note: If you remove or update all WWNs to which command devices are attached and select the **Apply changes to attached volumes** check box, the command device configuration will be disabled.

If you are removing WWNs, update the server without selecting the **Apply changes to attached volumes** check box, and then use Storage Navigator to make changes to keep the command device configuration enabled.

If you are updating an existing WWNs, add a WWN to the server first and then remove the existing WWN.

- c. (Optionally) You can add and edit user-defined names for WWNs.
- d. You can select a reference WWN to add the new WWN into host groups to which the reference WWN belongs.



Note: If reference WWN is specified, **Apply changes to attached volumes** must also be verified.

- e. Unselect the **Apply changes to attached volumes** check box if you do not want to make the following changes:
 - If a WWN has been updated, all host groups with the old WWN are updated with the new one. All zones created in Ops Center Administrator have the old WWN removed and new address populated.
 - If a WWN is deleted, it will be removed from any host group that has it. If the host group is left with no WWNs, the entire host group is deleted. All zones created in Ops Center Administrator have the deleted WWN removed.
 - If you add a reference WWN, host groups are edited. If you do not add a reference WWN, no host groups are edited. You can use Ops Center Administrator to add paths using the new LUN.
 - Changes to name, description, IP address, or OS type do not impact host groups.
3. Edit any of the following for a server with the iSCSI protocol:
 - **Server Name**
 - **Description**
 - **IP Address**
 - **OS Type**
 - **Modify Existing iSCSI Names**
 - **Add New iSCSI Names**
 - **Enable or Disable CHAP Settings.**
 - a. You can select a reference iSCSI name to add the modified or new iSCSI name into host groups to which the reference iSCSI belongs.
 - b. (Optionally) Click **iSCSI List** to select and update existing iSCSI names.



Note: If you remove or update all iSCSI initiator names that command devices are attached to and select the **Apply changes to attached volumes** check box, the command device configuration will be disabled.

If you are removing iSCSI initiator names, update the server without selecting the **Apply changes to attached volumes** check box, and then use Storage Navigator if you want to keep the command device configuration enabled.

If you are updating existing iSCSI initiator names, add the iSCSI initiator names to the server first and then remove the existing iSCSI initiator names.

- c. (Optionally) You can add and edit user-defined names for iSCSI names.
- d. Select **Update CHAP Credential** to add a new CHAP user.

- e. Unselect the **Apply changes to attached volumes** check box if you do not want to make the following changes:
- If an iSCSI initiator name has been updated, all iSCSI targets with the old iSCSI name are updated with the new one.
 - If an iSCSI initiator name is deleted, it will be removed from any iSCSI target that has it. If the iSCSI target is left with no iSCSI initiator names, the entire iSCSI target is deleted.
 - If an iSCSI initiator name is added, no iSCSI targets are edited. You can use Ops Center Administrator to add paths using the new LUN.
 - If a CHAP user information changes, Ops Center Administrator will update the information on the storage system.
 - Changes to name, description, IP address, or OS type do not impact iSCSI targets.

Deleting servers

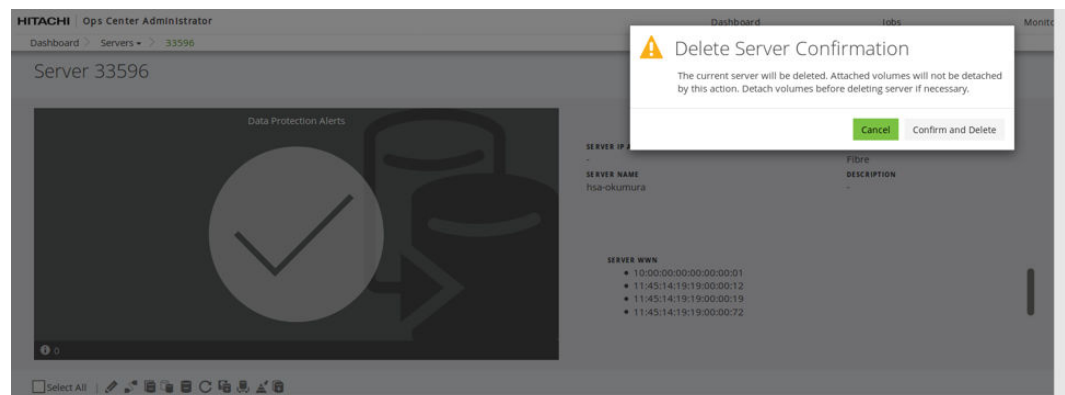
You can delete one or more existing servers.

Procedure

1. From the dashboard, click **Servers** to open the inventory of servers.
2. Select the servers that you want to delete, then click the trashcan icon.



Note: If at least one or more servers have an attached volume (the volume count is 1 or more), the following dialog box appears.



Editing LUN paths

You can update LUN paths for volumes that are attached to a server.

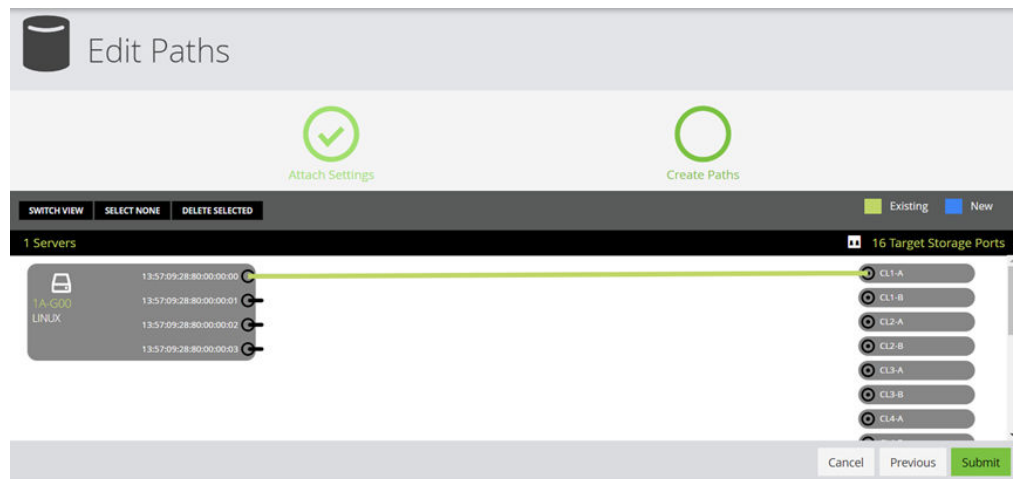
If LUN paths become outdated, you can update them through the detail window for the server.

Procedure

1. On the dashboard, click **Servers**, then click the server that needs to be updated to open the detail window for the server.
2. Select one or more volumes for which you want to edit LUN paths, then click **Edit LUN Path** to open the **Edit Paths** window.

LABEL	ID	SIZE	LUN
hsa-fs	26 (00:00:1A)	1.00 GiB	LUN (optional)

3. The **Host Mode** is set by default to the server operating system. You can make a selection if needed.
The server OS Type is provided when the server is added to Ops Center Administrator.
4. The prepopulated **Host Mode Option** depends on the **Host Mode** selection. The default Host Mode Option can be changed manually.
Default values are set only for **VMWARE EX** and **WIN EX** host modes. The default for all other Host Modes is none.
Ops Center Administrator identifies all host groups containing any of the server WWNs. If all of those host groups have the same host mode and host mode options, those settings are prepopulated with the same settings in the host groups.
5. **Auto Create Zone** is set to No by default. You can set it to Yes to automatically create zones.
6. Click **Next**.
7. In the **Create Paths** window, you can view servers and their WWNs, along with ports on the storage system.



Ops Center Administrator will scan for existing host groups on the storage system and try to reuse them by default. The following options are available for managing LUN paths:

- Existing paths are populated as follows: all existing host groups with one or more server WWNs and the exact same host mode and host mode options selected on the **Attach Settings** window are populated as paths.
- To prevent the volume from being added to an existing path, click the path to highlight it, then click **Delete Selected**.
- To manually create a path, click a WWN and a port to connect them with a blue connector line. Click the connector again and then click **Delete Selected** to delete the connection.
- Optionally, you can click **Switch View** to use tables to create paths by selecting Server Ports to attach to Storage Ports.

For VSP 5000 series storage systems, you can click a port to display a list of ports with their DKC locations, CTL locations, and redundancy levels.

Port Redundancy Level

NO.	PORT ID	DKC LOCATI...	CTL LO...	REDUNDANC...
1	CL4-A	DKC-1	CTL11	DKC
2	CL2-A	DKC-1	CTL11	DKC
3	CL6-A	DKC-1	CTL11	DKC
4	CL8-A	DKC-1	CTL11	DKC
5	CL1-J	DKC-2	CTL21	DKC
6	CL3-J	DKC-2	CTL21	DKC
7	CL5-J	DKC-2	CTL21	DKC
8	CL7-J	DKC-2	CTL21	DKC
9	CL2-J	DKC-3	CTL31	DKC
10	CL4-J	DKC-3	CTL31	DKC

Total Ports: 15

OK

- When you are satisfied with the paths, click **Submit**.

Reusing existing LUN paths

You can select an existing server and reuse its LUN paths in another server.

Before you begin

The selected server must have the same operating system as the server whose LUN paths you want to use.

Procedure

- In the **Servers** window, select a server, then click **Use Existing LUN Paths** to open the **Use Existing LUN Paths** window.

Use Existing LUN Paths

Select server Configure LUN Paths

STORAGE SYSTEM: RAID900_9_VSP5500HB#10.196.191.52_e (39303)

Select an original server

SERVER ID	SERVER LABEL	SERVER IP ADDRESS	PROTOCOL	REPLICATION TYPE
<input type="radio"/> 1	dummyGAD4portA	19.90.5.1	Fibre	-
<input type="radio"/> 3	dummyGAD2portC	19.90.5.3	Fibre	Clone, Snap on Snap
<input type="radio"/> 4	dummyGAD2portD	19.90.5.4	Fibre	-
<input type="radio"/> 7	zoning_server1	15.15.15.15	Fibre	-

Showing of Servers

Replication Type

- ☐ Snap on Snap
- ☐ Snap Clone
- ☐ Snap
- ☐ Clone

Protection Type

- ☐ Protected
- ☐ Partial
- ☐ Unprotected

Cancel Next

2. Select a server, then click **Next** to configure paths.

Use Existing LUN Paths

Select server Configure LUN Paths

SWITCH VIEW SELECT NONE DELETE SELECTED

Existing New

1 Servers 1 Storage Ports

Original Server

dummyGAD2portD WIN

Selected Server(s)

zoning_server2 host-mode-WIN

CL1-B

Cancel Previous Submit

Optionally, you can click **Switch View** to use tables to create paths by selecting Server Ports to attach to Storage Ports.

Server Ports

Filter

2 selected

SERVER NAME	WWN (FC) INITIATOR NAME	WWN (iSCSI) INITIATOR NAME
zoning_server2	10:00:00:90:00:00:00:02	---
zoning_server2	10:00:00:90:00:00:00:03	---
zoning_server2	10:00:00:90:00:00:00:04	---
zoning_server2	10:00:00:90:00:00:00:05	---

Loaded all 4 items

Storage Ports

Filter

1 selected

PORT ID
CL1-B

Loaded all 1 items

Paths

Filter

0 selected

SERVER NAME	WWN (FC) INITIATOR NAME	WWN (iSCSI) INITIATOR NAME	PORT ID
-------------	-------------------------	----------------------------	---------

Loaded all 0 items

Cancel OK

3. After you have configured the paths, click **Submit**.

Managing server groups

Use server groups to simplify management of servers.

Adding a server group

Add server groups to manage more than one server at a time.



Note: A server group can only use a single protocol, either Fibre host or iSCSI host.

Procedure

1. On the **Servers** window, click **Server Group**, then click the plus sign (+) to open the **Add Server Group** window.

2. Add the server group name. A description is optional.
3. Select the servers you want to add to the server group, then click **Add**. All the servers in a single server group must use the same protocol. Click **Submit**.

Updating a server group

You can add and remove servers from a group or update the name and description.

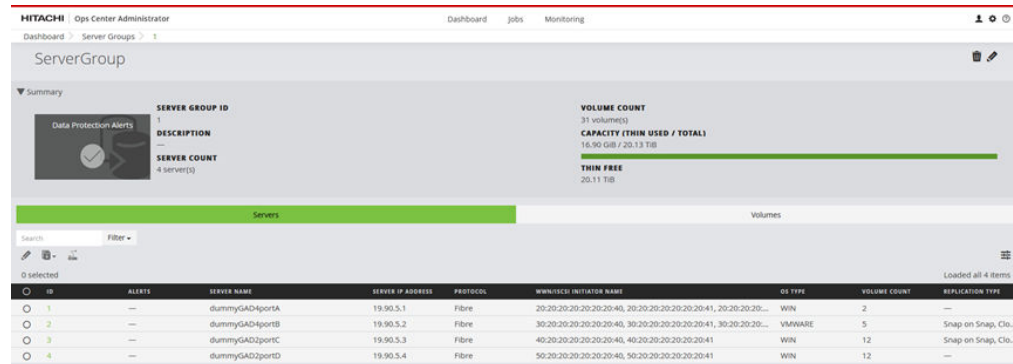
Procedure

1. Access the **Updating a Server Group** window by opening the **Servers** window. Click the **Server Groups** tab.
2. Select a server group, then click **Edit**.
3. You can update the name and description. Add or remove servers to the server group as needed, then click **Submit**.

Server group details

Use the server group details window to manage servers and attached volumes in the server group.

Access the server group details window by clicking Server Group ID in the Server Group tab of the Servers window.



The following server group functions can be performed on this window:

- Select a server group, then click Edit to access the Update Server Group window for a single server group, where you can add or remove servers, and edit the name or description.
- Select one or more server group, then click Attach Volumes to open the menu for Attach Existing Volumes, Create, Attach and Protect Volumes with Local Replication or with High Availability.

Managing volumes

A volume is a single accessible storage area, created in a pool within a parity group on a storage system. Create a volume and manage its properties with Ops Center Administrator.

Volumes inventory

The Volumes window enables you to filter, sort, and edit volumes for a single storage system, and to select volumes for data protection. The window shows tiered, thin, and snapshot volumes. The tiered volumes are DP-VOLs created from Dynamic Tiering pools. The thin volumes are DP-VOLs created from Dynamic Provisioning pools. The snapshot volumes are V-VOLs for Thin Image S-VOLs.



Note: The following volumes are not displayed:

- The emulation type is not OPEN_V
- Volumes using Cache Residency Manager cache
- Volumes encrypted by Encryption License Key feature
- Pool volumes
- NAS boot volumes
- Deduplication system data volumes
- Thin/tiered volumes with data direct mapping enabled
- Volumes related to the Namespace for FC-NVMe

Access the inventory of volumes for a storage system to gain insight into volume size, data protection, and utilization. You can also create volumes, attach existing volumes and perform data protection functions.

ID	VIRTUAL ID	S/N	VIRTUAL S/N	LABEL	POOL ID	ATTRIBUTE	STATUS	CAPACITY SAVING	TOTAL	USAGE	PROVISIONING STATUS
13 (00:00:00)	13 (00:00:00)	39303	39303	hadyos20200304...	6	Thin	Normal	No	150.00 MiB	0%	Unmanaged
14 (00:00:00)	14 (00:00:00)	39303	39303	hadyos20200304...	6	Thin	Normal	No	150.00 MiB	0%	Unmanaged
15 (00:00:00)	15 (00:00:00)	39303	39303	hadyos20200304...	6	Thin	Normal	No	150.00 MiB	0%	Unmanaged
16 (00:00:10)	16 (00:00:10)	39303	39303	vol1-1-1	7	Thin	Normal	No	512.00 MiB	0%	Unmanaged
20 (00:00:14)	20 (00:00:14)	39303	39303	hadyos20200304...	6	Thin	Normal	No	150.00 MiB	0%	Unmanaged
21 (00:00:15)	21 (00:00:15)	39303	39303	vol1-2-1	7	Thin	Normal	No	512.00 MiB	0%	Unmanaged
22 (00:00:16)	22 (00:00:16)	39303	39303	vol1-2-2	7	Thin	Normal	No	512.00 MiB	0%	Unmanaged
23 (00:00:17)	23 (00:00:17)	39303	39303	vol1-3-1	7	Thin	Normal	No	512.00 MiB	0%	Unmanaged
24 (00:00:18)	24 (00:00:18)	39303	39303	vol1-3-2	7	Thin	Normal	No	512.00 MiB	0%	Unmanaged
25 (00:00:19)	25 (00:00:19)	39303	39303	vol2-2-1	7	Thin	Normal	No	1.00 GiB	0%	Unmanaged
27 (00:00:18)	27 (00:00:18)	39303	39303	vol2-4-1	7	Thin	Normal	No	512.00 MiB	0%	Unmanaged
28 (00:00:1C)	28 (00:00:1C)	39303	39303	vol2-5-1	7	Thin	Normal	No	512.00 MiB	0%	Unmanaged
29 (00:00:10)	29 (00:00:10)	39303	39303	gs42ggsvol	7	Thin	Normal	No	10.00 GiB	0%	Unmanaged
30 (00:00:1E)	30 (00:00:1E)	39303	39303	gs42ggsvol	7	Thin	Normal	No	10.00 GiB	0%	Unmanaged
31 (00:00:1F)	31 (00:00:1F)	39303	39303	gs42ggsvol	7	Thin	Normal	No	10.00 GiB	0%	Unmanaged

View volume details in either a tile view or a list view. Capacity utilization is represented by a color bar:

- Green: under 70%.
- Orange: from 70% to 80%.
- Red: over 80%.



Note: Volumes labeled HSA-reserved are created automatically when file pools are created and no procedures can be performed on them in Ops Center Administrator.



Note: The HA Consistency Group column does not display its value if a displayed volume is in a storage system without SVP.

The following actions are available on this window:

- To interrupt shredding processes, click Interrupt Shredding and confirm.
- If there are Data Protection Alerts, click the tile to view the alerts in the Monitoring tab.
- Click the plus sign (+) to add volumes on the **Create Volumes** window.
- Click a volume to view details and attach, update, or delete the volume.
- Select a volume, then click Edit to update the volume by editing the name or updating the Capacity Saving setting and command device settings. You can also expand the capacity.
- You can also edit the names of multiple volumes and enable or disable command device at the same time if you select multiple volumes.

Select one or more volumes to perform one of the following actions:

- To delete volumes, click Delete. When you delete a volume, it is de-provisioned and removed from the storage system.



Note: If deduplication or compression has been enabled on the volume and the volume has data, it cannot be deleted in Ops Center Administrator. Access Storage Navigator to delete the volume.

- To attach volumes to one or more servers on the **Attach Volumes** window, click Attach Volumes.
- To detach volumes from servers, click Detach Volumes.



Note: The command device configuration will be disabled if all the attached LUN paths are removed.

- To apply data protection on the **Protect Volumes** window, click Protect Volumes.



Note: Command devices cannot be protected. Do not perform this action for command devices.



Note: Ops Center Administrator can protect volumes with high availability in existing replication group with following conditions:

- P-VOL STATUS and S-VOL STATUS of all paired volumes in the replication group is PAIR.
- Status of the corresponding replication in Ops Center Protector is OK.

- To unprotect a volume, select it, then click Unprotect.



Note: Ops Center Administrator can unprotect volumes with high availability in the replication group with following conditions:

- P-VOL STATUS and S-VOL STATUS of all paired volumes in the replication group is PAIR.
- Status of the corresponding replication in Ops Center Protector is OK.



Note: A volume in an "External" replication group cannot be unprotected.

- To restore a secondary volume to the primary volume, click Restore to open the **Restore Volume** window.
- To attach the volume to a storage system to prepare for volume migration, click Attach to Storage.
- To migrate volumes, click Migrate Volumes.



Note: Command devices cannot be migrated. Do not perform this action for command devices.

- To detach volumes from storage, click Detach from Target Storage.



Note: The command device configuration will be disabled if all the attached LUN paths are removed.

- To shred volumes that have been migrated, click Shred Volumes.
- Select a tiered volume, then click Update Tiering Policy to update the tiering policy.
- To attach like volumes, click Attach Like Volumes.

Click Open the **Logical Devices** window in Storage Navigator to open the **Volumes inventory** window and:

- Create a Command Device from a volume
- Restore a blocked volume

For more information, see the Storage Navigator online help.



Note: Any changes you make in Storage Navigator might not be reflected in Ops Center Administrator for a few minutes.

About capacity saving

Capacity saving reduces the used capacity of volumes.

Ops Center Administrator enables you to use capacity saving when you create or update a volume.

A capacity saving volume is a dynamically provisioned volume (DP-Vol) that is capable of compression or both deduplication and compression.

Ops Center Administrator offers the following capacity saving options:

- Compression
Compression allows you to reduce the used capacity of your volume.
- Deduplication and compression
Deduplication identifies duplicate data blocks to remove redundancy.
- No
No data compression or deduplication is done.

Creating volumes

Create volumes for a registered storage system so you can attach them to a server.

When you create volumes in Ops Center Administrator you can:

- Create multiple volumes of the same size or different sizes at the same time.
- Select the specific pool for volume creation or let Ops Center Administrator automatically select the best pool based on utilization.
- Specify a common label and starting label suffix for identical volumes that are the same size and have the same pool requirement.
- Set up compression, or deduplication and compression for the volume.

Before you begin

- Create parity groups.
- Create pools.
- Add servers (optional).

Procedure

1. On the Ops Center Administrator dashboard, click **Storage Systems** to see the inventory of storage systems and capacity information.
2. Click a storage system to see its configuration of servers, pools, ports, volumes, and parity groups.
3. Click **Volumes** to see the inventory of configured volumes for the storage system.
4. Click the plus sign (+) to open the **Create Volumes** window.

Create Volumes

STORAGE SYSTEM
10054
SUBSCRIBED CAPACITY
0% (100.00 GiB)
VIRTUAL STORAGE MACHINE
no-virtualize
Capacity
1.32 PiB Total

VOLUME LABEL: [Volume] LABEL SUFFIX: [0] NUMBER OF VOLUMES: [1] SIZE: [1] GiB

VOLUME ID / RANGE: Auto From To

POOL TYPE: Tiered POOL TIER: - POOL: Auto Selected TIERING POLICY: 0: All (default: Select Pool t...)

CAPACITY SAVING: No

COMMAND DEVICE SETTING: Disable COMMAND DEVICE SECURITY: Disable USER AUTHENTICATION: Enable DEVICE GROUP SETTING: Disable

Cancel Submit

5. Configure volumes for the specified storage system.
 - a. To add the volume to a virtual storage machine, select one in the **Virtual Storage Machine** menu.
 - b. Enter the volume label and select a suffix for it.
 - c. Select the number of volumes.
 - d. Select the size.
 - e. Select the volume unit: **GiB**, **TiB**, or **Blocks**.
 - f. Select the pool type: **Tiered** or **Thin**.
 - g. For a Thin pool, select the **Pool Tier**.
If the storage system has available capacity from external storage, you can also select the **External** tier.
 - h. (Optional) Select the pool from the list of available pools. The default selection is **Auto Selected**, which means that Ops Center Administrator selects the best pool for provisioning the volume based on utilization and tier requirements.



Note: DDM pools are not available and cannot be used to create volumes.

- i. You can select **Compression**, **Deduplication and Compression**, or **No** from the **Capacity Saving** drop down list.



Note: You can set capacity saving for volumes based on tiered pools for the following storage systems:

- VSP 5000 series
- VSP G1x00, F1500 with microcode version 80-05-4x or later.



Note: If you choose **Deduplication and Compression** and later want to update the volume to **Compression** you must first disable **Capacity Saving**.

- j. (Optional) Select a **Volume / ID Range** to specify one for the new volume for granular management of volumes. For instance, for ease of management, you can assign a certain range of IDs to certain departments in your organization. ID ranges can be specified in decimal or hexadecimal. The default selection is **Auto**, which means that Ops Center Administrator automatically selects the ID for the volume.
 - k. (Optional) For a Tiered Pool, select the **Tiering Policy** from the list. Tiering policy choices available in the list depend on the choice of pool that was made in the previous step. Tiering policy choice is not available for auto-selected pools.
 - l. You can set the volumes as command devices by selecting **Enable** from the **Command Device Setting** drop down list.
6. Click the plus sign to move the configured volume to the lower portion of the **Create Volumes** window.
A blank row appears at the top of the window where you can create more volumes.
 7. Click **Submit**.

Result

A job is started to create the volumes and add them to the volume inventory for the storage system.

Next steps

- Verify the status of the volume-creation job by clicking Jobs.
- Attach volumes to a server.

Attaching existing volumes to a server

If there are volumes that have not been attached to servers, you can use this procedure to attach them.

Before you begin

- Enable security on ports.
- Add servers.
- Create volumes.

Procedure

- To select volumes and servers, do one of the following:
 - On the **Volumes** inventory window for a storage system, select one or more volumes and click **Attach Volumes** to select servers on the **Attach Volumes** window.
 - On the **Servers** inventory window, select one or more servers, then click **Attach Volumes** to select volumes on the **Attach Volumes** window.
 - On the detail window for a pool, select one or more volumes, then click **Attach Volumes** to select servers on the **Attach Volumes** window.

