
High Availability User Guide

Legal notices

© 2019 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 Corporation (collectively "Hitachi"). Licensee may make copies of the Materials provided that any such copy is: (i) created as an essential step in utilization of the Software as licensed and is used in no other manner; or (ii) used for archival purposes. Licensee may not make any other copies of the Materials. "Materials" mean text, data, photographs, graphics, audio, video and documents.

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

Some of the features described in the Materials might not be currently available. Refer to the most recent product announcement for information about feature and product availability, or contact Hitachi Vantara Corporation 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 Corporation.

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

1. Acquiring the relevant consents as may be required under local privacy laws or otherwise from authorized employees and other individuals; 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 is a registered trademark of Hitachi, Ltd., in the United States and other countries.

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

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

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

Contents

Preface.....	7
Document conventions.....	7
Conventions for storage capacity values.....	8
Conventions for storage capacity values for Storage Navigator.....	9
Accessing product documentation.....	10
Getting help.....	10
Comments.....	11
Chapter 1: Overview.....	12
Overview of configuring for high availability.....	12
Overview of global-active device.....	12
About global-active device.....	12
Global-active device solutions.....	13
Fault-tolerant storage infrastructure.....	13
Failover clustering without storage impact.....	14
Server load balancing without storage impact.....	14
System configurations for GAD solutions.....	16
Global-active device and global storage virtualization.....	19
About the virtual ID.....	19
GAD status monitoring.....	20
GAD status.....	20
GAD status transitions.....	22
Pair status.....	23
I/O modes.....	24
Relationship between GAD status, pair status, and I/O mode.....	25
Global-active device and server I/O.....	27
Server I/O (GAD status: Mirrored).....	27
Server I/O (GAD status: Mirroring or Quorum disk blocked).....	27
Server I/O when the GAD status is Suspended.....	29
Server I/O when the GAD status is Blocked.....	30
Quorum disk and server I/O.....	30
I/O stoppage detected in the counterpart system.....	32
I/O stoppage not detected in the counterpart system.....	32
Server I/Os and data mirroring with blocked quorum disk or without quorum disk volumes.....	32

Quorum disk status.....	36
Initial copy and differential copy.....	37
GAD consistency groups.....	38
Operations on GAD pairs by consistency group.....	38
Suspension of GAD pairs by consistency group.....	39
Use cases for consistency groups.....	41
GAD consistency group statuses.....	41
Global-active device components.....	42
User interfaces for global-active device operations.....	44
Storage Advisor Embedded.....	44
Device Manager - Storage Navigator.....	44
Storage Advisor.....	44
Data Instance Director.....	45
Chapter 2: System requirements.....	46
Requirements and restrictions.....	46
Interoperability requirements.....	59
Volume types that can be used for GAD.....	59
Dynamic Provisioning / Dynamic Tiering / Active flash.....	62
ShadowImage.....	63
Limitations when sharing GAD and ShadowImage volumes.....	64
SI operations and GAD pair status.....	64
GAD operations and SI pair status.....	67
Thin Image.....	69
Limitations for using both GAD and Thin Image.....	70
Thin Image operations and GAD status.....	71
GAD operations and Thin Image pair status.....	73
Use cases for pairing GAD volumes with SI or HTI.....	75
Universal Replicator.....	75
GAD pair and Universal Replicator pair interoperability.....	77
Data Retention Utility.....	82
GAD status and I/O allowance by access attribute.....	82
Volume Migration (VSP 5000 series).....	84
Restrictions for using GAD with Volume Migration.....	84
GAD status and Volume Migration pair operations.....	84
Volume Migration pair statuses and GAD pair operations (VSP 5000 series).....	87
LUN Manager.....	89
Virtual Partition Manager.....	89
Nondisruptive migration (VSP F1500 and VSP G1x00).....	89
Restrictions for using GAD with nondisruptive migration.....	90
GAD status and nondisruptive migration pair operations.....	91

Volume Shredder.....	95
Performance Monitor.....	95
Server I/Os added to Performance Monitor.....	96
Port I/Os added to Performance Monitor.....	100
Connection types.....	100
Direct connection.....	101
Connection using switches.....	103
Connection using channel extenders.....	106

Chapter 3: High availability provisioning for Virtual Storage Platform G/F350, G/F370, G/F700, G/F900..... 108

Prerequisites to provisioning for high availability.....	108
Configuration workflow for high availability.....	109
Adding the external storage system for the quorum disk.....	109
Register global-active device license.....	110
Configuring remote paths.....	110
Configuring quorum disks.....	111
Configuring a quorum disk by creating an external volume.....	111
Configuring a quorum disk by selecting an external volume.....	113
Adding block storage to Data Instance Director.....	114
Hitachi Block Storage Node Wizard.....	114
Adding a storage system in Storage Advisor.....	129
Adding servers.....	130
Registering Data Instance Director.....	132
Creating virtual storage machines.....	133
Create, attach, and protect volumes with high availability.....	134
Attach volumes and select secondary servers.....	136
Select volume protection options.....	137
Configuring connections to the primary and secondary sites.....	138
Protect existing volumes with high availability.....	139

Chapter 4: High availability provisioning for VSP G1x00, F1500, VSP Gx00 and VSP Fx00..... 143

Prerequisites to provisioning for high availability.....	143
Configuration workflow for high availability.....	144
Register global-active device license.....	144
Adding the external storage system for the quorum disk.....	145
Connecting the primary and secondary storage systems.....	146
Defining port attributes for the remote connection.....	146
Adding a remote connection.....	146
Creating the quorum disk.....	148
Defining port attributes on VSP G1000,VSP G1500, and VSP F1500 for connecting the primary and secondary storage systems.....	148

Discover and add external volumes.....	148
Adding the quorum disk.....	149
Adding block storage to Data Instance Director.....	149
Hitachi Block Storage Node Wizard.....	150
Adding a storage system in Storage Advisor.....	164
Adding servers.....	165
Registering Data Instance Director.....	167
Creating virtual storage machines.....	168
Create, attach, and protect volumes with high availability.....	169
Attach volumes and select secondary servers.....	171
Select volume protection options.....	172
Configuring connections to the primary and secondary sites.....	173
Protect existing volumes with high availability.....	174
Monitoring and pair operations in Data Instance Director.....	177
Chapter 5: Monitoring and pair operations in Data Instance Director.....	178
Hitachi Block Snapshot or Replication Mount Wizard.....	180
Mount Wizard - Select Microsoft SQL Restore Options.....	188
Mount Wizard - Select Oracle Restore Options.....	190
Hitachi Block Replication Swap Wizard.....	195
Access Control Transfer Permissions Dialog.....	196
Chapter 6: Managing remote paths and quorum disks.....	198
Managing remote paths.....	198
Adding remote paths	198
Removing remote paths.....	199
Removing remote path groups.....	200
Managing quorum disks	200
Adding external paths to a quorum disk.....	200
Changing the name of the external volume of a quorum disk.....	201
Removing external paths to a quorum disk.....	201
Disabling the quorum disk settings and deleting the corresponding external volume.....	202
Disabling the quorum disk settings.....	202

Preface

This document describes and provides instructions for using the provisioning software to configure and perform operations on Hitachi Virtual Storage Platform G1000 systems.

Please read this document carefully to understand how to use these products, and maintain a copy for your reference.





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: <pre>pairdisplay -g group</pre> (For exceptions to this convention for variables, see the entry for angle brackets.)
Monospace	Indicates text that is displayed on screen or entered by the user. Example: <code>pairdisplay -g oradb</code>
< > angle brackets	Indicates variables in the following scenarios: <ul style="list-style-type: none">Variables are not clearly separated from the surrounding text or from other variables. Example: <pre>Status-<report-name><file-version>.csv</pre>Variables in headings.
[] square brackets	Indicates optional values. Example: [a b] indicates that you can choose a, b, or nothing.

Convention	Description
{ } braces	Indicates required or expected values. Example: { a b } indicates that you must choose either a or b.
vertical bar	Indicates that you have a choice between two or more options or arguments. Examples: [a b] indicates that you can choose a, b, or nothing. { a b } indicates that you must choose either a or b.

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

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

Conventions for storage capacity values

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

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

Logical 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 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

Conventions for storage capacity values for Storage Navigator

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 KB Open-systems: ▪ OPEN-V: 960 KB ▪ Others: 720 KB
1 KB	1,024 (2^{10}) bytes
1 MB	1,024 KB or $1,024^2$ bytes
1 GB	1,024 MB or $1,024^3$ bytes
1 TB	1,024 GB or $1,024^4$ bytes
1 PB	1,024 TB or $1,024^5$ bytes
1 EB	1,024 PB or $1,024^6$ bytes

Accessing product documentation

Product user documentation is available on Hitachi Vantara Support Connect: <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

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

[Hitachi Vantara Community](#) 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 Corporation.

Thank you!

Chapter 1: Overview

Overview of configuring for high availability

You can configure your Hitachi storage for high availability using the web-based user interfaces of Storage Advisor Data Instance Director, and Storage Advisor Embedded or Storage Navigator.

High availability is enabled by Hitachi global-active device technology.

High availability for VSP F1500, VSP G1x00, VSP Fx00 models, and VSP Gx00 models can be configured using Storage Navigator, Data Instance Director, and Storage Advisor.

High availability for VSP F350, F370, F700, F900 and VSP G350, G370, G700, G900 models can be configured using Storage Advisor Embedded, Data Instance Director, and Storage Advisor.

Overview of global-active device

An overview of the global-active device feature helps you to understand its components and capabilities.

About global-active device

Global-active device (GAD) enables you to create and maintain synchronous, remote copies of data volumes.

A virtual storage machine is configured in the primary and secondary storage systems using the actual information of the primary storage system, and the global-active device primary and secondary volumes are assigned the same virtual LDEV number in the virtual storage machine. This enables the host to see the pair volumes as a single volume on a single storage system, and both volumes receive the same data from the host.


A quorum disk, which can be located in a third and external storage system, is used to monitor the GAD pair volumes. The quorum disk acts as a heartbeat for the GAD pair, with both storage systems accessing the quorum disk to check on each other. A communication failure between systems results in a series of checks with the quorum disk to identify the problem for the system able to receive host updates.

Alternate path software on the host runs in the Active/Active configuration. While this configuration works well at campus distances, at metro distances a multi-pathing software with ALUA (asymmetric logical unit access) or Hitachi Dynamic Link Manager is required to support preferred/nonpreferred paths and ensure that the shortest path is used.

If the host cannot access the primary volume (P-VOL) or secondary volume (S-VOL), host I/O is redirected by the alternate path software to the appropriate volume without any impact to the host applications.

Global-active device provides the following benefits:


- Continuous server I/O when a failure prevents access to a data volume
- Server failover and failback without storage impact
- Migrating virtual machines without storage impact



Global-active device solutions


Fault-tolerant storage infrastructure

If a failure prevents host access to a volume in a GAD pair, read and write I/O can continue to the pair volume in the other storage system to provide continuous server I/O to the data volume.



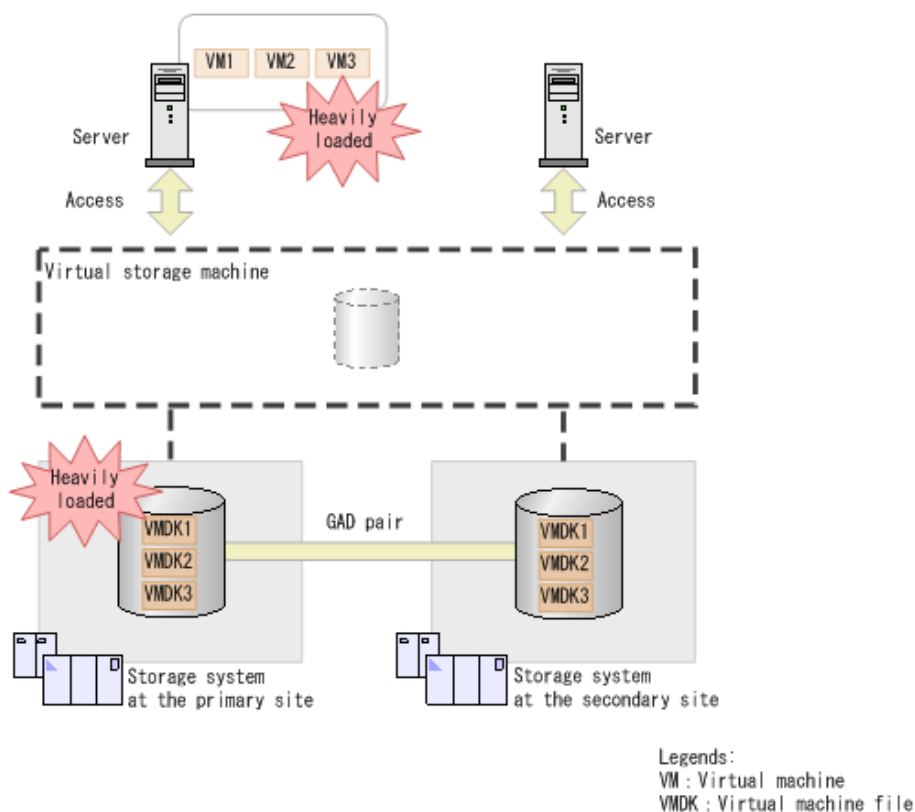
Failover clustering without storage impact

In a server-cluster configuration with global-active device, the cluster software is used to perform server failover and failback operations, and the global-active device pairs do not need to be suspended or resynchronized.

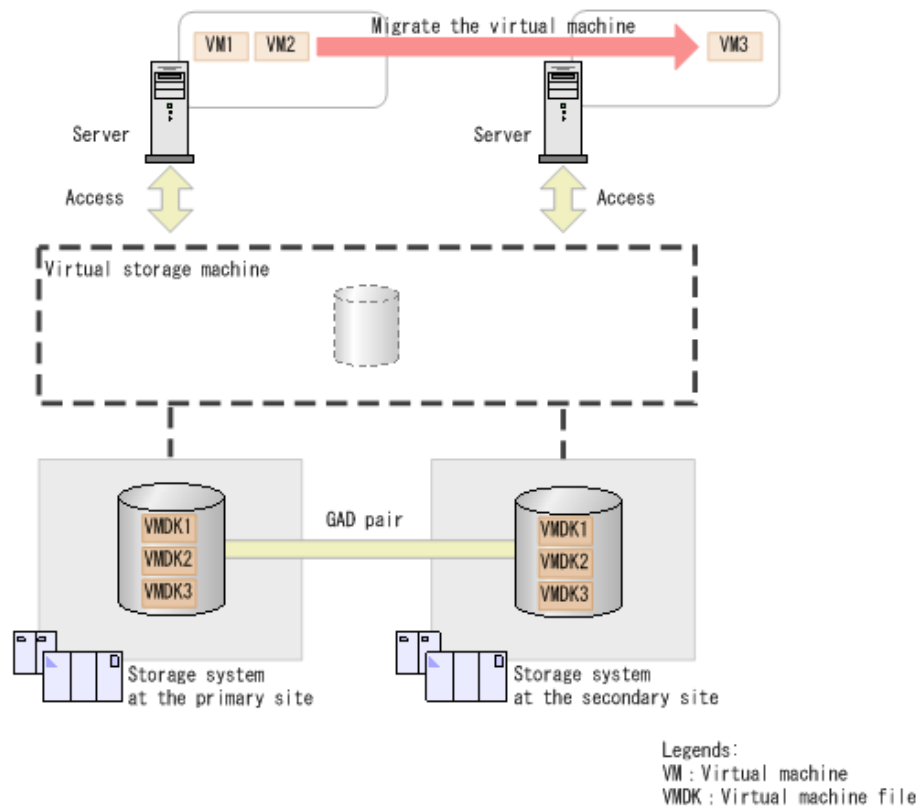


Server load balancing without storage impact

When the load on a virtual machine increases global-active device enables you to migrate the virtual machine to the paired server without performing any operations on the storage systems.



As shown in this example, the server virtualization function is used to migrate virtual machine VM3 from the primary-site server to the secondary-site server. Because the GAD primary and secondary volumes contain the same data, you do not need to migrate any data between the storage systems.



System configurations for GAD solutions

You have the option of implementing three different system configurations: a single-server configuration, a server-cluster configuration, and a cross-path configuration. The system configuration depends on the GAD solution that you are implementing.

The following table lists the GAD solutions and specifies the system configuration for each solution.



Caution: When you register GAD pairs to a consistency group, you should use the cross-path configuration. If GAD pairs in the Mirrored status are suspended due to a path failure between the primary site and the secondary site in the following condition, some GAD pairs might be able to be accessed only from the server at the primary site, and other GAD pairs might be able to be accessed only from the server at the secondary site.


- GAD pairs both in the Mirrored status and in the Mirroring status are in the consistency group.
- GAD pairs both in the Mirrored status and in the Suspended status are in the consistency group.

When you use the cross-path configuration that enables both servers at the primary and secondary sites to access both volumes at the primary and secondary sites, the servers can continue to access the GAD volumes even in this situation. If you use a configuration other than the cross-path configuration, the servers cannot access the GAD volumes.

GAD solution	Software		System configuration
	Alternate path software	Cluster software	
Continuous server I/O (if a failure occurs in a storage system)	Required	Not required	Single-server configuration
Failover and failback on the servers without using the storage systems	Not required	Required	Server-cluster configuration
Migration of a virtual machine of a server without using the storage systems	Not required	Required	Server-cluster configuration
Both of the following: <ul style="list-style-type: none"> Continuous server I/O (if a failure occurs in a storage system) Failover and failback on the servers without using the storage systems 	Required	Required	Cross-path configuration


Single-server configuration

In a single-server configuration, the primary and secondary storage systems connect to the host server at the primary site. If a failure occurs in one storage system, you can use alternate path software to switch server I/O to the other site.




Server-cluster configuration

In a server-cluster configuration, servers are located at both the primary and secondary sites. The primary storage system connects to the primary-site server, and the secondary storage system connects to the secondary-site server. The cluster software is used for failover and failback. When I/O on the virtual machine of one server increases, you can migrate the virtual machine to the paired server to balance the load.



Cross-path configuration

In a cross-path configuration, primary-site and secondary-site servers are connected to both the primary and secondary storage systems. If a failure occurs in one storage system, alternate path software is used to switch server I/O to the paired site. The cluster software is used for failover and failback.



Global-active device and global storage virtualization

GAD operations are based on the global storage virtualization function. When virtual information is sent to the server in response to the SCSI Inquiry command, the server views multiple storage systems as multiple paths to a single storage system.

The global storage virtualization function is enabled when you install the license for Resource Partition Manager, which is provided with the Storage Virtualization Operating System (SVOS). For more information about Resource Partition Manager, see the *Provisioning Guide* for the storage system.

About the virtual ID

The server is able to identify multiple storage systems as a single virtual storage machine when the resources listed below are virtualized and the virtual identification (virtual ID) information is set. You can set virtual IDs on resource groups and on individual volumes, as described in the following table.

Virtual information required by the server	Resource on which virtual IDs are set
Serial number	Resource group
Product	Resource group

Virtual information required by the server	Resource on which virtual IDs are set
LDEV ID*	Volume
Emulation type	Volume
Number of concatenated LUs of LUN Expansion (LUSE)	Volume
SSID	Volume
* A volume whose virtual LDEV ID has been deleted cannot accept I/O from a server. The virtual LDEV ID is temporarily deleted on a volume to be used as a GAD S-VOL because, when the pair is created, the P-VOL's physical LDEV ID is set as the S-VOL's virtual LDEV ID.	

When using global storage virtualization you can set the following:

- The same serial number or product as the virtual ID for more than one resource group
- Up to eight types of virtual IDs for resource groups in a single storage system
- Virtual IDs for a maximum of 1,023 resource groups (excluding resource group #0)
- Virtual IDs for a maximum of 65,279 volumes

For instructions on setting virtual IDs, see the *Command Control Interface Command Reference*.

GAD status monitoring

GAD operations are managed based on the following information: Pair status, I/O mode of the P-VOL and S-VOL, and GAD status, which is a combination of pair status and I/O mode

GAD status

It is important to be able to understand what the meaning of a GAD status is and what that status tells you about the GAD pair.

The following table lists and describes the GAD statuses.

GAD status	Description	Data redundancy	Updated volume	Volume with latest data
Simplex	The volume is not a pair volume.	No	Not applicable	Not applicable
Mirroring	The pair is changing to Mirrored status.	No	P-VOL and S-VOL	P-VOL

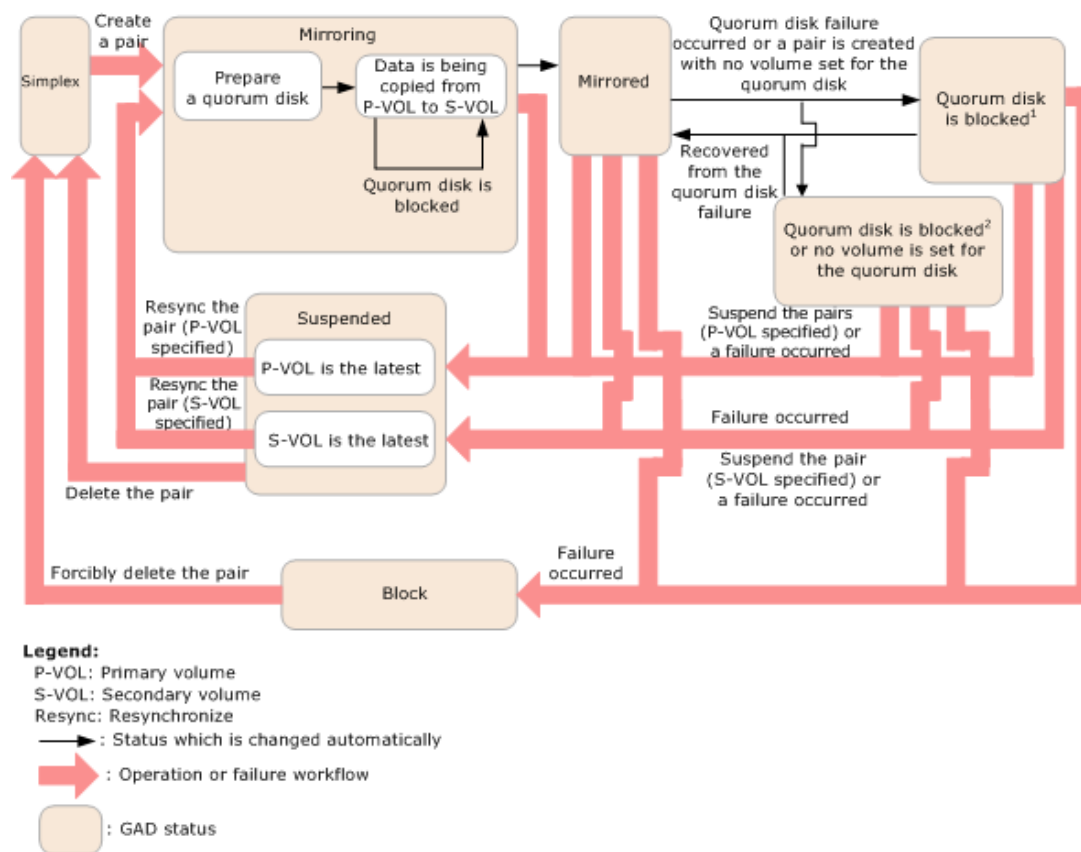
GAD status	Description	Data redundancy	Updated volume	Volume with latest data
	<p>This status is issued when you do the following:</p> <ul style="list-style-type: none"> Prepare a quorum disk. Copy data from the P-VOL to the S-VOL. 			
Mirrored	The pair is operating normally.	Yes	P-VOL and S-VOL	P-VOL and S-VOL
Quorum disk blocked	Quorum disk is blocked, but the data is mirrored. For a pair created, resynchronized, or swap resynchronized on microcode version 80-04-2x or earlier for VSP G1x00 and VSP F1500, I/O from the server to the S-VOL cannot be accepted.	Yes	P-VOL and SVOL	P-VOL and SVOL
Suspended	<p>The pair is suspended. I/O from the server is sent to the volume with the latest data.</p> <p>When a failure occurs or the pair is suspended, the status changes to Suspended.</p> <p>For a pair created, resynchronized, or swap resynchronized, the status changes to Suspended after the time specified for Read Response Guaranteed Time When Quorum Monitoring Stopped elapses.</p>	No	P-VOL or S-VOL	P-VOL or S-VOL
Blocked	<p>I/O is not accepted by either pair volume. This status occurs when:</p> <ul style="list-style-type: none"> Both the P-VOL and S-VOL have the latest data. If the pair is forcibly deleted, I/O can be restarted in either of the volumes. A failure occurs in the primary or secondary storage system, and I/O to the volume in the paired system is also stopped. 	No	None	P-VOL and S-VOL

GAD status	Description	Data redundancy	Updated volume	Volume with latest data
	If more than one failure occurs at the same time, the GAD status changes to Blocked.			

GAD status transitions

The GAD status changes depending on the pair operation and failure.

The following illustration shows the GAD pair status transitions.



If you resynchronize a pair specifying the P-VOL, I/O continues on the P-VOL. If you resynchronize a pair specifying the S-VOL, data flow switches from the S-VOL to the P-VOL, and then I/O continues on the new P-VOL.

If you suspend a pair specifying the P-VOL, I/O continues to the P-VOL. If you suspend a pair specifying the S-VOL, I/O continues to the S-VOL.

Pair status

The pair status provides information about the current state of a global-active device pair.

The following table lists and describes the pair statuses, which indicate the current state of a global-active device pair. As shown in the following table, the pair status terms displayed by the user interfaces are slightly different.

Pair status		Description
CCI	HDvM - SN	
SMPL	SMPL	The volume is not paired.
COPY	INIT/COPY	The initial copy or pair resynchronization is in progress (including creation of a GAD pair that does not perform data copy). A quorum disk is being prepared.
	COPY	The initial copy is in progress; data is being copied from the P-VOL to the S-VOL (including creation of a GAD pair that does not perform data copy).
PAIR	PAIR	The pair is synchronized.
PSUS	PSUS*	The pair was suspended by the user. This status appears on the P-VOL.
PSUE	PSUE*	The pair was suspended due to a failure.
SSUS	SSUS*	The pair was suspended by the user, and update of the S-VOL is interrupted. This status appears on the S-VOL.
SSWS	SSWS*	The pair was suspended either by the user or due to a failure, and update of the P-VOL is interrupted. This status appears on the S-VOL.
* When a GAD pair is suspended, you can view the suspend type on the View Pair Properties window.		

GAD suspend types

When a GAD pair is suspended, the suspend type is displayed in the Status field of the View Pair Properties window. The suspend type is not displayed by CCI.

The following table lists and describes the GAD suspend types.

Suspend type	Volume	Description
Primary Volume by Operator	P-VOL	The user suspended the pair from the primary storage system. The S-VOL suspend type is "by MCU".

Suspend type	Volume	Description
Secondary Volume by Operator	P-VOL S-VOL	The user suspended the pair from the secondary storage system.
by MCU	S-VOL	The secondary storage system received a request from the primary storage system to suspend the pair. The P-VOL suspend type is Primary Volume by Operator or Secondary Volume by Operator.
by RCU	P-VOL	The primary storage system detected an error condition at the secondary storage system, which caused the primary storage system to suspend the pair. The S-VOL suspend type is Secondary Volume Failure.
Secondary Volume Failure	P-VOL S-VOL	The primary storage system detected an error during communication with the secondary storage system, or an I/O error during update copy. In this case, the S-VOL suspend type is usually Secondary Volume Failure. This suspend type is also used when the number of paths falls below the minimum number of paths setting on the Add Remote Connection window.
MCU IMPL	P-VOL S-VOL	The primary storage system could not find valid control information in its nonvolatile memory during IMPL. This condition occurs only if the primary storage system is without power for more than 48 hours (that is, power failure and fully discharged backup batteries).
Initial Copy Failed	P-VOL S-VOL	The pair was suspended before the initial copy operation was complete. The data on the S-VOL is not identical to the data on the P-VOL.

I/O modes

You should understand the I/O actions on the P-VOL and the S-VOL of a GAD pair.

The following table lists and describes the GAD I/O modes. As shown in the following table, the I/O mode terms displayed by the user interfaces are slightly different.

I/O mode			Read processing	Write processing
I/O mode	CCI ¹	HDvM - SN		
Mirror (RL)	L/M	Mirror (Read Local)	Sends data from the storage system that received a read request to the server.	Writes data to the P-VOL and then the S-VOL.
Local	L/L	Local	Sends data from the storage system that received a read request to the server.	Writes data to the volume on the storage system that received a write request.
Block ²	B/B	Block	Rejected (Replies to illegal requests).	Rejected (Replies to illegal requests).
Notes: <ol style="list-style-type: none"> 1. In CCI, the I/O mode is displayed as <read processing>/<write processing> in which L indicates Local, M indicates Mirror, and B indicates Block (for example, L/L indicates Local read processing and Local write processing). 2. For volumes whose I/O mode is Block, a response indicating that the LU is undefined is returned to the Report LUN and Inquiry commands. Therefore, servers cannot identify a volume whose I/O mode is Block, or the path of this volume is blocked. 				

Relationship between GAD status, pair status, and I/O mode

You should understand the relationship between the GAD status, pair status, and I/O mode to be informed about your GAD pairs.

The following table lists the GAD statuses and describes the relationship between the GAD status, pair status, and I/O mode. "N" indicates that pair status or I/O mode cannot be identified due to a failure in the storage system.

GAD status	When to suspend	P-VOL		S-VOL		Volume that has the latest data
		Pair status	I/O mode	Pair status	I/O mode	
Simplex	Not applicable	SMPL	Not applicable	SMPL	Not applicable	Not applicable
Mirroring	Not applicable	INIT	Mirror(RL)	INIT	Block	P-VOL

GAD status	When to suspend	P-VOL		S-VOL		Volume that has the latest data
		Pair status	I/O mode	Pair status	I/O mode	
	Not applicable	COPY	Mirror(RL)	COPY	Block	P-VOL
Mirrored	Not applicable	PAIR	Mirror(RL)	PAIR	Mirror(RL)	P-VOL and S-VOL
Quorum disk blocked	Not applicable	PAIR	Mirror(RL)	PAIR	Block or Mirror(RL) 1	P-VOL and S-VOL
Suspended	Pair operation	PSUS	Local	SSUS	Block	P-VOL
		PSUE ²	Local	PSUE	Block	P-VOL
		PSUE ²	Local	SMPL	Not applicable	P-VOL
		PSUE ²	Local	N	N	P-VOL
	Pair operation	PSUS	Block	SSWS	Local	S-VOL
		PSUE	Block	SSWS ²	Local	S-VOL
		SMPL	Not applicable	SSWS ²	Local	S-VOL
		N	N	SSWS ²	Local	S-VOL
Blocked	Not applicable	PSUE	Block	PSUE	Block	P-VOL and S-VOL
	Not applicable	PSUE	Block	N	N	P-VOL and S-VOL
	Not applicable	N	N	PSUE	Block	P-VOL and S-VOL
Notes: <ol style="list-style-type: none"> 1. For microcode version 80-04-2x or earlier for VSP 5000 series and firmware version 83-03-3x or earlier for VSP Gx00 models, the status is Block. For microcode version 80-05-0x or later for VSP 5000 series, and for VSP G350, G370, G700, G900 and VSP F350, F370, F700, F900, the status is Mirror(RL). 2. If the server does not issue the write I/O, the pair status might be PAIR, depending on the failure location. 						


Global-active device and server I/O

I/O requests from the server to a GAD pair volume are managed according to the volume's I/O mode. The GAD status determines the I/O mode of the P-VOL and S-VOL of a pair.


Server I/O (GAD status: Mirrored)

When the GAD status is Mirrored, the I/O mode of the P-VOL and S-VOL is Mirror (RL).

As shown in the following figure, a write request sent to a GAD volume is written to both pair volumes, and then a write-completed response is returned to the host.



Read requests are read from the volume connected to the server and then sent to the server. There is no communication between the primary and secondary storage systems.




Server I/O (GAD status: Mirroring or Quorum disk blocked)


When the GAD status is Mirroring or Quorum disk blocked, the I/O mode for the P-VOL is Mirror(RL), and the I/O mode for the S-VOL is Block. The I/O mode and the I/O flow vary depending on the microcode or firmware version.

Behavior for microcode 80-04-2x or earlier for VSP 5000 series and firmware version 83-03-3x or earlier for VSP Gx00 models

Write requests are written to both pair volumes, and then the write-completed response is returned to the server. Because the S-VOL's I/O mode is Block, it does not accept I/O from the server, but the data written to the P-VOL is also written to the S-VOL by the primary storage system, as shown in the following figure.




Read requests are read by the P-VOL and then sent to the server. There is no communication between the primary and secondary storage systems.




Behavior for microcode version 80-05-0x or later for VSP 5000 series, or the firmware version 83-04-0x for VSP G/F350, G/F370, G/F700, G/F900

When the GAD status is Mirroring or Quorum disk blocked, the I/O mode for the P-VOL is Mirror(RL), and the I/O mode for the S-VOL is Mirror(RL).

Write requests are written to both pair volumes and then the write-completed response is returned to the server.



Read requests are read by the P-VOL or S-VOL and then sent to the server.



Server I/O when the GAD status is Suspended

When the GAD status is Suspended, the I/O mode differs depending on where the latest data is.


When the GAD status is Suspended and the latest data is on the P-VOL, the I/O mode is as follows:

- P-VOL: Local
- S-VOL: Block


When the latest data is on the S-VOL, the I/O mode is as follows:

- P-VOL: Block
- S-VOL: Local

When the latest data is on the P-VOL, write requests are written to the P-VOL, and then the write-completed response is returned to the host, as shown in the following figure. The S-VOL's I/O mode is Block, so it does not accept I/O from the server, and the P-VOL's I/O mode is Local, so the data written to the P-VOL is not written to the S-VOL.



Read requests are read by the P-VOL and then sent to the host. There is no communication between the primary and secondary storage systems.



Server I/O when the GAD status is Blocked


When the GAD status is Blocked, the I/O mode of the P-VOL and S-VOL is Block. Neither volume accepts read/write processing.

Quorum disk and server I/O

The quorum disk is used to determine the storage system on which server I/O should continue when a path or storage system failure occurs.

The quorum disk is a volume virtualized from an external storage system. The primary and secondary storage systems check the quorum disk for the physical path statuses.

When the primary and secondary storage systems cannot communicate, the storage systems take the following actions:



1. The primary storage system cannot communicate over the data path and writes this status to the quorum disk.
2. When the secondary storage system detects from the quorum disk that storage systems cannot communicate over the data path, it stops accepting read/write.
3. The secondary storage system communicates to the quorum disk that it cannot accept read/write.
4. When the primary storage system detects that the secondary storage system cannot accept read/write, the primary storage system suspends the pair. Read/write continues to the primary storage system.

If the primary storage system cannot detect from the quorum disk that the secondary storage system cannot accept I/O within five seconds of a communication stoppage, the primary storage system suspends the pair and I/O continues.

If both systems simultaneously write to the quorum disk that communication has stopped, this communication stoppage is considered to be written by the system with the smaller serial number.

In addition, you can create a GAD pair without setting a volume in an external storage system as the quorum disk volume.

The GAD configuration without a volume set for the quorum disk supports the following migration situations.

- Using GAD to migrate data from VSP F400, F600, F800, VSP G200, G400, G600, G800 to VSP F350, F370, F700, F900, VSP G350, G370, G700, G900
- Using GAD to migrate data from VSP F1500, VSP G1x00 to another VSP F1500, VSP G1x00

I/O stoppage detected in the counterpart system

When a stoppage is detected within 5 seconds in the counterpart system, the pair volume that will continue to receive read/write after the stoppage is determined based on the pair status.

- When the pair status is PAIR, read/write continues to the volume that wrote the communication stoppage to the quorum disk.
- When the pair status is INIT/COPY, read/write continues to the P-VOL. Read/write to the S-VOL remains stopped.
- When the pair status is PSUS, PSUE, SSWS, or SSUS, read/write continues to the volume whose I/O mode is Local. Read/write is stopped to the volume whose I/O mode is Block.

I/O stoppage not detected in the counterpart system

When a stoppage is not detected within 5 seconds in the counterpart system, the pair volume whose system wrote the communication stoppage to the quorum disk will continue to receive read/write after the stoppage.

Read/write processing depends on the pair status and I/O mode of the volume that did not detect the write as follows:

- When the pair status is PAIR, read/write continues.
- When the pair status is INIT/COPY, read/write continues to the P-VOL.

Read/write to the S-VOL remains stopped.

- When the pair status is PSUS, PSUE, SSWS, or SSUS, read/write continues to the volume whose I/O mode is Local.

Read/write is stopped to the volume whose I/O mode is Block. In addition, server I/O does not continue to the volume that should have notified the quorum disk, but did not, that it cannot accept I/O, because either a storage system failure occurred or the quorum disk is no longer accessible.

Server I/Os and data mirroring with blocked quorum disk or without quorum disk volumes

You should understand the server I/Os and data mirroring that occur when a failure occurs on the quorum disk.

GAD pairs that meet the following requirements can continue operation using the S-VOL if the P-VOL is blocked:

- The option for setting no LDEVs for the quorum disk is enabled when the quorum disk is created.
- The pair is created using the quorum ID assigned when the quorum disk is created.

If the quorum disk is blocked, GAD pairs that meet the following requirements can keep the same data in the P-VOL and S-VOL. In addition, if the P-VOL is also blocked, the operation can continue in the S-VOL.

- For VSP F1500 and VSP G1x00, the microcode version of the primary and secondary storage systems is 80-05-xx or later.
- For VSP Fx00 models and VSP Gx00 models, the firmware version is 83-04-0x or later.
- The pair must be created, resynchronized, or swap resynchronized in the following microcode or firmware versions:
 - Microcode version 80-05-xx or later for VSP F1500 and VSP G1x00
 - Firmware version 83-04-0x or later for VSP G/F350, G/F370, G/F700, G/F900

If the quorum disk is blocked, GAD pairs that meet the following requirements can keep the same data in the P-VOL and S-VOL, but the operation stops if the P-VOL is blocked. To continue the operation, you must delete the GAD pair.

- For VSP F1500 and VSP G1x00, the microcode version of the primary and secondary storage systems is 80-02-4x or later.
- For VSP G/F350, G/F370, G/F700, G/F900, the firmware version is from 83-03-0x to 83-03-3x.
- The pair must be created, resynchronized, or swap resynchronized in the following microcode or firmware versions:
 - Microcode version 80-02-4x or later for VSP F1500 and VSP G1x00
 - Firmware version 83-03-0x to 83-03-3x for VSP G/F350, G/F370, G/F700, G/F900

For a pair that does not meet these requirements, the pair is suspended when the quorum disk is blocked. In this case, data duplication cannot be maintained. When the microcode of the primary and secondary storage systems is upgraded to the following level, the existing GAD pairs should be resynchronized or swap resynchronized:

- Microcode version 80-02-4x or later for VSP F1500 and VSP G1x00
- Firmware version 83-03-0x or later for VSP G/F350, G/F370, G/F700, G/F900

For a pair created, resynchronized, or swap resynchronized (80-05-0x or later for VSP 5000 series, 83-04-0x or later for VSP G350, G370, G700, G900 and VSP F350, F370, F700, F900)

When a pair is created, resynchronized, or swap resynchronized, server I/Os and data mirroring are as follows:

- **When the quorum disk is blocked and the pair status is PAIR** The primary and secondary storage systems communicate through remote paths. Because the P-VOL and S-VOL pair status and the I/O mode remain PAIR (Mirror(RL)), server I/Os continue in the P-VOL and the S-VOL. Data mirroring can be maintained through remote paths between the primary and secondary storage systems.
- **When the quorum disk is blocked and the pair status is INIT/COPY** Server I/Os continue in the P-VOL; however, that the pair might be suspended if the quorum disk is blocked immediately after the pair status changes to COPY.
- **When the pair is suspended (pair status is PSUS, PSUE, SSWS, or SSUS) and the quorum disk is blocked or when the pair is suspended** Server I/Os continue in the volume of which I/O mode is Local. I/Os to the volume of which I/O mode is Block remains stopped, and data mirroring remains suspended.
- **When the remote paths are disconnected after the quorum disk is blocked or when the remote paths are disconnected** After the quorum disk is blocked, the pair is suspended when the remote paths are disconnected. The P-VOL status and the I/O mode change to PSUE (Local), and the S-VOL status and the I/O mode change to PAIR (Block). Server I/Os continue in the P-VOL. The pair might be suspended and the status and the I/O mode of the P-VOL and the S-VOL might change to PSUE (Block) depending on the timing of the remote path disconnection after the quorum disk is blocked.

Before the pair status of the S-VOL and the I/O mode change to PAIR (Block), reading data might be delayed. If you want to minimize the delay, set a smaller value for Read Response Guaranteed Time When Quorum Monitoring Stopped. The time between the remote path disconnection and the pair suspension is also shortened.

When you want to restore the remote path quickly and do not want to suspend pairs immediately after the remote path is disconnected, set a larger value for Read Response Guaranteed Time When Quorum Monitoring Stopped. If you set a value larger than the server timeout time, a timeout might occur on the server.

The following table lists the recommended values for Read Response Guaranteed Time When Quorum Monitoring Stopped.

Setting value for Blocked Path Monitoring (sec)	Recommended setting value for Read Response Guaranteed Time When Quorum Monitoring Stopped
40 (Default)	40 (Default)
2 to 5	5*
6 to 25	6 to 25*

Setting value for Blocked Path Monitoring (sec)	Recommended setting value for Read Response Guaranteed Time When Quorum Monitoring Stopped
26 to 44	26 to 44
45	45
* A GAD pair might be suspended if remote path communication is blocked temporarily due to an MP or path failure. To avoid this, a value which is greater than the RIO MIH time or at least 25 seconds must be set for Read Response Guaranteed Time When Quorum Monitoring Stopped. Note, however, that reading data might delay up to the time set for Read Response Guaranteed Time When Quorum Monitoring Stopped.	

Setting the same value as the blocked path monitoring for Read Response Guaranteed Time When Quorum Monitoring Stopped is recommended. Until the pair status and I/O mode of the S-VOL change to PSUE (Block), delay of reading data can be maintained within the seconds set for Read Response Guaranteed Time When Quorum Monitoring Stopped. Note that if a value equal to or less than 5 seconds is set for the blocked path monitoring, set 5 for Read Response Guaranteed Time When Quorum Monitoring Stopped.

If a value equal to or greater than 46 seconds is set for Read Response Guaranteed Time When Quorum Monitoring Stopped, GAD pair suspension caused by a remote path failure might be avoided. When you set a value of 46 or a greater, make sure that the application timeout setting for server I/Os is greater than this value. Also, make sure that multiple remote paths are set (at least four paths are recommended). Reading data might be delayed until the time set for Read Response Guaranteed Time When Quorum Monitoring Stopped elapses.



Note: If a pair created on microcode version 80-04-2x or earlier for VSP 5000 series or firmware version 83-03-3x or earlier for VSP Gx00 models is resynchronized or swap resynchronized in storage systems on microcode version 80-05-0x or later for VSP 5000 series and firmware version 83-04-0x or later for VSP Gx00 models at the primary and secondary sites, the pair operates similarly as pairs created on microcode version 80-05-0x or later.

For a pair created, resynchronized, or swap resynchronized (80-02-4x to 80-04-2x for VSP 5000 series , 83-03-0x to 83-03-3x VSP Gx00 models)

When a pair is created, resynchronized, or swap resynchronized on microcode version 80-02-4x to 80-04-2x (both the primary and secondary storage system), I/O and data mirroring are as follows:

- **When the quorum disk is blocked and the pair status is PAIR:** The primary and secondary storage systems communicate through remote paths. The P-VOL pair status and I/O mode changes to PAIR (Mirror(RL)). The S-VOL pair status and I/O mode changes to PAIR (Block). Server I/O continues on the P-VOL. Data mirroring can be maintained by using the remote paths between the primary and secondary storage systems.
- **When the quorum disk is blocked and the pair status is INIT/COPY:** The pair is suspended, and data duplication is also suspended. Server I/O continues on the P-VOL.
- **When the quorum disk is blocked and the pair status is PSUS, PSUE, or SSUS:** Server I/O continues on the volume whose I/O mode is Local. I/O to the volume whose I/O mode is Block remains stopped. Data mirroring remains suspended.
- **When the remote paths are disconnected after the quorum disk is blocked:** After the quorum disk is blocked and the P-VOL status and I/O mode changes to PAIR (Mirror(RL)) and the S-VOL status and I/O mode changes to PAIR (Block), the pair is suspended and data mirroring is also suspended when the remote paths between the primary and secondary storage systems are detected to be disconnected. Server I/O continues on the P-VOL.

For a pair created, resynchronized, or swap resynchronized (80-02-3x or earlier for VSP 5000 series, 83-03-0x or earlier for VSP Gx00 models)

For a pair that was created when the microcode was 80-02-3x or earlier, I/O stoppage and data duplication are as follows:

- **When the quorum disk is blocked and the pair status is PAIR, INIT, or COPY:** The pair is suspended, and data duplication is also suspended. Server I/O continues on the P-VOL.
- **When the quorum disk is blocked and the pair status is PSUS, PSUE, or SSUS:** Server I/O continues on the volume whose I/O mode is Local. I/O to the volume whose I/O mode is Block remains stopped. Data duplication remains suspended.
- **When the remote paths are disconnected after the quorum disk is blocked:** When the quorum disk is blocked and the pair is suspended, and when the remote paths between the primary and secondary storage system are detected to be disconnected, server I/O continues on the P-VOL. Data duplication remains suspended.

Quorum disk status

You need to check the status of the quorum disk before you replace the external storage system currently used by the quorum disk while you keep GAD pairs.

You can check the quorum disk status using the `raidcom get quorum` command. For details, see the Command Control Interface Command Reference.

In microcode version 80-05-4x or later (VSP 5000 series) or in firmware version 83-04-4x or later (VSP Gx00 models and VSP Fx00 models), you can replace the external storage system currently used by the quorum disk with a new external storage system while keeping GAD pairs.


There are five statuses for the quorum disk.