Attach Volumes

Select Volume(s) that you want to Attach to the Server(s)

☐ Select All

VOLUME ID	VIRTUAL ID	S/N	NAME	POOL ID	STATUS	REPLICATION STATUS	CAPACITY SAVING	TOTAL	THIN USED	PROVISIONING STATUS
<input type="radio"/>	13 (00:00:00)	39303	hadyoz20200...	6	NORMAL	-	No	150.00 MiB	0 GiB	UNMANAGED
<input type="radio"/>	14 (00:00:0E)	39303	hadyoz20200...	6	NORMAL	-	No	150.00 MiB	0 GiB	UNMANAGED
<input type="radio"/>	15 (00:00:0F)	39303	hadyoz20200...	6	NORMAL	-	No	150.00 MiB	0 GiB	UNMANAGED
<input type="radio"/>	16 (00:00:10)	39303	vol1-1-1	7	NORMAL	-	No	512.00 MiB	0 GiB	UNMANAGED
<input type="radio"/>	20 (00:00:14)	39303	hadyoz20200...	6	NORMAL	-	No	150.00 MiB	0 GiB	UNMANAGED
<input type="radio"/>	21 (00:00:15)	39303	vol1-2-1	7	NORMAL	-	No	512.00 MiB	0 GiB	UNMANAGED

Showing 50 of 3393 Volumes

Free Space 0 - 1024 PiB

Size of Volume 0 - 1024 PiB

Utilization (%) 0 - 100

Volume Type ALL

Cancel Next

- Click **Next** to view options for attaching existing volumes.

Attach Volumes

STORAGE SYSTEM

RAID900_9_VSP5500H@10.196...

HOST MODE

WIN

HOST MODE OPTION

6 - TRIOLO

HOST GROUP NAME

[Optional]

MANDATE LUN ALIGNMENT

YES NO

AUTO CREATE ZONE

YES NO

VIRTUAL ID / RANGE

Auto

Selected Volumes

LABEL	ID	SIZE	LUN
hadyoz20200304-1login1	13 (00:00:00)	150.00 MiB	LUN (optional)
hadyoz20200304-1login1	14 (00:00:0E)	150.00 MiB	LUN (optional)

Cancel Previous Next

- The **Host Mode** is set by default to the server operating system. You can make a selection if needed.
The server OS Type is provided when the server is added to Ops Center Administrator.

4. The prepopulated **Host Mode Option** depends on the **Host Mode** selection. The default Host Mode Option can be changed manually.

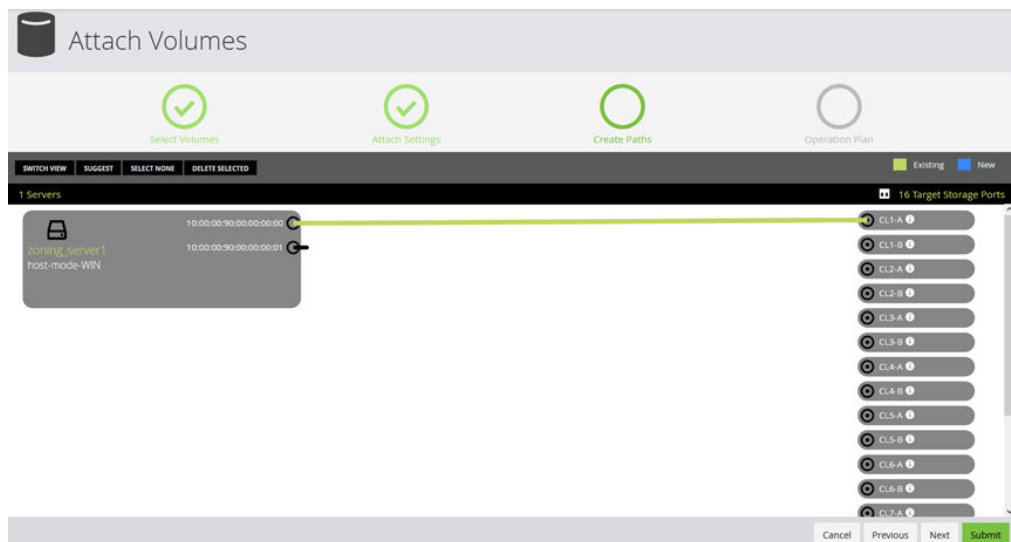
Ops Center Administrator identifies all host groups containing any of the server WWNs. If all of those host groups have the same host mode and host mode options, those settings are prepopulated with the same settings in the host groups.

Default values are set only for **VMWARE EX** and **WIN EX** host modes. The default for all other Host Modes is none.

5. Select **Mandate LUN Alignment**.

This option specifies whether to assign the same LUN number to multiple servers for a volume. If **Yes** is specified, the same LUN number is always assigned. If **No** is specified, the same value is not necessarily used. If the attachment is to only one server, this setting has no effect.

6. (Optional) Select a method and specify a virtual volume ID range in the **Virtual ID/Range** field.
7. The **Auto Create Zone** is set to **No** by default. You can set it to **Yes** to automatically create zones.
8. (Optional) Enter the LUN number to which you want to attach the volume.
9. Click **Next** to view options for creating and editing LUN paths.



10. In the **Create Paths** window, you can view servers and their WWNs, along with ports on the storage system. The following options are available for managing LUN paths:
For VSP 5000 series storage systems, you can click a port to display a list of ports with their DKC locations, CTL locations, and redundancy levels.

Port Redundancy Level				
NO.	PORT ID	DKC LOCATI...	CTL LO...	REDUNDANC...
1	CL4-A	DKC-1	CTL11	DKC
2	CL2-A	DKC-1	CTL11	DKC
3	CL6-A	DKC-1	CTL11	DKC
4	CL8-A	DKC-1	CTL11	DKC
5	CL1-J	DKC-2	CTL21	DKC
6	CL3-J	DKC-2	CTL21	DKC
7	CL5-J	DKC-2	CTL21	DKC
8	CL7-J	DKC-2	CTL21	DKC
9	CL2-J	DKC-3	CTL31	DKC
10	CL4-J	DKC-3	CTL31	DKC

Total Ports: 15

OK

Ops Center Administrator will scan for existing host groups on the storage array and try to reuse them by default. The following options are available for managing LUN paths:

- If you connect more than one server to the same port, the **Share a host group with all servers** box displays. Select the box to add the servers to a single host group.



Note: If host groups already exists on that port and single host group cannot be created, the check box will not appear.

- Existing paths are populated as follows: all existing host groups with one or more server WWN and the exact same host mode and host mode options selected on the Attach Settings window are populated as paths.

To prevent the volume from being added to an existing path, click the path to highlight it, then click **Delete Selected**.

- Click **Suggest** to populate automatically selected paths. By default, the least-used ports are selected.
 - To manually create a path, click a WWN and a port to connect them with a blue connector line. Click the connector again and then click **Delete Selected** to delete the connection. Suggest paths need for both server and storage ports to be logged into the fabric switches in the Ops Center Administrator inventory.
 - Optionally, you can click **Switch View** to use tables to create paths by selecting Server Ports to attach to Storage Ports.
11. Click **Next** to continue to the optional Operation Plan step, you can confirm settings specified in each window and can confirm the resources to be used in the job, including LUN candidates.



Note: Internet Explorer 11 does not support the Operation Plan page. Use another supported browser to access this feature.

Attach Volumes

Progress indicators: ✓ Select Volumes ✓ Attach Settings ✓ Create Paths ○ Operation Plan

Selected Servers

SERVER ID	SERVER NAME	SERVER IP ADDRESS	PROTOCOL	OS TYPE	VOLUME COUNT	REPLICATION TYPE
7	zoning_server1	15.15.15.15	Fibre	WIN	2	—

Selected Volumes

VOLUME ID	VIRTUAL ID	STORAGE SYST...	VIRTUAL STOR...	NAME	POOL ID	STATUS	REPLICATION T...	REPLICATION S...	CAPACITY SHEL...	TOTAL	THIN USED	PROVISIONING...
13 (00:00:00)	—	RAID900_9...	39303-VSP5...	hadyoz1202...	6	Normal	—	—	No	150.00 MiB	0.00 GiB	Unmanaged
14 (00:00:00)	—	RAID900_9...	39303-VSP5...	hadyoz1202...	6	Normal	—	—	No	150.00 MiB	0.00 GiB	Unmanaged

Attach Settings

PARAMETER	VALUE
Host Mode	WIN
Host Mode Options	6
Host Group Name	—
Mandate LUN Alignment	No
Auto Create Zone	No
Mandate Using Displayed Volume IDs	No

Buttons: Cancel Previous **Submit**

12. You can specify the LUN range for volumes in the **LUN Settings** pop up window which is launched by clicking LUN Settings button above the Planned Path Configuration table.
13. You can make it mandatory to use the LUNs displayed in the table by specifying Yes to **Mandate Using Displayed LUN**.
14. When you are satisfied with all the settings, click **Submit**.

Attaching like volumes

Allocate volumes to servers using existing volume settings. Attach like volumes for local replication or for high availability.

Procedure

1. Select a volume in one of the following ways:
 - On the **Server** detail window.
 - On the **Volumes** window.
 - Begin on a **Volume** detail window.
 - On the details window for a pool.
 - On the **Volumes** window in a virtual storage machine.

2. Click **Attach Like Volumes** and select one of the following:

- Select **Attach Like Volumes with Local Replication** to access the **Create, Attach and Protect Volumes** window.

The volume settings and paths will be copied from the selected volume. You can make changes as needed.

- Select **Attach Like Volumes with High Availability** to access the **Create, Attach and Protect Volumes with High Availability** window.

The volume settings and paths will be copied from the selected volume. You can make changes as needed.

See "Create, attach, and protect volumes with High Availability" in the *Hitachi Ops Center Administrator High Availability User Guide* for more information.

Updating a volume

You can expand and rename any volume and enable or disable capacity saving (deduplication and compression), command device setting, and ALUA mode.

You can also rename volumes and enable or disable command devices at the same time if you select multiple volumes.

Expanding protected volumes is supported for the following storage systems:

- VSP 5000 series
- VSP E series
- VSP G/F350, G/F370, G/F700, G/F900
- VSP G1x00, F1500



Note: From the Ops Center Protector block device advanced system settings, you can set the storage system prerequisite advanced system settings. This does not apply to VSP G1x00, F1500 storage systems.

Before you begin

To expand a protected volume, see the relevant hardware manual to confirm that all prerequisite conditions are met.

Procedure

1. To update a volume, navigate to the volume details by using one of the following options:
 - From the dashboard, click **Storage Systems**, and then click a storage system tile to view the resources. Click **Volumes**, and then click the volume tile for the volume that you want to update.
 - From the dashboard, click **Servers**, and then click a server tile to view the volumes. Click a volume tile for the volume that you want to expand.
2. In the **Volume Detail** window, click **Edit** to open the **Update Volume** window.



Note: To reset a volume to the default settings, click **Reset**.

3. You can rename the volume and change the size by clicking the up and down arrows next to the volume size.
4. Under **Change Compression Type**, select **Compression** or **Deduplication and Compression** to change the type. To disable compression, select **No**.



Note: If **Deduplication and Compression** is enabled and you want to update to **Compression**, you must first update to the option to **No** and then update it again.

5. You can update command device settings by selecting **Change Value** under **Change Command Device Settings**. To keep current settings, select **Keep Current Value**.



Note:

- You can enable or disable command devices.
- When changing the command device setting, you must specify the command device security, user authentication, and device group setting.
- If selected command devices have different settings or command devices belongs to a storage system without an SVP, then command device security, user authentication, and device group setting will be displayed as empty value. When updating at least one of these attributes, set all three attributes.
- If you have protected, T10 PI, or ALUA volumes, then you cannot enable them as command devices. The ALUA setting and command device settings cannot be updated at the same time.

6. Click **Submit**.



Note: If you have registered Ops Center Protector 7.2.5 or later, you do not need to complete a pair operation (split pair/resync pair) in advance to expand the HA pair volume.



Note: In order to use new features of Ops Center Protector after an upgrade installation, wait 4 hours or restart the Ops Center Administrator service so that the version information of Ops Center Protector is correctly updated.

To restart the service, see "Restart the service" in [Managing the Ops Center Administrator service with Podman \(on page 189\)](#) or [Managing the Ops Center Administrator service with Docker \(on page 188\)](#).



Note: For HA paired volumes, regardless of whether you change the volume size, this operation expands the capacity if the paired volume capacity is smaller than the target volume capacity. The HA paired volumes will have the same capacity.

Result

A job is started to update the volume.

Volume details

Access volume details in one of these ways:

- On the dashboard, click Storage Systems, then click a storage system tile. Click Volumes to open the **Volumes** window then click a volume in the list or tile view.
- On the detail window for a pool, click a volume in the list or tile view.



The volume details window exposes pool details including Pool ID, type, and tier of the pool. The volumes attached to servers are listed under Attached Servers.

The Provisioning Status can be any of the following:

- Attached: The volume is attached to a server.
- Unattached: The volume is not attached to a server.
- Unmanaged: The volume has only a LUN path associated. For example, a volume might be attached to a server that is not known to Ops Center Administrator.

Details for protected volumes also include details of all replication groups to which the volumes belong.

If the volume is a command device, then command device attributes (command device security, user authentication and device group setting) are shown.



Note: These three attributes are shown only if the command devices belongs to a storage system with an SVP.

On a volume detail window, you can do the following:

- Attach the volume to a server by clicking Attach Volumes.
- Edit the volume configuration by clicking Update Volumes to open the Update Volumes window.

When you click Submit, a notification appears stating that job has started to expand the volume.

- For a tiered volume, click Update Tiering Policy to access the Update Tiering Policy window.
- Delete the volume.



Note: If deduplication or compression has been enabled on the volume and the volume has data, it cannot be deleted in Ops Center Administrator. Access Storage Navigator to delete the volume.

- Protect the volume by clicking Protect Volumes to open the **Protect Volumes** window.



Note: Command devices cannot be protected. Do not perform this action for command devices.



Note: Ops Center Administrator can protect volumes with high availability in existing replication group with following conditions:

- P-VOL STATUS and S-VOL STATUS of all paired volumes in the replication group is PAIR.
- Status of the corresponding replication in Ops Center Protector is OK.

- Unprotect a volume by clicking Unprotect Volumes.



Note: Ops Center Administrator can unprotect volumes with high availability in the replication group with following conditions:

- P-VOL STATUS and S-VOL STATUS of all paired volumes in the replication group is PAIR.
- Status of the corresponding replication in Ops Center Protector is OK.

- Click Migrate Volumes to access the **Migrate Volumes** window.



Note: Command devices cannot be migrated. Do not perform this action for command devices.

- Click Attach Like Volumes to create, attach, and protect volumes with local replication or with High Availability.
- Expand any or all replication groups to which the volume belongs and view details of the volume pairs. In the table, you can click an S-Vol ID to open the detail window for the secondary volume.
- Restore a secondary volume to a primary volume by selecting the pair, then clicking Restore Volume Pair.
- Perform the following functions on a replication group:
 - Edit the replication group by clicking Edit.
 - Suspend the replication group by clicking Suspend Replication Group.
 - Resume the replication group by clicking Resume Replication Group.
 - Delete the replication group by clicking Delete.



Note: Replication group functions are not available for external replication groups. These are volume pairs that were created outside Ops Center Administrator.

Updating the tiering policy for a volume

You can choose a new tiering policy for a single volume.

Procedure

1. On the details window for a volume, click **Update Tiering Policy**.
2. On the **Update Tiering Policy** window, select a tier level from the list, then click **Submit**.

Shredding volumes

You can shred volumes using a shredding pattern that you create or a pre-defined pattern. You can view and find the volumes that are shredding and you can interrupt (stop) shredding processes.

You must be a storage administrator to perform shredding.

Ops Center Administrator allows you to interrupt (stop) the running shredding processes. To perform an urgent procedure while shredding is running, first interrupt the shredding job.



Note: Ops Center Administrator limits the number of volumes you can shred at any one time to 300, because of the high load to the storage system and the long processing time.

Before you begin

Ensure that:

- The volumes are unattached and are not snapshots.
- The volumes do not have capacity savings.

Procedure

1. From the dashboard, click **Storage System**, click a storage system tile to open the detail window, then click **Volumes**.
2. Select volumes to be shredded.
3. Click **Shred Volumes**.

4. Choose **Basic** or **Advanced**. The **Basic** method performs one pass. Choose **Advanced** to select the number of passes and enter a pattern for each pass.

Result

The volume has the status SHREDDING in the Volumes window.

Managing host groups

Create a host group on a storage system and manage its properties with Ops Center Administrator.

Host groups inventory

The Host Group inventory window displays host groups of a storage system and enables you to create, edit, and delete host groups.



Note: Host groups that belong to NVMe mode ports are not displayed.

Accessing the host groups inventory window

To access the **Host Groups** window, click **Storage Systems** in the Dashboard. Then select a storage system from the S/N column. From the storage system details window, click the **Host Groups/ISCSI Targets** tile.

The **Host Groups** inventory window appears:

ID	HOST GROUP NAME	TARGET ALIAS	PORT ID	PROTOCOL	FCID	TARGET NAME	HOST MODE	NUMBER OF LUNs	NUMBER OF WINDOWS	VIRTUAL LUN	VIRTUAL MODEL	HOST MODE OPTIONS
CL1-A-0	TA-G00		CL1-A	Fibre	---		LINUX	0	0	39303	VSP S100H, 5500H	---
CL1-A-1	htm-en3		CL1-A	Fibre	---		LINUX	9	2	39303	VSP S100H, 5500H	---
CL1-A-2	SSP-HTNM-F800		CL1-A	Fibre	---		LINUX	8	65	39303	VSP S100H, 5500H	---
CL1-A-3	test		CL1-A	Fibre	---		CHVMS	8	64	39303	VSP S100H, 5500H	2, 6, 7, 12, 13, 14, ...
CL1-A-4	hadycc20200304-1logn1		CL1-A	Fibre	---		AIX	4	1	39303	VSP S100H, 5500H	---
CL1-A-5	svt-1		CL1-A	Fibre	---		WIN	1	1	39303	VSP S100H, 5500H	---
CL1-A-6	gsd2g-gps-hgt		CL1-A	Fibre	---		LINUX	3	3	39303	VSP S100H, 5500H	---
CL1-A-7	gsd2g-gps-hgtst1		CL1-A	Fibre	---		LINUX	305	0	10057	VSP F1500 and VS...	---
CL1-A-8	gsd2g-gps-b-qps		CL1-A	Fibre	---		VMWARE EX	512	2	39303	VSP S100H, 5500H	54, 63
CL1-A-9	ham3x-rgv032		CL1-A	Fibre	---		LINUX	1	1	39303	VSP S100H, 5500H	---
CL1-A-10	gatekeeper		CL1-A	Fibre	---		LINUX	0	1	39303	VSP S100H, 5500H	---
CL1-A-11	zoning_server3		CL1-A	Fibre	---		WIN	0	1	39303	VSP S100H, 5500H	---
CL1-A-12	REST_TEST		CL1-A	Fibre	---		LINUX	0	0	10051	VSP F1500 and VS...	---
CL1-A-13	hisa-kokubun-1		CL1-A	Fibre	---		LINUX	0	0	39303	VSP S100H, 5500H	---
CL1-A-14	ham3x-rgv063		CL1-A	Fibre	---		LINUX	1	1	39303	VSP S100H, 5500H	---
CL1-A-15	ham3x-rgv064		CL1-A	Fibre	---		LINUX	1	1	39303	VSP S100H, 5500H	---
CL1-A-16	server5		CL1-A	Fibre	---		HP UX	1	2	39303	VSP S100H, 5500H	---
CL1-A-17	server6		CL1-A	Fibre	---		HP UX	1	1	39303	VSP S100H, 5500H	---

You can filter the displayed host groups by protocol, host mode, and host mode options. You can select multiple host mode options.

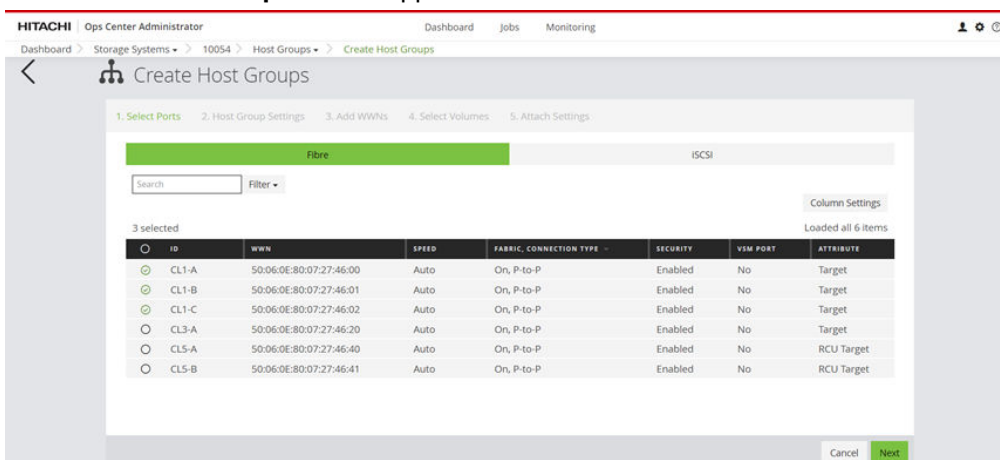
Creating host groups

Create a host group as follows:

Procedure

1. From the **Dashboard**, click **Storage Systems**.

2. Select a storage system from the **S/N** column, then click the **Host Groups/ISCSI Targets** tile. The **Host Groups** window appears.
3. Click the (+) sign to create a host group. The **Create Host Groups** wizard appears.



4. Follow the steps in the **Create Host Groups** wizard.
5. After providing the required information, click **Submit**.

Editing host groups

Edit host groups as follows:

Procedure

1. From the **Dashboard**, click **Storage Systems**.
2. Select a storage system from the list.
3. Click the **Host Groups/ISCSI Targets** tile. The **Host Groups** window appears.
4. Select a host group from the list, then click the Edit (✎) icon to edit a host group. The **Edit Host Group** window appears.

HITACHI | Ops Center Administrator Dashboard Jobs Monitoring

Dashboard > Storage Systems > 20055 > Host Groups > CL1-A-0 > Edit Host Group

Edit Host Group

HOST GROUP NAME
1A-G00

HOST MODE
LINUX

HOST MODE OPTIONS
Nothing selected

PREFERRED PATH

Preferred Non Preferred

Cancel Submit

From Edit Host Group, you can edit:

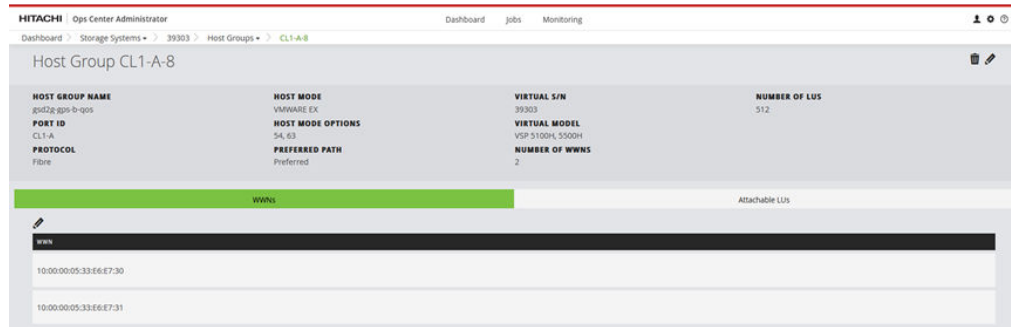
- Host Group Name
- Host Mode
- Host Mode Options
- Preferred Path

5. After you make changes, click **Submit**.

Host group details

From the **Host Groups** window of a storage system, you can click a host group listed in the ID column to view more detailed information.

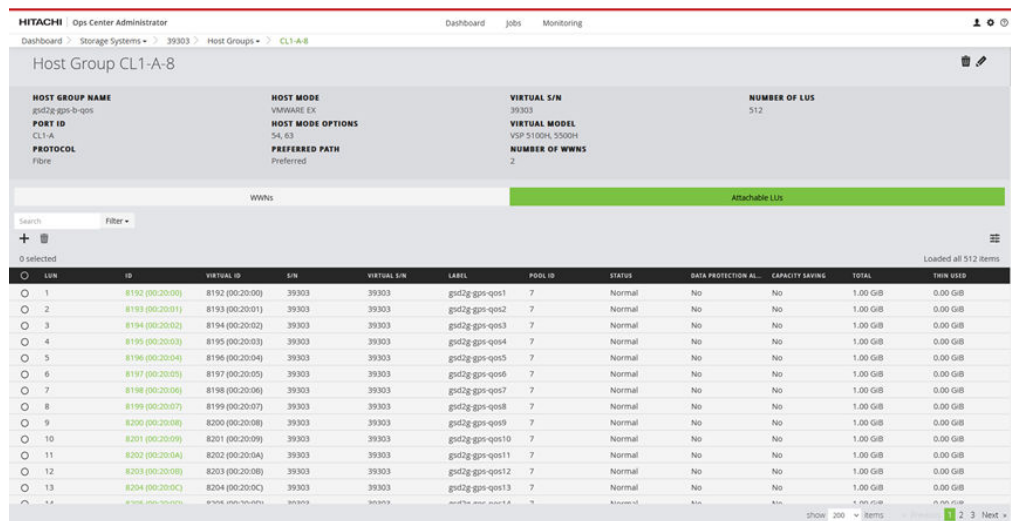
The WWNs tab displays a list of WWNs along with host group details.



The LUs tab displays the logical units along with host group details.



Note: The command device configuration will be disabled if all the attached LUN paths are removed.



Adding and deleting WWNs from a host group

The WWNs tab displays a list of WWNs along with detailed information on each host group. The LUs tab displays the logical units of a host group. You can use the toggle to switch from the WWNs tab to the LUs tab.

Add or delete WWNs from host groups as follows:

Procedure

1. From the **Dashboard**, click **Storage Systems**.
2. Select a storage system from the **S/N** column, then click the **Host Groups/ISCSI Targets** tile. The **Host Groups** window appears.
3. Select a host group from the list.
4. Click the **Edit** icon to open the **Edit WWNs** window.
5. Add WWNs by clicking **+ Add WWN** or you can click **WWN List** to select from a list of available WWNs.


You can also specify WWNs to delete from a host group.