Quorum disk status	Display by CCI	Description
Normal	NORMAL	The quorum disk is operating normally.
Transitioning	TRANSITIONING	The status of the quorum disk is being changed.
Blocked	BLOCKED	The quorum disk is blocked.
Replacing	REPLACING	The quorum disk is being replaced.
Failed	FAILED	The primary and secondary storage systems are connected to different quorum disks. Specify the external volume again, so that they can be connected to the same quorum disk, and reconfigure the quorum disk.

Initial copy and differential copy

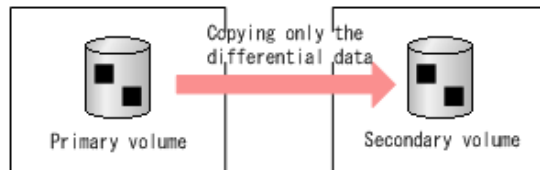
There are two types of GAD copy operations that synchronize the data on the P-VOL and S-VOL of a pair, initial copy and differential copy.

For an initial copy operation, all data in the P-VOL is copied to the S-VOL, which ensures that the data in the two volumes is consistent. The initial copy is executed when the GAD status changes from Simplex to Mirrored.



Differential copy For a differential copy operation, only the differential data between the P-VOL and the S-VOL is copied. Differential copy is used when the GAD status changes from Suspended to Mirrored.


When a GAD pair is suspended, the storage systems record the update locations and manage the differential data. The following figure shows the differential copy operation for a pair in which the P-VOL received server I/O while the pair was suspended. If the S-VOL receives server I/O while a pair is suspended, the differential data is copied from the S-VOL to the P-VOL.



GAD consistency groups

You can manage multiple GAD pairs as a group by using consistency groups.

The GAD pairs in a GAD 3DC delta resync (GAD+UR) configuration must be registered to a consistency group.



Registering GAD pairs to consistency groups enables you to perform operations on all GAD pairs in a consistency group at the same time. In addition, when a failure occurs, the GAD pairs are suspended by consistency group (concurrent suspension).

For details about storage system support (microcode) for consistency groups, see [Requirements and restrictions](#).

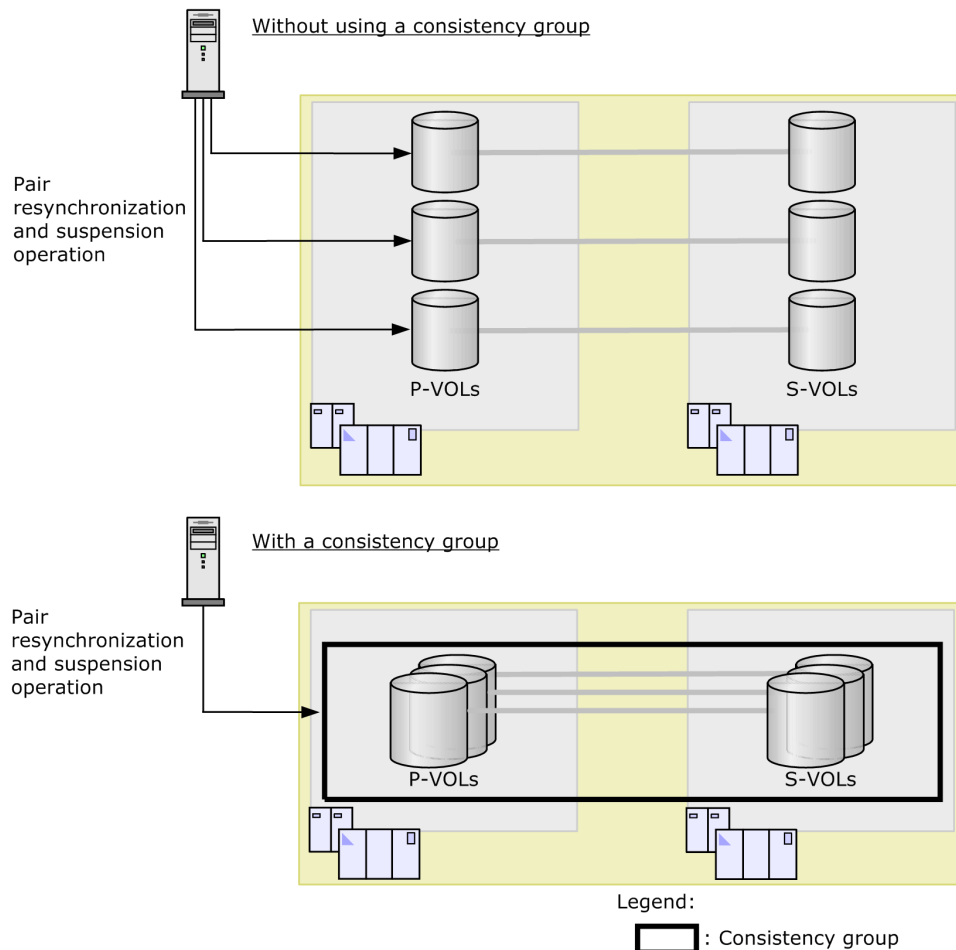


Note: When you register GAD pairs to a consistency group, you should use the cross-path configuration. For details, see [System configuration for GAD solutions \(on page 16\)](#).

Operations on GAD pairs by consistency group

By registering multiple GAD pairs to a consistency group, you can resynchronize or suspend the GAD pairs by consistency group.

You can resynchronize all GAD pairs registered to a consistency group by performing a single pair resynchronization operation. In addition, you can suspend all GAD pairs registered to a consistency group by performing a single pair suspension operation.



For details about storage system support (microcode) for consistency groups, see [Requirements and restrictions \(on page 46\)](#).

Suspension of GAD pairs by consistency group

When a failure occurs, suspension of GAD pairs by consistency group guarantees data consistency among primary volumes if the I/O mode of a primary volume changes to Block, or among secondary volumes if the I/O mode of a secondary volume changes to Block.

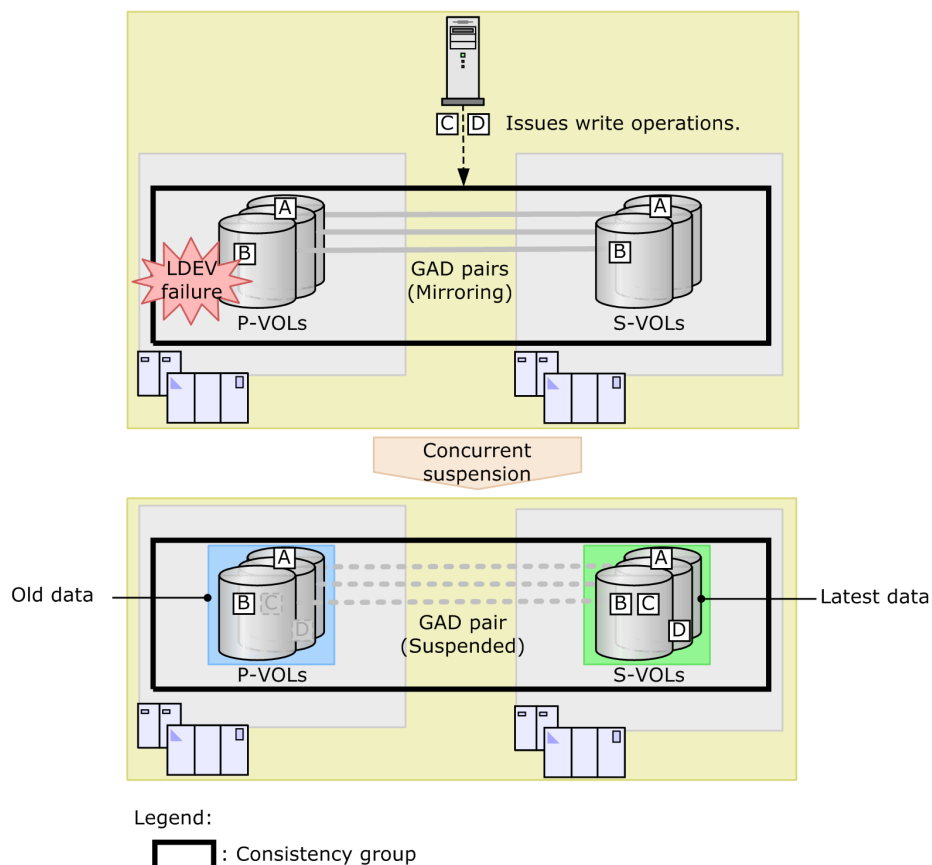
If some GAD pairs in a consistency group are suspended due to a failure, all GAD pairs in the consistency group to which the suspended GAD pairs are registered change to the suspended state. This is called concurrent suspension.

- The volumes that have the most recent data are aggregated to a single storage system.

If a failure occurs in some pairs, and all GAD pairs registered to a consistency group are in the Suspended state, the volumes that have the most recent data are aggregated to the storage system at either the primary site or the secondary site.

- Data consistency is guaranteed before and after the suspension of the GAD pairs.

If all GAD pairs registered to a consistency group are in the Suspended state, only the volumes (of either the primary or the secondary site) that have the most recent data will receive I/O from the server. The volumes of the other site will stop receiving I/O from the server (including I/O for volumes where no failure occurred). In addition, processing to write data will also stop. This ensures data consistency before and after the GAD pair suspension in the volumes that stopped receiving I/O from the server.



For example, a server issues write operations A to D. After the storage system receives write operation B, all GAD pairs registered to the consistency group change to the Suspended state because of an LDEV failure in the primary volume. In such a case, write operations A and B received before the GAD pairs changed to the Suspended state were completed for both the primary and secondary volume. Write operations C and D received after the GAD pairs changed to the Suspended state were completed only for the secondary volume.

Therefore, the volumes that have the most recent data are aggregated to the storage system at the secondary site.

For details about storage system support (microcode) for consistency groups, see [Requirements and restrictions \(on page 46\)](#).

Use cases for consistency groups

You can use GAD consistency groups for many use cases, for example, batch failover or resuming operations by using consistent backup data.

GAD consistency group statuses

You can view the status of a consistency group by using Device Manager - Storage Navigator.

The following table describes the statuses of GAD consistency groups.

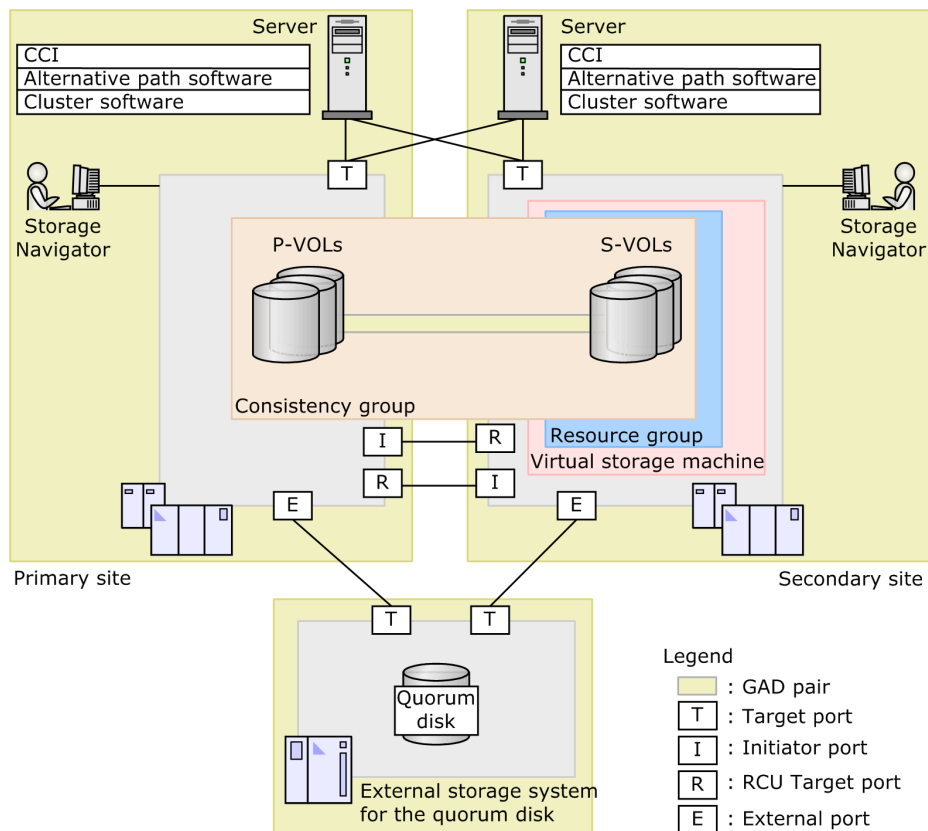
Status	Description
SMPL	All volumes in the consistency group are not used as GAD pair volumes.
INIT/COPY	The initial copy or pair resynchronization of all GAD pairs in the consistency group is in progress (including creation of a GAD pair that does not perform data copy). A quorum disk is being prepared.
COPY	The initial copy of all GAD pairs in the consistency group is in progress; data is being copied from the P-VOL to the S-VOL (including creation of a GAD pair that does not perform data copy).
PAIR	All GAD pairs in the consistency group are synchronized, including pairs whose quorum disk is blocked. The data is duplicated.
PSUS	All GAD pairs in the consistency group were suspended by the user. This status appears when the volumes in the consistency group on the local storage system are P-VOLs.
PSUE	All GAD pairs in the consistency group were suspended due to a failure.
SSUS	All GAD pairs in the consistency group were suspended by the user, and update of the S-VOL is interrupted. This status appears when the volumes in the consistency group on the local storage system are S-VOLs.
SSWS	All GAD pairs in the consistency group were suspended either by the user or due to a failure, and update of the P-VOL is interrupted. This status appears when the volumes in the consistency group on the local storage system are S-VOLs.
Suspending	GAD pair suspension processing is being performed by consistency group.

Status	Description
Resynchronizing	GAD pair resynchronization processing is being performed by consistency group.
Mixed	More than one pair status exists in the consistency group.
Unknown	The consistency group status cannot be obtained.
Blank	The consistency group is not used.

Global-active device components

A typical global-active device system consists of storage systems, paired volumes, a consistency group, a quorum disk, a virtual storage machine, paths and ports, alternate path software, and cluster software.

The following illustration shows the components of a typical global-active device system.



Storage systems

Both of the primary and secondary storage systems should be the same model type, but they do not have to be the same model. For example,

- If the primary storage system is a VSP G1x00 or VSP F1500, the secondary storage system can be a VSP G1x00 or VSP F1500.
- If the primary storage system is a VSP G1x00 or VSP F1500, the secondary storage system can be a VSP G350, G370, G700, G900, VSP F350, F370, F700, F900.
- If the primary storage system is a VSP G350, G370, G700, G900 storage system, the secondary storage system can be a VSP G350, G370, G700, G900.
- If the primary storage system is a VSP F350, F370, F700, F900 storage system, the secondary storage system can be a VSP F350, F370, F700, F900.

An external storage system, which is connected to the primary and secondary storage systems using Universal Volume Manager, is required for the quorum disk.

Paired volumes

A global-active device pair consists of a P-VOL in the primary storage system and an S-VOL in the secondary storage system. For model connectivity support requirements, see [System requirements \(on page 46\)](#).

Consistency group

A consistency group consists of multiple global-active device pairs. By registering GAD pairs to a consistency group, you can resynchronize or suspend the GAD pairs by consistency group.

For details about storage system support (microcode) for consistency groups, see [Requirements and restrictions \(on page 46\)](#).

Quorum disk

The quorum disk, required for global-active device, is used to determine the storage system on which server I/O should continue when a storage system or path failure occurs. The quorum disk is virtualized from an external storage system that is connected to both the primary and secondary storage systems.

Virtual storage machine

A virtual storage machine (VSM) is configured in the secondary storage system with the same model and serial number as the (actual) primary storage system. The servers treat the virtual storage machine and the storage system at the primary site as one virtual storage machine.

You can create GAD pairs using volumes in virtual storage machines. When you want to create a GAD pair using volumes in VSMs, the VSM for the volume in the secondary site must have the same model and serial number as the VSM for the volume in the primary site.

Paths and ports

GAD operations are carried out between hosts and primary and secondary storage systems that are connected by data paths composed of one or more physical links.

The data path, also referred to as the remote connection, connects ports on the primary storage system to ports on the secondary storage system. The ports have attributes that enable them to send and receive data. One data path connection is required, but you should use two or more independent connections for hardware redundancy.



Note: You do not need to set the port attributes (Initiator, RCU Target, Target External) on VSP G/F350, G/F370, G/F700, G/F900 models.

Alternate path software

Alternate path software is used to set redundant paths from servers to volumes and to distribute host workload evenly across the data paths. Alternate path software is required for the single-server and cross-path GAD system configurations.

Cluster software

Cluster software is used to configure a system with multiple servers and to switch operations to another server when a server failure occurs. Cluster software is required when two servers are in a global-active device server-cluster system configuration.

User interfaces for global-active device operations

Global-active device operations are performed using the management software and optionally, the command-line interface (CLI) software, for the storage system.

Storage Advisor Embedded

For VSP G/F350, G/F370, G/F700, G/F900: firmware 88-03-0x or later, use HSAE to configure the remote paths and the quorum disk required for using global-active device. These settings must be configured on the primary and secondary storage systems in the global-active device environment.

Device Manager - Storage Navigator

For VSP G200, G/F400, G/F600, G/F800 with firmware version 83-05-2x or later and VSP G1x00 and VSP F1500 with firmware version 80-06-6x or later, use Storage Navigator to configure the remote paths and the quorum disk required for using global-active device. These settings must be configured on the primary and secondary storage systems in the global-active device environment.

Storage Advisor

Use Storage Advisor to configure and manage virtual storage machine (VSM), GAD pairs, and monitor and manage your global-active device environment.

Data Instance Director

Use Data Instance Director to perform pair operations in a global-active device environment.

Chapter 2: System requirements

You should understand all requirements and restrictions for global-active device operations.

Requirements and restrictions

The following table lists the requirements and restrictions for global-active device operations.

Item	Requirements and restrictions
Primary and secondary storage systems	<ul style="list-style-type: none">▪ Model: The following information is current as of the publish date. See the TC, UR, GAD Replication intermix matrix at https://support.hitachivantara.com/en_us/interoperability.html to find a complete and current list of support information.▪ Microcode:<ul style="list-style-type: none">• VSP G1000: DKCMAIN 80-05-0x or later• VSP G1500: DKCMAIN 80-05-0x or later• VSP F1500: DKCMAIN 80-05-0x or later• VSP G350, G370, G700, G900, VSP F350, F370, F700, F900: Any version compatible with the models.• VSP F400, VSP F600, VSP F800 and VSP G200, VSP G400, VSP G600, VSP G800: DKCMAIN 83-04-0x or later• When you connect VSP G1x00, VSP F1500 and VSP G350, G370, G700, G900, VSP F350, F370, F700, F900:<ul style="list-style-type: none">▪ VSP G1x00, VSP F1500: DKCMAIN 80-06-4x or later▪ VSP G350, G370, G700, G900, VSP F350, F370, F700, F900: DKCMAIN 88-02-0x or later▪ Global-active device license: The global-active device feature must be installed and enabled on the primary and secondary storage systems.


Item	Requirements and restrictions
	<ul style="list-style-type: none"> ▪ Controller emulation type: The controller emulation type of the primary and secondary storage systems must be same. ▪ Shared memory: <ul style="list-style-type: none"> • VSP G200VSP G350, G370, G700, G900 and VSP F350, F370, F700, F900: You can use GAD only with shared memory in the base part. Adding shared memory expands the capacity of the pairs being created. • VSP G400, G600, G800, VSP F400, F600, F800: Additional shared memory is required in the primary and secondary storage systems. For information on adding shared memory and setting the GAD dedicated area, contact customer support. • VSP G1000, VSP G1500, and VSP F1500: Additional shared memory is required in the primary and secondary storage systems.

Item	Requirements and restrictions
GAD 3DC delta resync (GAD+UR)	<ul style="list-style-type: none"> ▪ The following intermix configurations are supported. No other configurations are supported. <ul style="list-style-type: none"> • VSP 5000 series GAD primary storage system, VSP 5000 series GAD secondary storage system, VSP 5000 series UR storage system. • VSP G900 GAD primary storage system, VSP G900 GAD secondary storage system, VSP G900 UR storage system. • VSP F900 GAD primary storage system, VSP F900 GAD secondary storage system, VSP F900 UR storage system. • VSP G800 GAD primary storage system, VSP G800 GAD secondary storage system, VSP G800 UR storage system. • VSP F800 GAD primary storage system, VSP F800 GAD secondary storage system, VSP F800 UR storage system. • VSP 5000 series GAD primary storage system, VSP 5000 series GAD secondary storage system, VSP G800 or VSP F800 UR storage system. ▪ Microcode/firmware: <ul style="list-style-type: none"> • When only VSP 5000 series storage systems are used: DKCMAIN microcode version 80-02-4x or later • When only VSP G800 storage systems are used: DKCMAIN firmware version 83-03-xx or later • When only VSP F800 storage systems are used: DKCMAIN firmware version 83-04-2x or later • When both VSP 5000 series and VSP G800 storage systems are used VSP 5000 series DKCMAIN microcode version 80-04-xx or later VSP G800: DKCMAIN firmware version 83-03-xx or later • When both VSP G1000, VSP G1500, VSP F1500, and VSP F800 storage systems are used: VSP 5000 series: DKCMAIN microcode version 80-04-xx or later VSP G800: DKCMAIN firmware version 83-04-2x or later

Item	Requirements and restrictions
	<ul style="list-style-type: none"> GAD pairs in a GAD+UR configuration must be registered to a consistency group.
External storage systems or servers (for quorum disk)	<ul style="list-style-type: none"> The storage system must be supported for attachment using Universal Volume Manager. For details, see the <i>Hitachi Universal Volume Manager User Guide</i> for the storage system. The maximum distance between the external storage system and the primary site and secondary site is 1,500 km.
Licensed capacity	<ul style="list-style-type: none"> The page size assigned to the virtual volume is counted as a licensed capacity for GAD; however, for a volume with capacity saving enabled, the GAD licensed capacity is the capacity before savings. If the actual licensed capacity exceeds the available licensed capacity, GAD can be used as usual for 30 days. After 30 days, only pair split and pair delete operations are allowed.
Host server platforms	<ul style="list-style-type: none"> AIX HP-UX OpenVMS Red Hat Enterprise Linux Solaris SuSE Linux VMware ESX Windows Server <p>For more information, refer to the Hitachi Vantara interoperability matrix: https://support.hitachivantara.com/en_us/interoperability.html</p>
Maximum number of storage systems that can be connected	One storage system can create pairs with a maximum of 15 storage systems.

Item	Requirements and restrictions
SCSI commands	<ul style="list-style-type: none"> ▪ The Thin Provisioning function of Windows Server 2012 is supported. ▪ The SCSI-2 Reserve command, the SCSI-3 Persistent Reserve command, and the VAAI command are supported. ▪ The reservation information is duplicated when the Reserve command or the Persistent Reserve command is received, or when the initial copy or resync copy starts.
Physical paths connecting the primary and secondary storage systems	<ul style="list-style-type: none"> ▪ Maximum number of physical paths: 8 ▪ Maximum distance between the primary and secondary storage systems: 500 km ▪ The maximum value of the round-trip delay, including the delay due to the failure of an interface device, is 20 ms. However, you must meet the following conditions to connect storage systems over the distance of 100 km or more. <ul style="list-style-type: none"> • The primary and secondary storage systems are connected by Fibre Channel interfaces. • The DKCMAIN microcode version of the primary and secondary storage systems is 80-04-21-00/00 or later for VSP G1000, VSP G1500, VSP F1500. • The DKCMAIN firmware version of the primary and secondary storage systems is 83-03-21-x0/01 or later for VSP Gx00 models and VSP Fx00 models VSP G200, G400, G600, G800, VSP F400, F600, F800. • The line speed between the primary and secondary storage systems is 1 Gbps or more. • The host mode option 51 is set to ON. ▪ Port type: Fibre Channel, VSP G200, G400, G600, G800 83-01-2x and later, VSP F400, F600, F800 83-04-2x and later, VSP 5000 series 80-03-3x and later) with direct, switch, or channel extenders. For details, see Connection types. ▪ Port attribute (VSP 5000 series): The ports that connect the primary and secondary storage systems must be configured as Initiator ports and RCU target ports.
Remote paths and path groups	<ul style="list-style-type: none"> ▪ Maximum number of remote paths per path group: 8 ▪ Maximum number of path groups per storage system: 64 (sum of the path groups used by TC, UR, and URz) ▪ Path group ID: 0-255.

Item	Requirements and restrictions
	<ul style="list-style-type: none"> ▪ Protocol: All remote paths in a path group must be the same protocol, either Fibre Channel. Remote paths for Fibre Channel cannot coexist within the same path group. ▪ The path group is specified during the create pair operation and cannot be changed by resynchronization. ▪ The remote path must be set by each path group of the storage systems at the primary site and the secondary site. <p>You can also use multiple path groups with the same combination of the storage systems at the primary and the secondary sites.</p> <ul style="list-style-type: none"> ▪ When using the System connection type and not the CU connection type (specified on the Add Remote Connection window), specify different paths and path groups for TrueCopy, Universal Replicator, and Universal Replicator for Mainframe secondary storage systems.
Virtual storage machines (VSMs)	<ul style="list-style-type: none"> ▪ Maximum number of VSMs per storage system: 7 ▪ Maximum number of GAD volumes per VSM: <ul style="list-style-type: none"> • VSP G200: 2,048 • VSP G400, G600, VSP F400, F600: 4,096 • VSP G350, VSP F350VSP G800, VSP F800: 16,384 • VSP G370, VSP F370: 32,768 • VSP G700, VSP F700: 49,152 • VSP 5000 series, VSP G900, VSP F900: 65,280 ▪ You can create GAD pairs using volumes in virtual storage machines. When you want to create a GAD pair using volumes in VSMs, the VSM for the volume in the secondary site must have the same model and serial number as the VSM for the volume in the primary site.

Item	Requirements and restrictions
	<ul style="list-style-type: none"> ▪ (VSP Gx00 models, VSP Fx00 models) When a resource group (virtual storage machine) in a storage system at the secondary site has the same virtual LDEV ID as the P-VOL, you cannot create a GAD pair. In addition, when a volume is not created and only LDEV IDs exist, virtual LDEV IDs must be deleted. ▪ The VSP 5000 series microcode version for both the primary and secondary storage systems must be 80-02-01 or later to create a GAD pair using a P-VOL that already has a virtual LDEV ID assigned and is managed by a user-defined virtual storage machine. This requirement is not necessary for GAD pairs created by adding secondary storage resources to the primary storage default virtual storage machine.
Resource groups (VSP 5000 series)	<ul style="list-style-type: none"> ▪ DKCMAIN 80-02-xx or later: A volume in a resource group that was migrated from a VSP or USP V/VM storage system VSP 5000 series can be used as a GAD volume. ▪ DKCMAIN 80-01-xx or earlier: A volume in a resource group that was migrated from a VSP or USP V/VM storage system VSP 5000 series cannot be used as a GAD volume. <div data-bbox="695 1087 1393 1199">  Note: Storage Advisor supports DP volumes. Basic volumes, including external volumes, are not supported. </div>

Item	Requirements and restrictions
Maximum number of GAD pairs	<ul style="list-style-type: none"> ▪ When all pairs are created with DP-VOLs and external volumes (calculated by subtracting the number of quorum disks (at least one) from the maximum number of virtual volumes that can be defined in a storage system): <ul style="list-style-type: none"> • VSP G200: 2,046 • VSP G400, VSP G600, VSP F400, F600: 4,094 • VSP G800, VSP F800: 14,079 • VSP G350, VSP F350: 16,382 • VSP G370, VSP F370: 32,766 • VSP G700, VSP F700: 49,150 • VSP 5000 series, VSP G900, VSP F900: 63,231 ▪ When CCI is used in the in-band method and all pairs are created with DP-VOLs or external volumes, and one virtual volume or external volume is used as a command device, and a volume is set for the quorum disk: <ul style="list-style-type: none"> • VSP G350, VSP F350: 16,381 • VSP G370, VSP F370: 32,765 • VSP G700, VSP F700: 49,149 • VSP 5000 series, VSP G900, VSP F900: 63,230 ▪ When CCI is used in the in-band method and all pairs are created with DP-VOLs or external volumes, one normal volume (VSP Gx00 models, VSP Fx00 models) or internal volume (VSP 5000 series) is used as a command device, and a volume is set for the quorum disk: <ul style="list-style-type: none"> • VSP G350, VSP F350: 16,381 • VSP G370, VSP F370: 32,765 • VSP G700, VSP F700: 49,149 • VSP 5000 series, VSP G900, VSP F900: 63,231 ▪ When CCI is used in the in-band method and all pairs are created with DP-VOLs or external volumes, one virtual volume (DP-VOL) or external volume is used as a command device: <ul style="list-style-type: none"> • VSP F1500 and VSP G1x00: 63,231

Item	Requirements and restrictions
	<ul style="list-style-type: none"> ▪ When CCI is used in the in-band method and all pairs are created with DP-VOLs or external volumes, one internal volume is used as a command device: <ul style="list-style-type: none"> • VSP F1500 and VSP G1x00: 63,232 ▪ When all pairs are created with internal volumes (calculated by subtracting the number of quorum disks (at least one) from the maximum number of internal volumes that can be defined in a storage system): <ul style="list-style-type: none"> • VSP G350, VSP F350: 16,383 • VSP G370, VSP F370: 32,767 • VSP G700, VSP F700: 49,151 • VSP 5000 series, VSP G900, VSP F900: 65,279 ▪ When CCI is used in the in-band method and all pairs are created with internal volumes, and one virtual volume or external volume for VSP 5000 series is used as a command device, and a volume is set for the quorum disk: <ul style="list-style-type: none"> • VSP G350, VSP F350: 16,381 • VSP G370, VSP F370: 32,765 • VSP G700, VSP F700: 49,149 • VSP 5000 series, VSP G900, VSP F900: 65,277 ▪ When CCI is used in the in-band method and all pairs are created with internal volumes, and one normal volume (VSP Gx00 models, VSP Fx00 models) or internal volume (VSP 5000 series) is used as a command device, and a volume is set for the quorum disk: <ul style="list-style-type: none"> • VSP G350, VSP F350: 16,382 • VSP G370, VSP F370: 32,766 • VSP G700, VSP F700: 49,150 • VSP 5000 series, VSP G900, VSP F900: 65,278 ▪ When CCI is used in the in-band method and all pairs are created with internal volumes, and one virtual volume or external volume for VSP F1500 and VSP G1x00 is used as a command device: <ul style="list-style-type: none"> • VSP F1500 and VSP G1x00: 65,278

Item	Requirements and restrictions
	<ul style="list-style-type: none"> ▪ When CCI is used in the in-band method and all pairs are created with internal volumes, and one normal volume (VSP Gx00 models, VSP Fx00 models) or internal volume (VSP 5000 series) is used as a command device: <ul style="list-style-type: none"> • VSP F1500 and VSP G1x00: 65,279 ▪ Virtual storage machine: same as the maximum number of pairs for the storage system model. ▪ For details about calculating the maximum number of pairs based on the number of cylinders used in volumes or the number of bitmap areas used in volumes, see Maximum number of GAD pairs.
Pair volumes	<ul style="list-style-type: none"> ▪ Provisioning type: The following provisioning types are supported for the GAD pair volumes. The provisioning type of the P-VOL and S-VOL must be same. For example, if the P-VOL is a DP-VOL, the S-VOL must also be a DP-VOL. <ul style="list-style-type: none"> • Dynamic Provisioning virtual volumes (DP-VOLs) For DP-VOLs, you can only create a GAD pair when both DP-VOLs do not have the Data Direct Mapping attribute or when both DP-VOLs have the Data Direct Mapping attribute. You cannot create a GAD pair when the Data Direct Mapping attribute is enabled for one DP-VOL but not for the other. • Internal volumes • External volumes ▪ Emulation type: OPEN-V. ▪ Volume size: The P-VOL and S-VOL must be equal in size. ▪ Maximum volume size: <p>DP-VOL: same as the maximum size of a DP-VOL. For details, see the <i>Provisioning Guide</i> for the storage system.</p> <p>(VSP 5000 series)</p> <ul style="list-style-type: none"> ▪ Internal volume: 3,145,663 MB (6,442,317,824 blocks) ▪ External volume: 4,194,304 MB (8,589,934,592 blocks) <p>(VSP G350, VSP G370, VSP G700, VSP G900, VSP F350, VSP F370, VSP F700, VSP F900)</p> <ul style="list-style-type: none"> ▪ Internal volume: 256 TB ▪ External volume: 256 TB ▪ SAN boot: You can use GAD pair volumes for SAN boot.

Item	Requirements and restrictions
	<ul style="list-style-type: none"> ▪ Virtual LDEV ID: The same virtual LDEV ID as the P-VOL must not exist in the resource group of the secondary storage system (virtual storage machine). You cannot create a GAD pair when the same virtual LDEV ID as the P-VOL exists in the resource group of the secondary storage system (virtual storage machine). To use the P-VOL, you must delete the virtual LDEV ID in the resource group of the secondary storage system. You must delete the virtual LDEV ID even if the volume is not created and only the LDEV ID exists. ▪ Dynamic volume expansion: You cannot dynamically expand a GAD pair volume. If you need to expand a GAD pair volume, you must delete the pair, expand the volume, and then re-create the pair. ▪ T10 PI: The same value must be set for the T10 PI attribute of the P-VOL and the S-VOL. ▪ A volume (LDEV) from a parity group with accelerated compression enabled cannot be used directly as a GAD pair volume. Such volumes must be used as pool volumes for an HDP or HDT pool.
Quorum disks	<ul style="list-style-type: none"> ▪ Maximum number of quorum disks: 32 per storage system in the primary storage system and secondary storage system. ▪ Quorum disk ID: Specify a value from 0 to 31.

Item	Requirements and restrictions
	<ul style="list-style-type: none"> ▪ Maximum number of pairs per quorum disk: <ul style="list-style-type: none"> • VSP G200: 2,046 when you create all pairs with DP-VOLs and external volumes, and 2,047 when you create all pairs with internal volumes. • VSP G400, VSP G600, VSP F400, F600: 4,094 when you create all pairs with DP-VOLs and external volumes, and 4,095 when you create all pairs with internal volumes. • VSP G800, VSP F800: 14,079 when you create all pairs with DP-VOLs and external volumes, and 16,383 when you create all pairs with internal volumes. • VSP G350, VSP F350: 16,382 when you create all pairs with DP-VOLs or external volumes, and 16,383 when you create all pairs with internal volumes. • VSP G370, VSP F370: 32,766 when you create all pairs with DP-VOLs or external volumes, and 32,767 when you create all pairs with internal volumes. • VSP G700, VSP F700: 49,150 when you create all pairs with DP-VOLs or external volumes, and 49,151 when you create all pairs with internal volumes. • VSP G900, VSP F900: 63,232 when you create all pairs with DP-VOLs or external volumes, and 65,279 when you create all pairs with internal volumes. • VSP 5000 series: <ul style="list-style-type: none"> 63,231 when you create all pairs with DP-VOLs or external volumes. 65,279 when you create all pairs with internal volumes. ▪ Emulation type (VSP 5000 series): OPEN-V ▪ Minimum size: 12,292 MB (25,174,016 blocks) ▪ Maximum size: same as the maximum limit for an external volume supported by Universal Volume Manager: 4 TB. ▪ One external volume group must be mapped to one external volume.

Item	Requirements and restrictions
	<ul style="list-style-type: none"> ▪ Interoperability: A GAD quorum disk cannot also be used as a quorum disk with High Availability Manager for VSP, HUS VM, and USP V/VM. ▪ Requirements for the external storage system volume: <ul style="list-style-type: none"> • The T10 PI attribute must not be enabled. • The Data Direct Mapping attribute must not be set. • The volume must not belong to the system resource group.
Consistency groups	<ul style="list-style-type: none"> ▪ The following firmware is required for CTG support on VSP Gx00 models: DKCMAIN version 83-03-xx or later ▪ Maximum number of consistency groups per storage system: <ul style="list-style-type: none"> • VSP G200: 16 (CTG ID 0-15) • VSP G400, G600, VSP F400, F600: 64 (CTG ID 0-63) • VSP G350, VSP G370, VSP G700, VSP F350, VSP F370, VSP F700, VSP G800, VSP F800: 128 (CTG ID 0-127) • VSP 5000 series, VSP G900, VSP F900: 256 (CTG ID 0-255) ▪ Maximum number of GAD pairs per consistency group: <ul style="list-style-type: none"> • VSP G200: 2,047 • VSP G400, G600, VSP F400, F600: 4,095 • VSP F350, F370, F700, F900 and VSP G350, G370, G700, G900, VSP G800, VSP F800: 8,192 • VSP 5000 series: 8,192 ▪ Quorum disk ID: The same quorum disk ID must be set for all GAD pairs in a single consistency group. ▪ VSM: GAD pairs in the same consistency group must be created on the same virtual storage machine. ▪ CTG ID: If consistency groups have the same ID but their physical storage systems are different, they are treated as different consistency groups. You can use the same consistency group ID for groups on different storage systems.
Alternate path software	<p>Alternate path software is required for the single-server GAD configuration and the cross-path GAD configuration (two servers). When ALUA is used in the cross-path configuration, use the same models at both the primary and secondary sites.</p>

Item	Requirements and restrictions
	<p>The following microcode/firmware is required in the storage systems at the primary and secondary sites: DKCMAIN version 80-03-31-00/00 or later for VSP 5000 series, and DKCMAIN version 83-03-0x-xx/xx or later for VSP G200, G400, G600, G800, or 83-04-2x-xx/xx or later for VSP F400, F600, F800 for VSP 5000 series.</p> <p>Refer to the Hitachi Vantara interoperability matrix: https://support.hitachivantara.com/en_us/interoperability.html</p>
Cluster software	<p>Cluster software is required for the server-cluster and cross-path GAD configurations.</p> <p>Refer to the Hitachi Vantara interoperability matrix: https://support.hitachivantara.com/en_us/interoperability.html</p>
User interfaces	<ul style="list-style-type: none"> ▪ Command Control Interface: <ul style="list-style-type: none"> • VSP G200, G400, G600, G800VSP G350, G370, G700, G900 and VSP F350, F370, F700, F900: 01-45-03/02 or later • VSP F400, F600, F800: 01-39-03 or later • VSP 5000 series: 01-32-03/00 or later

Interoperability requirements

You should understand how global-active device (GAD) operates with other features of the VSP G1000, VSP G1500, VSP F1500, VSP G350, VSP G370, VSP G700, VSP G900, VSP F350, VSP F370, VSP F700, VSP F900 storage systems. It is possible that volume types that are used for features other than GAD can be used as GAD P-VOL or S-VOL.

Volume types that can be used for GAD

Volumes used for other than GAD can be or cannot be used with GAD P-VOL and S-VOL.

The following table explains whether the volume types used for features other than GAD can be used with the GAD P-VOL or S-VOL.

Volume type	Used as GAD P-VOL?	Used as GAD S-VOL?	Used as quorum disk?
Dynamic Provisioning / Dynamic Tiering / Active flash			
Virtual volume	Yes ¹	Yes ¹	No
Pool volume	No	No	No

Volume type	Used as GAD P-VOL?	Used as GAD S-VOL?	Used as quorum disk?
V-VOL with capacity saving enabled	Yes	Yes	No
Deduplication system data volume	No	No	No
ShadowImage / Thin Image ²			
P-VOL	Yes	Yes	No
S-VOL	No	No	No
TrueCopy			
P-VOL	No	No	No
S-VOL	No	No	No
Universal Replicator			
P-VOL	Yes	Yes ³	No
S-VOL	No	No	No
Journal volume	No	No	No
Universal Volume Manager			
External volume	Yes ¹	Yes ¹	Yes
Data Retention Utility			
Volume with access attribute	Yes	Yes ⁴	No
Volume Migration			
Source volume	No (VSP Gx00 models and VSP Fx00 models) Yes (VSP 5000 series)	No (VSP Gx00 models and VSP Fx00 models) Yes (VSP 5000 series)	No
Target volume	No	No	No
Cache Residency Manager (VSP 5000 series)			

Volume type	Used as GAD P-VOL?	Used as GAD S-VOL?	Used as quorum disk?
The volume on which Cache Residency Manager is set	No	No	No
Hitachi Virtual LUN			
Virtual LUN volume	Yes	Yes	Yes ⁵
LUN Manager			
The volume on which paths are defined	Yes	Yes	No
Volume on which paths are not defined	No	No	Yes
CCI command device			
Command device	No	No	No
Remote command device	No	No	No
Encryption License Key			
Volume whose parity groups have been encrypted	Yes	Yes	You can use an encrypted volume in the external storage system as a quorum disk. ⁶
FMD Encryption License Key			
Volume whose parity groups have been encrypted	Yes	Yes	You can use an encrypted volume in the external storage system as a quorum disk. ⁶
Nondisruptive migration			

Volume type	Used as GAD P-VOL?	Used as GAD S-VOL?	Used as quorum disk?
Volume which is being migrated	Yes	No	No
Notes: <ol style="list-style-type: none"> 1. A DP-VOL that uses an external volume as its pool volume can be used as a GAD P-VOL or S-VOL. 2. For the node volume or the leaf volume of Thin Image, see the description of the S-VOL, not the P-VOL. 3. GAD S-VOL is used as UR delta resync pair P-VOL. 4. If you set the S-VOL Disable attribute of Data Retention Utility to the GAD S-VOL, GAD pair operations using CCI are restricted. Release the S-VOL Disable attribute on the GAD S-VOL, and then perform the GAD pair operations. 5. Quorum disks can be set only on external volumes that have been configured so that one external volume group is mapped to one external volume. 6. You cannot encrypt a nonencrypted quorum disk in the external storage system from the primary or secondary storage system. 			

Dynamic Provisioning / Dynamic Tiering / Active flash

Dynamic Provisioning, Dynamic Tiering, and active flash virtual volumes (DP-VOLs) can be used as GAD pair volumes.

A V-VOL with capacity saving enabled can be used as a P-VOL or S-VOL of a GAD pair. A deduplication system data volume cannot be used as a P-VOL or S-VOL of a GAD pair.




Note:

- Data compressed or deduplicated by the capacity saving function is copied to a volume after compression and deduplication are released, that is, the capacity saving function is not performed immediately for copied data. Therefore, before creating or resynchronizing a GAD pair, make sure that the available capacity in the copy destination volume is greater than the used capacity in the copy origination volume before capacity saving. For details, see the *Provisioning Guide for Open Systems*.
- If you create a GAD pair using a volume for which the capacity saving function is used, compressed or deduplicated data is copied. Because of this, copy or I/O performance might be degraded.
- When the capacity saving function is used, management information is stored in a pool. As a result, there might be a difference in the number of used pages or licensed capacity between a P-VOL and an S-VOL.


ShadowImage

You can use the GAD P-VOL and S-VOL as a ShadowImage P-VOL.


You can create up to three ShadowImage pairs respectively on the GAD primary and secondary storage systems.



Because the server recognizes a GAD pair as one volume, it sees the volume as paired with six ShadowImage volumes.



You can create three additional, cascaded SI pairs using the SI S-VOLs. This means that up to nine SI pairs can be created with the GAD P-VOL, and nine SI pairs can be created with the GAD S-VOL.



Note:

- Pairs in an SI consistency group must reside in the same storage system. Because of this, the SI pairs that are associated with both the GAD P-VOL and the S-VOL cannot be registered to the same consistency group.
- When you use GAD pair volumes to create an SI pair, you must specify the physical LDEV ID, not the virtual LDEV ID.

Limitations when sharing GAD and ShadowImage volumes

Any operation that deletes the virtual LDEV ID of a volume used as a ShadowImage volume cannot be performed.

When a GAD pair is deleted with the P-VOL specified, the virtual LDEV ID of the S-VOL is deleted. If you delete the pair with the S-VOL specified, the virtual LDEV ID of the P-VOL is deleted. When the virtual LDEV ID is deleted, the server does not recognize the volume.

SI operations and GAD pair status

The ability to perform a ShadowImage pair operation depends on the SI pair status and GAD pair status.

The following tables show SI pair operations and whether they can be performed (Yes, No) with the listed GAD status. The information assumes the required SI status for the operation.

The Virtual LDEV ID column shows whether the volume has a virtual LDEV ID or not (Yes, No).

Table 1 SI operations when GAD status is Simplex

GAD pair status	Virtual LDEV ID	I/O		ShadowImage pair operation				
		Read	Write	Create pairs	Split pairs	Resync pairs	Restore pairs	Delete pairs/ Suspend copy
SMPL	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
	No	No	No	Yes	Yes	Yes	Yes	Yes
	No, but the GAD reserve attribute is set	No	No	No	No	No	No	Yes

Table 2 SI operations when GAD status is Mirroring

GAD pair status	I/O mode	Pair location	I/O		ShadowImage pair operation				
			Read	Write	Create pairs	Split pairs	Resync pairs	Restore pairs	Delete pairs/ Suspend copy
INIT/ COPY	Mirror (RL)	Primary	Yes	Yes	Yes	Yes	Yes	No ¹	Yes
	Block	Secondary	No	No	Yes	No ²	No ²	No ^{1, 3}	Yes
COPY	Mirror (RL)	Primary	Yes	Yes	Yes	Yes	Yes	No ¹	Yes
	Block	Secondary	No	No	Yes	No ²	No ²	No ^{1, 3}	Yes

Notes:

1. Cannot be used because GAD pairs are not suspended.
2. Cannot be used because S-VOL data is not fixed.
3. Cannot be used because the volume at the GAD copy destination is the same as the volume at the ShadowImage copy destination.

Table 3 ShadowImage operations when GAD status is Mirrored

GAD pair status	I/O mode	Pair location	I/O		ShadowImage pair operation				
			Read	Write	Create pairs	Split pairs	Resync pairs	Restore pairs	Delete pairs/ Suspend copy
PAIR	Mirror (RL)	Primary	Yes	Yes	Yes	Yes	Yes	No*	Yes
		Secondary	Yes	Yes	Yes	Yes	Yes	No*	Yes
* Cannot be used because GAD pairs are not suspended, and also because the volume at the GAD copy destination is the same as the volume at the ShadowImage copy destination.									