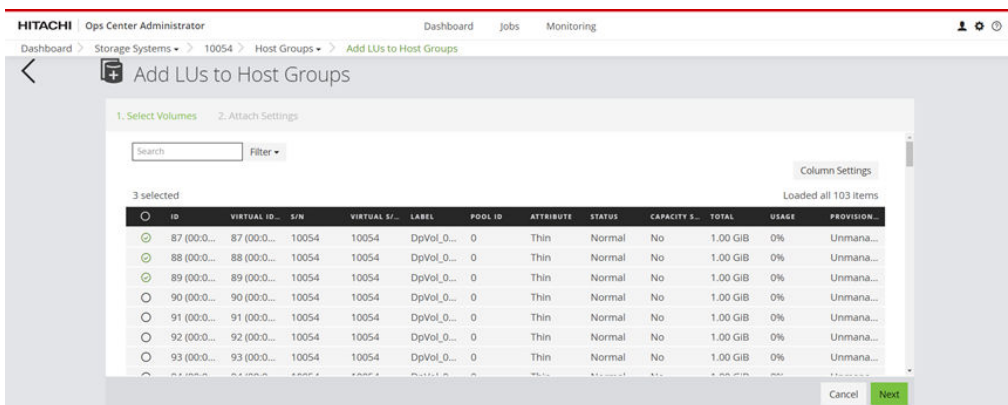
6. Click **Submit**.

Adding volumes to host groups as logical units

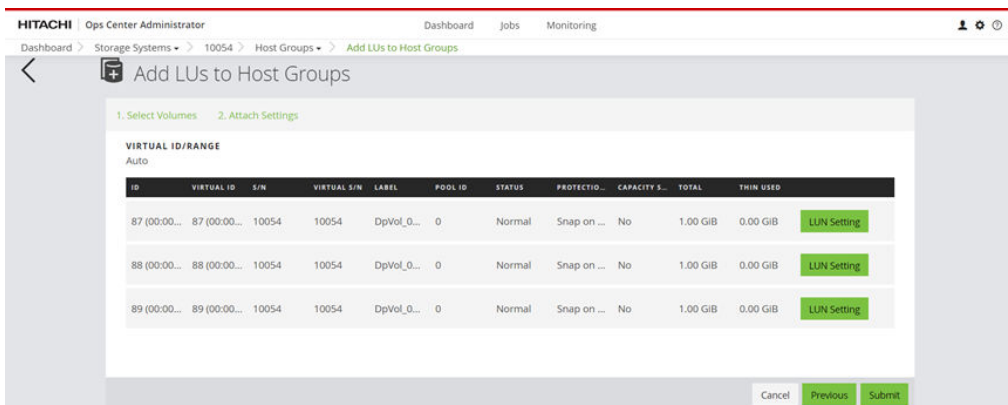
Add volumes to host groups as logical units (LUs) as follows:

Procedure


1. From the **Dashboard**, click **Storage Systems**.
2. Select a storage system from the list.
3. Click the **Host Groups/ISCSI Targets** tile.
The **Host Groups** window appears.
4. Select a host group from the list.
5. Click the Add LUs to Host Groups () icon to open the wizard.



6. Select one or more volumes to add to a host group.
7. Click **Next** to move to **Attach Settings**.



8. To edit the LUN setting for a volume, click **LUN Setting**.



9. Click **Submit**.

Chapter 4: Configure file storage

To configure file storage, begin by creating file pools and virtual file servers. You need both to create file systems.

Then create file systems and use them to create shares and exports that you can use to offer storage to users.

Managing file pools

Ops Center Administrator uses the following best practices in file pool creation:

- For creation of the related block pool, use of RAID6 with a minimum of four parity groups for SAS or one parity group for SSD, FMD, FMD DC2, SSD NVMe, or SSD(RI). s



Note: You should use this best practice. If capacity is chosen from non-conforming parity groups during file pool creation, a warning displays.

- Creation of a Thin pool that comprises each selected tier (either one or two).
- Minimum disk requirements for file pool creation:
 - 24 disks for SAS 10K, 15K
 - 48 disks for SAS 7.2K
 - 8 disks for FMD, FMD DC2, SSD, SSD NVMe, SSD(RI)


When a user creates a file pool in Ops Center Administrator, the system automatically does the following:

1. Identifies parity groups.
2. Creates a block pool.
3. Creates volumes and attaches them to NAS modules.
4. Configures volumes as system drives.
5. Creates a file pool.



File pool inventory

Access the File Pools window from the detail window for a storage system.



Use this window to view file pool information in a tile or list view, and to delete or expand existing pools, as well as to create new pools.

The colored bars display used capacity, physical capacity allocated for use on the underlying block storage, and the file over-commit capacity. The over-commit capacity reflects the amount by which the file pool is over-subscribed.

Each tile or list row shows whether the file pool is tiered, its File Pool Label, and capacity utilization.

The following actions are available:

- Create a new file pool by clicking the plus (+) sign to open the Create File Pool window.
- Delete one or more file pools by selecting them, then clicking Delete.
- Add capacity by selecting a file pool, then clicking Expand to open the Expand File Pool window.
- Open the File Pool detail window for the file pool by clicking a File Pool Label.
- Click a DP Pool [ID] link in a tile to open the detail window for the underlying block storage pool.

Creating file pools

Create file pools to create file systems.

Before you begin

The minimum disk requirements are:

- 24 disks for SAS 10K, 15K
- 48 disks for SAS 7.2K
- 8 disks for FMD, FMD DC2, SSD, FMD HDE, SSD NVMe, or SSD(RI)

You must also have a minimum number of parity groups. See the following table for a summary of both disk and parity group requirements:

Tier	Minimum number of disks	Minimum number of parity groups
Diamond	8	1
Platinum	8	1
Gold	24	4
Silver	24	4
Bronze	48	6

Both requirements must be met to create a file pool using Ops Center Administrator.

Examples of valid combinations:

- 1 or more 6D+2P SSD parity groups (Platinum)
- 2 or more 3D+1P SSD parity groups (Platinum)
- 4 or more 6D+2P SAS 15k Parity Groups (Gold)
- 4 or more 6D+2P SAS 10k parity groups (Silver)
- 6 or more 6D+2P SAS 7.2k parity groups (Bronze)

Examples of unacceptable combinations:

- 1 3D+1P SSD parity group (Platinum; meets only the 1 parity group requirement, not the 8-disk requirement.)
- 2 14D+2P SAS 10k parity groups (Silver; meets only the 24-disk requirement, not the 4-parity group requirement.)



Note: Ops Center Administrator only supports file pool creation from DP volumes and not from basic volumes created outside Ops Center Administrator.



Note: You cannot create a file pool with encrypted parity groups with Ops Center Administrator. Use another management tool to create a file pool with encrypted parity groups.

Procedure

1. On the **File Pools** window, click the plus (+) sign to open the **Create File Pool** window.
2. Enter a label for the file pool.
3. In the **Select Capacity for New File Pool** pane, choose storage from one or two tiers. For a tiered pool, select Platinum or Diamond and one other tier.



Note: Best practices suggest that file over-commit capacity should be an estimated 200%. Best practices also include using RAID 6, with a minimum of four parity groups for SAS drives or one parity group for SSD, FMD, FMD DC2, FMD HDE, SSD NVMe, or SSD(RI) drives. If another RAID layout is selected, a warning displays.

The information gauge displays the total capacity, with the physical block capacity and estimated over-commit capacity also detailed.

The **Estimated Max Potential**: in IOPS displays after each type of capacity is selected.

Table 3 Tier Definitions

Pool Tier	Description
Diamond	SCM NVMe
Platinum	SSD, FMD, FMD DC2, FMD HDE, SSD NVMe, SSD(RI)
Gold	SAS 15 k
Silver	SAS 10 k
Bronze	SAS 7.2 k

4. (Optional) Adjust the Utilization Thresholds.

5. Click **Submit**.

A job is added to create the file pool.

File pool details

Access the File Pool details window by clicking a File Pool Label in the File Pools window.

Use this window to view details about a file pool and to manage file systems in the pool.

You can search for a specific file system by searching for its File System Label in the search box. You can also use the Free Space and File Systems Size filters to find specific file systems, or filter by virtual file server.

View file system information in either a tile view or a list view, including capacity utilization and whether the file system is mounted.

Capacity utilization is represented by a color bar in each tile.

The following actions are available:

- Delete the file pool by clicking Delete.
- Expand the file pool by clicking Expand to open the Expand File Pool window.
- View details of an underlying block pool by clicking an DP pool link in the file pool tile.
- Add a file system by clicking the plus (+) sign to open the Create File System window.
- Delete one or more file systems by selecting them, then clicking Delete.
- Edit a file system by selecting it, then clicking Edit to open the Update File System window where you can change the label or expand the file system.
- Select a file system and mount or unmount it.

Expanding a file pool

File pool capacity might need to be expanded when the used capacity approaches the capacity of the underlying block pool.

You can also relabel a file pool or change utilization thresholds.

To access the Expand File Pool window do one of the following:

- In the File Pools window, select a file pool, then click Expand.
- In any file pool detail window, click Expand.

Procedure

1. To expand the file pool, select capacity from one of the tiers.
An untiered file pool cannot be expanded into a tiered file pool.
2. To change the file pool label, enter it in the **File Pool Label** field.
3. To change utilization thresholds, use the sliders.
4. Click **Submit**.
The job is added to the **Jobs** window.

Managing virtual file servers

This module shows you how to manage virtual file servers.

Virtual file server inventory

Access the Virtual File Servers window from the dashboard or from the detail window for a storage system.

Use this window to view information about virtual file servers in a tile or list view. Use the filters to show only offline or online virtual file servers, or filter by blade number.

The following actions are available:

- Open the Create Virtual File Server window by clicking the plus (+) sign.
- Open the Virtual File Server detail window, where you can delete, edit, and enable or disable the virtual file server by clicking the virtual file server label in any tile or list row.
- Open the storage system detail window by clicking the Storage System [ID] link in any tile or list row.
- Delete one or more virtual file servers by selecting them, then clicking Delete.
- Change the label for a virtual file server by selecting it, then clicking Edit to access the Update Virtual File Server window.
- You can also Delete, Enable, or Disable by selecting a virtual file server, then clicking the applicable icon.

Creating a virtual file server

Create virtual file servers so you can create file systems.

Before you begin

- Configure ports.
- Obtain the required IP address and subnet mask information for the virtual file server.

Procedure

1. On the **Virtual File Servers** window, click the plus sign (+) to open the **Create Virtual File Server** window.
2. Enter a label for the virtual file server.
3. Enter the IP address (ipv4 or ipv6) for the virtual file server.
4. Enter the subnet mask for the virtual file server.
5. Select the cluster node Gigabit Ethernet port to which the IP address for the virtual file server is assigned.
6. Click **Submit**.
A job is added on the **Jobs** window to create the virtual file server.

Virtual file server details

Access the Virtual File Server detail window from the Virtual File Servers window.

Use this window to manage a virtual file server's file systems.

The following actions are available:

- Open the storage system detail window by clicking a Storage System [ID].
- Select a file system in either the list or tile view and do one of the following:
 - Delete the file system by clicking Delete.
 - Expand a file system or change the label on the Update File System window by clicking Edit.
 - Click Mount or Unmount to mount or unmount the file system.
 - Open the File System detail window by clicking a file system.
 - Open the detail window for the underlying file pool by clicking the File Pool link in a tile.

Managing file systems

File system inventory

Access the File Systems window from the detail window for a storage system. Use this window to view information about file systems in a list or tile view, and to manage file systems.

You can search for a specific file system by searching for its file system label in the search box. You can also use the Free Space and File Systems Size filters to find specific file systems, or filter by virtual file server.

View file system details in either a tile view or a list view, including capacity utilization and whether the file system is mounted.

Capacity utilization is represented by a color bar in each tile.

The following actions are available:

- View details of a file system by clicking the file system label in a tile or row to open the File System detail window.
- Add a file system by clicking the plus (+) sign to open the Create File System window.
- Delete one or more file systems by selecting them, then clicking Delete.
- Edit a file system by selecting it, then clicking Edit to open the Update File System window.
- Select a file system and Mount or Unmount it.
- View details of the underlying pool by clicking a File Pool label in a tile.
- View details of the related virtual file server by clicking the VFS label in a tile.

Creating file systems

Create file systems so that you can create shares on Windows and exports on Linux.

Before you begin

Create file pools and virtual file servers.

Procedure

1. Click the plus sign (+) on the **File Systems** window to open the **Create File System** window.
2. Select a virtual file server.
 - a. Enter a label up to 255 alphanumeric characters.
 - b. Select a format:
 - **4K**: This choice is best for small block random applications (Virtual Server / Virtual Desktops / Databases).
 - **32K**: This choice is best for large block sequential applications (Video, Media, Images) or when the file system is hosted on NL-SAS drives.
3. Enter capacity up to 1PiB.
4. Select a pool, then click **Submit**.
A job is added to the **Jobs** window to create the file system.

File system details

Access the File System detail window from the File Systems window.

Use this window to create and manage shares and exports. You can also view information about a file system and access the associated file pool and virtual file server.

The following actions are available:

- Delete the file system by clicking Delete.
- Update label and capacity by clicking Edit to open the Update File System window.
- Mount or Unmount the file system by clicking Mount or Unmount.
- Access the file pool where the file system was created by clicking the File Pool [ID] link in the tile.
- Access the virtual file server where the file system was created by clicking the virtual file server label in the tile.
- Create a share or export by clicking the plus sign (+) to open the Create Shares/Exports window.
- Delete a share or export by selecting it, then clicking Delete.
- Update a share or export by selecting it and clicking Edit to open the Update Share or Update Export window.

Updating a file system

Update a file system to re-label or expand it.

Procedure

1. On the **File Systems** window, select a file system, then click **Edit** to open the **Update File System** window.
2. You can change the file system label, the capacity, or both.
 - To edit the label, enter changes in the label field.
 - To change the capacity, enter it and select a new unit of measure, if necessary.
3. Click **Submit**.
The new job is added to the **Jobs** window.

Managing shares and exports

Shares enable file sharing in Windows and exports enable file sharing in Linux.

Shares and exports inventory

Access the inventory of shares and exports in the detail window for a storage system to view information about existing shares and exports in a tile or list view.

Use this window to add, edit, and delete shares and exports. You can also access related resources.

The following actions are available:

- To add a share or export, click the plus sign (+) to open the Create Shares/Exports window.
- To delete a share or export, select a tile, then click Delete.
- To edit a share or export, select a tile, then click Edit to open the Update Export window or the Update Share window.
- Click File System label in a tile to open the detail window for the file system where the share or export was created.
- Click VFS label in a tile to open the Virtual File Server detail window where the file system was created.

Creating shares and exports

Create shares and exports to offer storage capacity to users.

Procedure

1. Access the **Create Shares/Exports** window by clicking the plus sign (+) on the **Shares/Exports** window.
2. Enter a **File System Path** and a **Share/Export Label**.
3. Select a **File System**.
4. Create a share by selecting **Windows** or create an export by selecting **Linux**.
5. Click **Submit**.
A job to create the share or export is added to the **Jobs** window.

Export details

Access the detail window for an export to review internal and external paths and to edit the export.

The following controls are available:

- Delete the export by clicking Delete.
- Edit the export by clicking Edit to open the Update Export window.
- Access the file system where the export was created by clicking File System.
- Access the virtual file server where the file system was created by clicking VFS.

Updating an export

You can update an export to change the file system path and access configuration.

Procedure

1. Access the **Update Export** window by clicking **Edit** in an Export detail window.
2. (Optional) Enter a new path in the **File System Path** field.
3. (Optional) In the **Access Configuration** field, enter IP addresses, host names, or the NIS netgroups of the clients who can access the NFS export (up to 5957 characters).

If the system has been set up to work with a name server, you can enter the NIS netgroup to which the clients belong, or the client's computer name rather than its IP address (not case-sensitive).

You can also specify the required flavors of NFS security in a colon-separated list using the option (sec=<list>).

The supported flavors are:

- none: Connect as a null user
- sys: The traditional security flavor used by NFS, users are not authenticated by the server
- krb5: Kerberos authentication
- krb5i: Kerberos authentication with per-messaging integrity
- krb5p: Kerberos authentication with per-message privacy

For example: 10.1.*.*(sec=sys:krb5:krb5i)

See the mount-point-access-configuration man window for further information.

4. Click **Submit** to save any changes.
A job is added to the **Jobs** window to update the export.

Share details

Access the **Share details** window to review internal and external paths and to edit the export.

The following controls are available:

- Delete the share by clicking Delete.
- Edit the share clicking Edit to open the Update Share page.
- Access the file system where the export was created by clicking the File System label.
- Access the virtual file server where the file system was created by clicking the VFS label.
- Add an existing Account Domain user group to the share by clicking the plus sign (+) to open the Create Groups page.

Adding existing account domain groups to a share

You can add an existing account domain group to a share. This enables you to change users permissions when they access the share.

Before you begin

The account domain group and CIFS setup must already exist.

Procedure

1. On the detail page for a share, click the plus sign (+) to open the **Create Groups** page.
2. In the **Group Name** field, type in a user group name from the account domain.
3. Choose permissions; **Full Control**, **Change**, or **Read**.

4. Click **Submit**.

A job is started to add the group to the share.

Updating a share

You can update a share to change the file system path, permissions, and access configuration.

Procedure

1. Access the **Update Share** page by clicking **Edit** in a Share detail page.
2. (Optional) Enter a new path in the **Modify File System Path** field.
3. (Optional) In the **Access Configuration** field, enter IP addresses of the clients who can access the share (up to 5,957 characters allowed in this field).

What to type	Means	Example
Blank or * Partial addresses using wildcards.	All clients can access the share.	10.168.*.* Clients with matching addresses can access the share.
Specific addresses	Only clients with the specified IP address can access the share.	10.168.20.2
Specific address range	Only clients with an IP address within the specified IP address range (10.168.20.0 to 10.168.20.255) can access the share.	10.168.20.0/16
Partial addresses using wildcards	Clients with matching addresses can access the share.	10.168.*.*
Permissions		
What to type	Permissions granted	
(rw)	read-write	
(ro)	read only	



Note: The order in which you specify the entries is important. Take the following two lines:

- *(ro)
- 10.1.2.38(rw)

The first grants read-only access to all clients, whereas the second grants read/write access to the specified client. The second line is redundant, however, as the first line matches all clients. You must transpose the lines to grant write access to 10.1.2.38. Examples:

- 10.1.2.38(ro) grants read-only access to the client whose IP address is 10.1.2.38.
- 10.1.2.0/24(ro) grants read-only access to all clients with an IP address in the range 10.1.2.0 to 10.1.2.255
- 10.1.*(rw) grants read-write access to all the matching clients.

4. (Optional) Select the supplied **Everyone** group and select permissions.
5. Click **Submit** to save any changes.
A job is added to the **Jobs** page to update the share.

Chapter 5: Protecting data

Data protection refers to replication of volumes. Ops Center Administrator has local and remote replication options.

Data protection concepts

You can protect volumes with local replication or with remote replication.



Note: Command devices cannot be protected. Do not perform this action for command devices.

Local replication

You can protect volumes locally on the same storage system and choose to protect a single volume, a group of volumes, or all volumes attached to a server.

Clone Now

The Clone Now method creates a copy of a volume in a point in time. This method creates a full copy of a primary volume and automatically suspends the copy procedure. Any further read-write functions of the primary volume will not be updated to the secondary volumes.

Snap

The Snap method takes incremental snapshots of volumes based on a user-created schedule. The retention policy (maximum number of snapshots) is from 1 to 1024, based on user input. When the retention policy is exceeded, the oldest snapshots are deleted when newer snapshots are created.



Note: When a user removes a storage system and adds it again, all the protected volumes using Clone Now will remain untouched, but any snapshot schedules will be suspended. The user will need to resume scheduling after adding the storage system.

Snap on Snap

Snap on Snap is a snapshot that can support a cascaded configuration. They can be created inside or outside Ops Center Administrator. If they are created outside Ops Center Administrator, you cannot perform functions on the replication group, which is labeled "External".



Note: The cascaded configuration is not supported within Ops Center Administrator.

Snap Clone

Snap Clone is a full copy of a snapshot (P-VOL) and creates a real volume.



Note: You cannot create Snap Clone pairs using Ops Center Administrator.

Remote replication

Use global-active device to help ensure reliability.

Selecting volumes to protect

You can select a server, create volumes and attach them to the server, and then (optionally) protect the volumes.

Alternatively, you can create volumes separately and attach them to servers later. After attaching volumes to servers, you can select volumes to protect.

Procedure

1. To select volumes to protect, do one of the following:
 - Select a server, create volumes for the server, and attach the volumes to the server.
 - Select a server to view volumes already attached to the server. You can filter for unprotected volumes.
 - On the dashboard, click **Storage Systems**, then click a storage system tile, then click **Volumes** to open the **Volumes** window.
 - Open the detail window for any volume that is attached to a server.
2. Select volumes and then click **Protect Volumes** and select **Protect Volumes with Local Replication** or **Protect Volumes with High Availability**.
3. For local replication, choose a type of data replication, **Clone Now** or **Snap**.



Note: The **Clone Now** method is available for volumes that are already existing and attached to a server. It is not available when the create-and-attach workflow is used to create, attach, and protect volumes in a single window.

If you use the workflow provided to create, attach, and protect volumes, you can choose **None** as a data replication type to avoid protecting volumes.

Protecting volumes with local replication

You can protect volumes by cloning or creating snapshots.

Protecting volumes by cloning

After creating and attaching volumes, you can implement data protection by creating a clone

The Clone Now method can be useful if maintenance or testing of a volume is needed. You can clone the primary volume and perform testing when the clone is complete. If any data is lost, you can retrieve it from the secondary volume.

You can create up to three clones for a volume. Using Ops Center Administrator, you can do this performing the cloning function on a set of volumes three times.



Note: The Clone Now method is available for volumes that exist and are attached to a server. It is not available when the create-and-attach workflow is used to create, attach, and protect volumes in a single window.

Procedure

1. In the **Protect Volumes** window, select the volumes you want to protect, then click **Clone Now**.
2. Choose whether to use **Consistency**.
Using Consistency means that copy functions will run on all pairs in the group simultaneously. Consistency, schedule and retention policy are all filters for replication groups, so your choices might affect the availability of replication groups.
3. Select a replication group from the options in the **Replication Group Name** list or enter a new name.
Up to 26 alphanumeric characters, and hyphens (-) and underscores (_), are allowed. Spaces are not allowed.
4. Click **Submit** to clone the volumes.

Result

A job is created and a summary message displays.

Protecting volumes by creating snapshots

After creating and attaching volumes, you can implement data protection by creating incremental snapshots.

You can use the Snapshot method if you don't need continuous copying but want to schedule the copy depending on expected changes in the data.



Note: For volumes on the following storage systems, you can create snapshots by accessing Storage Advisor Embedded:

- VSP E series
- VSP G/F350, G/F370, G/F700, G/F900

The snapshot creation process creates the secondary volume on a snap pool. If there is more than one snap pool, the least-used snap pool is used.



Note: To create snapshots outside of Ops Center Administrator for a volume that is in Ops Center Administrator, use a new snapshot group.

Procedure

1. In the **Protect Volumes** window, select the volumes you want to protect, then click **Snapshot**.
2. Enter a number from 1 to 1024 in the **Snapshot Retention Policy (Max 1024)** field to specify the number of snapshots to be created.
When the specified number of snapshots is reached, older snapshots will be deleted as new ones are added. The maximum number of snapshots that can be retained is 1024.
3. You can enable **Consistency** to perform data protection on all pairs in the group simultaneously.
Consistency acts as a filter for replication groups, so your choice might affect the availability of replication groups.
4. Set an hourly, daily, weekly, or monthly schedule. A volume cannot have more than one schedule. The schedule is based on UTC time.
 - **Hourly**: Select a start time in minutes after the hour, and the number of hours between snapshots to start creating snapshots at the selected time and interval.
 - **Daily**: Select a start time, available in five minutes increments, to start a snapshot every day at the selected time.
 - **Weekly**: Select a start time and a day to create a snapshot every week at the selected date and time. You can select more than one day to create more than one weekly snapshot.
 - **Monthly**: Select a start time and the day of the month (from 1 to 31) to create a snapshot one time a month at the selected day and time.
5. In **Volumes to be Protected**, select a **Replication Group** for each volume or select **Use New** and enter a name for the new replication group in the **Replication Group Name**. Up to 26 alphanumeric characters, also hyphens (-) and underscores (_), are allowed. Spaces are not allowed.
6. (Optional) Select the pool from the list of available pools. The default selection is **Auto Selected**, which means that Ops Center Administrator selects the best pool based on utilization and tier requirements.
If you choose a Thin (DP) pool or a Snap (HTI) pool, the replication technology used will be Snap on Snap, unless there is already a Snap volume created in an earlier version of Ops Center Administrator. In that case, the replication technology used will continue to be Snap.
7. Click **Submit**.

Result

A summary message is displayed, and snapshots are created according to the schedule. The replication group can be viewed on the **Replication Groups** window.

Protecting existing volumes with high availability

You can protect existing volumes with high availability.

Before you begin

- Register Ops Center Protector in Ops Center Administrator.
- Set up your environment for high availability.



Note: If you want to align the same volume ID between a primary volume with a secondary volume of a high availability pair, the secondary volume ID must belong to the metadvice or the same virtual storage machine as the primary storage system. If you want to use the same virtual storage machine, ensure it has the same resource group.

Procedure

1. Access the **Volumes** window and select one or more attached volumes. They cannot belong to different VSMs. Select **Protect Volumes with High Availability**.