Table 4 SI operations when status of GAD pairs created, resynchronized, or swap resynchronized is Quorum disk blocked

GAD pair status	I/O mode	Pair location	I/O		ShadowImage pair operation				
			Read	Write	Create pairs	Split pairs	Resync pairs	Restore pairs	Delete pairs/ Suspend copy
Microcode version 80-04-2x or earlier (VSP 5000 series) or firmware version 83-03-3x or earlier VSP Gx00 models									
PAIR	Mirror(RL)	Primary	Yes	Yes	Yes	Yes	Yes	No*	Yes
	Block	Secondary	No	No	Yes	Yes	Yes	No*	Yes
* Cannot be used because GAD pairs are not suspended, and also because the volume at the GAD copy destination is the same as the volume at the ShadowImage copy destination.									

Table 5 ShadowImage operations when GAD status is Suspended

GAD pair status	I/O mode	Pair location	I/O		ShadowImage pair operation				
			Read	Write	Create pairs	Split pairs	Resync pairs	Restore pairs	Delete pairs/ Suspend copy
PSUS	Local	Primary	Yes	Yes	Yes	Yes	Yes	Yes*	Yes

GAD pair status	I/O mode	Pair location	I/O		ShadowImage pair operation				
			Read	Write	Create pairs	Split pairs	Resync pairs	Restore pairs	Delete pairs/ Suspend copy
	Block	Primary	No	No	Yes	Yes	Yes	No	Yes
PSUE	Local	Primary	Yes	Yes	Yes	Yes	Yes	Yes*	Yes
	Block	Primary	No	No	Yes	Yes	Yes	No	Yes
		Secondary	No	No	Yes	Yes	Yes	No	Yes
SSUS	Block	Secondary	No	No	Yes	Yes	Yes	No	Yes
SSWS	Local	Secondary	Yes	Yes	Yes	Yes	Yes	Yes*	Yes
* Quick Restore cannot be executed. For VSP 5000 series, if the microcode is 80-05-7x or later, Quick Restore can be executed.									

Table 6 ShadowImage operations when GAD status is Blocked

GAD pair status	I/O mode	Pair location	I/O		SI pair operations				
			Read	Write	Create pairs	Split pairs	Resync pairs	Restore pairs	Delete pairs/ Suspend copy
PSUE	Block	Primary	No	No	Yes	Yes	Yes	No	Yes
		Secondary	No	No	Yes	Yes	Yes	No	Yes

GAD operations and SI pair status

The ability to perform a GAD pair operation depends on GAD pair status and SI pair status.

The following tables show GAD operations and whether they can be performed (Yes, No) with the listed SI status. The information assumes the required GAD status for the operation.

Table 7 GAD operations and SI pair statuses, when GAD P-VOL is shared

SI pair status	GAD pair operations							
	Create pairs	Suspend pairs		Delete pairs			Resync pairs	
		P-VOL selected	S-VOL selected	P-VOL selected ¹	S-VOL selected ²	Forced deletion	P-VOL selected	S-VOL selected
SMPL(PD)	Yes	Yes	Yes	Yes	No ³	Yes	Yes	Yes
COPY	Yes	Yes	Yes	Yes	No ³	Yes	Yes	Yes
PAIR	Yes	Yes	Yes	Yes	No ³	Yes	Yes	Yes
COPY(SP)	Yes	Yes	Yes	Yes	No ³	Yes	Yes	Yes
PSUS(SP)	Yes	Yes	Yes	Yes	No ³	Yes	Yes	Yes
PSUS	Yes	Yes	Yes	Yes	No ³	Yes	Yes	Yes
COPY(RS)	Yes	Yes	Yes	Yes	No ³	Yes	Yes	Yes
COPY(RS-R)	No ⁴	impossible	impossible	Yes	No ³	Yes	No ⁴	No ⁴
PSUE	Yes	Yes	Yes	Yes	No ³	Yes	Yes	Yes
Notes: <ol style="list-style-type: none"> 1. You can delete a GAD pair by specifying the P-VOL, only when the I/O mode is Local and the GAD pair status of the P-VOL is PSUS or PSUE. 2. You can delete a GAD pair by specifying the S-VOL, only when the I/O mode is Local and the GAD pair status of the S-VOL is SSWS. 3. Cannot be used because, when you delete a GAD pair specifying the S-VOL, the P-VOL's virtual LDEV ID is also deleted, which makes it unusable as the SI P-VOL. 4. To continue SI restore copy, the GAD pairs must be suspended. 								

Table 8 GAD operations and SI pair statuses, when GAD S-VOL is shared

SI pair status	GAD pair operations							
	Create pairs	Suspend pairs		Delete pairs			Resync pairs	
		P-VOL selected	S-VOL selected	P-VOL selected ¹	S-VOL selected ²	Forced deletion	P-VOL selected	S-VOL selected
SMPL(PD)	No ³	Yes	Yes	No ⁴	Yes	Yes	Yes	Yes
COPY	No ³	Yes	Yes	No ⁴	Yes	Yes	Yes	Yes

SI pair status	GAD pair operations							
	Create pairs	Suspend pairs		Delete pairs			Resync pairs	
		P-VOL selected	S-VOL selected	P-VOL selected ¹	S-VOL selected ²	Forced deletion	P-VOL selected	S-VOL selected
PAIR	No ³	Yes	Yes	No ⁴	Yes	Yes	Yes	Yes
COPY(SP)	No ³	Yes	Yes	No ⁴	Yes	Yes	Yes	Yes
PSUS(SP)	No ³	Yes	Yes	No ⁴	Yes	Yes	Yes	Yes
PSUS	No ³	Yes	Yes	No ⁴	Yes	Yes	Yes	Yes
COPY(RS)	No ³	Yes	Yes	No ⁴	Yes	Yes	Yes	Yes
COPY(RS-R)	No ^{3, 5}	impossible	impossible	No ⁴	Yes	Yes	No ^{5, 6}	No ⁶
PSUE	No ³	Yes	Yes	No ⁴	Yes	Yes	Yes	Yes

Notes:


1. You can delete a GAD pair by specifying the P-VOL, only when the I/O mode is Local and the GAD pair status of the P-VOL is PSUS or PSUE.
2. You can delete a GAD pair by specifying the S-VOL, only when the I/O mode is Local and the GAD pair status of the S-VOL is SSWS.
3. When a GAD pair is created, the GAD reserve attribute is assigned to the volume that will become the S-VOL, which removes the virtual LDEV ID of this volume, making it unusable as an SI pair volume.

The GAD reserve attribute is set, and the virtual LDEV ID is deleted for the volume that will become the GAD S-VOL, making it unusable as an SI volume.
4. Cannot be used because, when you delete a GAD pair specifying the S-VOL, the P-VOL's virtual LDEV ID is also deleted, which makes it unusable as the SI P-VOL.
5. Cannot be used because the volume at the GAD copy destination is the same as the volume at the ShadowImage copy destination.
6. To continue ShadowImage restore copy, GAD pairs must be suspended.


Thin Image

You can use a GAD P-VOL or S-VOL as a Thin Image (HTI) P-VOL.

You can create up to 1,024 Thin Image pairs using a GAD P-VOL, and up to 1,024 Thin Image pairs using a GAD S-VOL.



Because the server recognizes the GAD pair as one volume, it sees the volume as paired with 2,048 HTI volumes.



Note:

- Pairs in an HTI consistency group and snapshot group must reside in the same storage system. Because of this, the HTI pairs that are associated with both the GAD P-VOL and S-VOL cannot be registered to the same consistency group or snapshot group.
- When you use GAD pair volumes to create a Thin Image pair, specify the physical LDEV ID, not the virtual LDEV ID.

Limitations for using both GAD and Thin Image

Any operation that deletes the virtual LDEV ID of a volume used as a Thin Image volume cannot be performed.

When a GAD pair is deleted with the P-VOL specified, the virtual S-VOL's LDEV ID is deleted. If you delete the pair with the S-VOL specified, the P-VOL's virtual LDEV ID is deleted. When the virtual LDEV ID is deleted, the server does not recognize the volume, making it unusable as a Thin Image volume.

Thin Image operations and GAD status

The ability to perform a Thin Image pair operation depends on the HTI pair status and the GAD pair status.

The following tables show HTI operations and whether they can be performed (Yes, No) with the listed GAD status. The information assumes the required HTI status for the operation.

The Virtual LDEV ID column shows whether the volume has a virtual LDEV ID or not (Yes, No).

Table 9 Thin Image operations when GAD status is Simplex

GAD pair status	Virtual LDEV ID	I/O		Thin Image pair operation				
		Read	Write	Create pairs	Split pairs	Resync pairs	Restore pairs	Delete pairs
SMPL	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
	No	No	No	Yes	Yes	Yes	Yes	Yes
	No, but the GAD reserve attribute is set	No	No	No	No	No	No	Yes

Table 10 Thin Image operations when GAD status is Mirroring

GAD pair status	I/O mode	Pair location	I/O		Thin Image pair operation				
			Read	Write	Create pairs	Split pairs	Resync pairs	Restore pairs	Delete pairs
INIT/ COPY	Mirror (RL)	Primary	Yes	Yes	Yes	Yes	Yes	No ¹	Yes
	Block	Secondary	No	No	No	No ²	No ²	No ^{1,3}	Yes
COPY	Mirror (RL)	Primary	Yes	Yes	Yes	Yes	Yes	No ¹	Yes
	Block	Secondary	No	No	No	No ²	No ²	No ^{1,3}	Yes

Notes:

1. Cannot be used because GAD pairs are not suspended.
2. Cannot be used because the data is being copied and the volume data is not fixed yet.
3. Cannot be used because the volume at the GAD copy destination is the same as the volume at the Thin Image copy destination.

Table 11 Thin Image operations when GAD status is Mirrored

GAD pair status	I/O mode	Pair location	I/O		Thin Image pair operation				
			Read	Write	Create pairs	Split pairs	Resync pairs	Restore pairs	Delete pairs
PAIR	Mirror (RL)	Primary	Yes	Yes	Yes	Yes	Yes	No*	Yes
		Secondary	Yes	Yes	Yes	Yes	Yes	No*	Yes
* Cannot be used because GAD pairs are not suspended, and also because the volume at the GAD copy destination is the same as the volume at the Thin Image copy destination.									

Table 12 Thin Image operations when status of GAD pairs created, resynchronized, or swap resynchronized is Quorum disk blocked

GAD pair status	I/O mode	Pair location	I/O		Thin Image pair operation				
			Read	Write	Create pairs	Split pairs	Resync pairs	Restore pairs	Delete pairs/ Suspended copy
Microcode version 80-04-2x or earlier (VSP 5000 series) or firmware version 83-03-3x or earlier (VSP Gx00 models)									
PAIR	Mirror (RL)	Primary	Yes	Yes	Yes	Yes	Yes	No*	Yes
	Block	Secondary	No	No	Yes	Yes	Yes	No*	Yes
Microcode version 80-05-0x or later (VSP 5000 series), firmware version 88-01-0x or later (VSP G350, G370, G700, G900 and VSP F350, F370, F700, F900), firmware version 83-04-0x or later (VSP Gx00 models), or 83-04-2x or later (VSP Fx00 models)									
PAIR	Mirror (RL)	Primary	Yes	Yes	Yes	Yes	Yes	No*	Yes
	Mirror (RL)	Secondary	Yes	Yes	Yes	Yes	Yes	No*	Yes
* Cannot be used because GAD pairs are not suspended, and also because the volume at the GAD copy destination is the same as the volume at the Thin Image copy destination.									

Table 13 Thin Image operations when GAD status is Suspended

GAD pair status	I/O mode	Pair location	I/O		Thin Image pair operation				
			Read	Write	Create pairs	Split pairs	Resync pairs	Restore pairs	Delete pairs
PSUS	Local	Primary	Yes	Yes	Yes	Yes	Yes	Yes	Yes
	Block	Primary	No	No	Yes	Yes	Yes	No	Yes
PSUE	Local	Primary	Yes	Yes	Yes	Yes	Yes	Yes	Yes
	Block	Primary	No	No	Yes	Yes	Yes	No	Yes
		Secondary	No	No	Yes	Yes	Yes	No	Yes
SSUS	Block	Secondary	No	No	Yes	Yes	Yes	No	Yes
SSWS	Local	Secondary	Yes	Yes	Yes	Yes	Yes	No	Yes

Table 14 Thin Image operations when GAD status is Blocked

GAD pair status	I/O mode	Pair location	I/O		Thin Image pair operation				
			Read	Write	Create pairs	Split pairs	Resync pairs	Restore pairs	Delete pairs
PSUE	Block	Primary	No	No	Yes	Yes	Yes	No	Yes
		Secondary	No	No	Yes	Yes	Yes	No	Yes

GAD operations and Thin Image pair status

The ability to perform a GAD pair operation depends on the GAD pair status and the HTI pair status.

The following tables show GAD operations and whether they can be performed (Yes, No) with the listed HTI status. The information assumes the required GAD status for the operation.

Table 15 GAD operations and HTI pair status, when the GAD P-VOL is shared

TI pair status	GAD pair operations							
	Create GAD Pairs	Suspend Pairs		Delete Pairs			Resync Pairs	
		P-VOL specified	S-VOL specified	P-VOL specified ¹	S-VOL specified ²	Forced deletion	P-VOL specified	S-VOL specified
SMPL(PD)	Yes	Yes	Yes	Yes	No ³	Yes	Yes	Yes
COPY	Yes	Yes	Yes	Yes	No ³	Yes	Yes	Yes
PAIR	Yes	Yes	Yes	Yes	No ³	Yes	Yes	Yes
PSUS	Yes	Yes	Yes	Yes	No ³	Yes	Yes	Yes
COPY(RS-R)	No ⁴	impossible	impossible	Yes	No ³	Yes	No ⁴	No ⁴
PSUE	Yes	Yes	Yes	Yes	No ³	Yes	Yes	Yes
Notes: <ol style="list-style-type: none"> 1. You can delete a GAD pair by specifying the P-VOL, only when the I/O mode is Local and the GAD pair status of the P-VOL is PSUS or PSUE. 2. You can delete a GAD pair by specifying the S-VOL, only when the I/O mode is Local and the GAD pair status of the S-VOL is SSWS. 3. Cannot be used because, when you delete a GAD pair specifying the S-VOL, the P-VOL's virtual LDEV ID is also deleted, which makes it unusable as the HTI P-VOL. 4. To continue resynchronizing the HTI pair, you must split the GAD pair. 								

Table 16 GAD operations and HTI pair status, when the GAD S-VOL is shared

TI pair status	GAD pair operations							
	Create GAD Pairs	Suspend Pairs		Delete Pairs			Resync Pairs	
		P-VOL specified	S-VOL specified	P-VOL specified ¹	S-VOL specified ²	Forced deletion	P-VOL specified	S-VOL specified
SMPL(PD)	No ³	Yes	Yes	No ⁴	Yes	Yes	Yes	Yes
COPY	No ³	Yes	Yes	No ⁴	Yes	Yes	Yes	Yes
PAIR	No ³	Yes	Yes	No ⁴	Yes	Yes	Yes	Yes
PSUS	No ³	Yes	Yes	No ⁴	Yes	Yes	Yes	Yes
COPY(RS-R)	No ^{3, 5}	No	No	No ⁴	Yes	Yes	No ^{5, 6}	No ⁶

TI pair status	GAD pair operations							
	Create GAD Pairs	Suspend Pairs		Delete Pairs			Resync Pairs	
		P-VOL specified	S-VOL specified	P-VOL specified ¹	S-VOL specified ²	Forced deletion	P-VOL specified	S-VOL specified
PSUE	No ³	Yes	Yes	No ⁴	Yes	Yes	Yes	Yes

Notes:

1. You can delete a GAD pair by specifying the primary volume, only when the I/O mode is Local and the GAD pair status of the primary volume is PSUS or PSUE.
2. You can delete a GAD pair by specifying the secondary volume, only when the I/O mode is Local and the GAD pair status of the secondary volume is SSWS.
3. To create a GAD pair, you must assign the GAD reserve attribute to the volume used as an S-VOL. Because the virtual LDEV ID of the volume to which the GAD reserve attribute is assigned is deleted, you cannot create a GAD pair by specifying the volume shared with HTI as the S-VOL of the pair.
4. Cannot be used because, when you delete a GAD PAIR specifying the P-VOL, the S-VOL's virtual LDEV ID is also deleted, which makes it unusable as an HTI P-VOL.
5. Cannot be used because the GAD pair's target volume is the same as the HTI pair's target volume.
6. To continue resynchronizing the HTI pair, you must split the GAD pair.

Use cases for pairing GAD volumes with SI or HTI

Backing up GAD pair volumes with ShadowImage (SI) or Thin Image (HTI) provides further protection for GAD data.

Further protection for GAD data is provided in the following ways:

- When the GAD pair is resynchronized, pair status changes to COPY. While in this status, S-VOL consistency is temporarily lost. You can protect data when in COPY status by pairing the S-VOL with SI or HTI before resynchronizing the GAD pair.
- Though data in a blocked GAD pair is inconsistent, host activity can continue with the P-VOL or S-VOL. Therefore, before correcting the failure by forcibly deleting the pair, you should pair the volumes with SI or HTI.
- The SI and HTI pairs can then be copied, and the copies used for other purposes.

Universal Replicator

You can combine GAD and Universal Replicator to create a configuration that can continue to operate in the event of a multi-site failure.

In a GAD system, the server accesses the primary and secondary sites simultaneously and shares the same data between the two sites (at campus distance). If a failure occurs at one site, you can continue operations at the other site. However, if a failure occurs at both sites, for example due to a large-scale disaster, you will not be able to continue operations with the data redundancy provided only by GAD.

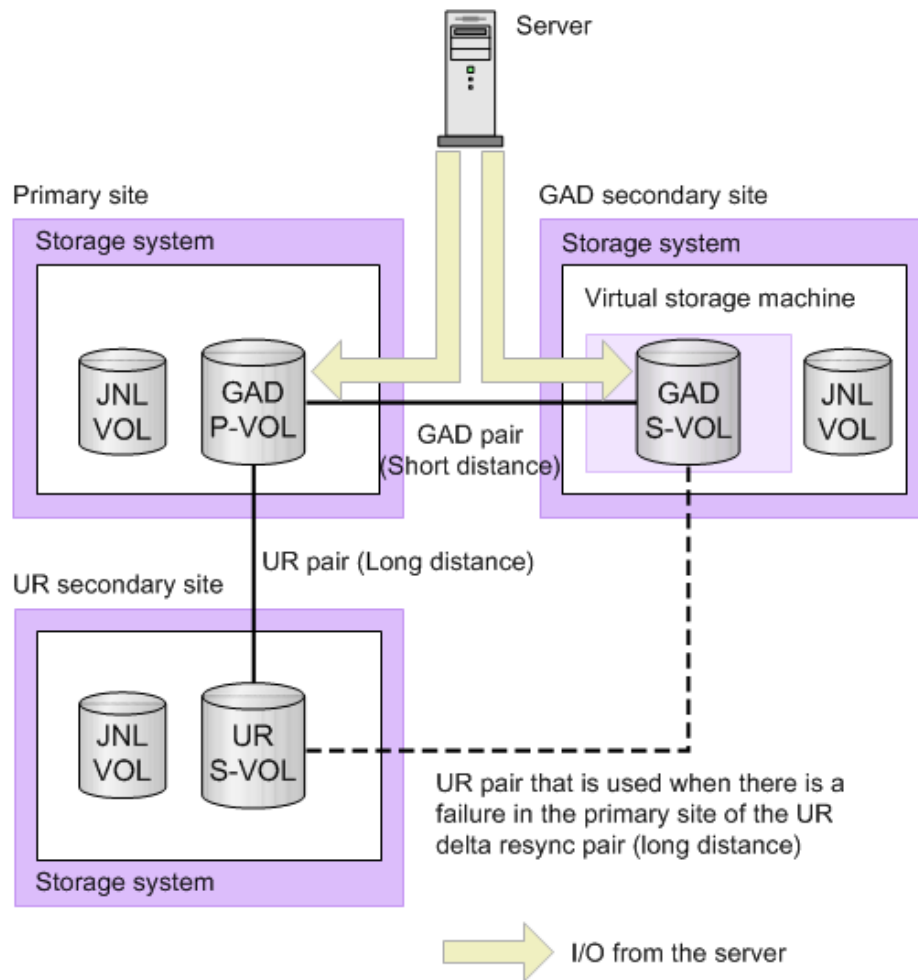
To manage this situation, you can implement a 3-data-center (3DC) configuration by combining GAD and Universal Replicator (UR). This configuration is called a GAD 3DC delta resync (GAD+UR) configuration. If there is a failure at both the primary site and the GAD secondary site, the GAD+UR configuration enables you to continue operations using the UR secondary site (at metro distance).



Note:

- When all storage systems in a GAD+UR configuration are VSP 5000 series, the DKCMAIN microcode version must be 80-02-4-xx/xx or later. If both VSP G1000, VSP G1500, VSP F1500 and VSP G800 or VSP F800 exist, the DKCMAIN firmware version for VSP 5000 series must be 80-04-xx-xx/xx or later; and for VSP G800, the DKCMAIN firmware version must be 83-03-xx-xx/xx or later, or for VSP F800, the DKCMAIN firmware version must be 83-04-2x-xx/xx or later. You cannot combine GAD and UR using other storage system models.
- The CCI remote command device is not required in GAD+UR configurations.
- Volumes in a 3DC GAD+UR configuration cannot be shared with the following volumes:
 - UR volumes in a system with multiple storage systems in the primary and secondary sites
 - UR volumes in a 3DC multi-target configuration with TC and UR
 - UR volumes in a 3DC multi-target configuration with 3 UR sites
 - UR volumes in a 3DC cascade configuration with 3 UR sites
- 3DC delta resync configuration with GAD and UR are supported only by VSP G/F900, VSP 5000 series.
- 3DC delta resync configuration with GAD and UR is supported only by VSP F370, VSP F700, VSP F900, VSP G370, VSP G700, VSP G900, VSP 5000 series.

The following figure shows the required configuration for GAD+UR operations. You must use this configuration when combining GAD with UR.



In a GAD+UR configuration:

- The P-VOL of the GAD pair functions as the P-VOL of the UR pair.
- The S-VOL of the GAD pair functions as the P-VOL of the UR delta resync pair.
- The UR S-VOL has two mirror IDs, one for the UR pair, and one for the UR delta resync pair.
- The UR delta resync pair consists of the GAD S-VOL at the GAD secondary site and the UR S-VOL at the UR secondary site. This UR delta resync pair allows you to synchronize the UR pair S-VOL by using the journal data at the GAD secondary site in the event of a failure at the primary site.
- When combining GAD with UR, create the GAD pair first, and then create the UR pair. The GAD pair cannot be used with a UR pair that was created first.

GAD pair and Universal Replicator pair interoperability

You should know how GAD pairs and Universal Replicator pairs work together for different pair statuses for GAD pairs or UR pairs.

The following table shows GAD pair operations and whether they can be performed (Yes or No) with the listed UR pair status.