SERVER NAME	SERVER ID	OS TYPE	ATTACHED VOLUME IDS
1A-G00_006	735	Linux	200 (00:00:C8)

2. The **Host Mode** is set by default to the server operating system. You can make a selection if needed.
The server OS Type is provided when the server is added to Ops Center Administrator.
3. The prepopulated **Host Mode Option** depends on the **Host Mode** selection. The default **Host Mode Option** can be changed manually.
Default values are set only for **VMWARE EX** and **WIN EX** host modes. The default for all other host modes is none.
Ops Center Administrator identifies all host groups containing any of the server WWNs. If all host groups have the same host mode and host mode options, those settings are prepopulated with the same settings in the host groups.
4. Select **Mandate LUN Alignment**.
This option specifies whether to assign the same LUN number to multiple servers for volumes containing primary and secondary HA pairs. If **Yes** is specified, the same LUN number is always assigned. If **No** is specified, Administrator tries to use the same LUN if it is available. If it is not available, Administrator selects available LUNs separately for each primary and secondary volume.
5. The **Auto Create Zone** is set to **No** by default. You can set it to **Yes** to automatically create zones.

6. **ALUA Enabled** is set to **Yes**. If it is set to **No**, you must rediscover volume or restart the OS after job completion.
7. Click **Next** to select secondary servers, if applicable.
8. Click **Next** to choose a secondary storage system and replication group, and quorum disk.

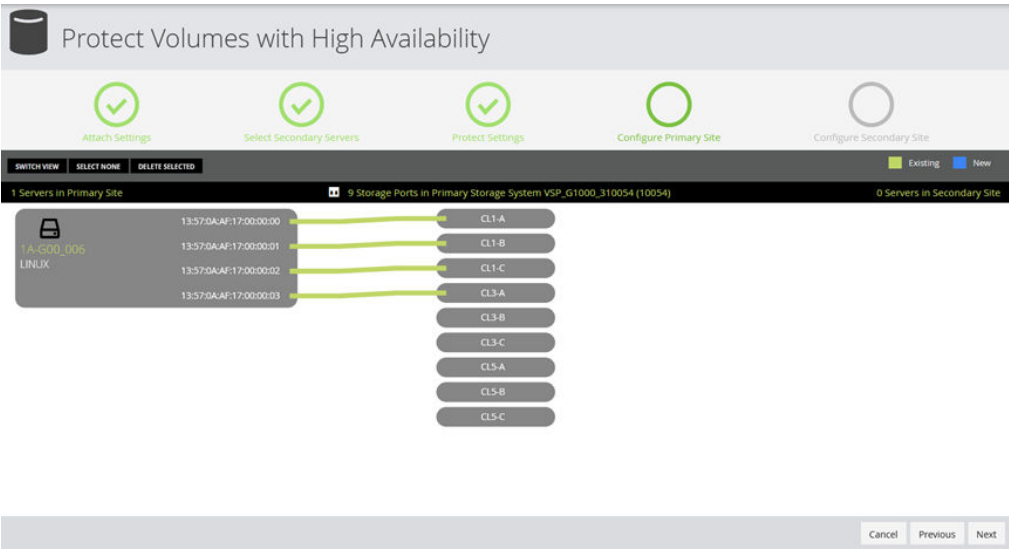
9. In the **Protect Volumes** window, select a secondary storage system.
10. Choose whether to use an existing replication group.
11. Select the **Consistency** setting. By default, this is set to **Yes**. Using consistency means that copy processes run on all pairs in the group simultaneously.
12. Select a replication group from the options in the **Replication Group** list or enter a new name. Up to 26 alphanumeric characters, and hyphens (-) and underscores (_), are allowed. Spaces are not allowed. If you select an existing replication group, the **Storage Pool of Secondary Storage System for S-Vol** and **Quorum Disk** are selected automatically.



Note: Ops Center Administrator can protect volumes with high availability in existing replication group with following conditions:

- P-VOL STATUS and S-VOL STATUS of all paired volumes in the replication group is PAIR.
- Status of the corresponding replication in Ops Center Protector is OK.

13. Select the pool to use for S-Vols and select a quorum disk.
14. Click **Next** to configure the primary site.
15. In the **Configure Primary Site** window, you can view the path of the primary server. If you need to change it, you must edit the LUN path before applying data protection.



Optionally, you can click **Switch View** to use tables to create paths by selecting Server Ports to attach to Storage Ports.

For VSP 5000 series storage systems, you can click a port to display a list of ports with their DKC locations, CTL locations, and redundancy levels.

Port Redundancy Level

NO.	PORT ID	DKC LOCATI...	CTL LO...	REDUNDANC...
1	CL4-A	DKC-1	CTL11	DKC
2	CL2-A	DKC-1	CTL11	DKC
3	CL6-A	DKC-1	CTL11	DKC
4	CL8-A	DKC-1	CTL11	DKC
5	CL1-J	DKC-2	CTL21	DKC
6	CL3-J	DKC-2	CTL21	DKC
7	CL5-J	DKC-2	CTL21	DKC
8	CL7-J	DKC-2	CTL21	DKC
9	CL2-J	DKC-3	CTL31	DKC
10	CL4-J	DKC-3	CTL31	DKC

Total Ports: 15

OK

16. Click **Next** to configure ports in the secondary storage system to the server in the primary site, and to the server in the secondary site, if there is one.
17. Click **Submit** to create a job to protect volumes.
18. You can monitor the job in the **Jobs** window.



Note: If the job does not complete successfully, access Ops Center Protector to remove the related resources (**Block Host Node**, **Policy**, and **Data Flow**) with the same name as the selected replication group. When using an existing replication group, remove added **P-Vol** from **Block Host Node**.

To edit a Block Host Node created by Ops Center Administrator in Ops Center Protector, specify LDEV IDs in decimal format per line and not in hex format or range format.

Viewing replication groups

Ops Center Administrator enables you to view all replication groups in a single window.

You can also access replication details for a volume to view and manage all replication group membership and volume pairs.

Ops Center Administrator detects and displays all volume pairs, whether they were created in Ops Center Administrator.

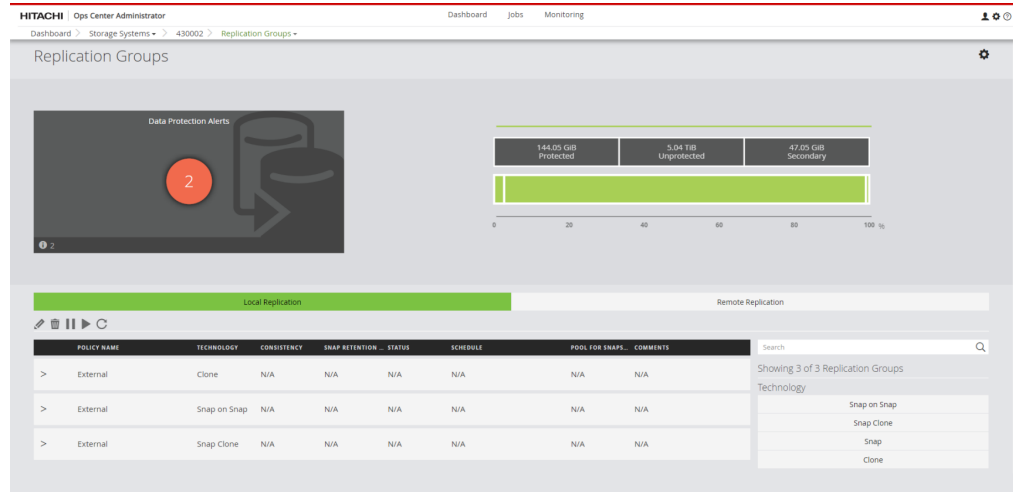
Replication groups apply the policies that control how volumes are replicated. Their parameters include:

- Type of technology:
 - High Availability for remote replication.
 - For local replication:
 - Snap on Snap
 - Snap Clone
 - Snap
 - Clone
- Whether consistency is used
- Snapshot schedule
- Snapshot retention

Replication groups are created when volumes are replicated.

Replication groups inventory

Access the inventory of replication groups from the detail window for a storage system.



This window enables you to view all replication groups, sorted by technology:

- Snap on Snap
- Snap Clone
- Snap
- Clone

Replication groups created outside Ops Center Administrator are listed as External.

You can view all volume pairs, including those created outside of Ops Center Administrator.



Note: Replication group functions other than Restore are not available for External replication groups. If a copy group is renamed outside Ops Center Administrator, the replication group will become External. Copy groups cannot be renamed using Ops Center Administrator.



Note: You can import HA pairs from Ops Center Protector. See [Importing HA replication pairs from Ops Center Protector \(on page 140\)](#) for more information.

Use this window to:

- View replication policies applicable to a volume and all S-VOL details.
- View all replication policies and all volume pair details.
- Suspend or resume a snapshot schedule.

Data replication can be suspended and resumed. Data will be missing for the period of suspension.

- Suspend or resume a clone.

Data replication can be resumed to update cloned data and then suspended to keep a point-in-time full copy.

- Remove a volume from a replication policy. Delete a replication group to remove the policy from all volumes in the group.
- Restore a volume pair by restoring a secondary volume to a primary volume.

The following actions are available after selecting a replication group:

- Edit the replication group by clicking Edit to open the Edit Replication Group window and edit the label or remove volumes from the replication group.
- Suspend the replication group by clicking Suspend Replication Group. You cannot suspend a replication group that is in COPY status.
- Resume the replication group by clicking Resume Replication Group.
- Delete the replication group by clicking Delete.

To view volume pair details, expand one or more replication groups.

- Access detail windows for each P-VOL and each S-VOL by clicking a P-VOL ID or an S-VOL ID in the expanded view of a replication group.
- Before using the Restore function, stop all I/O activity to the volume. To restore a volume pair, select it, then click Restore Volume Pair.



WARNING: The Restore function will overwrite the volume with backup data. The data on current volume will be destroyed.

Use the gear icon at the top of the **Replication Groups** window to open the following windows:

- Click Open the Local Replication window in Storage Navigator to view local replication groups information for clones and snapshots.
- Click Open the Remote Replication window in Storage Navigator to view remote replication groups information for clones and snapshots.
- Click Open the Remote Connection window in Storage Navigator to view remote connection groups information for clones and snapshots.
- Click Open the Journals window in Storage Navigator.
- Click Launch Ops Center Protector.

For more information, see the Storage Navigator online help.



Note: Any changes you make in Storage Navigator might not be reflected in Ops Center Administrator for a few minutes.

Editing replication groups

You can edit a replication group to adjust the schedule for snapshots or to remove volumes from the replication group.



Note: External replication groups cannot be edited. These include volume pairs that are created outside Ops Center Administrator.



Note: You can import HA pairs from Ops Center Protector. See [Importing HA replication pairs from Ops Center Protector \(on page 140\)](#) for more information.

Procedure

1. Access the **Edit Replication Group** window from the **Replication Groups** window.
2. You can edit the following:
 - Name
 - Comments
 - For snapshots, the schedule and number of snapshots. The schedule is based on UTC time.
 - (Optional) Select the pool from the list of available pools. The default selection is **Auto Selected**, which means that Ops Center Administrator selects the best pool based on utilization and tier requirements.



Note: You cannot edit the name and comments of imported replication groups.

3. To remove a primary volume from the replication group, select the corresponding row. You can choose whether to remove the corresponding secondary volume.



Note: Ops Center Administrator can unprotect volumes with high availability in the replication group with following conditions:

- P-VOL STATUS and S-VOL STATUS of all paired volumes in the replication group is PAIR.
- Status of the corresponding replication in Ops Center Protector is OK.

4. Click **Submit**.
A job is added to the Jobs window to apply the changes to the replication group.

Importing HA replication pairs from Ops Center Protector

Before you begin

Before importing the HA replication pairs, make sure of the following:

- Both primary and secondary storage systems of the high availability (HA) pairs are onboarded on Ops Center Administrator.
- The data flow for HA pairs is created with the following conditions in Ops Center Protector:
 - The node for the primary storage is Hitachi Block Host.
 - The logical devices for the primary volumes are defined in the Included Logical Devices of the Hitachi Block Host.
The logical device must be defined with the *LDEV_ID*. Ops Center Administrator cannot use logical devices defined with *Host Group ID*.
- Specify Use Block Host selections in the Specify Hitachi Block Storage classification attributes of the policy.
- Configure the following in Secondary Volume Host Groups in the Dataflow:
 - Specify Automatically Provisioned Host Group.
 - Do not specify any host groups in Optionally specify on or more host groups on the destination storage system.
- Replication is marked as Active Full Copy in Ops Center Protector.
- The HA pair must use a Quorum disk.

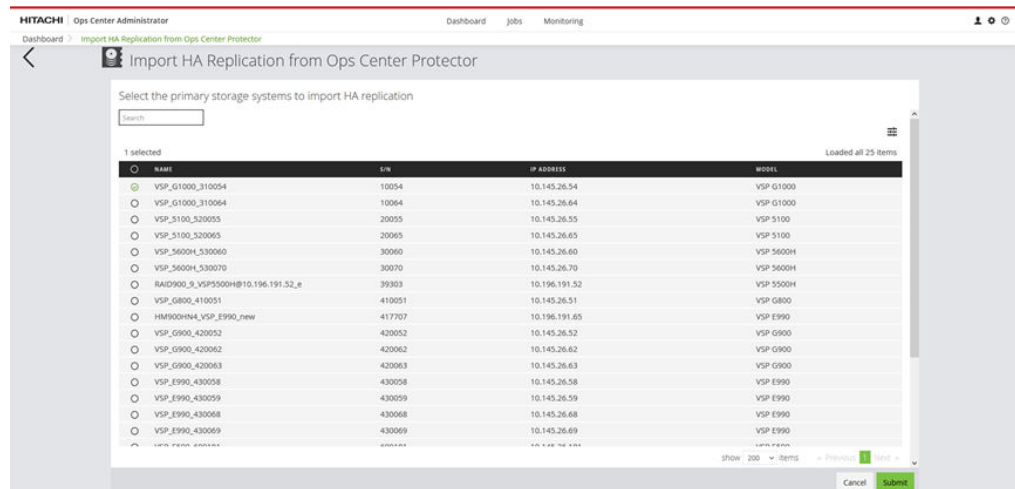


Note:

- If you add or remove pairs in the replication group, the logical device range will be regenerated and the updated logical device specification will be listed as a list of single logical device specification and the Excluded Logical Devices specification will be retained.
- Do not specify optional host groups in Secondary Volume Host Groups on Ops Center Protector after you import the replication. If you specify optional host groups, unexpected paths between the hosts and volumes may be created when adding or removing pairs in the replication group.

Procedure

1. From the **Settings** menu, click **Import HA Replication**.
2. Select the primary storage systems of the HA pair on which you want to import HA replication.



3. Click **Submit**.

Chapter 6: Manage virtual storage machines

Ops Center Administrator allows you to create virtual storage machines and add volumes and host groups to them.

You can also add defined and undefined resources to an existing virtual storage machine and remove defined and undefined resources from an existing virtual storage machine.

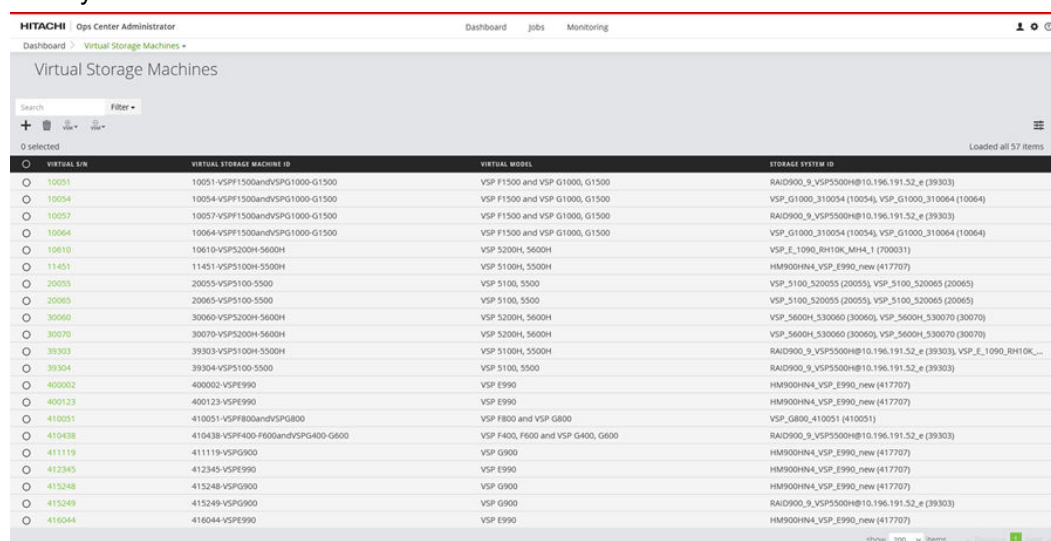


Note: To manage VSMs for high availability setup, click Go to VSM Inventory in the High Availability Setup window.

Virtual storage machine inventory

You can access the Virtual Storage Machines window to view all virtual storage machines and perform procedures on them.

Click Virtual Storage Machines on the dashboard to open the **Virtual Storage Machines** inventory.



VIRTUAL MACHINE ID	VIRTUAL STORAGE MACHINE ID	VIRTUAL MODEL	STORAGE SYSTEM ID
10051	10051-VSPF1500andVSPG1000-G1500	VSP F1500 and VSP G1000, G1500	RAID900_9_VSP5500HB10.196.191.52_e (39303)
10054	10054-VSPF1500andVSPG1000-G1500	VSP F1500 and VSP G1000, G1500	VSP_G1000_310054 (10054), VSP_G1000_310064 (10064)
10057	10057-VSPF1500andVSPG1000-G1500	VSP F1500 and VSP G1000, G1500	RAID900_9_VSP5500HB10.196.191.52_e (39303)
10064	10064-VSPF1500andVSPG1000-G1500	VSP F1500 and VSP G1000, G1500	VSP_G1000_310054 (10054), VSP_G1000_310064 (10064)
10610	10610-VSP5200H-5600H	VSP 5200H, 5600H	VSP_E_1090_RH10K_MHA_1 (700031)
11451	11451-VSP5100H-5500H	VSP 5100H, 5500H	HMM00HNA_VSP_E990_new (417707)
20055	20055-VSP5100-5500	VSP 5100, 5500	VSP_5100_520055 (20055), VSP_5100_520065 (20065)
20065	20065-VSP5100-5500	VSP 5100, 5500	VSP_5100_520055 (20055), VSP_5100_520065 (20065)
30060	30060-VSP5200H-5600H	VSP 5200H, 5600H	VSP_5600H_530060 (30060), VSP_5600H_530070 (30070)
30070	30070-VSP5200H-5600H	VSP 5200H, 5600H	VSP_5600H_530060 (30060), VSP_5600H_530070 (30070)
39303	39303-VSP5100H-5500H	VSP 5100H, 5500H	RAID900_9_VSP5500HB10.196.191.52_e (39303), VSP_E_1090_RH10K...
39304	39304-VSP5100-5500	VSP 5100, 5500	RAID900_9_VSP5500HB10.196.191.52_e (39303)
400002	400002-VSP E990	VSP E990	HMM00HNA_VSP_E990_new (417707)
400123	400123-VSP E990	VSP E990	HMM00HNA_VSP_E990_new (417707)
410051	410051-VSPF800andVSPG800	VSP F800 and VSP G800	VSP_G800_410051 (410051)
410438	410438-VSPF400-F600andVSPG400-G600	VSP F400, F600 and VSP G400, G600	RAID900_9_VSP5500HB10.196.191.52_e (39303)
411119	411119-VSPG900	VSP G900	HMM00HNA_VSP_E990_new (417707)
412345	412345-VSP E990	VSP E990	HMM00HNA_VSP_E990_new (417707)
415248	415248-VSPG900	VSP G900	HMM00HNA_VSP_E990_new (417707)
415249	415249-VSPG900	VSP G900	RAID900_9_VSP5500HB10.196.191.52_e (39303)
416044	416044-VSP E990	VSP E990	HMM00HNA_VSP_E990_new (417707)

The following actions are available on this window:

- Click a storage system ID to open the virtual storage machine details window.
- Select one or more virtual storage machines, then click Delete to permanently remove them.

All resources will be removed from the virtual storage machine and the virtual storage machine will be deleted. Removed resources will be moved to the meta resource group. A virtual storage machine with attached volumes or HA pairs cannot be deleted.

- Select a virtual storage machine, then click Remove VSM to open the menu and do one of the following:
 - Click Remove Defined Volumes from VSM to open the **Remove Defined Volumes from VSM** window and remove volumes from the virtual storage machine.
 - Click Remove Undefined Resources from VSM to open the **Remove Undefined Resources from VSM** window and remove resources from the virtual storage machine.
- Select a virtual storage machine, then click Add VSM to open the menu and do one of the following:
 - Click Move Volumes to VSM to open the **Move Volumes to VSM** window and move volumes from a storage system to a virtual storage machine.
 - Click Add Resources to VSM to open the **Add Undefined Resources to VSM** window and add resources from one or more storage systems to a virtual storage machine.

Creating virtual storage machines

You can use Ops Center Administrator to create virtual storage machines.

Each virtual storage machine must have a model number and a serial number. After assigning those, you can add volumes and host groups.

Procedure

1. On the dashboard, click **Virtual Storage Machines** and then click the plus sign (+) on the **Virtual storage machines** window.
2. On the **Create VSM** window, select the storage systems you want to use.

Create VSM

INTEND TO USE THE SAME MODEL/SERIAL NUMBER OF EXISTING STORAGE?

Yes No

VIRTUAL MODEL: Nothing selected

VIRTUAL SERIAL NUMBER: The range is between:

Select Storage Systems

Please select one or more storage systems to create VSM

☐ Select All

NAME	S/N	IP ADDRESS	MODEL
<input checked="" type="radio"/> VSP_G1000_310054	10054	10.145.26.54	VSP G1000
<input type="radio"/> VSP_G1000_310064	10064	10.145.26.64	VSP G1000
<input type="radio"/> VSP_5100_520055	20055	10.145.26.55	VSP 5100
<input type="radio"/> VSP_5100_520065	20065	10.145.26.65	VSP 5100
<input type="radio"/> VSP_5600H_530060	30060	10.145.26.60	VSP 5600H
<input type="radio"/> VSP_5600H_530070	30070	10.145.26.70	VSP 5600H

Showing 26 of 26 Storage Systems

Free Space: 0 - 1024 PiB

Size of Array: 0 - 1024 PiB

Migration Tasks Exist: ALL

False True

Cancel Next

3. Choose whether to use the model and serial number of one of the selected storage systems for the new virtual storage machine.
 - If you choose **YES**, the entire meta_resource group of the storage system will be the virtual storage machine.
 - Click **NO** to select a virtual model and a virtual serial number.



Note: If a resource group already exists, it will be used and a new one will not be created.

4. Click **Next** to add volumes to the virtual storage machine.

You can add volumes from any or all selected storage systems, or skip this and the next step and add volumes later.
5. For each storage system you selected, specify undefined volumes you want to add by choosing one of the following from the **Specify Volumes By** drop down:
 - Select **Number of Volumes** and enter a number in the **Number of Volumes** field beneath it.
 - Select **Number of volumes and volume ID range** and enter the **Number of Volumes** and **Volume ID/Range** in the fields beneath it.
 - Select **Volume ID range (add all volumes within the range)** and enter the **Volume ID/Range** in the field beneath it.
6. Click **Next** to add host groups. Add host groups from one storage system at a time, then click the plus sign (+) to add them to the list.
7. Click **Submit** to create a job to add the virtual storage machine.

Removing undefined resources from a virtual storage machine

You can remove undefined resources if you no longer need them.

Procedure

1. On the **Virtual Storage Machines** window, click the VSM icon (⊖) and select **Remove Undefined Resource** from VSM from the menu.
2. On the **Remove Undefined Resources from VSM** window, select a storage system from the menu.

3. For each storage system you selected, specify undefined volumes you want to remove by choosing one of the following from the **Specify Volumes By** drop down:
 - Select **Number of Volumes** and enter a number in the **Number of Volumes** field beneath it.
 - Select **Number of volumes and volume ID range** and enter the **Number of Volumes** and **Volume ID/Range** in the fields beneath it.
 - Select **Volume ID range (remove all volumes within the range)** and enter the **Volume ID/Range** in the field beneath it.
4. Click the plus sign (+) to display your selections in the lower half of the window. Add more selections if needed.
5. Click **Next**.
6. Remove host groups using the same procedure as for volumes.
7. Click **Submit** to start a job to remove the undefined resources.

Moving defined volumes to a virtual storage machine

You can move defined volumes from a storage system that is in a virtual storage machine to the virtual storage machine.

Procedure

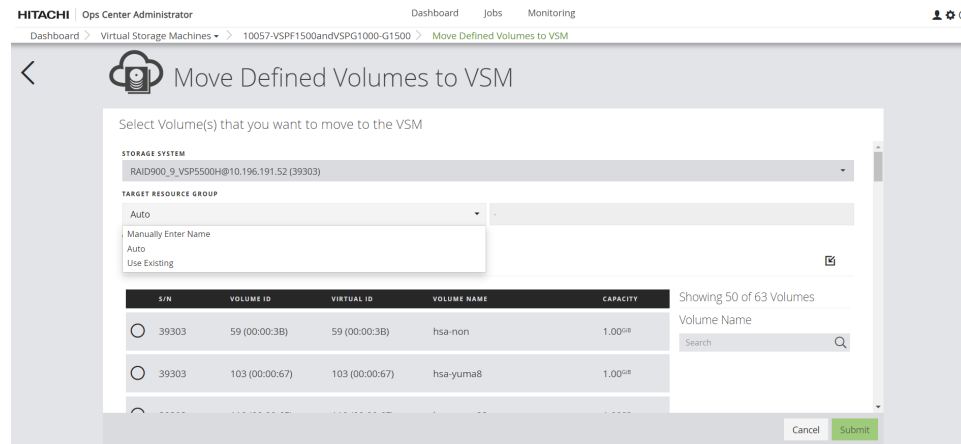
1. Select a virtual storage machine in the **Virtual Storage Machines** window.
2. Select **Add VSM** then click **Move Defined Volumes to VSM**.



Note: You can access the **Move Defined Volumes to VSM** window from the **Volumes** tab of the virtual storage machine detail window.

3. Select the volumes to be moved. You can specify the resource group from the **TARGET RESOURCE GROUP** menu. By default, the resource group name is assigned automatically.


- Click **Submit** to start a job to move the volumes.

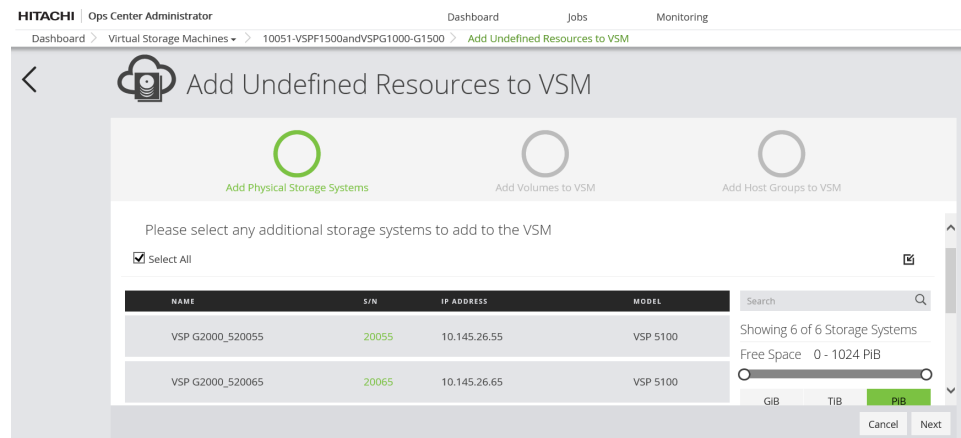


Adding resources to a virtual storage machine

You can add additional resources to a virtual storage machine.

Procedure

- To access the **Add Undefined Resources to VSM** window, select a virtual storage machine on the **Virtual Storage Machines** window.
- Click the VSM icon () and select **Add Undefined Resources to VSM**.



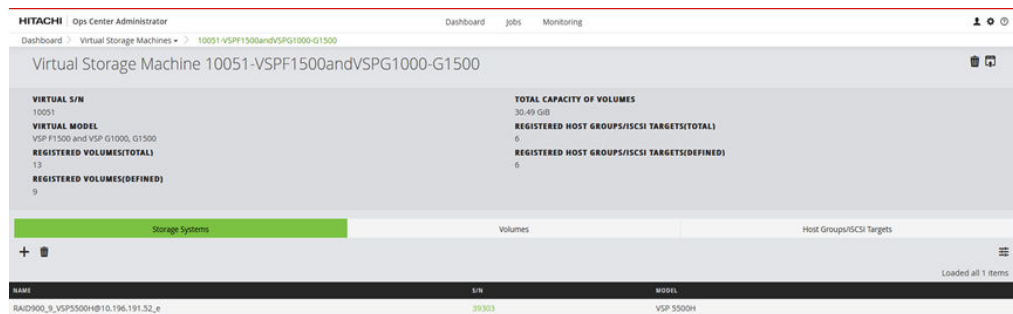
- Optionally, select one or more storage systems, then click **Next** to add volumes.
- For each storage system you selected, specify undefined volumes you want to add by choosing one of the following from the **Specify Volumes By** drop down:
 - Select **Number of Volumes** and enter a number in the **Number of Volumes** field beneath it.
 - Select **Number of volumes and volume ID range** and enter the **Number of Volumes** and **Volume ID/Range** in the fields beneath it.
 - Select **Volume ID range (add all volumes within the range)** and enter the **Volume ID/Range** in the field beneath it.

5. Click the plus sign (+) to add your selections to the lower window. When you are finished adding selections, click **Next** to add host groups.
At least one undefined volume must remain in each storage system.
6. Add host groups for each required storage system and port combination. Choose one of the following from the **SPECIFY HOST GROUPS BY** menu:
 - Select **Number of Host Groups**, then click the plus (+) sign to add them to the list.
 - Select **Host Group ID**, then click **HOST GROUP LIST** to specify the host group ID.
7. When all selections are complete, click **Submit** to start a job to add the undefined resources.

Virtual storage machine details

Review the details for a single virtual storage machine.

Click on a virtual storage machine in the **Virtual Storage Machines** inventory window to view details for the virtual storage machine.



The following actions are available:

- Click the Delete icon to delete the virtual storage machine.
- Click the Open HA Pairs icon (🔗) to view the global-active device pairs on this virtual storage machine.
- Click a storage system serial number in the list of storage systems under the Storage Systems tab to view the physical storage machine detail.
- Click the Storage Systems tab, Volumes tab, or Host Groups/iSCSI Targets tab to review the resources across physical storage systems and manage resources on the virtual storage machine.

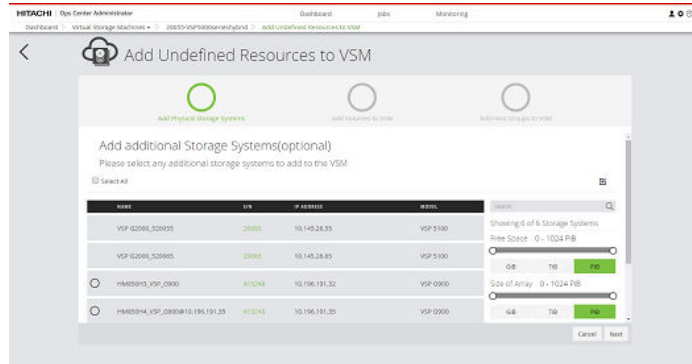
Adding storage systems to a virtual storage machine

You can add physical storage systems to a virtual storage machine (VSM) as follows:

Procedure

1. On the Ops Center Administrator dashboard, click **Virtual Storage Machines**.

2. Click the virtual storage machine listed by ID and model number) on which you want to add a storage system.
3. Under the **Storage Systems** tab, click the **Add Undefined Resources to VSM** (+) icon.
The **Add Undefined Resources to VSM** window appears.



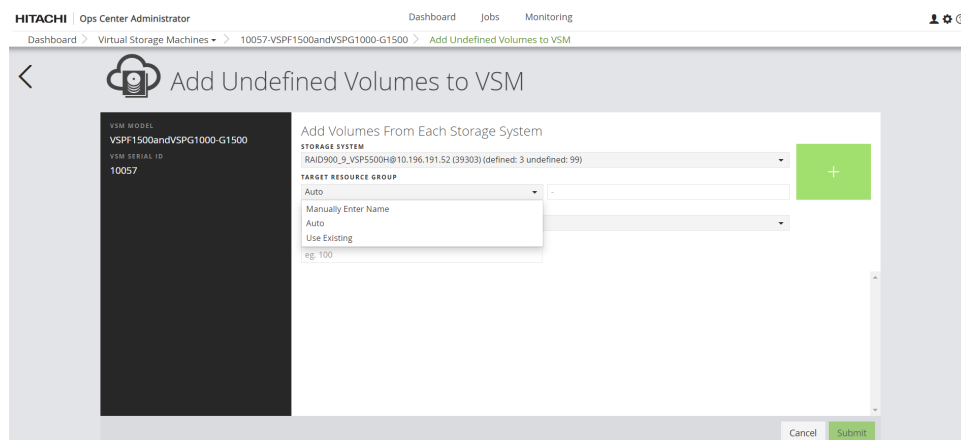
4. Select one or more storage systems to add to the VSM.
5. You can add volumes or host groups to the VSM, or just click **Next** twice, and then click **Submit**.

Adding undefined volumes to a virtual storage machine

You can add undefined volumes from a storage system to a virtual storage machine (VSM) as follows:

Procedure

1. On the Ops Center Administrator dashboard, click **Virtual Storage Machines**.
2. Click the virtual storage machine (listed by ID and model number) on which you want to add a volume.
3. Under the **Volumes** tab, click the **Add Undefined Volumes to VSM** (+) icon. The **Add Undefined Volumes to VSM** window appears.
4. Specify the resource group from the **TARGET RESOURCE GROUP** menu. By default, the resource group name is assigned automatically.



5. Select the storage system from which you want to add a volume to the VSM from the **STORAGE SYSTEMS** menu.
6. Specify the volumes that you want to add:
 - **Number of volumes**
 - **Volume ID range (add all volumes within the range)**
 - **Number of volumes and volume ID range**
7. Click **Next**. You can then choose to add host groups from a storage system to the VSM or click **Submit**.

Adding undefined host groups to a virtual storage machine

You can add undefined host groups and iSCSI targets from storage systems to a virtual storage machine (VSM) as follows:

Procedure

1. On the Ops Center Administrator dashboard, click **Virtual Storage Machines**.
2. Click the virtual storage machine (listed by ID and model number) on which you want to add a volume.
3. Under the **Host Groups/iSCSI Targets** tab, click the **Add Undefined Host Groups to VSM** (+) icon.

The **Add Undefined Host Groups to VSM** window appears.



Note: If the **Add Undefined Host Groups to VSM** window does not appear, either add storage systems and volumes to the VSM first, or click **Next** until it appears.

4. Select the storage system from which you want to add a volume to the VSM from the **STORAGE SYSTEMS** menu.
5. Specify the resource group from the **TARGET RESOURCE GROUP** menu. By default, the resource group name is assigned automatically.

Add Undefined Host Groups to VSM

VSM MODEL: VSPF1500 and VSPG1000-G1500
VSM SERIAL ID: 10057

SPECIFY HOST GROUPS BY: Number of host groups

Add Host Group From Each Storage System

STORAGE SYSTEM: HMB0052_VSP_G200@10.196.191.20 (430002)
PORT ID: CL4-B (defined: 0 undefined: 0)
TARGET RESOURCE GROUP: Auto
NUMBER OF HOST GROUPS: eg. 100

STORAGE SYSTEM	PORT ID	TARGET RESOURCE GROUP	NUMBER OF HOST GROUPS

Cancel Submit

6. In the **Add Undefined Host Groups to VSM** window, you can specify the host group you are adding either by **Number of hosts groups** or by **Host Group ID selection**.
7. Click **Submit**.

Removing storage systems from a virtual storage machine

You can remove physical storage systems from a virtual storage machine (VSM) as follows:



Note: If there are any existing registered volumes or host groups, Ops Center Administrator tries to remove them before removing the storage systems.

Procedure

1. On the Ops Center Administrator dashboard, click **Virtual Storage Machines**.
2. Click the virtual storage system ID number of the VSM from which you want to remove a storage system.
3. Under the **Storage Systems** tab, click the minus icon (⊖) to remove one or more storage systems from the VSM.

The **Remove Storage Systems from VSM** window appears.

Dashboard > Virtual Storage Machines > 10051-VSPF1500andVSPG1000-G1500 > Remove Storage Systems from VSM

Remove Storage Systems from VSM

Select Storage System(s) that you want to remove from the VSM

☒ Select All

NAME	S/N	MODEL
<input checked="" type="checkbox"/> VSP G2000_520055	20055	VSP 5100
<input checked="" type="checkbox"/> VSP G2000_520065	20065	VSP 5100
<input checked="" type="checkbox"/> HM850H4_VSP_G900@10.196.191.35	415249	VSP G900

Cancel Submit

4. Select the storage systems that you want to remove from the VSM, then click **Submit**.

Removing defined volumes from a virtual storage machine

You can remove defined volumes from a virtual storage machine (VSM) as follows:

Procedure

1. On the Ops Center Administrator dashboard, click **Virtual Storage Machines**.
2. Click the virtual storage machine (listed by ID and model number) from which you want to remove a volume.
3. Under the **Volumes** tab, click the icon and select **Remove Defined Volumes from VSM** from the menu. The **Remove Defined Volumes from VSM** window appears.

Remove Defined Volumes from VSM

Select Volume(s) that you want to remove from the VSM

☐ Select All

S/N	VOLUME ID	VIRTUAL ID	VOLUME NAME	CAPACITY	RESOURCE GROUP
410438	631 (00:02:77)	65535 (00:FF:FF)	-	321.00 ^{MB}	HDID_VSM_M85_430002
410438	632 (00:02:78)	65535 (00:FF:FF)	-	1.00 ^{MB}	HDID_VSM_M85_430002
410438	637 (00:02:7D)	637 (00:02:7D)	-	1.00 ^{MB}	HDID_VSM_M85_430002
410438	234 (00:00:EA)	65535 (00:FF:FF)	hsa-uehara-gad-hdid	6.50 ^{MB}	HDID_VSM_M85_430002
410438	2023 (00:07:E7)	undefined	vn305	2.00 ^{MB}	VSP G200@430002
410438	2022 (00:07:E6)	undefined	vn304	2.00 ^{MB}	VSP G200@430002
410438	2021 (00:07:E5)	undefined	vn303	2.00 ^{MB}	VSP G200@430002
410438	2020 (00:07:E4)	undefined	vn302	2.00 ^{MB}	VSP G200@430002

Showing 20 of 20 Volumes

S/N
Search

Volume Name
Search

Resource Group
Search

Cancel Submit

- Specify the volumes that you want to remove and then click **Submit**.

Removing undefined volumes from a virtual storage machine

You can remove undefined volumes from a virtual storage machine (VSM) as follows:

Procedure

- On the Ops Center Administrator dashboard, click **Virtual Storage Machines**.
- Click the virtual storage machine (listed by ID and model number) from which you want to remove a volume.
- Under the **Volumes** tab, click the **Remove Undefined Volumes from VSM** (⊖) icon. The **Remove Undefined Volumes from VSM** window appears.

Remove Undefined Volumes from VSM

VSM MODEL
VSP5000seriesAFA

VSM SERIAL ID
20055

Remove Volumes from VSM

STORAGE SYSTEM
VSP_5100_520065 (20065) (defined: 0 undefined: 20)

SOURCE RESOURCE GROUP
Auto (defined: 0 undefined: 20)

SPECIFY VOLUMES BY
Number of volumes

NUMBER OF VOLUMES
eg. 100

Cancel Submit

- Select the storage system from which you want to remove a volume from the VSM from the **STORAGE SYSTEMS** menu.
- Specify the resource group from the **SOURCE RESOURCE GROUP** menu.

6. Specify the volumes that you want to remove by:
 - **Number of volumes**
 - **Volume ID range (remove all volumes within the range)**
 - **Number of volumes and volume ID range**
7. Click **Submit**.

Removing host groups from a virtual storage machine

You can remove defined or undefined host groups (FC) and iSCSI targets (iSCSI) from a virtual storage machine (VSM) as follows:

Procedure

1. On the Ops Center Administrator dashboard, click **Virtual Storage Machines**.
2. Click the virtual storage machine (listed by ID and model number) on which you want to add a volume.
3. Under the **Host Groups/iSCSI Targets** tab, click the **Remove Undefined Host Groups from VSM** (⊖) icon. The **Remove Undefined Host Groups from VSM** window appears.

4. In the **Remove Undefined Host Groups from VSM** window, you can specify the host group you are adding either by the Number of hosts groups or by the Host Group ID selection.
5. Select the storage system from which you want to remove a host group from the VSM from the **STORAGE SYSTEMS** menu.
6. Specify the host groups that you want to remove by::
 - **Number of host groups**
 - **Number of host groups and host group ID range**
 - **Host Group ID range (add all host groups within the range)**

7. Click **Submit**.

Chapter 7: Migrate volumes

Volume migration is moving a volume and its associated data from a source location to a destination.



Note: Command devices cannot be migrated. Do not perform this action for command devices.

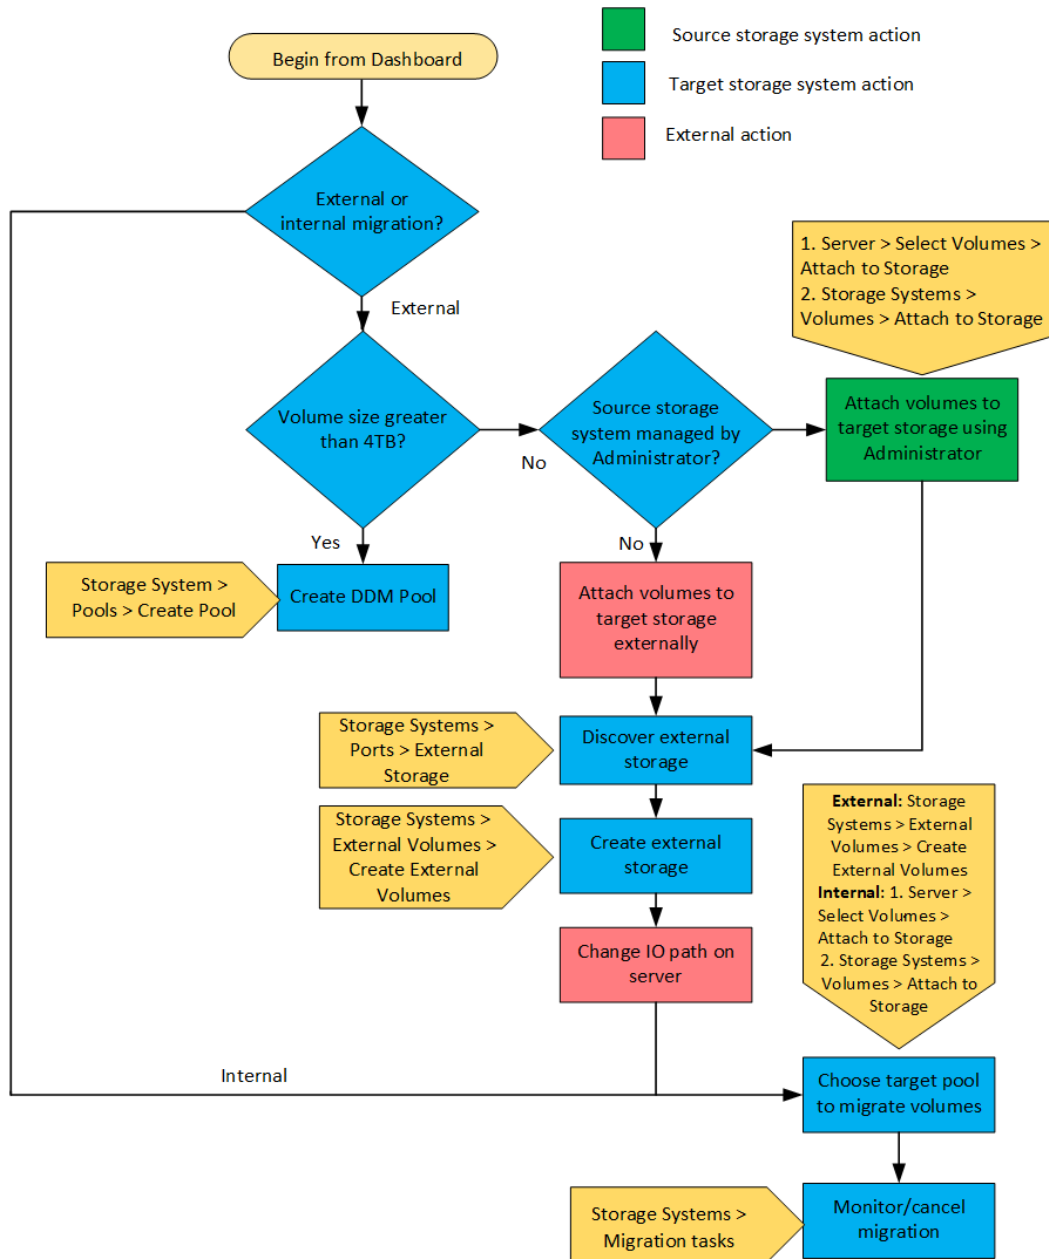
Volume migration workflow

Ops Center Administrator gives you an end-to-end workflow for migrating volumes from one pool to a different pool within one storage system (internal migration) and an end-to-end workflow for migrating volumes from one storage system to another (external migration).

Volume migration involves moving data from a set of source volumes to a set of destination volumes. 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.

For internal migration, see the following: [Migrating volumes \(on page 162\)](#),

The following figure shows the high-level workflow of volume migration in Ops Center Administrator.



External volume migration supports two options; in each case the volumes must be attached to the target storage as the first step in migration:

1. Both the source and target storage systems are managed in Ops Center Administrator.
2. The source storage system is one that is not supported by Ops Center Administrator. In this case, the volume cannot be attached using Ops Center Administrator.

Volume virtualization is a required step in external migration and (optionally) attaching volumes. It creates external parity groups, creates external volumes on the external parity groups, and attaches the volumes. It creates a pointer to the original volume from the target storage system.

External migration of volumes over 4TiB

To migrate an external volume greater than 4TiB, first create a data direct mapping (DDM) pool.

DDM pools are flagged **DDM Enabled: YES** in the **Pools** window and in the pool details windows.

To add a DDM pool, navigate to the Advanced tab of the Create Pool window and select Yes under Intend to enable DDM?

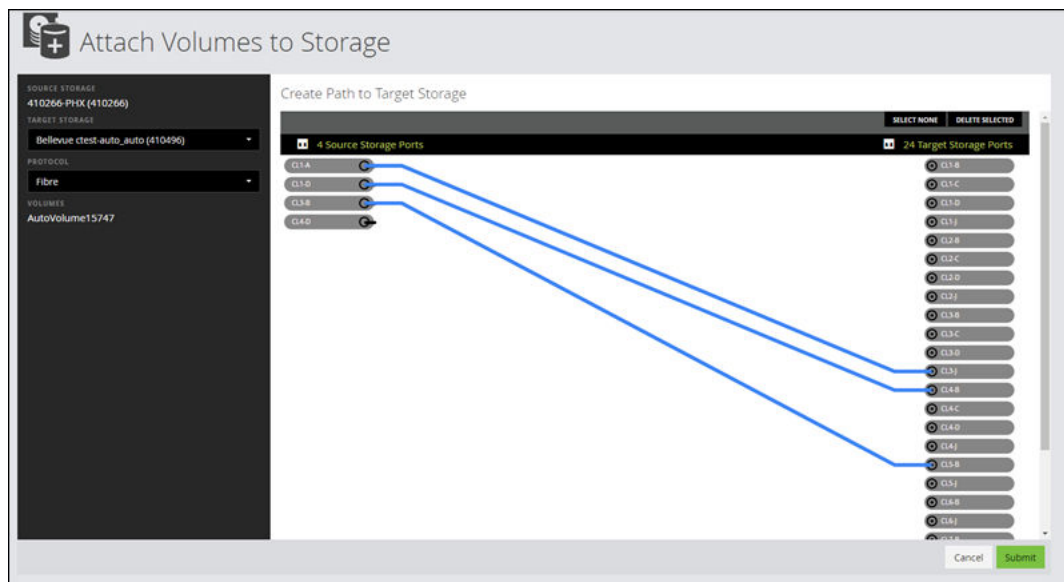
For DDM pool sizing, refer to "Estimation of the required capacity of pool-VOLs with system area in a pool with data direct mapping enabled" section in your storage system's Provisioning Guide.

Attaching volumes to storage when the source is managed in Ops Center Administrator (for external migration)

As a part of external volume migration with the source storage system managed in Ops Center Administrator, the volumes can be attached to the target storage system.

Procedure

1. On the Ops Center Administrator dashboard, click **Storage Systems** to see the inventory of storage systems and capacity information.
2. To select volumes to attach to storage, do one of the following:
 - Select a storage system, then click **Volumes** to see the configured volumes for the storage system.
 - Click **Servers** on the dashboard and select a server.
3. Select one or more volumes, then click **Attach to Storage**.



4. Select the protocol to display the ports for that protocol.

For VSP 5000 series storage systems, you can click a port to display a list of ports with their DKC locations, CTL locations, and redundancy levels.

Port Redundancy Level				
NO.	PORT ID	DKC LOCATI...	CTL LO...	REDUNDANC...
1	CL4-A	DKC-1	CTL11	DKC
2	CL2-A	DKC-1	CTL11	DKC
3	CL6-A	DKC-1	CTL11	DKC
4	CL8-A	DKC-1	CTL11	DKC
5	CL1-J	DKC-2	CTL21	DKC
6	CL3-J	DKC-2	CTL21	DKC
7	CL5-J	DKC-2	CTL21	DKC
8	CL7-J	DKC-2	CTL21	DKC
9	CL2-J	DKC-3	CTL31	DKC
10	CL4-J	DKC-3	CTL31	DKC

Total Ports: 15

OK

5. Click a source storage port and then click a target storage port to create a path.
6. Click **Submit**.

Result

A job is started in Ops Center Administrator to create a path from the volume to the target storage system.

Discovering external storage

Discover external storage on the target storage system so you can virtualize volumes.

Before you begin

Volumes must be attached to the target storage system.

Procedure

1. Open the **Storage System** window for the target storage system.
2. Click **External Volumes**.
3. Click **Discover External Storage** and select the ports that are attached to the source storage system.



Note: The bidirectional port attribute includes all port attributes (such as Target and External). Ports with this attribute are available on some of storage system models.