Table 17 GAD operations and UR pair statuses

UR pair status	Attribute of the target UR pair	GAD pair operation					
		Create pairs ¹	Split pairs	Suspend pairs	Resynchronize pairs		Delete pairs
			P-VOL specified	S-VOL specified	P-VOL specified	S-VOL specified	P-VOL or S-VOL specified
COPY	P-VOL	No	Yes	impossible	Yes	Yes	No ²
	S-VOL	No	No ³	impossible	No	impossible	No ²
PAIR	P-VOL	No	Yes	impossible	Yes	Yes	No ²
	S-VOL	No	No ³	impossible	No	impossible	No ²
PSUS	P-VOL	No	Yes	impossible	Yes	Yes	No ²
PSUE	P-VOL	No	Yes	impossible	Yes	Yes	No ²
	S-VOL	No	No ³	impossible	No	impossible	No ²
SSUS	S-VOL	No	No ³	impossible	No	impossible	No ²
SSWS	S-VOL	No	No ³	impossible	No	impossible	No ²
HOLD	P-VOL	No	impossible	Yes	impossible	impossible	No ²
	S-VOL	No	impossible	impossible	impossible	impossible	No ²
HLDE	P-VOL	No	impossible	Yes	impossible	impossible	No ²
	S-VOL	No	impossible	impossible	impossible	impossible	No ²
Note: 1. When creating a GAD and UR configuration, create a GAD pair first.							

UR pair status	Attribute of the target UR pair	GAD pair operation					
		Create pairs ¹	Split pairs	Suspend pairs	Resynchronize pairs		Delete pairs
			P-VOL specified	S-VOL specified	P-VOL specified	S-VOL specified	P-VOL or S-VOL specified
<p>2. When deleting a GAD pair, delete the UR pair and the UR delta resync pair first.</p> <p>3. The UR S-VOL is also used as the GAD P-VOL only when the UR S-VOL is duplicated using GAD. In this case, the GAD P-VOL status must be PSUS, therefore, the pair split terminates abnormally.</p>							

The following table shows UR pair operations and whether they can be performed (Yes or No) with the listed GAD pair status.

Table 18 UR pair operations and GAD pair statuses

GAD pair status	I/O mode	Attribute of the target GAD pair	UR pair operation				
			Create pairs	Split pairs	Suspend pairs	Resynchronize pairs	
				P-VOL specified	S-VOL specified	P-VOL specified	S-VOL specified
INIT/COPY	Mirror (RL)	P-VOL	No	Yes	impossible	Yes ¹	impossible
	Block	S-VOL	No	impossible	impossible	impossible	impossible
COPY	Mirror (RL)	P-VOL	No	Yes	impossible	Yes ¹	impossible
	Block	S-VOL	No	impossible	impossible	impossible	impossible
PAIR	Mirror (RL)	P-VOL	Yes ²	Yes	impossible	Yes ¹	impossible
		S-VOL	No	impossible	impossible	impossible	impossible
PSUS	Local	P-VOL	No	Yes	Yes ³	Yes ¹	Yes ³
	Block	P-VOL	No	Yes	Yes ³	No	impossible

GAD pair status	I/O mode	Attribute of the target GAD pair	UR pair operation				
			Create pairs	Split pairs	Suspend pairs	Resynchronize pairs	
				P-VOL specified	S-VOL specified	P-VOL specified	S-VOL specified
PSUE	Local	P-VOL	No	Yes	Yes ³	Yes ¹	Yes ³
	Block	P-VOL	No	Yes	Yes ³	No	impossible
SSUS	Block	S-VOL	No	impossible	impossible	impossible	impossible
SSWS	Local	S-VOL	No	Yes	impossible	Yes	impossible
Note: <ol style="list-style-type: none"> 1. Can be performed only when the UR delta resync pair status is HLDE. 2. GAD pairs can be used with UR pairs only when the volume specified as the UR S-VOL is a UR delta resync S-VOL. 3. The UR S-VOL is also used as the GAD P-VOL only when the UR S-VOL is duplicated using GAD. 							

The following table shows UR delta resync pair operations and whether they can be performed (Yes or No) with the listed GAD pair status.

Table 19 UR delta resync pair operations and GAD pair statuses

GAD pair status	I/O mode	Attribute of the target GAD pair	UR delta resync pair operation					
			Create UR delta resync pair	Delta resync	Delete pairs			
					P-VOL specified	S-VOL specified	UR delta resync P-VOL	UR delta resync S-VOL
INIT/COPY	Mirror (RL)	P-VOL	No	No ¹	Yes	impossible	impossible	impossible
	Block	S-VOL	No	No ¹	impossible	impossible	impossible	impossible
COPY	Mirror (RL)	P-VOL	No	No ¹	Yes	impossible	impossible	impossible

GAD pair status	I/O mode	Attribute of the target GAD pair	UR delta resync pair operation					
			Create UR delta resync pair	Delta resync	Delete pairs			
					P-VOL specified	S-VOL specified	UR delta resync P-VOL	UR delta resync S-VOL
	Block	S-VOL	No	No ¹	impossible	impossible	impossible	impossible
PAIR	Mirror (RL)	P-VOL	No	No ¹	Yes	impossible	impossible	impossible
		S-VOL	Yes	No ¹	impossible	impossible	Yes ⁴	impossible
PSUS	Local	P-VOL	No	No ¹	Yes	Yes ³	Yes ⁴	impossible
	Block	P-VOL	Yes ²	No ¹	Yes	Yes ³	Yes ⁴	impossible
PSUE	Local	P-VOL	No	No ¹	Yes	Yes ³	Yes ⁴	impossible
	Block	P-VOL	Yes ²	No ¹	Yes	Yes ³	Yes ⁴	impossible
SSUS	Block	S-VOL	No	No ¹	impossible	impossible	Yes ⁴	impossible
SSWS	Local	S-VOL	No	No ¹	Yes	impossible	Yes ⁴	impossible
Note: <ol style="list-style-type: none"> 1. Cannot be performed by users. The storage system performs the operation automatically. 2. If a failure occurs at the primary site, the operation can be performed only when the volume specified as the UR delta resync S-VOL is the UR S-VOL. 3. The UR S-VOL is also used as the GAD P-VOL only when the UR S-VOL is duplicated using GAD. 4. Deleting a UR delta resync pair deletes the UR pair. 								

Data Retention Utility

You can create a GAD pair using volumes that have been assigned the Data Retention Utility access attribute.

- When you create or resynchronize a GAD pair, the access attribute set for the P-VOL is copied to the S-VOL.
- If you change the access attribute when GAD status is Mirrored or Mirroring, make sure to set the access attribute to both the P-VOL and S-VOLs.
- Server I/O can be controlled, depending on GAD status and the access attribute.
- If you set the Data Retention Utility S-VOL Disable attribute on the GAD S-VOL, GAD pair operations using CCI are restricted. Release the S-VOL Disable attribute from the S-VOL, then perform CCI operations.

GAD status and I/O allowance by access attribute

Even when the access attribute is assigned to a GAD volume, the initial copy and pair resynchronization operations are not controlled. The following table shows whether server I/O is allowed for the listed GAD status and access attribute.

GAD statuses	Access attribute		I/O	
	P-VOL	S-VOL	P-VOL	S-VOL
Mirrored	Read/Write	Read/Write	Ends normally	Ends normally
	Read Only or Protect	Read/Write	Depends on the attribute*	Ends normally
	Read/Write	Read Only or Protect	Ends normally	Depends on the attribute*
	Read Only or Protect	Read Only or Protect	Depends on the attribute*	Depends on the attribute*
Quorum disk blocked (status of GAD pairs created, resynchronized, or swap resynchronized on 80-04-2x or earlier for VSP 5000 series and 83-03-3x or earlier for VSP Gx00 models)	Read/Write	Read/Write	Ends normally	Rejected
	Read Only or Protect	Read/Write	Depends on the attribute*	Rejected
	Read/Write	Read Only or Protect	Ends normally	Rejected
	Read Only or Protect	Read Only or Protect	Depends on the attribute*	Rejected
Quorum disk blocked (status of	Read/Write	Read/Write	Ends normally	Ends normally

GAD statuses	Access attribute		I/O	
	P-VOL	S-VOL	P-VOL	S-VOL
GAD pairs created, resynchronized, or swap resynchronized on 80-05-0x or later for VSP 5000 series, 88-01-0x or later for VSP G350, G370, G700, G900 and VSP F350, F370, F700, F900	Read Only or Protect	Read/Write	Depends on the attribute*	Ends normally
	Read/Write	Read Only or Protect	Ends normally	Depends on the attribute*
	Read Only or Protect	Read Only or Protect	Depends on the attribute*	Depends on the attribute*
Mirroring Suspended (when the I/O mode of the primary volume is Local and the I/O mode of the secondary volume is Block)	Read/Write	Read/Write	Ends normally	Rejected
	Read Only or Protect	Read/Write	Depends on the attribute*	Rejected
	Read/Write	Read Only or Protect	Ends normally	Rejected
	Read Only or Protect	Read Only or Protect	Depends on the attribute*	Rejected
Suspended (when the I/O mode of the primary volume is Block and the I/O mode of the secondary volume is Local)	Read/Write	Read/Write	Rejected	Ends normally
	Read Only or Protect	Read/Write	Rejected	Ends normally
	Read/Write	Read Only or Protect	Rejected	Depends on the attribute*
	Read Only or Protect	Read Only or Protect	Rejected	Depends on the attribute*
Block	Read/Write	Read/Write	Rejected	Rejected
	Read Only or Protect	Read/Write	Rejected	Rejected
	Read/Write	Read Only or Protect	Rejected	Rejected
	Read Only or Protect	Read Only or Protect	Rejected	Rejected
* If the attribute is Read Only, Read is allowed but not Write. If the attribute is Protect, Read and Write are not allowed.				

Volume Migration (VSP 5000 series)

You can use Volume Migration to move volumes in an overloaded drive to a non-overloaded drive online.

Volume Migration moves volumes by specifying the P-VOL and S-VOL of a GAD pair in an overloaded drive as the migration source volume which then moves the volumes to a non-overloaded drive.

Restrictions for using GAD with Volume Migration

There are restrictions of which you should be aware before you use Volume Migration to move a GAD pair volume.

- The GAD pair must be split before you specify the volumes as the migration source volumes for Volume Migration.
- Provisioning types of the GAD P-VOL and S-VOL must be the same. Ensure that the provisioning types of the GAD P-VOL and S-VOL are still the same after the migration by Volume Migration.
- When a GAD pair volume is shared by an SI pair volume, you cannot specify it as the source volume during Quick Restore of the SI pair. When Quick Restore completes, start using Volume Migration.
- For Volume Migration operation when using GAD with nondisruptive migration, see [GAD status and nondisruptive migration pair operations \(on page 91\)](#).
- When combining GAD with Volume Migration, create the GAD pair first, and then operate Volume Migration. The GAD pair cannot be paired with the volumes for which Volume Migration is being operated.

GAD status and Volume Migration pair operations

The ability of Volume Migration to create or cancel migration plans depends on the status of the GAD pair.

The following table describes if you can create or cancel migration plans when GAD volumes are not mirrored.

GAD pair status	Virtual LDEV ID	I/O from the server		Volume Migration operation	
		Read	Write	Create migration plans	Cancel migration plans
SMPL	Available	Y	Y	Y	Y
	Not available	N	N	N	Y

GAD pair status	Virtual LDEV ID	I/O from the server		Volume Migration operation	
		Read	Write	Create migration plans	Cancel migration plans
	Not available (Virtual attribute: GAD reserve)	N	N	N	Y
Legend Y: Can be performed N: Cannot be performed					

The following table describes if you can create or cancel migration plans when GAD volumes are being mirrored.

GAD pair status	I/O mode	Pair location	I/O from the server		Volume Migration operation	
			Read	Write	Create migration plans	Cancel migration plans
INIT/COPY	Mirror(RL)	Primary	Y	Y	N	Y
	Block	Secondary	N	N	N	Y
Legend Y: Can be performed N: Cannot be performed						

The following table describes if you can create or cancel migration plans when GAD volumes are mirrored.

GAD pair status	I/O mode	Pair location	I/O from the server		Volume Migration operation	
			Read	Write	Create migration plans	Cancel migration plans
PAIR	Mirror(RL)	Primary	Y	Y	N	Y
		Secondary	Y	Y	N	Y
Legend						
Y: Can be performed						
N: Cannot be performed						

The following table describes if you can create or cancel migration plans when the GAD pair is suspended.

GAD pair status	I/O mode	Pair location	I/O from the server		Volume Migration operation	
			Read	Write	Create migration plans	Cancel migration plans
PSUS/ PSUE	Local	Primary	Y	Y	Y	Y
	Block	Any	N	N	Y	Y
SSUS	Block	Secondary	N	N	Y	Y
SSWS	Local	Secondary	Y	Y	Y	Y
Legend Y: Can be performed N: Cannot be performed						

The following table describes if you can create or cancel migration plans when the GAD pair is blocked.

GAD pair status	I/O mode	Pair location	I/O from the server		Volume Migration operation	
			Read	Write	Create migration plans	Cancel migration plans
PSUE	Block	Primary	N	N	Y	Y
		Secondary	N	N	Y	Y
Legend						
Y: Can be performed						
N: Cannot be performed						

Volume Migration pair statuses and GAD pair operations (VSP 5000 series)

You should understand what GAD operations you can perform given the Volume Migration pair status and whether the GAD P-VOL or S-VOL is shared by Volume Migration.

The following table describes operations you can perform according to the Volume Migration pair status when the GAD P-VOL is shared by Volume Migration.

Volume Migration pair status	GAD pair operation							
	Create	Suspend		Delete			Resynchronize	
		P-VOL specified	S-VOL specified	P-VOL specified ¹	S-VOL specified ²	Force delete	P-VOL specified	S-VOL specified
SMPL(PD)	No ⁴	Yes ⁵	Yes ⁵	Yes	No ³	Yes	No	No
COPY	No ⁴	Yes ⁵	Yes ⁵	Yes	No ³	Yes	No	No
PSUS	No ⁴	Yes ⁵	Yes ⁵	Yes	No ³	Yes	No	No
PSUE	No ⁴	Yes ⁵	Yes ⁵	Yes	No ³	Yes	No	No

Note:

1. You can delete a GAD pair by specifying the P-VOL, only when the I/O mode is Local and the GAD pair status of the P-VOL is PSUS or PSUE.
2. You can delete a GAD pair by specifying the S-VOL, only when the I/O mode is Local and the GAD pair status of the S-VOL is SSWS.
3. Cannot be performed because the virtual LDEV ID of the GAD P-VOL (volume for Volume Migration) is deleted if you specify the S-VOL to delete the GAD pair.
4. To perform Volume Migration, you must suspend the GAD pair.

Volume Migrati on pair status	GAD pair operation							
	Create	Suspend		Delete			Resynchronize	
		P-VOL specifi ed	S-VOL specifi ed	P-VOL specified 1	S-VOL specified 2	Force delete	P-VOL specifi ed	S-VOL specifi ed
5. To perform Volume Migration, you must suspend the GAD pair. Therefore, you can suspend the GAD pair, however the status of the GAD pair does not change because the GAD pair is already suspended.								

The following table describes operations you can perform according to the Volume Migration pair status when the GAD S-VOL is shared by Volume Migration.

Volume Migration pair status	GAD pair operation							
	Create	Suspend		Delete			Resynchronize	
		P-VOL specified	S-VOL specified	P-VOL specified ¹	S-VOL specified ²	Force delete	P-VOL specified	S-VOL specified
SMPL(PD)	No ³	Yes ⁵	Yes ⁵	No ⁴	Yes	Yes	No	No
COPY	No ³	Yes ⁵	Yes ⁵	No ⁴	Yes	Yes	No	No
PSUS	No ³	Yes ⁵	Yes ⁵	No ⁴	Yes	Yes	No	No
PSUE	No ³	Yes ⁵	Yes ⁵	No ⁴	Yes	Yes	No	No
Note: <ol style="list-style-type: none"> 1. You can delete a GAD pair by specifying the primary volume, only when the I/O mode is Local and the GAD pair status of the primary volume is PSUS or PSUE. 2. You can delete a GAD pair by specifying the secondary volume, only when the I/O mode is Local and the GAD pair status of the secondary volume is SSWS. 3. To create a GAD pair, you must assign the GAD reserve attribute to the volume used as an S-VOL. Because the virtual LDEV ID of the volume to which the GAD reserve attribute is assigned is deleted, you cannot create a GAD pair by specifying the volume used for Volume Migration as the S-VOL of the pair. 4. Cannot be performed because the virtual LDEV ID of the GAD S-VOL (volume for Volume Migration) is deleted if you specify the P-VOL to delete the GAD pair. 5. To perform Volume Migration, you must suspend the GAD pair. Therefore, you can suspend the GAD pair, however the status of the GAD pair does not change because the GAD pair is already suspended. 								

LUN Manager

Use the volumes for which LU paths have been set to create a GAD pair. You can add LU paths to or delete LU paths from GAD pair volumes. However, you cannot delete the last LU path because at least one LU path must be set for GAD pair volumes.

A volume for which no LU path has been set cannot be used as a GAD pair volume.



Caution: When you remove the path that is defined on an LDEV with the GAD reserve attribute, the path removal might fail if the number of LDEVs whose path is to be removed is too large.

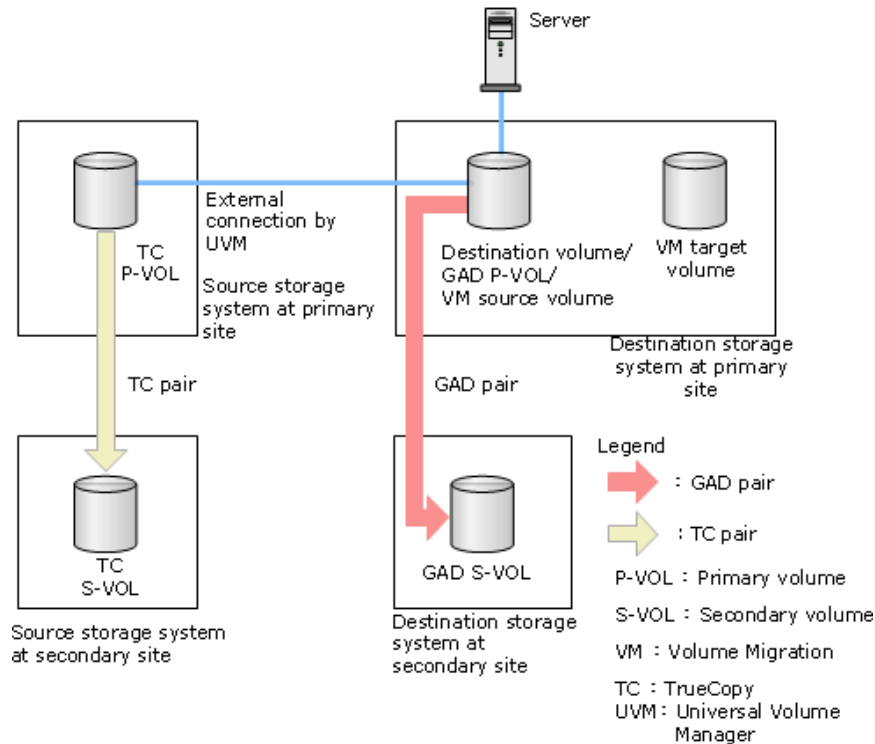
When a port has a path defined for an LDEV with the GAD reserve attribute and you need to configure the port (delete a host group, set a command device, or change the host mode, topology, AL-PL, or transfer speed), the operation might fail if you configure more than one port at a time. For ports that have a path defined on an LDEV with the GAD reserve attribute, perform these operations on one port at a time.

Virtual Partition Manager

GAD pair volumes and quorum disks can migrate across CLPRs.

Nondisruptive migration (VSP F1500 and VSP G1x00)

You can create a GAD pair using a volume being migrated by nondisruptive migration to migrate your data without changing system configuration, even if your system is designed for disaster recovery. The following figure shows an example of such data migration. For example, If you create a GAD pair using a volume being migrated while maintaining the configuration of the disaster recovery system using TrueCopy, you can maintain the disaster recovery system by GAD soon after migration.



For details, see the *Nondisruptive Migration User Guide*.

Restrictions for using GAD with nondisruptive migration

You should be aware of restrictions if you want to use GAD with nondisruptive migration.

- When creating a GAD pair, you cannot specify a volume being migrated by nondisruptive migration as its S-VOL.
- Before creating a GAD pair, you need to change the cache mode of the destination volume to Write Sync. Use CCI when changing the cache mode.
- I/O operations from the host to the GAD S-VOL cannot be performed until Volume Migration is completed. If you want to perform the operation which allows the host to recognize the GAD S-VOL, perform this operation after Volume Migration is completed.
- You cannot swap resynchronize GAD pairs. If you want to swap resynchronize them, delete Volume Migration pairs after Volume Migration ends.
- Before you start Volume Migration, set the same provisioning type for both the target volume of Volume Migration and the GAD S-VOL.
- 3-data-center (3DC) configurations using both GAD and Universal Replicator (UR) cannot be used for destination volumes. To use a 3DC configuration combining GAD and UR, configure the 3DC system after volumes are completely migrated by Volume Migration.

GAD status and nondisruptive migration pair operations

The following table shows the cache modes of nondisruptive migration and GAD pair operations.

Operation target	Virtual LDEV ID	GAD pair operation	
		P-VOL	S-VOL
Cache modes for external volume groups of the destination storage system	TM (Cache Through)	N	N
	SM (Write Sync)	Y	N
Legend Y: Can be performed N: Cannot be performed			

The following table shows the possibility of GAD pair operation and cache mode change in nondisruptive migration.

Operation target	Cache modes to which GAD P-VOL can be changed		
	TM (Cache Through)	SM	EM/DM
GAD P-VOL	N	Y	N
Legend Y: Can be performed N: Cannot be performed			

The following table shows the possibility of Volume Migration pair operations (when a GAD pair operating with nondisruptive migration is in the initial status).

GAD pair status	Virtual LDEV ID	I/O from the server		Volume Migration operation	
		Read	Write	Migrate volumes	Stop migrating volumes
SMPL	Specified	Y	Y	Y	Y
	None	N	N	N	Y

GAD pair status	Virtual LDEV ID	I/O from the server		Volume Migration operation	
		Read	Write	Migrate volumes	Stop migrating volumes
	None (GAD reserve is specified as the virtual attribute)	N	N	N	Y
Legend Y: Can be performed N: Cannot be performed					

The following table shows the possibility of Volume Migration pair operations (when a GAD pair operating with nondisruptive migration is being mirrored).

GAD pair status	I/O mode	Pair location	I/O from the server		Volume Migration operation	
			Read	Write	Migrate volumes	Stop migrating volumes
INIT/COPY	Mirror (RL)	P-VOL	Y	Y	N	Y
	Block	S-VOL	N	N	N	Y
Legend Y: Can be performed N: Cannot be performed						

The following table shows the possibility of Volume Migration pair operations (when a GAD pair operating with nondisruptive migration is already mirrored).

GAD pair status	I/O mode	Pair location	I/O from the server		Volume Migration operation	
			Read	Write	Migrate volumes	Stop migrating volumes
PAIR	Mirror (RL)	P-VOL	Y	Y	Y	Y

GAD pair status	I/O mode	Pair location	I/O from the server		Volume Migration operation	
			Read	Write	Migrate volumes	Stop migrating volumes
	Block	S-VOL	N	N	N	Y
Legend Y: Can be performed N: Cannot be performed						

The following table shows the possibility of Volume Migration pair operations (when a GAD pair operating with nondisruptive migration is suspended).