ID	WWN	SPEED	FIBRE CONNECTION TYPE	SECURITY	VSW PORT	ATTRIBUTE
CL1-A	50:06:0E:80:08:99:87:00	Auto	On, P-to-P	Enabled	Yes	Bidirectional
CL1-B	50:06:0E:80:08:99:87:01	Auto	On, P-to-P	Enabled	No	Target
CL2-A	50:06:0E:80:08:99:87:10	Auto	On, P-to-P	Enabled	No	Target
CL2-B	50:06:0E:80:08:99:87:11	Auto	On, P-to-P	Enabled	No	Target
CL3-A	50:06:0E:80:08:99:87:20	Auto	On, P-to-P	Disabled	Yes	Bidirectional
CL3-B	50:06:0E:80:08:99:87:21	Auto	On, P-to-P	Disabled	No	Bidirectional
CL4-A	50:06:0E:80:08:99:87:30	Auto	On, P-to-P	Disabled	No	Target
CL4-B	50:06:0E:80:08:99:87:31	Auto	On, P-to-P	Disabled	No	Bidirectional
CL5-A	50:06:0E:80:08:99:87:40	Auto	On, P-to-P	Enabled	No	Target
CL5-B	50:06:0E:80:08:99:87:41	Auto	On, P-to-P	Enabled	No	Target
CL6-A	50:06:0E:80:08:99:87:50	Auto	On, P-to-P	Disabled	No	Target
CL6-B	50:06:0E:80:08:99:87:51	Auto	On, P-to-P	Disabled	No	Target
CL7-A	50:06:0E:80:08:99:87:60	Auto	On, P-to-P	Enabled	No	Target
CL7-B	50:06:0E:80:08:99:87:61	Auto	On, P-to-P	Enabled	No	Target
CL8-A	50:06:0E:80:08:99:87:70	Auto	On, P-to-P	Enabled	No	Target

4. Click the iSCSI tab to discover iSCSI ports.

ID	IQN NAME	IPEN	IPFQDN (LOCAL)	IPFQDN (LOCAL)	SECURITY	VSW PORT	ATTRIBUTE
CL1-C	iqn.1994-04.jp.co.hitachi:rsd.r90.08998...	192.168.252.34	==	fe80::	Disabled	Yes	Target
CL1-D	iqn.1994-04.jp.co.hitachi:rsd.r90.08998...	192.168.252.36	==	fe80::	Disabled	No	Target
CL2-C	iqn.1994-04.jp.co.hitachi:rsd.r90.08998...	192.168.252.38	==	fe80::	Disabled	No	Target
CL2-D	iqn.1994-04.jp.co.hitachi:rsd.r90.08998...	192.168.252.40	==	fe80::	Disabled	No	Target
CL3-C	iqn.1994-04.jp.co.hitachi:rsd.r90.08998...	192.168.252.35	==	fe80::	Disabled	No	Target
CL3-D	iqn.1994-04.jp.co.hitachi:rsd.r90.08998...	192.168.252.37	==	fe80::	Disabled	No	Target
CL4-C	iqn.1994-04.jp.co.hitachi:rsd.r90.08998...	192.168.252.39	==	fe80::	Disabled	No	Target
CL4-D	iqn.1994-04.jp.co.hitachi:rsd.r90.08998...	192.168.252.41	==	fe80::	Disabled	No	Target

5. Click **Submit**.

Next steps

Discover and virtualize volumes.

Discovering and virtualizing volumes

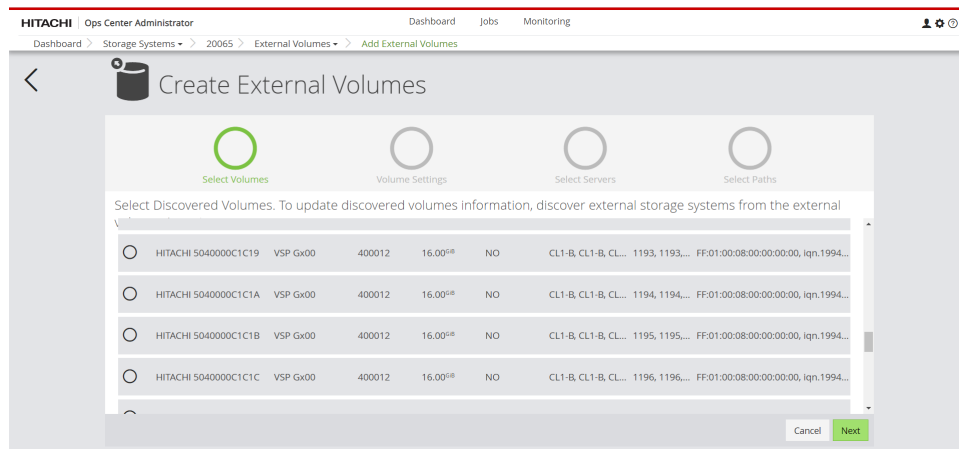
Before you begin

When using external migration, volumes must be virtualized before they can be migrated.

Perform the steps to attach volumes to storage and discover external storage.

Procedure

1. Navigate to **External Volumes** on the target storage system.
2. Click the plus (+) sign to open the **Create External Volumes** window. A list of discovered LUNs appears.
3. Select the discovered LUNs, then click **Next**.

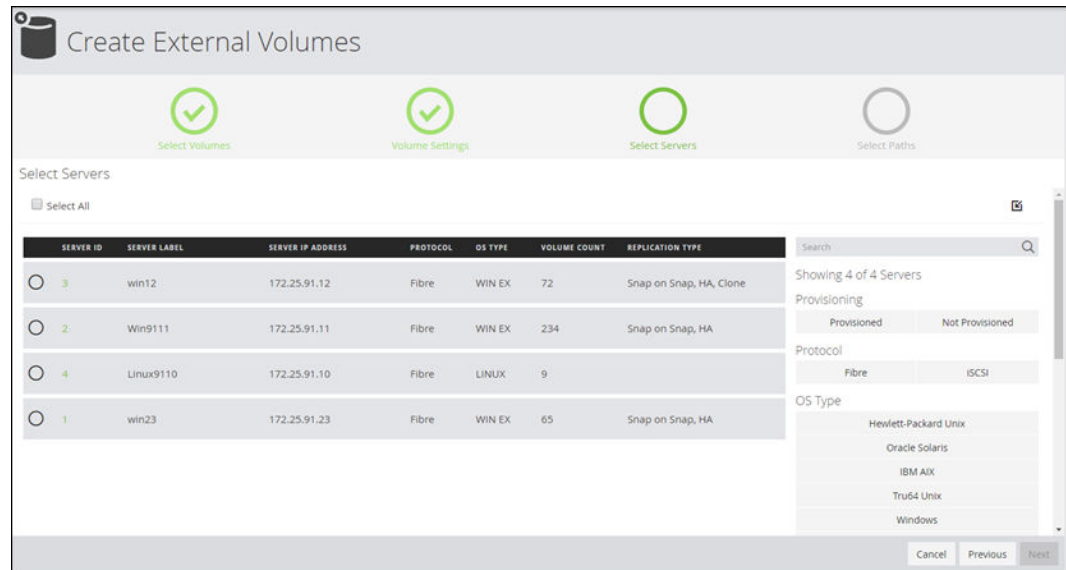


4. In the **Volume Settings** window, you can use **Volume Settings** to select a **Volume ID / Range**. Use the default selection **Auto**, to allow automatic selection of the volume. You can select either a **Decimal** or **Hexadecimal** range and enter the endpoints of the range. Click **Next**.
5. You can confirm which Volume IDs will be used by clicking **View Proposed Volume ID Selection**.
To mandate using volume IDs shown in pop-up dialog, select **Yes** under **Mandate Using Displayed Volume ID**.

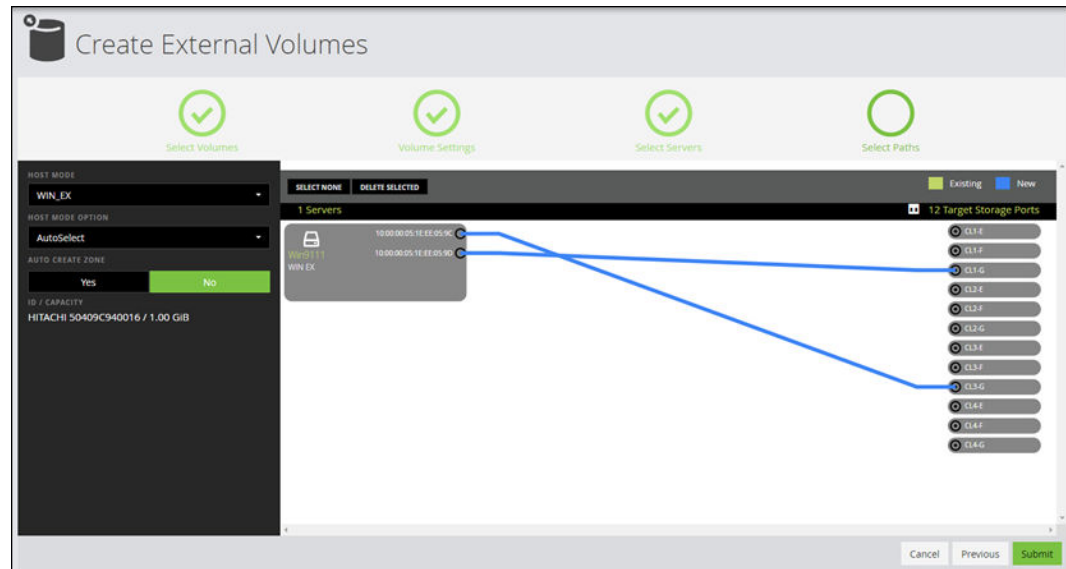


Note: The **Mandate Using Displayed Volume ID** selection will be reset to **No** if you move back to the previous step.

6. Select servers to attach external volumes to be created, then click **Next**. To create external volumes without attaching to servers, click **Next** without selecting servers.



7. You can view servers and their WWNs, along with ports on the storage system.



For VSP 5000 series storage systems, you can click a port to display a list of ports with their DKC locations, CTL locations, and redundancy levels.


NO.	PORT ID	DKC LOCATI...	CTL LO...	REDUNDANC...
1	CL4-A	DKC-1	CTL11	DKC
2	CL2-A	DKC-1	CTL11	DKC
3	CL6-A	DKC-1	CTL11	DKC
4	CL8-A	DKC-1	CTL11	DKC
5	CL1-J	DKC-2	CTL21	DKC
6	CL3-J	DKC-2	CTL21	DKC
7	CL5-J	DKC-2	CTL21	DKC
8	CL7-J	DKC-2	CTL21	DKC
9	CL2-J	DKC-3	CTL31	DKC
10	CL4-J	DKC-3	CTL31	DKC

Total Ports: 15

OK

The following options are available for managing LUN paths:


- If you connect more than one server to the same port, the **Share a host group with all servers** box displays. Select the box to add the servers to a single host group.

-  **Note:** If host groups already exists on that port and single host group cannot be created, the check box will not appear.

To prevent the volume from being added to an existing path, click the path to highlight it, then click **Delete Selected**.

Existing paths are populated as follows: all existing host groups with one or more server WWN and the exact same host mode and host mode options selected in the Attach Settings window are populated as paths.

- Click **Suggest** to populate automatically selected paths. By default, the least-used ports are selected. Suggest paths needs both server and storage ports to be logged into the fabric switches in the Ops Center Administrator inventory.

-  **Note:** For VSP 5000 series storage systems, the first port suggested has the fewest LUN definitions and the second port suggested has the second fewest LUN definitions.

- To manually create a path, click a WWN and a port to connect them with a blue connector line. Click the connector again and then click **Delete Selected** to delete the connection.

8. Click a server and then click a port to create a path between server and port, then click **Submit**.

Next steps

1. After virtualization completes successfully, go to the server and locate and mount the newly attached external volume. This changes the I/O from the older volume(LUN) to the new LUN. This action is required before data can be migrated and will ensure that no data is lost in the process.
2. Select volumes on the **External Volumes** window and migrate them.

Migrating volumes

You can select volumes and migrate them immediately or on a schedule.

You must be a storage administrator to migrate volumes.



Note: Ops Center Administrator limits the number of volumes you can migrate at any one time to 300, because of the high load to the storage system and the long processing time.

Before you begin

For internal migration, ensure the following:

- All selected volumes belong to a single storage system that is managed in Ops Center Administrator.
- All selected volumes will be migrated to a different pool of the same storage system.
- All selected volumes are attached to servers.
- The volumes are not protected. However, a clone pair can be migrated if the P-VOL status is PSUS and the S-VOLs status is SSUS.
- The storage system has at least one eligible pool (Thin or Tiered) with enough space to migrate volumes.

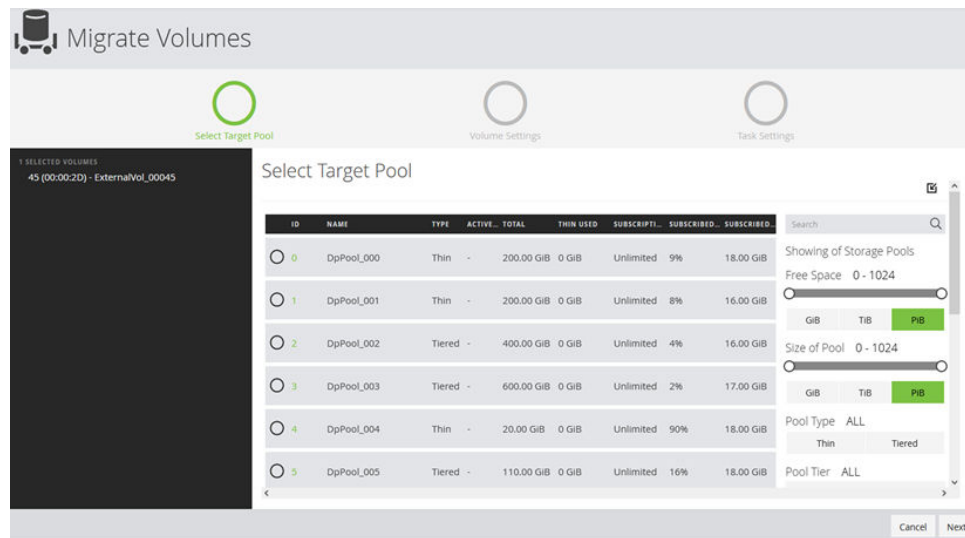
For external migration, ensure the following:

- The target storage system must be onboarded in Ops Center Administrator.
- The source volumes must be virtualized and selectable as external volumes.
- The target storage system must have at least one eligible pool (Thin or Tiered) with enough space to migrate volumes.

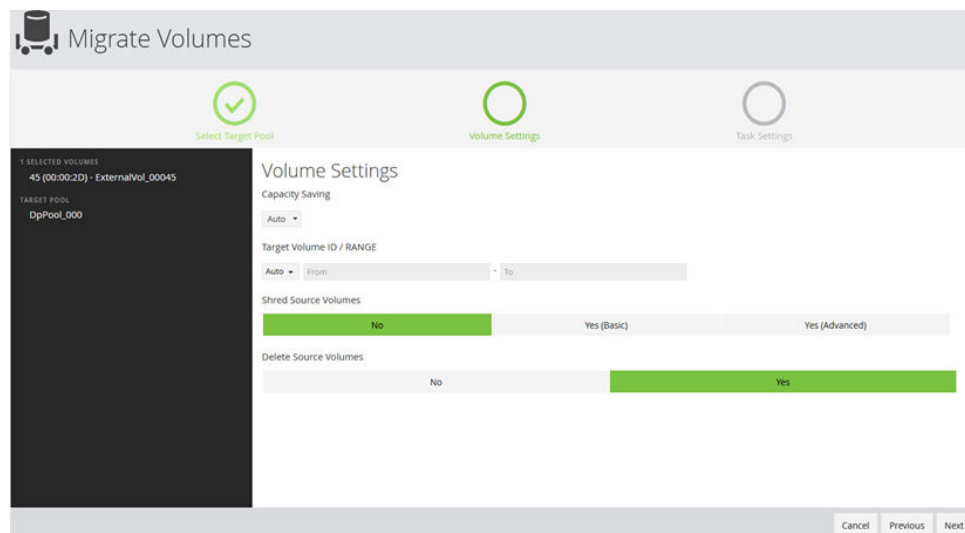
Procedure

1. For internal migration, navigate to one of the following to select volumes:
 - The Volumes window
 - The details window for a volume
 - The details window for a server
 - The details window for a pool

2. For external migration, navigate to the **External Volumes** window.
3. Select volumes to migrate, then click **Migrate volumes**.



4. Select a pool that has enough space to migrate the volume, then click **Next**.



5. In **Volume Settings** you can choose a type of **Capacity Savings** and select a **Volume ID / Range**. You can also set up cleanup tasks by selecting **Shred Source Volumes** and **Delete Source Volumes**.

When you select these options, Ops Center Administrator will shred and/or delete source volumes after a successful migration

The default selection for **Volume ID / Range** is **Auto**, which allows automatic selection. You can select either a **Decimal** or **Hexadecimal** range and enter the endpoints of the range.

The default selection for **Shred Source Volumes** is **No** and the default selection for **Delete Source Volumes** is **Yes**.

Click **Next**.

6. In **Task Settings**, choose to **Migrate Now** or **Migrate Later**:
 - **Migrate Now** needs a **Migration Task Name**.
 - **Migrate Later** needs a **Migration Task Name**, a **Start Time**, and a **Start Date**. You can change the schedule or discard it after submitting the task.
7. Click **Submit** and confirm that you want to migrate the volume.

Managing external volumes

You can access the External Volumes window to add, migrate, or delete external volumes.

Procedure

1. Navigate to the detail window for a storage system, then click **External Volumes**. The External Volumes window shows external volumes and thin/tiered volumes with data direct mapping enabled. The external volumes are volumes on other storage systems mapped to the storage system. The thin/tiered volumes are DP-VOLs created from Dynamic Provisioning pools or Dynamic Tiering pools.



Note: The following external volumes and thin/tiered volumes with data direct mapping enabled are not displayed:

- The emulation type is not OPEN_V
- Remote command devices
- Volumes using Cache Residency Manager cache
- Pool volumes
- NAS boot volumes
- Quorum disks
- Volumes related to the Namespace for FC-NVMe

ID	NAME	POOL ID	TOTAL	PROVISIONING STATUS	STATUS	ASSIGNED TO MIGRATION	EXTERNAL PARITY GROUP ID	MAPPED VOLUME ID
69 (00:00:45)	---	---	1.00 GiB	Unattached	BLOCKED	No	1-319	---
72 (00:00:48)	test_chida_	---	60.00 GiB	Unattached	NORMAL	No	1-51	---
76 (00:00:4C)	galahad-external-R900#10	---	3.00 GiB	Unattached	NORMAL	No	1-52	---
80 (00:00:50)	---	---	100.00 GiB	Unattached	BLOCKED	No	1-6	---
97 (00:00:61)	12345678901234567890...	---	2.00 GiB	Unattached	NORMAL	No	1-5	---
99 (00:00:63)	---	---	88.00 GiB	Unattached	NORMAL	No	1-74	---
102 (00:00:66)	hsa-test1	---	1.00 GiB	Unattached	NORMAL	No	1-58	---
116 (00:00:74)	---	---	10.00 GiB	Unattached	NORMAL	No	1-314	---
124 (00:00:7C)	hsa-test2	---	1.00 GiB	Unattached	NORMAL	No	1-317	---
126 (00:00:7E)	hsa-test3	---	1.00 GiB	Unattached	NORMAL	No	1-318	---
201 (00:00:C6)	hsa-test4	---	1.00 GiB	Unattached	NORMAL	No	1-320	---
202 (00:00:CA)	hsa-test5	---	1.00 GiB	Unattached	NORMAL	No	1-321	---
214 (00:00:D6)	hsa-test6	---	1.00 GiB	Unattached	NORMAL	No	1-322	---
1923 (00:07:83)	hsa-non-ext	---	19.00 GiB	Unmanaged	BLOCKED	No	1-20	---
1925 (00:07:85)	test_chida_	---	100.00 GiB	Unattached	BLOCKED	No	1-57	---
1926 (00:07:86)	hsapool2	---	18.00 GiB	Unmanaged	BLOCKED	No	1-41	---
2173 (00:08:7D)	hsa-test1	---	3.00 GiB	Unmanaged	BLOCKED	No	1-13	---
2174 (00:08:7E)	hsa-test2	---	3.00 GiB	Unmanaged	BLOCKED	No	1-14	---
2175 (00:08:7F)	hsa-test3	---	3.00 GiB	Unmanaged	BLOCKED	No	1-15	---
2298 (00:08:FA)	GSDLG_DP_FREE_OTHER	---	4.00 GiB	Unattached	NORMAL	No	1-11	---
2543 (00:09:F1)	GSDLG_IN_FREE	---	4.00 GiB	Unattached	NORMAL	No	1-7	---

2. On the **External Volumes** window, the following actions are available:

- Click **Open the Logical Devices window in Hitachi Storage Navigator** to perform advanced migration functions in Storage Navigator.
- To create external volumes, click the plus sign (+) to open the **Create External Volumes** window.
- To migrate external volumes, select attached volumes, then click **Migrate Volumes**.



Note: The **Assigned to Migration** filter can display which volumes are already part of a migration plan. If a volume status is **Yes (Unmanaged)** the volume is assigned for migration in other storage management tools.

Managing migration tasks

Access the Migration Tasks window to modify or delete migration tasks, or to interrupt a migration.

Procedure

- Navigate to the detail window for a storage system, then click **Migration Tasks**.

Migration Tasks						
<input type="checkbox"/> Select All Interrupt Migration in Progress						
	ID	NAME	PAIR COUNT	SCHEDULED DATE	STATUS	JOB START
>	110	VolumeMigration-...	16	Apr 21, 2018 8:42:...	N/A	N/A
>	114	VolumeMigration-...	120	Apr 26, 2018 2:42:...	Scheduled	N/A
>	117	VolumeMigration-...	16	Apr 26, 2018 5:42:...	Scheduled	N/A
>	112	VolumeMigration-...	120	Apr 26, 2018 7:42:...	Scheduled	N/A
>	113	VolumeMigration-...	32	Apr 27, 2018 8:42:...	Scheduled	N/A
>	115	VolumeMigration-...	64	Apr 27, 2018 7:42:...	Scheduled	N/A
>	111	VolumeMigration-...	16	Apr 29, 2018 2:42:...	Scheduled	N/A
>	119	VolumeMigration-...	120	N/A	N/A	N/A
>	118	VolumeMigration-...	16	N/A	N/A	N/A

2. On the **Migration Tasks** window you can perform the following actions:

- Select a scheduled task, then click **Edit Migration Task** to open the **Edit Migration Task** window where you can update the pool and the schedule. You can update the capacity saving setting and volume ID range of the target volume as well as update the shred and delete source volumes cleanup options.
- Select an in-progress migration task, click **Interrupt Migration in Progress** and confirm that you want to stop the migration. The migration cannot be restarted after it has been interrupted.
- Select a migration task that is not in progress, click **Delete Migration Task** and confirm that you want to delete it.

For more information about deleting migration tasks, see [Deleting a volume migration task \(on page 170\)](#).

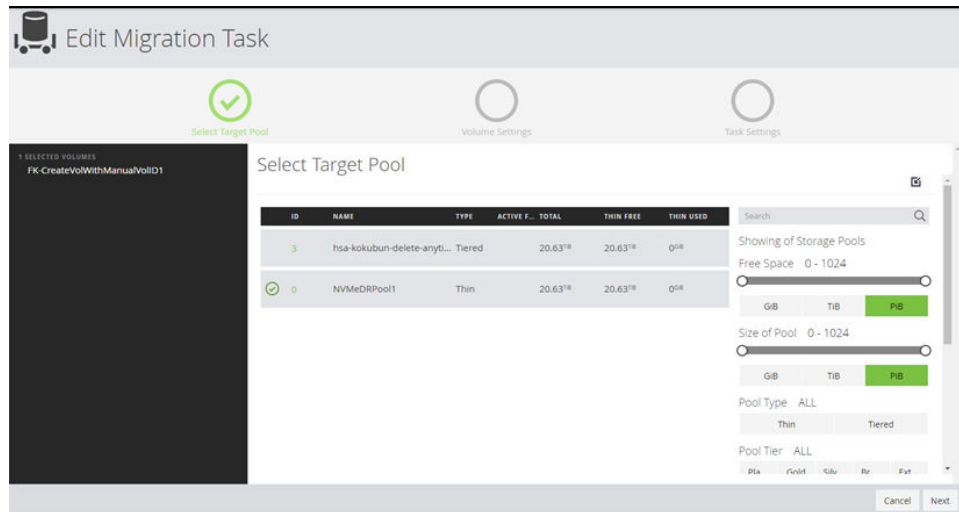
Editing a volume migration task

You can update the pool or the schedule for a planned migration.

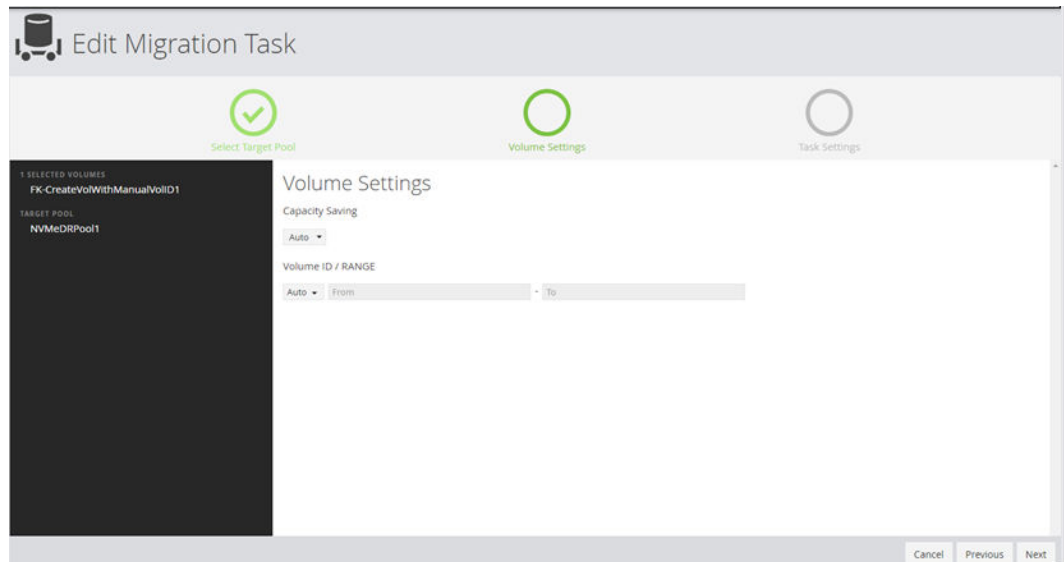
If a volume migration has been scheduled but is not in progress you change the selected pool or schedule.