GAD pair status	I/O mode	Pair location	I/O from the server		Volume Migration operation	
			Read	Write	Migrate volumes	Stop migrating volumes
PSUS/ PSUE	Local	P-VOL	Y	Y	Y	Y
	Block	Volume of your own choice (P-VOL or S-VOL)	N	N	Y	Y
SSUS	Block	S-VOL	N	N	Y	Y
SSWS	Local	S-VOL	Y	Y	Y	Y
Legend Y: Can be performed N: Cannot be performed						

The following table shows the possibility of Volume Migration pair operations (when a GAD pair operating with nondisruptive migration is blocked).

GAD pair status	I/O mode	Pair location	I/O from the server		Volume Migration operation	
			Read	Write	Migrate volumes	Stop migrating volumes
PSUE	Block	P-VOL	N	N	Y	Y
		S-VOL	N	N	Y	Y
Legend						
Y: Can be performed						
N: Cannot be performed						

The following table shows the GAD pair status in Volume Migration when GAD and nondisruptive migration are used together, and the possibility of Volume Migration pair operations (when the GAD P-VOL is a target volume of Volume Migration).

Volume Migration pair status	Is VM P-VOL shared with NDM?	Is VM S-VOL shared with NDM?	GAD pair operation							
			Create	Suspend		Delete			Resynchronize	
				P-VOL specified	S-VOL specified	P-VOL specified ¹	S-VOL specified ²	Force delete	P-VOL specified ³	S-VOL specified
SMP L (PD)	Yes	No ⁴	Y	Y	Y	Y	N	Y	N	N
COPY	Yes	No ⁴	N	Y	Y	Y	N	Y	N	N
PSUE	Yes	No ⁴	N	Y	Y	Y	N	Y	N	N
PSUS	No ⁴	Yes	N	Y	Y	Y	N	Y	N	N
Legend VM: Volume Migration NDM: nondisruptive migration Y: Can be performed N: Cannot be performed										

Volume Migration pair status	Is VM P-VOL shared with NDM ?	Is VM S-VOL shared with NDM?	GAD pair operation							
			Create	Suspend		Delete			Resynchronize	
				P-VOL specified	S-VOL specified	P-VOL specified ¹	S-VOL specified ²	Force delete	P-VOL specified ³	S-VOL specified
Note: <ol style="list-style-type: none">1. You can delete a GAD pair by specifying its P-VOL, only when the I/O mode is Local and the pair status of the P-VOL is PSUS or PSUE.2. You can delete a GAD pair by specifying its S-VOL, only when the I/O mode is Local and the pair status of the S-VOL is SSWS.3. If a GAD pair is suspended due to a failure, you can resynchronize it by deleting the Volume Migration pair, and then specifying the GAD P-VOL.4. You cannot use the S-VOL of Volume Migration as a volume of nondisruptive migration.										

Volume Shredder

GAD pair volumes and quorum disks cannot use Volume Shredder to delete data.

Performance Monitor

Performance Monitor can be used to collect performance information about GAD pair volumes and the quorum disk.

The amount of a port's I/O that can be added to Performance Monitor depends on the type of the volume to which I/O is issued, or on the volume's I/O mode.

For example, when the I/O mode of both GAD volumes is Mirror (RL), each time the server writes to the P-VOL, performance data is recorded for all of the following ports and volumes:

- Primary storage system port connected to the host (Target)
- Primary storage system port connected to the secondary storage system (Initiator)
- Secondary storage system port connected to the primary storage system (RCU Target)
- P-VOL
- S-VOL

When the I/O mode of both GAD volumes is Mirror (RL), each time the server reads the P-VOL data, performance data is recorded for only the primary storage system host (Target) port and the P-VOL.

Server I/Os added to Performance Monitor

The number of I/Os (reads and writes) to GAD volumes that is added to Performance Monitor depends on the GAD status, as shown the following tables.

Table 20 Writes to GAD volumes to be added to Performance Monitor

GAD status	P-VOL	S-VOL
Mirrored	The sum of the following values: <ul style="list-style-type: none"> Number of writes to the P-VOL Number of RIOs to the P-VOL from the S-VOL 	The sum of the following values: <ul style="list-style-type: none"> Number of reads from the S-VOL Number of RIOs to the S-VOL from the P-VOL
Quorum disk blocked (status of GAD pairs created, resynchronized, or swap resynchronized on 80-04-2x or earlier for VSP 5000 series, 88-01-0x or later for VSP G350, G370, G700, G900 and VSP F350, F370, F700, F900)	Number of writes to the P-VOL	Number of RIOs to the S-VOL from the P-VOL
Quorum disk blocked (status of GAD pairs created, resynchronized, or swap resynchronized on 80-05-0x or later for VSP 5000 series, 88-01-0x or later for VSP G350, G370, G700, G900 and VSP F350, F370, F700, F900)	Number of writes to the P-VOL	The sum of the following values: <ul style="list-style-type: none"> Number of writes from the server to the S-VOL Number of RIOs from the P-VOL to the S-VOL
Mirroring	Number of writes to the P-VOL	Number of RIOs to the S-VOL from the P-VOL

GAD status	P-VOL	S-VOL
Suspended (when the P-VOL has the latest information)	Number of writes to the P-VOL	Not counted*
Suspended (when the S-VOL has the latest information)	Not counted*	Number of writes to the S-VOL
Blocked	Not counted*	Not counted*
* Reads and writes by a server are illegal requests and cause an error. However, they could be counted as I/O.		

Table 21 Reads to GAD volumes to be added to Performance Monitor

GAD status	P-VOL	S-VOL
Mirrored	Number of reads from the P-VOL	Number of reads from the S-VOL
Quorum disk blocked (status of GAD pairs created, resynchronized, or swap resynchronized on 80-04-2x or earlier for VSP 5000 series, 88-01-0x or later for VSP G350, G370, G700, G900 and VSP F350, F370, F700, F900)	Number of reads from the P-VOL	Not counted*

GAD status	P-VOL	S-VOL
Quorum disk blocked (status of GAD pairs created, resynchronized, or swap resynchronized on 80-05-0x or later for VSP 5000 series, 88-01-0x or later for VSP G350, G370, G700, G900 and VSP F350, F370, F700, F900)	Number of reads from the P-VOL	Number of reads from the S-VOL
Mirroring	Number of reads from the P-VOL	Not counted*
Suspended (when the P-VOL has the latest information)	Number of reads from the P-VOL	Not counted*
Suspended (when the S-VOL has the latest information)	Not counted*	Number of reads from the S-VOL
Blocked	Not counted*	Not counted*
* Reads and writes from a server are illegal requests and cause an error. However, they could be counted as I/O.		

Table 22 Relation between the number of I/Os added to Performance Monitor and the number of server I/Os

GAD status	Number of writes	Number of reads
Mirrored	Approximately the same* as the number of writes to the P-VOL or S-VOL	The same as the total number of reads from the P-VOL and S-VOL

GAD status	Number of writes	Number of reads
Quorum disk blocked (status of GAD pairs created, resynchronized, or swap resynchronized on 80-04-2x or earlier for VSP 5000 series, 88-01-0x or later for VSP G350, G370, G700, G900 and VSP F350, F370, F700, F900)	The same as the number of writes to the P-VOL	The same as the number of reads from the P-VOL
Quorum disk blocked (status of GAD pairs created, resynchronized, or swap resynchronized on 80-05-0x or later for VSP G1000, VSP G1500, VSP F1500, 88-01-0x or later for VSP G350, G370, G700, G900 and VSP F350, F370, F700, F900)	The same as the number of writes to the P-VOL	The same as the total number of reads from the P-VOL and S-VOL
Mirroring	The same as the number of writes to the P-VOL	The same as the number of reads from the P-VOL
Suspended (P-VOL has latest data)	The same as the number of writes to the P-VOL	The same as the number of reads from the P-VOL
Suspended (S-VOL has latest data)	The same as the number of writes to the S-VOL	The same as the number of reads from the S-VOL
Blocked	Not counted	Not counted
* For writes by a server, RIOs might be divided before being issued. For this reason, this number might differ from the number of writes by a server.		

Port I/Os added to Performance Monitor

The number of I/Os (reads or writes) of the port added to Performance Monitor depends on the P-VOL or S-VOL (I/O destination), or on the I/O mode of the destination volume, as shown in the following table.

I/O destination volume I/O mode	I/O destination volume	Primary storage system			Secondary storage system		
		Target	Initiator	RCU Target	Target	Initiator	RCU Target
Mirror (RL)	P-VOL	Total writes and reads	Number of writes	Not added	Not added	Not added	Number of writes
	S-VOL	Not added	Not added	Number of writes	Total writes and reads	Number of writes	Not added
Local	P-VOL	Total writes and reads	Not added	Not added	Not added	Not added	Not added
	S-VOL	Not added	Not added	Not added	Total writes and reads	Not added	Not added
Block	P-VOL	Total writes and reads*	Not added	Not added	Not added	Not added	Not added
	S-VOL	Not added	Not added	Not added	Total writes and reads*	Not added	Not added
* Reads and writes by a server are illegal requests and cause an error. However, they might be counted as I/Os.							

Connection types


Three types of connections are supported for GAD physical paths: direct, switch, and channel extenders.

You can use Hitachi Command Suite or CCI to configure ports and topologies.

Establish bidirectional physical path connections from the primary to the secondary storage system and from the secondary to the primary storage system.

Direct connection

You can connect two storage systems directly to each other.



You can use the following host mode options (HMOs) to improve response time of host I/O by improving response time between the storage systems for distance direct connections (up to 10 km Long Wave) when the open package is used.

- HMO 49 (BB Credit Set Up Option1) (VSP 5000 series)
- HMO 50 (BB Credit Set Up Option2) (VSP 5000 series)
- HMO 51 (Round Trip Set Up Option)

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

The fabric and topology settings depend on the settings of packages, the protocol used for the connection between the storage systems, and the setting of HMO 51. The link speed that can be specified differs for each condition.

Package name	Protocol	HMO 51 setting	Fabric setting	Topology: remote replication ports	Link speed that can be specified
16FC8 (VSP 5000 series) 4HF8 (VSP G200, G400, G600, G800 VSP F400, F600, F800)	8 Gbps FC	OFF	OFF	FC-AL	<ul style="list-style-type: none"> ▪ 2 Gbps ▪ 4 Gbps ▪ 8 Gbps
	8 Gbps FC	ON	OFF	Point-to-Point	<ul style="list-style-type: none"> ▪ 2 Gbps ▪ 4 Gbps ▪ 8 Gbps
	8 Gbps FC	OFF	OFF	Point-to-Point	<ul style="list-style-type: none"> ▪ 2 Gbps ▪ 4 Gbps ▪ 8 Gbps

Package name	Protocol	HMO 51 setting	Fabric setting	Topology: remote replication ports	Link speed that can be specified
16FC16 (VSP 5000 series)	16 Gbps FC	OFF	OFF	FC-AL	<ul style="list-style-type: none"> 4 Gbps 8 Gbps
	16 Gbps FC	ON	OFF	Point-to-Point	<ul style="list-style-type: none"> 4 Gbps 8 Gbps 16 Gbps
		OFF	OFF	Point-to-Point	<ul style="list-style-type: none"> 4 Gbps 8 Gbps 16 Gbps
8FC16 (VSP 5000 series) 2HF16 (VSP G200, G400, G600, G800 VSP F400, F600, F800)	16 Gbps FC	OFF	OFF	FC-AL	<ul style="list-style-type: none"> 4 Gbps 8 Gbps
	16 Gbps FC	ON	OFF	Point-to-Point	<ul style="list-style-type: none"> 4 Gbps 8 Gbps 16 Gbps
		OFF	OFF	Point-to-Point	<ul style="list-style-type: none"> 4 Gbps 8 Gbps 16 Gbps
4HF32R* (VSP G200, G400, G600, G800, VSP F400, F600, F800)	16 Gbps FC	OFF	OFF	FC-AL	<ul style="list-style-type: none"> 4 Gbps 8 Gbps
		ON	OFF	Point-to-Point	<ul style="list-style-type: none"> 4 Gbps 8 Gbps 16 Gbps
		OFF	OFF	Point-to-Point	<ul style="list-style-type: none"> 4 Gbps 8 Gbps 16 Gbps
	32 Gbps FC	OFF	OFF	FC-AL	8 Gbps
		ON	OFF	Point-to-Point	<ul style="list-style-type: none"> 8 Gbps 16 Gbps 32 Gbps


Package name	Protocol	HMO 51 setting	Fabric setting	Topology: remote replication ports	Link speed that can be specified
		OFF	OFF	Point-to-Point	<ul style="list-style-type: none">▪ 8 Gbps▪ 16 Gbps▪ 32 Gbps
CHB(FC32G) (VSP G350, VSP G370, VSP G700, VSP G900, VSP F350, VSP F370, VSP F700, VSP F900)	32GbpsFC	OFF	OFF	FC-AL	<ul style="list-style-type: none">▪ 4 Gbps▪ 8 Gbps
		ON	OFF	FC-AL	<ul style="list-style-type: none">▪ 4 Gbps▪ 8 Gbps
		OFF	OFF	Point-to-Point	<ul style="list-style-type: none">▪ 16 Gbps▪ 32 Gbps
		ON	OFF	Point-to-Point	<ul style="list-style-type: none">▪ 16 Gbps▪ 32 Gbps
* 4HF32R (4 ports, FC 32 Gbps Ready Package) supports multiple transfer speed protocol. Depending on the mounted SFP parts, you can use either 16 Gbps or 32 Gbps protocol.					

Connection using switches

You can use host mode options to improve response times when switches are used for distance connections.



Note: You do not need to set the port attributes (Initiator, RCU Target, Target) on VSP G/F350, G/F370, G/F700, G/F900 models.



Switches from some vendors (for example, McData ED5000) require F_port.

You can use the following host mode options (HMOs) to improve response time of host I/O by improving response time between the storage systems when switches are used for distance connections (up to approximately 500 km with a round-trip response of 20 ms or less) and the open package is used.

- HMO 49 (BB Credit Set Up Option1) (VSP 5000 series)
- HMO 50 (BB Credit Set Up Option2) (VSP 5000 series)
- HMO 51 (Round Trip Set Up Option)

For details about HMOs, see the *Provisioning Guide* for the storage system.

The fabric and topology settings depend on the settings of packages, and protocol used for the connection between storage systems, and the HMO 51 setting. The link speed that can be specified differs on each condition.

Package name	Protocol	HMO 51 setting	Fabric setting	Topology: Initiator and RCU Target	Link speed that can be specified
16FC8 (VSP 5000 series) 4HF8 (VSP G200, G400, G600, G800, VSP F400, F600, F800)	8 Gbps FC	OFF	ON	Point-to-Point	<ul style="list-style-type: none"> ▪ 2 Gbps ▪ 4 Gbps ▪ 8 Gbps
	8 Gbps FC	ON	ON	Point-to-Point	<ul style="list-style-type: none"> ▪ 2 Gbps ▪ 4 Gbps ▪ 8 Gbps
	8 Gbps FC	OFF	OFF	Point-to-Point	<ul style="list-style-type: none"> ▪ 2 Gbps ▪ 4 Gbps ▪ 8 Gbps
8FC16 (VSP 5000 series) 2HF16 (VSP G200, G400, G600, G800, VSP F400, F600, F800)	16 Gbps FC	OFF	ON	Point-to-Point	<ul style="list-style-type: none"> ▪ 4 Gbps ▪ 8 Gbps ▪ 16 Gbps
	16 Gbps FC	ON	ON	Point-to-Point	<ul style="list-style-type: none"> ▪ 4 Gbps ▪ 8 Gbps ▪ 16 Gbps
	16 Gbps FC	OFF	OFF	Point-to-Point	<ul style="list-style-type: none"> ▪ 4 Gbps ▪ 8 Gbps ▪ 16 Gbps

Package name	Protocol	HMO 51 setting	Fabric setting	Topology: Initiator and RCU Target	Link speed that can be specified
16FC16 (VSP 5000 series)	16 Gbps FC	OFF	ON	Point-to-Point	<ul style="list-style-type: none"> ▪ 4 Gbps ▪ 8 Gbps ▪ 16 Gbps
	16 Gbps FC	ON	ON	Point-to-Point	<ul style="list-style-type: none"> ▪ 4 Gbps ▪ 8 Gbps ▪ 16 Gbps
	16 Gbps FC	OFF	OFF	Point-to-Point	<ul style="list-style-type: none"> ▪ 4 Gbps ▪ 8 Gbps ▪ 16 Gbps
4HF32R* (VSP G200, G400, G600, G800, VSP F400, F600, F800)	16 Gbps FC	OFF	ON	Point-to-Point	<ul style="list-style-type: none"> ▪ 4 Gbps ▪ 8 Gbps ▪ 16 Gbps
		ON	ON	Point-to-Point	<ul style="list-style-type: none"> ▪ 4 Gbps ▪ 8 Gbps ▪ 16 Gbps
		OFF	OFF	Point-to-Point	<ul style="list-style-type: none"> ▪ 4 Gbps ▪ 8 Gbps ▪ 16 Gbps
	32 Gbps FC	OFF	ON	Point-to-Point	<ul style="list-style-type: none"> ▪ 8 Gbps ▪ 16 Gbps ▪ 32 Gbps
		ON	ON	Point-to-Point	<ul style="list-style-type: none"> ▪ 8 Gbps ▪ 16 Gbps ▪ 32 Gbps
		OFF	OFF	Point-to-Point	<ul style="list-style-type: none"> ▪ 8 Gbps ▪ 16 Gbps ▪ 32 Gbps
		ON	ON	Point-to-Point	<ul style="list-style-type: none"> ▪ 8 Gbps ▪ 16 Gbps ▪ 32 Gbps
		OFF	OFF	Point-to-Point	<ul style="list-style-type: none"> ▪ 8 Gbps ▪ 16 Gbps ▪ 32 Gbps
16FE10 (VSP 5000 series)	10 Gbps FCoE	OFF	ON	Point-to-Point	10 Gbps

Package name	Protocol	HMO 51 setting	Fabric setting	Topology: Initiator and RCU Target	Link speed that can be specified
	10 Gbps FCoE	ON	ON	Point-to-Point	10 Gbps
CHB(FC32 G) VSP G350, VSP G370, VSP G700, VSP G900, VSP F350, VSP F370, VSP F700, VSP F900	32GbpsFC	OFF	ON	Point-to-Point	<ul style="list-style-type: none">▪ 4 Gbps▪ 8 Gbps▪ 16 Gbps▪ 32 Gbps
		ON	ON	Point-to-Point	<ul style="list-style-type: none">▪ 4 Gbps▪ 8 Gbps▪ 16 Gbps▪ 32 Gbps
* 4HF32R (4 ports, FC 32 Gbps Ready Package) supports multiple transfer speed protocol. Depending on the mounted SFP parts, you can use either 16 Gbps or 32 Gbps protocol.					


Connection using channel extenders

You should use channel extenders and switches for long-distance connections (up to 500 km and the round trip time is 20 ms or less).


Set Fabric to ON and topology to Point-to-Point for the remote replication ports (Initiator and RCU Target).



Note: You do not need to set the port attributes (Initiator, RCU Target, Target) on VSP G/F350, G/F370, G/F700, G/F900 models.



Legend:



Note:

- When the primary and secondary storage systems are connected using switches with a channel extender, and multiple data paths are configured, the capacity of data to be transmitted might concentrate on particular switches, depending on the configuration and the settings of switch routing. Contact customer support for more information.
- Make sure that your channel extenders can support remote I/O. For details, contact customer support.
- Create at least two independent physical paths (one per cluster) between the primary and secondary storage systems for hardware redundancy for this critical element.
- If you plan to use more than 4,000 pairs, when creating pairs you should restrict the number of pairs to 4,000 or less per physical path to distribute the load across multiple physical paths.

Chapter 3: High availability provisioning for Virtual Storage Platform G/F350, G/F370, G/F700, G/F900

Prerequisites to provisioning for high availability

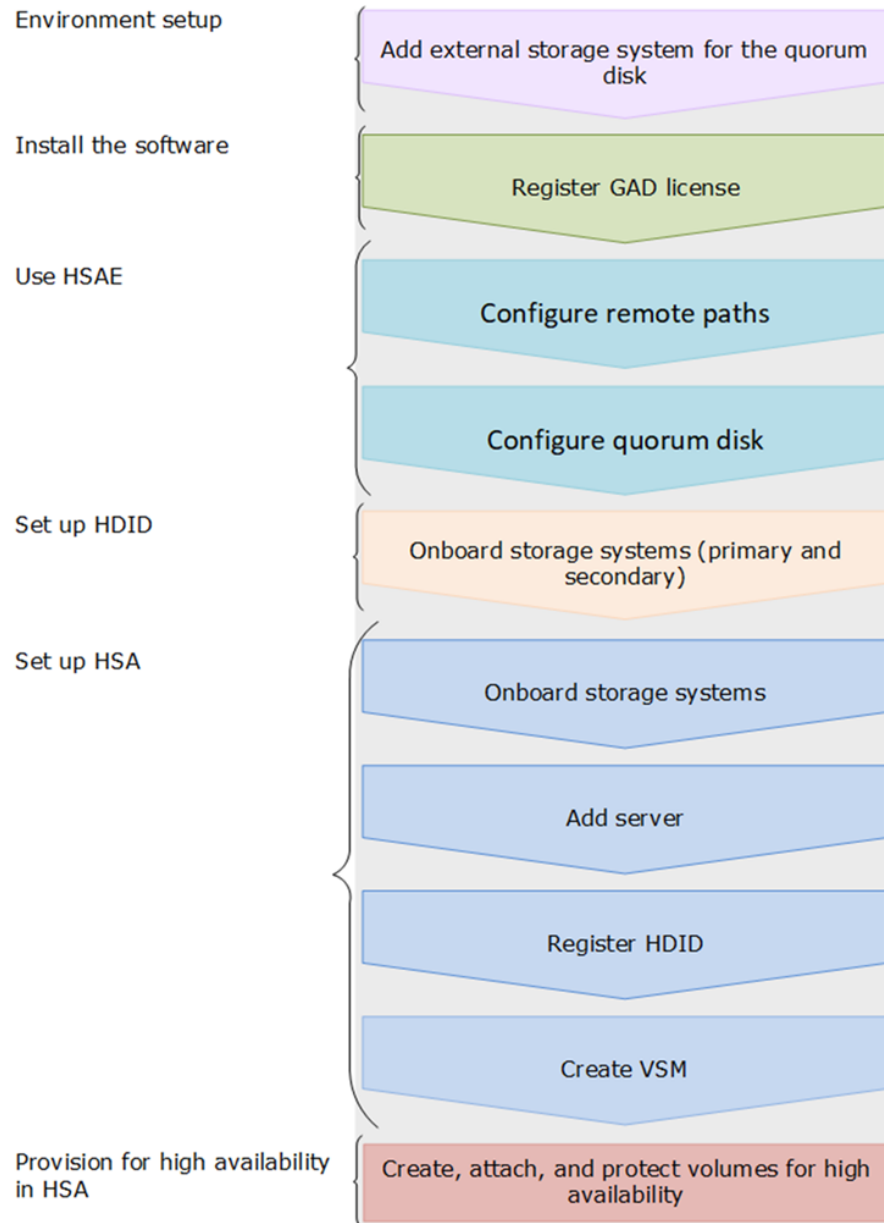
Review the conditions required to provision for high availability.