Procedure

1. On the storage system detail window, click **Migration Tasks**.
2. On the **Migration Tasks** window, select a scheduled task, then click **Edit Migration Task**.



3. You can choose a different pool by clicking the selected pool to unselect and then selecting another pool. Click **Next**.



4. In the **Volume Settings** window, you can update the **Capacity Saving** setting and **Volume ID / Range** of the target volume. You can select either a Decimal or Hexadecimal range and enter the endpoints of the range
You can also set up cleanup tasks by selecting **Shred Source Volumes** and **Delete Source Volumes**. By default, the volumes are deleted but not shredded.
5. In **Task Settings**, you can change the schedule or choose to migrate immediately.

Interrupting an in-progress volume migration

You can interrupt a volume migration that is in progress.

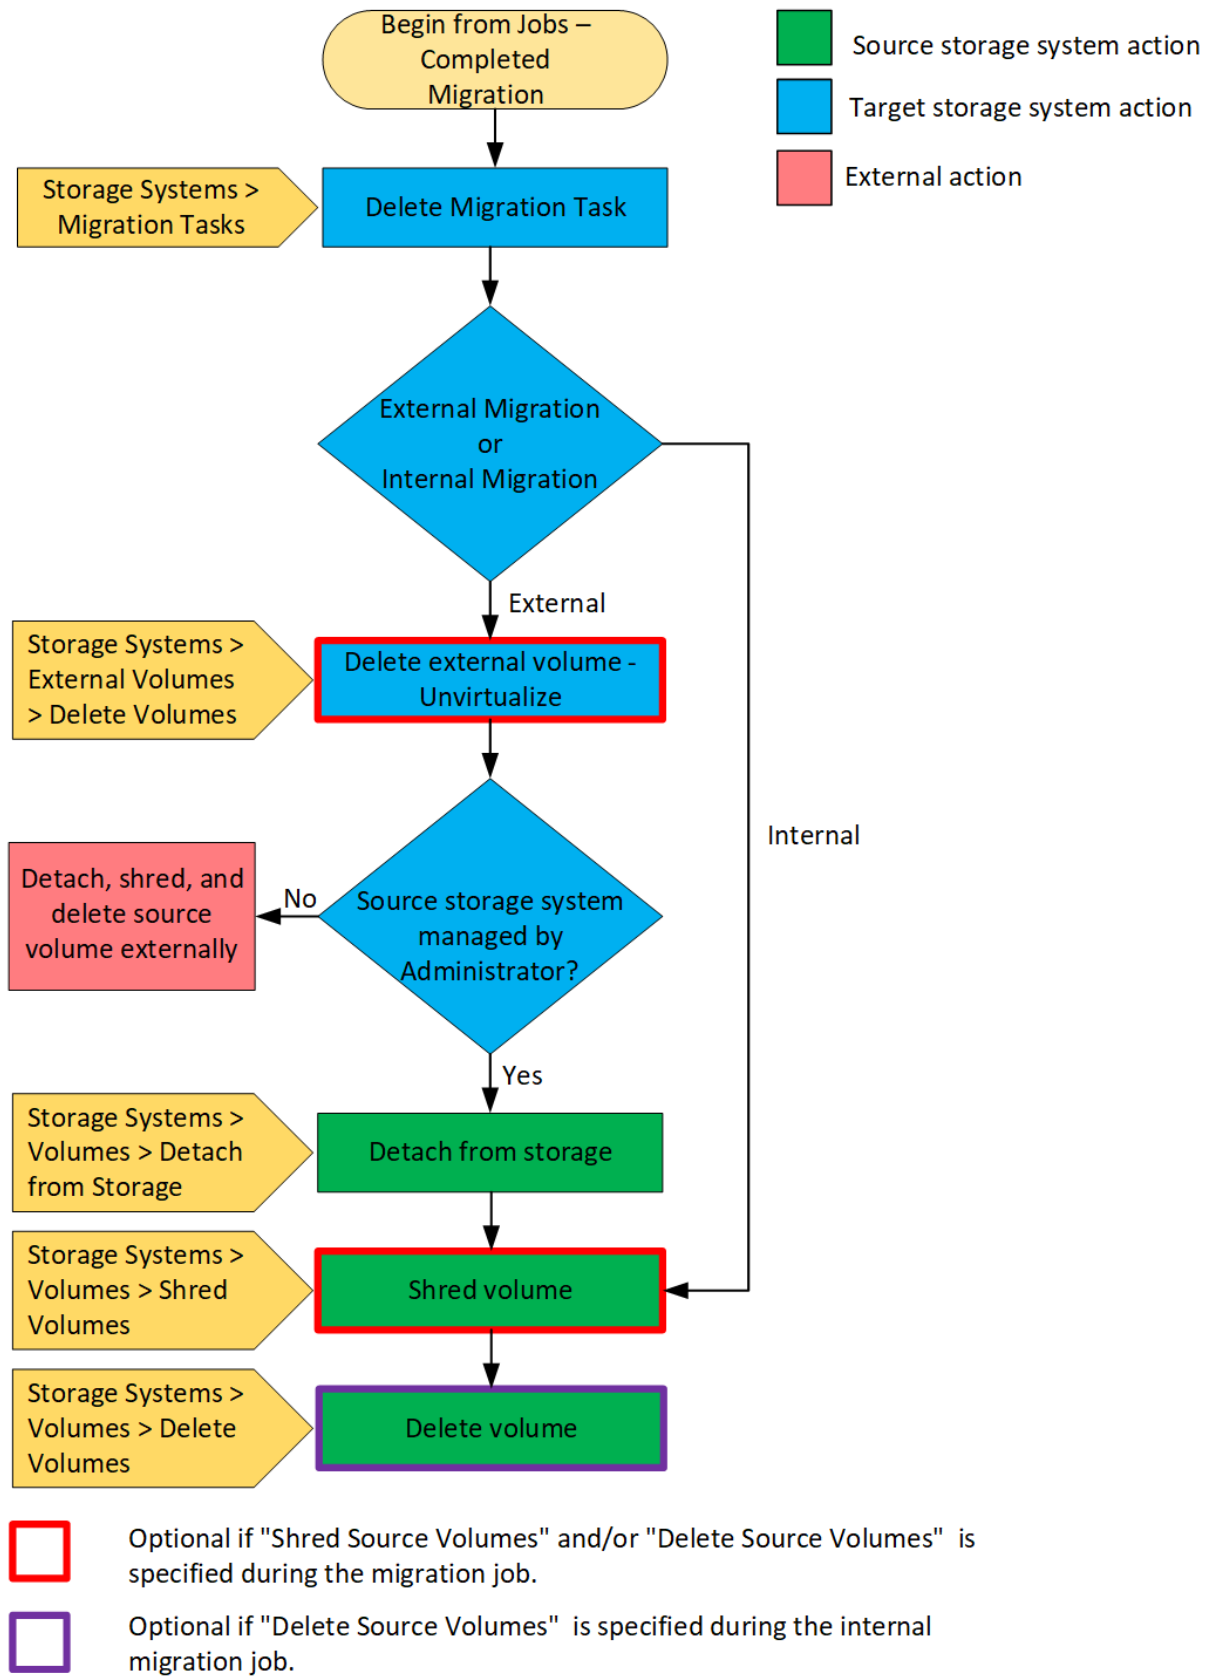
The length of time for volume migration depends on the number of volumes that will be migrated and the protocol. Fibre Channel is faster than iSCSI. If it is taking too long, you can interrupt it. After the migration is interrupted, you will need to discard the task by deleting the migration task. To try again, you will then need to create a new migration task.

Procedure

1. Click **Migration Tasks** on the storage system details window.
2. Select a migration task.
3. To stop a task in progress, click **Interrupt Migration in Progress** and confirm that you want to stop the migration.

Workflow for cleanup after migration

The following diagram shows the workflow for cleanup after migration.

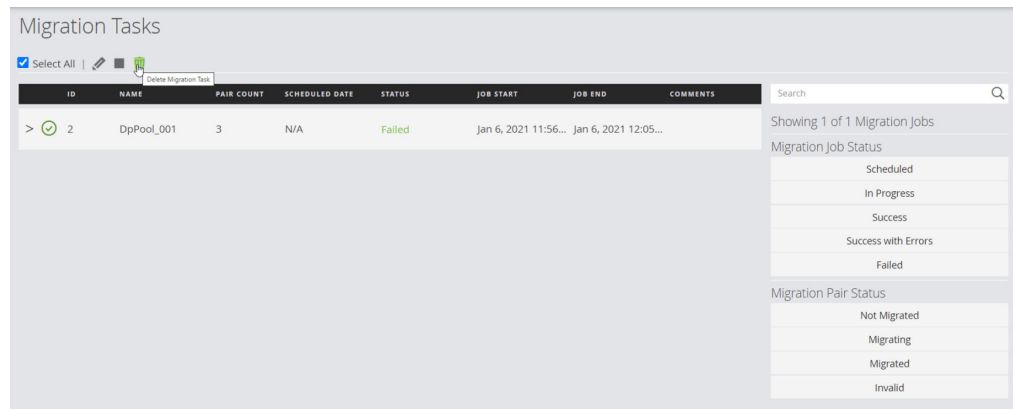


Deleting a volume migration task

You can delete a migration task that is not in-progress.

Procedure

1. Navigate to the detail window for a storage system, then click **Migration Tasks**.
2. On the **Migration Tasks** window, select a migration task, then click **Delete Migration Task**.



3. Confirm that you want to delete the migration task.

Unvirtualizing external volumes

After you have finished migration, you have an external link to an old volume that you no longer need. Unvirtualizing volumes is cleaning up the environment after migration.



Note: You can skip this task if you specify Delete Source Volumes in the Volume Settings page of the Migrate Volumes wizard.

Unvirtualization deletes the external paths and logical devices created during the virtualization flow. It also deletes the external parity groups and the paths from the external parity group to the source storage system. Unvirtualization detaches the external volume, removes the zone, if applicable. It also deletes the external volume, external parity group and paths. Unvirtualization is required to restore the system back to its original state.

Procedure


1. To unvirtualize, select the external volumes in the details window for the storage system.
2. Click the trash can. This also deletes the external parity groups and the paths from the external parity group to the source storage system.

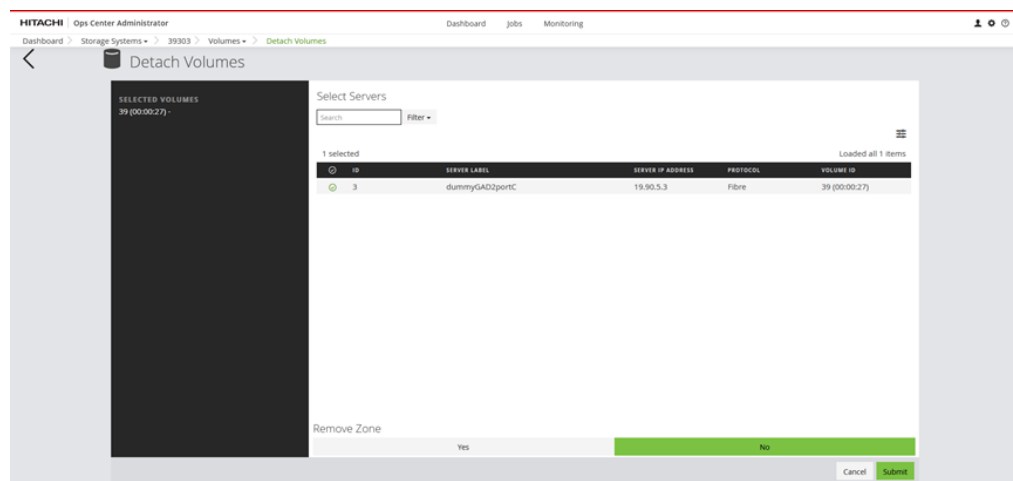
Detaching volumes from storage

Detaching volumes from storage means removing the connection between source and target storage systems.

Detaching volumes from storage deletes the paths from the source storage system to the target storage system. The paths were set when volumes were attached to storage. Detaching volumes from storage is required to restore the system back to its original state.

Procedure

1. Navigate to **Volumes**.
2. Select the attached volume(s).
3. Click the Detach Volumes icon ()



4. Optionally, click **Yes** to remove the zone.

Chapter 8: Monitor block storage

Hitachi Ops Center Administrator supports monitoring of capacity, data protection, jobs, and hardware.

Monitoring capacity

Monitor the capacity of block storage pools.

You can use Ops Center Administrator to monitor block storage capacity so you can tell when physical capacity usage is exceeding thresholds set for your pools.



Note: Capacity alerts for Snap (HTI) pools are not supported.

Alerts are represented by the number in the Capacity tile. If there is a check mark, there are no alerts. A red circle represents an Error and a yellow triangle represents a Warning. Only the highest level of alert is shown, so a red circle can include Errors and Warnings if there is more than one level of alert. The Error threshold is set in the "Utilization Threshold (High)" field when the pool is created and the Warning threshold is set in the "Utilization Threshold (Low)" field.

Capacity Alerts for all storage systems are shown on the dashboard and on the Monitoring tab. Click the Capacity tile on the Monitoring tab to view the pools that have alerts. In the expanded carousel, click Pools to view alert details. The number of Capacity Alerts cannot be more than 1 because Ops Center Administrator monitors at the component level, which is pools.

The alert clearance process runs every 20 minutes for each storage system that has pools with errors. Alerts are only cleared when all pools with errors in the storage system return to normal state. For example, if there are 5 pools with errors in one storage system, then all alerts will display until all the pool alerts are cleared.

Procedure

1. Monitoring alerts display for the following resources:
 - All storage systems: If **Capacity Alerts** on the dashboard has a red circle or a yellow triangle, click it to open the **Monitoring** tab with the number of **Pool** alerts exposed in the **Capacity** carousel. Alternatively, you can click the **Monitoring** tab and then click the **Capacity** tile.
 - Storage systems inventory: This option also displays all capacity alerts. Click **Storage Systems** on the dashboard and if the **Capacity Alerts** tile has a red circle or a yellow triangle, click it to open the carousel in the **Monitoring** tab.

- Storage system detail: In the storage systems inventory, click a storage system tile. If the **Capacity Alerts** tile has a red circle or yellow triangle, click it to open the tile in the **Monitoring** tab and display the **Pool** alerts for the selected storage system.
 - Pool inventory: On the storage system inventory window, click a storage system tile to view resources for the storage system. Click **Pools**. Then click the **Capacity Alerts** tile to open the alerts carousel for the storage system pools in the **Monitoring** tab.
2. In the **Capacity** alerts carousel, click **Pools** to populate the table with details of the alerts.



Note: Ref Code mapping to trap descriptions is available in a file included with the ISO image.

Next steps

If there are pools with alerts displaying the "Error" status, you can expand the pools to increase their capacity.

Monitoring hardware

You can monitor hardware devices to view and investigate warnings and critical alerts.

Ops Center Administrator monitors components in your storage systems and displays alerts on resource and storage system windows. You can also see how many components are in normal status.

Alerts are represented by a number in the Hardware tile or the Hardware Alerts tile. A red circle represents a critical alert and a yellow triangle represents a warning. Only critical alerts are shown unless there are no critical alerts, in which case warning alerts are displayed.

The alerts are received from the storage array: one alert per component type, for a maximum of eight.

The alert clearance process runs every 20 minutes for each storage system that has components with errors. Alerts are only cleared when all components of a given type with errors in the storage system return to normal state. The exception is disks, each of which can have alerts cleared, even if other disks remain in error. Alerts for ports and processors are cleared together, so alerts are cleared only when all ports and all processors are normal.

For example, if there are five fans with alerts in one storage system, they will be cleared only when each and every fan alert is cleared. However, for disks, each alert is cleared as and when disk alert clearance is available.

Ops Center Administrator monitors the following components:

- Disks
- Fans
- Batteries
- Cache
- Processors

- Power supplies
- Ports
- Shared memory

Ops Center Administrator does not report the following types of alerts:

- Alerts caused by the storage system's maintenance operation
- Alerts not categorized as system failure (e.g. recommend replacement)

Procedure

1. Hardware alerts display in the following windows. Alerts are displayed by type of component.
 - Dashboard: Click the **Hardware Alerts** tile or the **Monitoring** tab, then click the **Hardware** tile to display the types of components.
 - Storage system inventory: Includes all hardware alerts.
 - Storage system detail: Click a storage system tile, then click the **Hardware Alerts** tile to view hardware alerts for the selected storage system in the **Monitoring** tab.
 - Parity Group inventory: Click a storage system tile and then click **Parity Groups** to open the inventory. Click the **Hardware Alerts** tile to view alerts for disks where parity groups are located.



Note: Ref Code mapping to trap descriptions is available in a file included with the installation media.

2. Click a component type to display details in the table. Review the details to resolve the issue.

Monitoring data protection

Data protection monitoring enables you to monitor the status of each data protection task.

You can view jobs alerts for your data protections tasks in a single window. Data protection alerts can be viewed by storage system and by server, and are available for the aggregated storage systems.

Alerts for the following operations are available:

- Failed Clone Now operations
- Failed Snapshot operations

A red circle with a number in a Data Protection alerts tile indicates the number of primary volumes and servers on which data protection tasks have failed. You can click the tile to investigate failed tasks.

Procedure

1. View data protection alerts for the following resources:
 - All data protection alerts:
 - From the dashboard, click **Monitoring** or **Data Protection Alerts** to open the **Monitoring** tab.
 - Click **Storage Systems** on the dashboard and then click **Data Protection Alerts** to open the **Monitoring** tab.
 - A single storage system: Click **Storage Systems** and then click a storage system tile. Click the **Data Protection Alerts** tile to open the **Monitoring** tab.
 - Volumes attached to a single server: Click **Servers** from the dashboard, then click a **Server** tile. Then click **Data Protection Alerts** to open the **Monitoring** tab.

The volumes inventory for a server includes failed tasks marked by an exclamation point. Click an exclamation point to view all volumes with failed tasks in the **Monitoring** tab.
2. Click the **Data Protection** tile in the **Monitoring** tab to open the carousel and then click **Volumes** to display details that can help you identify the issue. You can click the exclamation point for any failed task to view details for the volume.
GAD pairs with the status **PSUE** are displayed.
3. Select a volume to perform one of the available functions:
 - Click **Delete** to delete the volume.
 - Click **Edit** to open [Update Volume window \(on page 105\)](#).
 - Click **Unprotect Volumes** to remove data protection.
 - Click **Restore** to open the **Restore Volume** window and restore secondary volumes to a primary volume.

Monitoring jobs

You can view the **Jobs** window to get an update on the status of a task, download the details of a job, perform a job search by keyword, user, or target resource IDs, or view the errors associated with a job.

In this window you see two types of jobs or tasks by switching the displayed tab. In Ops Center Administrator Jobs tab, you can verify submitted jobs and their statuses in Ops Center Administrator. In the System Tasks tab, you can verify tasks submitted in Storage Navigator which runs on the onboarded storage system's SVP.

Procedure



1. If there is a red circle with a number in the **Jobs Alert** tile, it reflects the number of failed jobs and jobs that succeeded with errors. Access the Jobs window in one of these ways:
 - From the dashboard, click **Jobs Alerts** to open the **Jobs** window, or click the **Jobs** tab.
 - On the dashboard, click **Storage Systems** to open the storage systems inventory. If the **Jobs Alert** tile has a red circle with a number, click it to view the Jobs window.
2. Click the **Export Jobs in CSV Format** icon (📄) to download the details of a job and its tasks.
3. Use the **Search** function to view job status and search by keyword, start date, end date, resource, or user.



STATUS	TITLE	USERNAME	START DATE	END DATE	SYSTEM JOB	TARGET RESOURCES
Success	Delete server	sysadmin	3/21/2022, 4:20:51 PM	3/21/2022, 4:20:53 PM	No	Servers: [322]
Success	Delete volume	sysadmin	3/21/2022, 4:20:40 PM	3/21/2022, 4:20:45 PM	No	Storage Systems: [4200...
Success	Delete volume	sysadmin	3/21/2022, 4:20:29 PM	3/21/2022, 4:20:35 PM	No	Storage Systems: [4200...
Success	Detach volume from server	sysadmin	3/21/2022, 4:20:18 PM	3/21/2022, 4:20:22 PM	No	Servers: [322], Storage ...
Success	Detach volume from server	sysadmin	3/21/2022, 4:20:08 PM	3/21/2022, 4:20:10 PM	No	Servers: [322], Storage ...
Success	Snapshot deletion for replication group id 6	sysadmin	3/21/2022, 4:19:57 PM	3/21/2022, 4:20:05 PM	No	Storage Systems: [4200...
Success	Create snapshot for replication group id 6 on Storage System 420053.	sysadmin	3/21/2022, 4:18:01 PM	3/21/2022, 4:18:49 PM	No	
Success	Create volumes, attach to servers and set up data protection	sysadmin	3/21/2022, 4:15:07 PM	3/21/2022, 4:15:40 PM	No	Servers: [322], Storage ...
Success	Onboarding servers	sysadmin	3/21/2022, 4:14:13 PM	3/21/2022, 4:14:14 PM	No	
Success	Deleting pool 8 on Storage System 420053.	sysadmin	3/21/2022, 4:13:29 PM	3/21/2022, 4:14:05 PM	No	Storage Systems: [4200...
Success	Delete server	sysadmin	3/21/2022, 4:13:18 PM	3/21/2022, 4:13:22 PM	No	Servers: [321]
Success	Delete volume	sysadmin	3/21/2022, 4:13:07 PM	3/21/2022, 4:13:11 PM	No	Storage Systems: [4200...
Success	Detach volume from server	sysadmin	3/21/2022, 4:12:55 PM	3/21/2022, 4:13:06 PM	No	Servers: [321], Storage ...
Success	Unprotecting volumes (250) from replication group AutoSnap-0759 on Storage System 420053.	sysadmin	3/21/2022, 4:12:41 PM	3/21/2022, 4:12:50 PM	No	Storage Systems: [4200...
Success	Update replication group 5 on storage system 420053.	sysadmin	3/21/2022, 4:12:30 PM	3/21/2022, 4:12:34 PM	No	Storage Systems: [4200...
Success	Update replication group 5 on storage system 420053.	sysadmin	3/21/2022, 4:12:03 PM	3/21/2022, 4:12:13 PM	No	Storage Systems: [4200...

Common search terms include:

- Volume
- Pool
- Storage
- Create
- Provision
- Expand
- Delete

The **Status** column displays the job status.

Icon	Description
 success	The job completed successfully.
 success with errors	The job completed successfully, but there were errors.

Icon	Description
 in progress	The job is still in progress.
 failed	The job failed.
*: If the job status does not change from "In Progress" for more than 20 minutes, verify that storage system is not locked.	

- Click a job to open the **Job Details** window where you can view any child jobs, reports, target resources, and request information associated with each job. Reports are the messages that the Ops Center Administrator service returns regarding the progress of each job. Child jobs serve to track jobs that start multiple, smaller tasks (such as creating multiple volumes).

Next steps

- If the job failed, review the job reports to determine the cause of the failure.
- If the job completed successfully with errors, click the job to view the error message.

Setting up to receive SNMP alerts

If you are using external management software, you can add an SNMP (Simple Network Management Protocol) manager to enable alerts about devices in registered storage systems. This option is not required to monitor hardware in Hitachi Ops Center Administrator.

When an SNMP manager is added, it collects information from SNMP agents that are installed in components of registered storage systems. The storage systems send traps to Ops Center Administrator using SNMP v3.



Note: Ops Center Administrator manages such alerts automatically, for processors, memory, disks, power supplies, batteries, ports, caches, pools, and fans.

Ops Center Administrator uses SNMP v3 to send traps to external management software as soon as they are received. The structure of the traps is defined in a MIB file and Ops Center Administrator uses its own MIB file to add information to traps received so that the trap is easier to understand. It is important to import the Ops Center Administrator file to access the additional information. The Ops Center Administrator MIB file is included in the product installation media.

The SNMP manager enables you to do the following in the external software:

- Determine whether there are any problems with storage system components or whether the status is normal.
- Identify the specific component that in warning or critical status.

The following types of components can be monitored:

- Processor
- Memory
- Power supply
- Battery
- Fan
- Physical drive or disk
- Port
- Cache
- Pool

Before you begin

SNMP management needs the following:

- External management software capable of receiving SNMP alerts.
- SNMP agents must be configured in the storage environment.
- The Ops Center Administrator MIB file has been imported to the SNMP manager.

Procedure

1. Click **Settings** and select **SNMP Settings** to open the **SNMP** window. Click the plus sign (+) to open the **SNMP Manager Management** window.
2. Enter the following details:
 - **Name:** up to 32 characters.



Note: If you are editing an existing SNMP manager, Name cannot be edited.

- **Username** of the SNMP manager
- **IP Address:** of the SNMP manager.
- **Port:** Positive integers from 0 to 65535. are allowed.
- (Optional) **Privacy Protocol:** You can choose from the following:
 - **NO_PRIV:** No privacy protocol.
 - **DES:** Data Encryption Standard.
 - **TRIPLE_DES:** Triple Data Encryption Algorithm
 - **AES-128:** Advanced Encryption Standard.
- **Privacy Password:** If you choose a privacy protocol, you must enter a password.
- **Auth Protocol:** If you choose a privacy protocol, you must also select an authentication protocol. The choices are **MD5**, **SHA**, or **NO_AUTH**.
- **Auth Password:** If you are using an authentication protocol, you must enter a password.


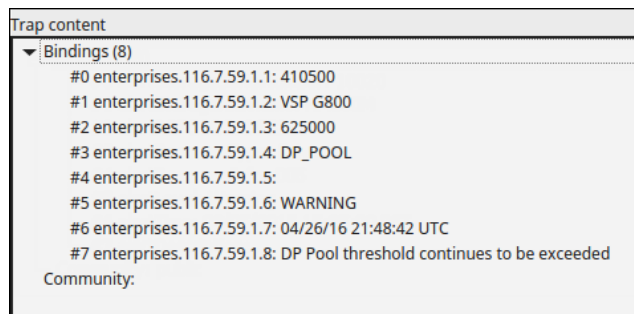
3. Click **Submit**.

Result

A job is created to add or update an SNMP manager.

Example

The following images show the difference between a trap with the MIB file imported and a trap without the MIB file in use.



Chapter 9: Administer Ops Center Administrator

Access Ops Center Administrator administration tasks in Virtual Appliance Manager or through Settings in the top navigation menu.

Administering security

Ops Center Administrator comes with a default local user that has permissions for all functions in Ops Center Administrator.

On a regular basis, you can change the password of the local user by accessing Change Local Password in Settings in the top navigation menu.

To provide more granular control to your users, you can use role-based access to control the functions that can be managed by different users. To do so, you must first add your AD server as the source of authentication. Ops Center Administrator supports adding one external AD server.

You can then assign Ops Center Administrator roles to user groups that exist in the Active Directory account domain. Members of the user groups will use the assigned user roles, and therefore, only the privileges attached to those user roles.



Note: You must use your User Principal Name (UPN) to log on. The UPN is a system user name in the format *username@domainname* (for example, *user@domain.example.com*).

Privileges and roles

Access to the storage system is restricted by roles. Roles determine what a user can and cannot do. The security administrator can grant access to users by assigning users roles with privileges that are attached to those specific roles.

System administrator

- Addition, administration, and deletion of servers, storage system, and fabric switches, SNMP manager, and tier management.
- Addition, administration, and deletion of parity groups.
- Addition, administration, and deletion of virtual file server resource groups.
- Read-only privileges to monitor everything in Ops Center Administrator.
- Registration of Ops Center Protector in Ops Center Administrator.

Storage administrator

- Addition, administration, and deletion of pools.
- Addition, administration, and deletion of volumes, including creating, attaching to hosts, and data protection.
- Addition, administration, and deletion of port configurations.
- Read-only privileges to monitor everything in Ops Center Administrator.

Security administrator

- Addition, administration, and deletion of remote account domains.
- User role assignment to groups.
- Read-only privileges to monitor everything in Ops Center Administrator.

Monitoring role

Read-only privileges to monitor everything in Ops Center Administrator.

Registering an AD domain using a DNS server

- Supported AD platforms: Currently, only Microsoft Active Directory is supported for AD authentication.
- AD domain registration is via DNS or by VAM:

Two methods can be used to register a Microsoft Active Directory account domain with Ops Center Administrator. The first uses the Ops Center Administrator security settings and locates AD servers based on DNS-based SRV records. The second method, using VAM, should be used only in the case that the first solution doesn't work. The one known reason why the first solution might fail is a large number of SRV records.

Before you begin

- Make sure that the DNS server in Ops Center Administrator points only to DNS servers that are able to resolve SRV records for your AD environment. In a Microsoft AD environment, Microsoft DNS servers for a given domain will be able to resolve the SRV records for your AD environment.
- The Security Administrator role is required.
- To enable SSL / TLS, obtain a Root Certification Authority Certificate and save it where you can access it during the procedure.

Procedure

1. Open a browser and enter the IP address of your Ops Center Administrator host machine. Log in to Ops Center Administrator.
2. Click **Settings** and select **Security Settings** to open the **Security** window.
3. Click the pencil icon to open the **Account Credentials** window .
4. Enter a valid user name in UPN format (for example: `FirstName.LastName@example.com`) and the user's password. Click **OK**.

5. In the **Account Domain** field, enter your Active Directory DNS domain name. Click **Submit** to create a job to add an account domain.
To enable SSL or TLS, select **Enable SSL/TLS** and then click **Import Certificate** to import a Root Certification Authority Certificate.
You can view the status on the job on the **Jobs** window.
6. When the job completes successfully, you can add Active Directory groups in **Group Name** in **Security Settings**.

Adding an account domain

Add an account domain so you can assign roles to user groups.



Note: When a fully qualified domain for a current active directory server is added, you can change the credentials on an ongoing basis and save the changes. To update the AD domain name, remove the existing Account Domain and add a new one.

Procedure

1. Click **Settings** and select **Security** to open the **Security** window.
2. In the **Account Domain** field, enter the AD domain name for the active AD server.
3. Enter the Active Directory server **Username** and **Password**. Ops Center Administrator uses these credentials to connect to the external domain and get the list of user groups in the Active Directory source. You can assign Ops Center Administrator roles to these user groups. Ops Center Administrator does not modify the user groups in the Active Directory source.
4. To enable SSL or TLS, select **Enable SSL/TLS** and then click **Import Certificate** to install an Active Directory CA Root Certificate.
5. Click **Submit**.
Changes take effect immediately.

Next steps

When the job is complete, you can assign user roles.

Assigning user roles

After adding an external AD source to Ops Center Administrator, you can assign Ops Center Administrator roles to user groups.

More than one role can be assigned to a user group. In this case the user will have permissions for multiple roles.

After a role is assigned to a user group, you can change the role assignment and save it.



Note: A user who has all of the following roles can run any task in Storage Navigator launched from Ops Center Administrator:

- System Administrator
- Storage Administrator
- Security Administrator

Before you begin

Add an account domain.

Procedure

1. Access the **Security** window by opening the settings menu and then selecting **Security Settings**.
2. In the **Group Name** field, type in a user group name from the account domain. Partial matches automatically populate the field.
3. In the **User Role** field, select an Ops Center Administrator user role:
 - **SecurityAdministrator:** Includes privileges needed for the following:
 - Addition, administration, and deletion of remote account domains.
 - User role assignment to groups.
 - Read-only privileges to monitor everything in Ops Center Administrator.
 - **SystemAdministrator:** Includes privileges needed for the following:
 - Addition, administration, and deletion of servers, storage, and fabric switches, onboarding of storage (block storage with or without NAS modules), SNMP manager, and tier management.
 - Addition, administration, and deletion of parity groups and port configurations.
 - Addition, administration, and deletion of virtual file server resource groups.
 - Read-only privileges to monitor everything in Ops Center Administrator.
 - **StorageAdministrator:** Includes privileges needed for storage provisioning and data protection:
 - Addition, administration, and deletion of pools.
 - Addition, administration, and deletion of volumes, including creating, attaching to servers, and data protection.
 - Read-only privileges to monitor everything in Ops Center Administrator.
 - **MonitoringRole:** Read-only privileges to monitor everything in Ops Center Administrator.
4. Click **Submit**.
The changes take effect immediately. Roles of the accounts related to the updated user group that are currently logged in will be updated from the next login.

Example

As an example, there might be a user group named "IT" in the account domain, whose members need to perform security functions in Ops Center Administrator. Assigning the User Role "SecurityAdministrator" to the Group Name "IT" gives security administration privileges to all members of "IT".

Registering an AD domain using the virtual appliance manager

- Supported AD platforms: Currently, only Microsoft Active Directory is supported for AD authentication.
- AD domain registration is via DNS or by VAM:

Two methods can be used to register a Microsoft Active Directory account domain with Ops Center Administrator. The first uses the Ops Center Administrator security settings and locates AD servers based on DNS-based SRV records. The second method, using VAM, should be used only in the case that the first solution doesn't work. The one known reason why the first solution might fail is a large number of SRV records.

Before you begin

The Security Administrator role is required.

Procedure

1. Open a browser and enter `https://ip-address/vam` in the address bar.
For Ops Center preconfigured media or application installer installations, log in by using the `sysadmin` credentials.
For stand-alone preconfigured media installations, log in by using the `service` credentials.
2. Click the **Domain** tab.
3. Enter a domain name, then click **Add**. A valid domain name is as follows:
 - No more than 255 alphanumeric characters. Periods (.) and hyphens (-) are allowed.
 - No more than 63 characters between dots.
 - Does not begin or end with a hyphen (-).
 - Domain names are not case sensitive.
4. Enter a TCP port.
5. Enter an IP address, then click the plus sign (+) to add it to the list of IP Addresses. Click **Submit**.
6. Log in to Ops Center Administrator and navigate to **Security Settings**. Enter the domain name with the suffix `.hid.local`.

Example:

- Domain to be registered: eng-hsa.com. User present in the domain: user_one.
- User registers eng-hsa.com with the previous steps in virtual appliance manager.
- On the Security Settings window enter the domain name: eng-hsa.com.hid.local , username: user_one, and password: user_one credentials.

You can enter Active Directory groups in **Group Name** and then map them to roles.

Updating an AD domain registered using the virtual appliance manager

Before you begin

You can update or delete a domain that was registered using VAM.

The Security Administrator role is required to delete a registered domain or an account domain.

Procedure

1. Log into the virtual appliance manager, then click **Domain**.
 - Update the port: Select a registered domain from **Registered Domain** list. Update the TCP port and save your change.
 - Update or delete an IP address:
 - Click the trash icon next to an IP address to delete it.
 - Click the pencil icon next to an IP address to update it.
 - Delete a registered domain:

Make sure the domain is not listed in the **Security Settings** in **Ops Center Administrator**. To remove it, you will need to enter your Security Administrator credentials.

In the **Registered Domain** list, select a registered domain, then click **Delete**. Enter your Security Administrator credentials in the **Please enter your credential** box.

2. Save your changes in the **Domain** window.

Adding a login warning banner

You can add a warning banner to display information about system access on the Ops Center Administrator and virtual appliance manager login pages.

Before you begin

The Security Administrator role is required.

Procedure

1. Open a browser and enter the IP address of your Ops Center Administrator host machine and log on to Administrator.

2. Click **Settings > Warning Banner Settings** to open the **Warning Banner Settings** window.

From here you can:

- Choose whether to display the warning banner.
- Customize the default warning message or compose your own.

3. Click **Submit**.

Result

A job is started to update the warning banner settings.

Next steps

- Verify the status of the job in the Jobs window.
- Verify that the setting update is reflected correctly on the Ops Center Administrator and virtual appliance manager login pages.

Backing up and restoring system settings

Use Virtual Appliance Manager to backup system settings before moving Ops Center Administrator to another instance.

The settings include configuration files, an email template file and an SQL file.

Procedure

1. Open a browser and enter `https://ip-address/vam/` to access the Virtual Appliance Manager and adjust the settings for your server environment.

For Ops Center preconfigured media or application installer installations, log in by using the `sysadmin` credentials.

For stand-alone preconfigured media installations, log in by using the `service` credentials.

2. Click **Backup** to download a tar.gz file containing system settings.
3. When you are ready to restore, log in to Virtual Appliance Manager.
4. Click **Restore** and either drag and drop the backup file or click the plus sign (+).
5. Click **Submit**. The restore procedure can take up to an hour.



Note: Do not cancel the restore procedure. Cancellation might corrupt the Ops Center Administrator instance. The progress bar can indicate 100% completion even though the procedure has not completed.

The Ops Center Administrator appliance restarts.

System elements preserved during backup

The following system elements are preserved as part of the backup:

- Element inventories: Storage Systems, Servers, and Fabric Switch inventories
- SNMP managers: Locations for forwarding SNMP traps
- Jobs: All jobs on the system
- Alerts: Monitoring alerts
- Tier names: Tier names for DT pools
- Security information: Local usernames and passwords, and integrated Active Directories
- Replication groups: All copy groups and replication groups and their associated snapshot schedules
- Virtual Appliance Manager settings: Connected NTP servers, log level settings, SSL certificate and service settings are preserved. Host settings are not preserved
- Migration tasks: All migration tasks and their associated migration pairs
- Storage launch settings: Option settings under storage launch settings

Downloading audit logs

Procedure

1. Open a browser and enter `https://ip-address/vam` in the address bar.

For Ops Center preconfigured media or application installer installations, log in by using the `sysadmin` credentials.

For stand-alone preconfigured media installations, log in by using the `service` credentials.

2. Click **Download Audit Logs**.

The system might take a couple of minutes to respond as it compiles the logs. The logs are downloaded as compressed tar archives.

Downloading maintenance information

You can use the CLI to download Ops Center Administrator maintenance information.

Procedure

1. Either open an SSH connection to the VM or open the VMware console and press `Alt + F2` to reach the console.
2. Log in as a root account.
3. In the terminal, enter the command: `/opt/rainier/bin/rainier-getlogs -dir [output-directory]`.

4. Verify the location of collected log file (`rainier-logs.tar.gz`) in the specified output folder, and copy the file to your client computer as required.

Managing the Ops Center Administrator service with Docker

In a Docker environment, you can manage the Ops Center Administrator service containers with the Docker command.

Check the service status

To check the Ops Center Administrator service status:

1. Log on to the host OS as a `root` user.
2. Enter the following command:

```
docker ps | grep rdocker:6000
```

Start the service

To start Ops Center Administrator from a stopped state:

1. Log on to the host OS as a `root` user.
2. Enter the following command:

```
docker ps -a | grep rdocker:6000 | docker start `awk '{print $1}'`
```

Stop the service

To stop Ops Center Administrator from a running state:

1. Log on to the host OS as a `root` user.
2. Enter the following command:

```
docker ps -a | grep rdocker:6000 | docker stop `awk '{print $1}'`
```

Restart the service

To restart Ops Center Administrator from a running state:

1. Log on to the host OS as a `root` user.
2. Enter the following command:

```
docker ps -a | grep rdocker:6000 | docker restart `awk '{print $1}'`
```

Managing the Ops Center Administrator service with Podman

In a Podman environment, you can manage the Ops Center Administrator service with the **systemctl** command. In your host OS, a service unit named `rainier` is registered for Ops Center Administrator. To control the service unit state, use the following commands.



Note: The following Podman features are not supported:

- Do not use the `podman pause` and `podman container checkpoint` commands.
- Do not enable the `podman-restart.service`.
- Do not change the `runroot` podman storage configuration after installation.

Check the service status

To check the Ops Center Administrator service status:

1. Log on to the host OS as a `root` user.
2. Enter the following command:

```
systemctl status rainier
```



Note: Status reported by `systemctl` may not report Ops Center Administrator availability correctly. If Ops Center Administrator is not available despite `systemctl status` reporting that it is in a running state, try to restart the service with the **systemctl** command.

Start the service

To start Ops Center Administrator from a stopped state:

1. Log on to the host OS as a `root` user.
2. Enter the following command:

```
systemctl start rainier
```

Stop the service

To stop Ops Center Administrator from a running state:

1. Log on to the host OS as a `root` user.
2. Enter the following command:

```
systemctl stop rainier
```



Note: The status after stopping Ops Center Administrator may appear as `failed` instead of `inactive` despite the fact that Ops Center Administrator stopped with no errors.

Restart the service

To restart Ops Center Administrator from a running state:

1. Log on to the host OS as a `root` user.
2. Enter the following command:

```
systemctl restart rainier
```

Managing the JDK in Ops Center Administrator environment

Ops Center Administrator uses the bundled JDK. If you want to use a newer version, you can replace it.

Viewing the JDK version used by Ops Center Administrator

To view the current JDK version:

Procedure

1. Open an SSH connection to the VM or the server, or open the VMware console and press **Alt+F2** to reach the console.
2. Log in as a root account.
3. To view the current JDK version, enter the command `/opt/rainier/bin/rainier-replace-jdk --current-java`.

Replacing the JDK used by Ops Center Administrator

To replace the JDK used by Ops Center Administrator:

**Note:**

- You cannot revert the changes. If you want to replace the JDK with the bundled JDK, re-install Ops Center Administrator. The JDK re-installation process includes backup, uninstall, fresh install, and then restore backup.
- When you upgrade Ops Center Administrator, the JDK is replaced with the bundled JDK. If necessary, replace the JDK again.

Before you begin

JDK prerequisites:

- Amazon Corretto 8
- File type must be Linux x64 Compressed Archive



Note: Make sure no one is using Ops Center Administrator and no jobs are running. Jobs running during the JDK replacement might fail.

Procedure

1. Open an SSH connection to the VM or the server, or open the VMware console and press **Alt+F2** to reach the console.
2. Log in as a root account.
3. Enter the command `/opt/rainier/bin/rainier-replace-jdk` with the following parameters to replace the file:
 - `--file` Path of the Linux x64 Compressed Archive JDK File.
 - `-u` User name of the Virtual Appliance Manager.
 - `--backup-dir` Path of an empty directory where the backup file is stored.
4. Enter the Virtual Appliance Manager password.
5. To confirm that the new JDK version was applied, enter `/opt/rainier/bin/rainier-replace-jdk --current-java`.

Example

```
# /opt/rainier/bin/rainier-replace-jdk --file /tmp/jdk/jdk-8u281-linux-x64.tar.gz -u vumuser --backup-dir /tmp/replacejdk_backup
```

Troubleshooting the JDK replacement

If the `rainier-replace-jdk` command fails during replacement of the JDK in the container, the system might not start the Ops Center Administrator service. Restore the setting by following the procedure and then replace the JDK.

1. Copy the `rainier-backup-yyyymmdd-hhmmss.tar.gz` file to the directory specified by the `--backup-dir` option.
2. Uninstall and then re-install Ops Center Administrator server.
3. Use the backup file to restore data. See *Backing up and restoring system settings* in the User Guide.

If you encounter the same problem, collect the following information and contact customer support:

- Support information stored in getlogs of the directory specified in the `--backup-dir` option.
- Latest file in `/var/logs/rainier-tools` with the following format: `replace-jdk_yyyyymmddhhmmss.log`.

Index

A

- accelerated data compression 43
- accessing
 - NAS Manager 74
 - Storage Navigator 73
- account domain
 - adding 182
- Active Directory 180
- AD 184
- AD Domain 185
- add array 22
- add login warning banner 185
- add server group 93
- add storage system 22
- add undefined host groups to a vsm 149
- add undefined volume to a virtual storage machine 148
- add undefined volume to a vsm 148
- adding a server group 93
- adding a volume to a host group 114
- adding resources to a virtual storage machine 146
- adding resources to a vsm 146
- adding storage system 147
- adding storage system to virtual storage machine 147
- adding storage systems to a VSM 147
- adding undefined host groups to a virtual storage machine 149
- adding volumes to host groups 114
- adding volumes to host groups/iSCSI targets 114
- administration 180
- administrator types 74
- advanced option 49
- alerts 172
- allocating volumes to servers 104
- attach like volumes 104
- attach volumes 156
- attach volumes to storage 156
- attaching existing volumes 100
- attaching existing volumes to a server 100
- attaching like volumes 104
- attaching volumes to server 100
- audience 8

B

- backup
 - elements preserved 187
- block storage
 - configuring 25

C

- capacity
 - monitoring 172
- cloning 128, 130
- copy group
 - renaming 137
- create a pool 49
- creating
 - external parity groups 41
 - volumes 98
- creating host groups 110
- creating host groups and iSCSI targets 110
- creating iSCSI targets 110

D

- dashboard
 - analyzing data 17
- data compression 43
- data protection
 - cloning 130
 - monitoring 174
 - snapshots 130
- data protection concepts 128
- deduplication 98
- delete
 - migration task 170
- delete hosts 88
- delete server 88
- deleting external parity group 42
- deleting servers 88
- detach
 - volume 171
- devices

devices (*continued*)

monitoring 173

discovering and virtualizing volumes 159

discovering external storage 40

discovering external storage 157

docker

managing state 188

E

editing host groups 111

editing iSCSI names 113

editing iSCSI targets 111

editing port settings 27

editing wwns 113

enabling file pools 116

enabling host group/iSCSI targets inventory 110

enabling host groups 110

EVS

details 121

inventory 120

EVSs

creating 121

export

server 81

exports

creating 124

details 124

updating 124

external migration 156

external parity group

deleting 42

initializing 41

external parity groups

creating 41

external parity groups inventory 38

external storage

discovering 40

external volume

cleaning up old volume 170

unvirtualizing 170

external volume migration 156

external volumes 164

F

fabric switches

adding 28

supported models 29

system requirements 29

file pool

details 119

file pools

creating 117

expanding 120

inventory 116

managing 116

file server

virtual 120

file storage

configuring 116

file system

inventory 121

file systems

creating 122

deleting 122

expanding 123

managing 121

mounting/unmounting 122

updating 123

viewing details 122

G

getting started 12

H

hardware 173

Hardware Alerts 173

high availability 131

host group 110

host group details 112

host group target details 112

host groups 112

host groups inventory 110

hosts

deleting 88

updating 85

hsa 3.1 128

I

import HA pairs 140

importing HA replication pairs from Ops Center Protector 140

importing high availability pairs 140

initializing external parity group 41

inventory 19, 38

iSCSI 27

iSCSI port settings 27

iSCSI target 110

iSCSI target details 112

J

- job status 175
- jobs
 - monitoring 175

L

- launch settings option
 - Storage Navigator 72
- log
 - download 187
- login
 - add warning banner 185
- LUN paths
 - editing 88

M

- manage
 - storage systems 64
- managing
 - docker state 188
 - podman state 189
- managing external volumes 164
- managing file pools 116
- managing host groups 110
- managing host groups and iSCSI targets 110
- managing iSCSI targets 110
- managing port security 25
- managing servers 75
- managing settings 25
- managing storage systems 12
- managing volumes 164
- migrate 162
- migrate volumes 154
- migrating information 79, 80
- migrating using scan host groups 80
- migrating volumes 162
- migration
 - cleanup 168
 - delete task 170
- monitoring
 - capacity 172
 - data protecting 174
 - hardware 173
 - jobs 175
- move volumes from storage system 145
- moving defined volumes 145
- moving defined volumes to a virtual storage machine 145
- moving volumes to a vsm 145

N

- NAS Manager
 - accessing 74
 - using for advanced file storage configuration 74

O

- ops center administrator 8
- Ops Center Administrator 156
- overview 12

P

- parity groups
 - creating 30
 - initializing 35
 - inventory 35
 - reserved for NAS firmware 35
 - status 35
- parity groups, creating
 - advanced method 33
 - basic method 31
- password
 - local, changing 14
- podman
 - managing state 189
- pool creation
 - basic method, calculation of pool sizes 49
- pools
 - creating 43
 - creating with advanced method 49
 - details 58
 - expanding 52
 - inventory 44
 - updating 52, 58
 - viewing details 52
- pools, creating 45
- ports
 - enabling security 25
- preface 8
- product version 8
- protect volumes 131
- protect volumes with high availability 131
- protecting volumes
 - local replication 129

R

- register 184
- remove host groups from virtual storage machine 152
- remove host groups from vsm 152

- remove storage from vsm 150
- remove storage systems 150
- remove storage systems from a virtual storage machine 150
- removing defined volumes from a virtual storage machine 150
- removing defined volumes from a vsm 150
- removing resources from a vsm 144
- removing storage systems 150
- removing storage systems from a vsm 150
- removing undefined volumes from a virtual storage machine 151
- removing undefined volumes from a vsm 151
- replication 128
- replication groups
 - editing 138
 - inventory 137
 - viewing 136
- resources 19

S

- scan host groups 79, 80
- Scan Host Groups 79
- scanning host groups 79
- security
 - administering 180
- server
 - details 83
 - export 81
- servers
 - adding 75
 - attach existing volumes 81
 - create and attach volumes 81
 - delete 81
 - edit 81
 - inventory 81
 - manage 75
- Settings menu 14
- shares
 - adding AD groups to 125
 - creating 124
 - details 125
 - updating 126
- shares and exports
 - inventory 123
 - managing 123
- shredding volumes 109
- snapshots 128, 130
- SNMP alerts
 - enabling 177

- SNMP manager
 - adding 177
 - updating 177
- spares
 - managing 35
- storage configuration
 - using 71
- storage management 8
- storage management solution 8
- Storage Navigator
 - accessing 73
- Storage Navigator for advanced storage configuration 71
- Storage Navigator launch settings option 72
- storage system
 - adding 22
 - configuration workflow 20
 - deleting 65
 - details 65
 - editing 65, 75
 - inventory 64
 - onboarding
 - workflow 20
 - updating 75
- storage systems
 - managing 64
- system elements 187
- system elements preserved during backup 187
- system log
 - download 187

T

- tiers
 - managing 74

U

- undefined volumes 148
- unified management
 - overview 13
- unvirtualizing
 - volume 170
- Update AD Domain 185
- update description of server group 93
- update name of server group 93
- update storage 75
- updating a host group 93
- updating a server group 93
- updating a storage system 75
- user roles
 - assigning 182

V

VAM 185

virtual

- file server 120

virtual appliance manager

- download system log 187

virtual file server

- deleting 120

- details 121

- enabling/disabling 120

- inventory 120

- renaming 121

virtual file servers

- creating 121

virtual storage machine 142, 147

virtual storage machine inventory 142

vms inventory 142

volume

- detach 171

- erase 109

- shred 109

volume details 128

volume migration 154

volumes

- cloning 130

- create and attach 59

- create and attach to servers 21

- creating 98

- details 107

- expanding 105

- inventory 95

- managing 94

- snapshots 130

- updating 105

vsm 142, 145

W

workflow

- configuring a storage system 20

- onboarding a storage system 20

Hitachi Vantara



Corporate Headquarters

2535 Augustine Drive

Santa Clara, CA 95054 USA

HitachiVantara.com | community.HitachiVantara.com

Contact Information

USA: 1-800-446-0744

Global: 1-858-547-4526

HitachiVantara.com/contact