Make sure that the following conditions are met before provisioning for high availability:

- Data Instance Director v6.6.1 or later is registered in Storage Advisor.
- The registered Storage Advisor satisfies the following conditions
 - Both primary and secondary storage systems are registered with Data Instance Director.
 - Global-active device licenses for both primary and secondary storage systems are installed in Data Instance Director.
- At least one Fibre Channel server must be registered.
- The primary and secondary storage systems must have the following configured:
 - A quorum disk. (The same disk and ID should be assigned in each storage system.)
 - Remote paths , which should be configured bidirectionally.
 - - If the VSM in the primary storage system is using the meta_resource, a VSM with at least one undefined host group must be configured in Storage Advisor for the secondary storage system.
 - If the VSM in the primary storage system is not using the meta_resource, a VSM with at least one undefined volume and one undefined host group must be configured for the primary storage system and, a VSM with at least one undefined host group must be configured for the secondary storage system in Storage Advisor.
- Make sure that both primary and secondary storage systems:
 - Are onboarded in Storage Advisor.
 - Have at least one Fibre port.
 - Have global-active device licenses.
 - Have enough room to create volumes and remote pairs including CTG.
 - Have a Thin or Tiered pool with enough capacity.

Configuration workflow for high availability

The following figure shows the workflow for configuring and provisioning for high availability.



Adding the external storage system for the quorum disk

Install an external storage system for the quorum disk .

The storage system must be supported by Universal Volume Manager for connection as external storage.

Register global-active device license

When installing a software product that enables you to use additional features, you first need to register the license key provided when you purchased that software product.


Before you begin

Prepare the license key code or the license key file for the software product to be installed.

Invoke the maintenance utility from Hitachi Storage Advisor Embedded.

Perform the following procedure on the primary and secondary storage system.

Procedure

1. In the navigation bar, click  (**Settings**), and then select **Licenses**.
2. In the maintenance utility, click **Install**.
3. Specify the license key code or the license key file, and register the license key.



Tip: For more information, refer to Help in the maintenance utility.

4. In the list of license keys, confirm that the status of the software product has changed to **Installed**.
5. In the maintenance utility, click **Log Out**.

Configuring remote paths

Use Storage Advisor Embedded to configure remote paths between the two storage systems that make up a global-active device environment. Because a remote path is required to send data to the connection-destination storage system, you need to configure remote paths for both storage systems so they can send data to each other.

Before you begin

- Connect physical paths between the two storage systems that make up the global-active device environment.
- Identify the following items:
 - The model and serial number of the connection-destination storage system
 - The ID of the path group
 - The port to be used on the connection-source storage system and the port of the connection-destination storage system

Procedure

1. In the navigation bar, click **Others > Remote Path Groups**.
2. Click the plus sign (+).

Create Remote Path Group

MODEL
VSP Fx00 and VSP Gx00

SERIAL NUMBER
424438

PATH GROUP ID
1

PROTOCOL
FC iSCSI

REMOTE PATHS TO ADD

PORT	REMOTE PORT
CL1-A	CL1-D

+

Cancel Submit

3. Specify the required items, and then configure the remote path.
4. Click the ID of the path group to open the details page. Make sure that the status of the remote path you configured is **Normal**.

**Note:**

If the status of the remote path is not **Normal**, see the *Global-Active Device User Guide*.

Configuring quorum disks

You can use Storage Advisor Embedded to create an external volume to configure a quorum disk, or to configure a quorum disk by using an external volume that has already been created.

Configuring a quorum disk by creating an external volume

Create an external volume in the two storage systems that make up the global-active device environment, and specify settings to use the external volume as a quorum disk.

Before you begin

- Ensure that the two storage systems that make up the global-active device environment and the external storage system in which the volume used as the quorum disk is configured are connected by physical paths.
- Ensure that the volume used as the quorum disk is assigned a port on each of the two storage systems that make up the global-active device environment.

- Identify the following items required to configure the external volume:
 - The port to be used for external connection
 - The model and serial number of the external storage system
 - The port information for the external storage system that is connected with a port for external connections:
 - When using Fibre Channel for the protocol you can specify WWN
 - The LUN of the volume used as the quorum disk

If you use a volume whose size exceeds 4 TiB, the size of the external volume that is created will be 4 TiB.
 - The name of the external volume
 - The ID of the external parity group
 - The ID of the external path group
- Identify the following information required to configure the quorum disk:
 - The ID of the quorum disk
 - The models and serial numbers of the paired storage system that shares the quorum disk

Procedure

1. In the navigation bar, click **Others > External Volumes**.
2. Click the plus sign (+).

Create External Volume

Select Ports for External Paths

PROTOCOL

FC | iSCSI

Type in keyword to search

☐ Select All 0 selected

SELECT	ID	WWN
<input type="radio"/>	CL1-A	50060e8012016000
<input type="radio"/>	CL2-A	50060e8012016010
<input type="radio"/>	CL3-A	50060e8012016020
<input type="radio"/>	CL4-A	50060e8012016030
<input type="radio"/>	CL5-A	50060e8012016040
<input type="radio"/>	CL6-A	50060e8012016050

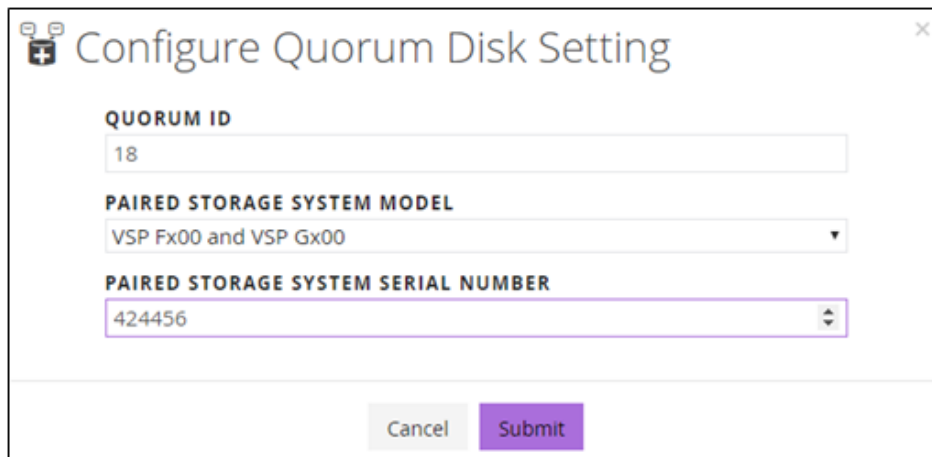
Cancel Next

3. Select the port for external connection, and then click **Next**.
4. Select the external path to be used, and then click **Next**.

5. Select the LUN of the volume used as the quorum disk, and specify the name of the external volume, the ID of the external parity group, and the ID of the external path group. Click **Submit**.

Information about the newly configured external volume is displayed in the list of external volumes.

6. Click  (**Configure Quorum Disk Setting**) for the external volume for which the quorum disk is to be configured.



Configure Quorum Disk Setting

QUORUM ID
18

PAIRED STORAGE SYSTEM MODEL
VSP Fx00 and VSP Gx00

PAIRED STORAGE SYSTEM SERIAL NUMBER
424456

Cancel Submit

7. Specify the required items, and then click **Submit**.

Configuring a quorum disk by selecting an external volume

Configure a quorum disk by using an existing external volume.

Before you begin

- Identify the following information required to configure the quorum disk:
 - The ID of the quorum disk
 - The models and serial numbers of the paired storage system that shares the quorum disk
 - The external volume name

Procedure

1. In the navigation bar, click **Others > Quorum Disks**.
2. Click the plus sign (+).

Configure Quorum Disk Setting

Specify Quorum Disk Setting

QUORUM ID
18

PAIRED STORAGE SYSTEM MODEL
VSP Fx00 and VSP Gx00

PAIRED STORAGE SYSTEM SERIAL NUMBER
424456

EXTERNAL VOLUME
☒ With External Volume
 ☐ Without External Volume

Cancel Next

3. Specify the required items.

To specify the external volume to be used, select **With External Volume**, and then click **Next**.

4. Select the external volume to be used, and then click **Submit**.

Adding block storage to Data Instance Director

Onboard the primary and secondary storage systems to Data Instance Director

Hitachi Block Storage Node Wizard

This wizard is launched when a new Hitachi Block Node is added to the Nodes Inventory.

To add a block node to the Nodes Inventory, click Nodes in the navigation sidebar. Click the plus sign (+) in the Nodes Inventory and select Storage in the Create Node page. Select Hitachi Block Device as the Storage Type to launch the Hitachi Block Storage Node Wizard.

Create Node - Hitachi Block Device

Specify Node name

Node Name

Must be between 2 and 64 characters, contain only letters, numbers, underscores, hyphens and full stops.

Resources or replication relationships created or adopted by HDID must only be managed, modified and deleted via HDID. Failure to do so will cause unpredictable consequences and is not supported unless specifically advised to do so by the HDID documentation.


☐ I confirm that I have read and understood this requirement

Cancel

Previous

Next

Figure 1 Hitachi Block Device Wizard - Specify Node Name

Control	Description
Node Name	Enter a name for the Hitachi Block storage node.
I confirm that ...	<div>This checkbox must be checked to proceed with the node configuration.</div> <div><div> Caution:</div><div>Resources or replication relationships created or adopted by HDID must only be managed, modified and deleted via HDID.</div><div>Failure to do so will cause unpredictable consequences and is not supported unless specifically advised to do so by the HDID documentation.</div></div>

Create Node - Hitachi Block Device

Allocate node to Access Control Resource Group

This node will be added to the 'default' resource group. Select an additional resource group as required.

Name	Description
<input type="radio"/> myResourceGroup	A user defined resource group

Cancel Previous **Next**

Figure 2 Hitachi Block Device Wizard - Allocate node to Access Control Resource Group

Control	Description
Resource Groups	Select the resource group(s) to which this node will be allocated for the purposes of RBAC. All nodes are automatically allocated to the 'default' resource group.

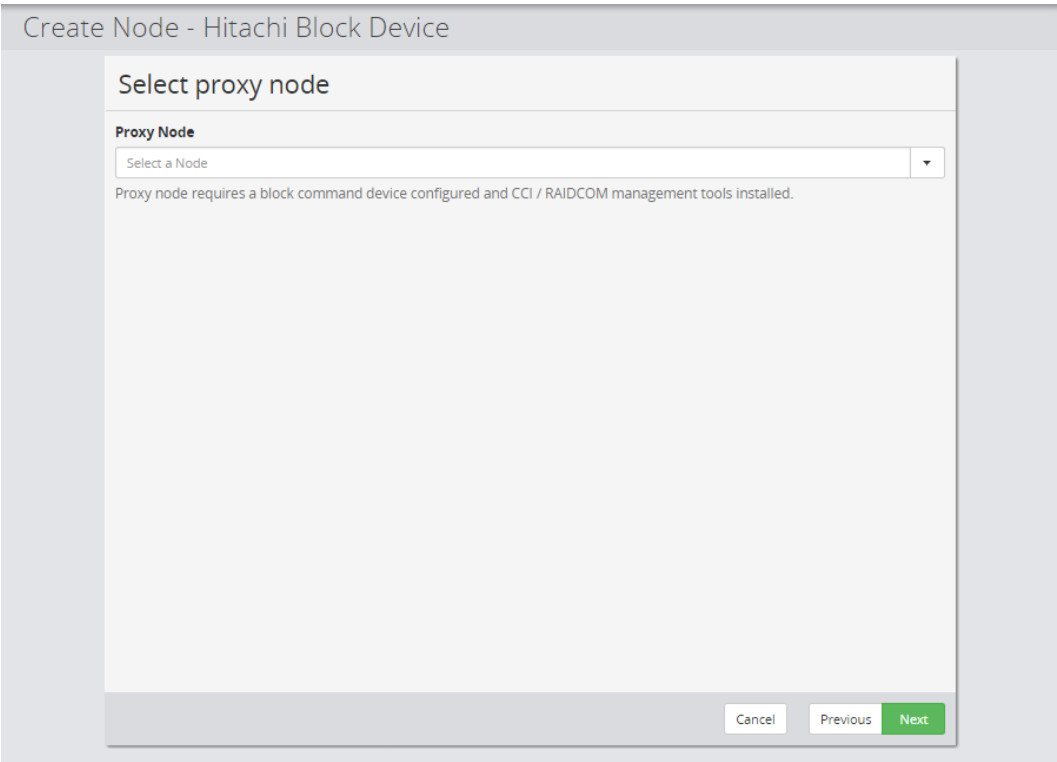




Figure 3 Hitachi Block Device Wizard - Select proxy node

Control	Description
Proxy Node	<p>Select an HDID node to act as a proxy.</p> <div> Caution: ISM nodes and their associated CMDs used to control storage devices must not be shared with other applications.</div> <div> Note: The proxy node is responsible for interfacing with the Block storage device. It can be a Windows or Linux machine with the HDID Client software installed and must be connected via a command device to the Block storage device. The command device must <i>only</i> have user authentication enabled. The proxy node must have supported version of CCI installed. Refer to the Command Control Interface documentation available at https://knowledge.hitachivantara.com.</div>

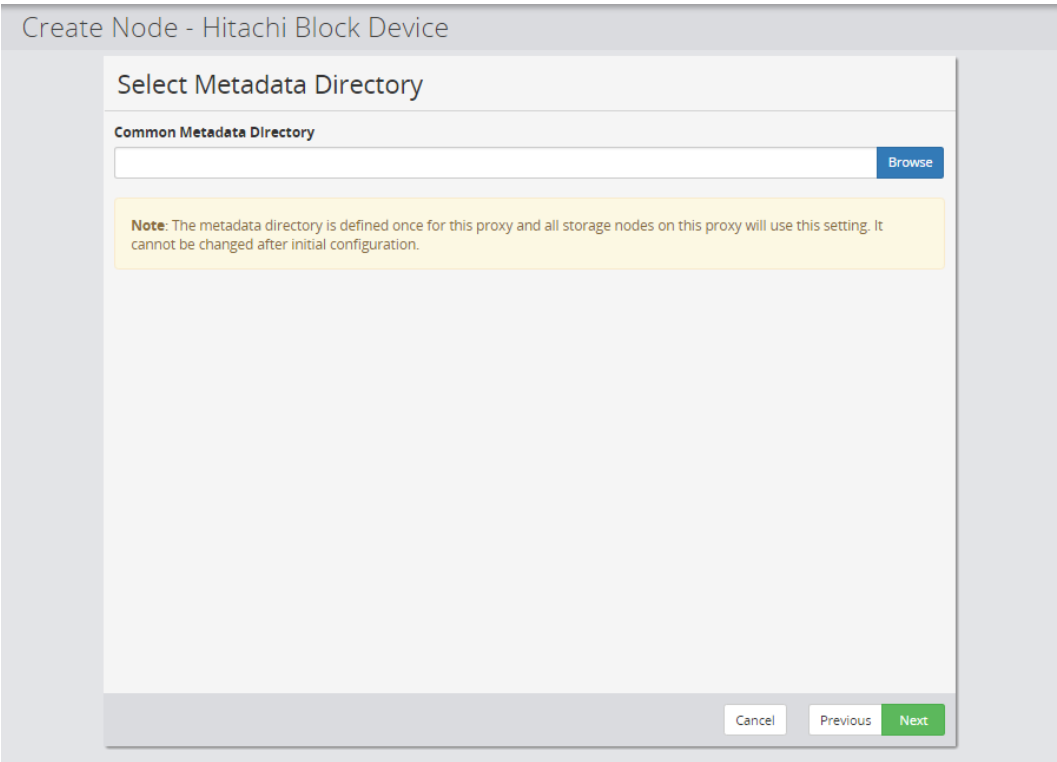



Figure 4 Hitachi Block Device Wizard - Select Metadata Directory

Control	Description
Common Metadata Directory	<div>Enter a directory on the Proxy node where HDID can place metadata files related to Block snapshots and replications. Click the Browse button to open the Path Dialog if required.</div> <div> Note: The metadata directory is defined once for this proxy and all storage nodes on this proxy will use this setting. It cannot be changed after initial configuration.</div>

Create Node - Hitachi Block Device

Specify Device

☒ Select from detected storage devices

Select a Storage Serial Number ▼

If the desired serial does not appear in this list then it may not have an available fibre command device on the selected proxy.


☐ Specify by IP

IP Address

Port Number

Cancel Previous Next

Figure 5 Hitachi Block Device Wizard - Specify Device

Control	Description
Select from detected storage devices	Select this option to specify the hardware storage device by serial number. A list of storage device serial numbers, available to the proxy node selected in the previous step, is displayed in the dropdown menu below.
Specify by IP	<p>Select this option to specify the hardware storage device using the IP address and port number of an IP command device on the storage device. Additional fibre and IP command devices can be added at a step later in the wizard.</p> <div> <p> Note:</p> <p>For HUS VM storage devices, use the IP address of the SVP.</p> <p>For VSP storage devices, use the IP address of CTL1 or CTL2. DO NOT use the IP address of the SVP.</p> </div>

Create Node - Hitachi Block Device

Specify credentials for device

Storage Device Serial Number

410297



Username

Password

The device account requires the following roles: Storage Administrator (Provisioning, Local Copy, Remote Copy) and Security Administrator (View Only, View and Modify).

Cancel Previous Next

Figure 6 Hitachi Block Device Wizard - Specify credentials for device

Control	Description
Storage Device Serial Number	Displays the serial number of the Block device specified in the previous step.
Username	<p>Enter the username for the Block device.</p> <p> Note: The username specified must be a member of the <i>Storage Administrator (Provisioning, Local Copy, Remote Copy)</i> and <i>Security Administrator (View Only, View and Modify)</i> on the Block device. If the Block device cannot be accessed or its credentials are invalid then the node will fail authorization. The configuration wizard can be reopened to correct any errors.</p>
Password	<p>Enter the password for the Block device.</p> <p> Note: The password for authorizing a Block device must contain only useable CCI command characters: A-Za-z0-9'-. / : @ \ _</p>

Create Node - Hitachi Block Device

Specify LDEV Provisioning Range

LDEV Range

☒ All
☐ User defined

Start

0x00

End

0x00

Cancel Previous **Next**




Figure 7 Hitachi Block Device Wizard - Specify LDEV Provisioning Range

Control	Description
All	Select this option if you want HDID to automatically detect the LDEV range from which snapshots and replications should be allocated.
User defined	Select this option if you want to manually specify the LDEV range from which snapshots and replications should be allocated.
Start	Enabled only if User defined is selected. Enter the lower limit of the LDEV range to use for allocation.
End	Enabled only if User defined is selected. Enter the upper limit of the LDEV range to use for allocation.

Create Node - Hitachi Block Device

Command Device Specification and Priority

If any command devices are specified HDID will only use these command devices. If no command devices are specified HDID will attempt to use any fibre based command device available to the proxy. IP based command devices must be specified for HDID to use them. HDID will attempt commands on command devices in the order they are specified, failing over to the next if there is an issue with a command device.

☐ Select All (0)   

Priority	Type	LDEV ID	IP Address	Port	Status
+					
<input type="radio"/> 1	Fibre	Any Available	-	-	

Cancel Previous **Next**



Figure 8 Hitachi Block Device Wizard - Command Device Specification and Priority




The user able to specify zero or more fibre or IP command devices in priority order.

If no command devices are specified, then HDID will attempt to control the hardware storage device via any fibre connected command device, available to the Proxy Node specified, in an order specified by HORCM.

If one or more command devices are specified, then HDID will attempt to control the hardware storage device via a command device in the order specified by the user. If the first command device fails, HDID will progress to the next. If all specified command devices fail then the operation fails. HDID will not attempt to use any command devices that are not specified, even if they are available.

for example, it is possible to specify a specific fibre command device, followed by any fibre command device, followed by a specific IP command device.

Control	Description
 Edit	Enabled when only one command device is selected. Launches the Configure Command Device wizard below to allow the settings to be edited.
 Increase Priority	Enabled when only one command device is selected. Increases the priority of the selected command device.

Control	Description
 Decrease Priority	Enabled when only one command device is selected. Decreases the priority of the selected command device.
 Delete	Enabled when one or more command devices are selected. Deletes selected command device.
 Add	Launches the Configure Command Device wizard below to guide you through setting up a fibre or IP command device.

Create Node - Hitachi Block Device

Configure Command Device

Fibre

IP

Fibre Command Device Options

☒ Use any available fibre command device

☐ Select from detected fibre command devices

Select a Fibre Command Device ▼

Cancel Discard Previous **Apply**

Figure 9 Hitachi Block Device Wizard - Configure Command Device - Fibre

Control	Description
Use any available fibre command device	Select this option to insert an entry in the command device list that allows HDID to use any available fibre command device.

Control	Description
Select from detected fibre command devices	Select this option to insert a specific fibre command device in the list. The detected fibre command devices are displayed in the dropdown menu below using their decimal LDEV ID.

Create Node - Hitachi Block Device

Configure Command Device

Fibre

IP

IP Command Device Options

IP Address

Port Number

Cancel Discard Previous Apply

Figure 10 Hitachi Block Device Wizard - Configure Command Device - IP



Note: When configuring IP command devices for VSP storage devices, we recommend adding one for CTL1 and one for CTL2, to maintain dual redundancy.



Control	Description
IP Address	<p>Enter the IP address of the command device to add to the list.</p> <p>Note:</p> <p>For HUS VM storage devices, use the IP address of the SVP.</p> <p>For VSP storage devices, use the IP address of CTL1 or CTL2. DO NOT use the IP address of the SVP.</p>

Control	Description
Port Number	Enter the port number of the command device.

Create Node - Hitachi Block Device

Specify LDEV Ranges for each VSM

Define a virtual LDEV range for each VSM serial you intend to use within HDID. The ranges are used to control the virtual LDEV IDs used for replications and snapshots (excluding GAD), and should be defined to not include the IDs of any GAD volumes. Failure to provide such a range (or providing an incorrect range) may result in ID clashes when attempting to set up GAD replications.

☐ Select All (0)  



Virtual Serial	Start	End
+		


Cancel Previous Next

Figure 11 Hitachi Block Wizard - Specify LDEV Ranges for each VSM



Note: GAD replications require P-VOLs and S-VOLs to have matching virtual serial numbers and virtual LDEV IDs. To avoid virtual LDEV ID collisions between GAD volumes and non-GAD S-VOLs (created by HDID for other types of replications and snapshots), it is possible to define virtual LDEV ID ranges to be used by those non-GAD operations. Virtual LDEV ranges can be specified for each VSM (Virtual Storage Machine) to be used.

Control	Description
 Edit	Enabled only if one Virtual LDEV range is selected. The Configure Virtual LDEV Range page of the wizard (see below) is displayed to enable a port to be added.
 Delete	Enabled only if one or more Virtual LDEV ranges are selected. Deletes the Virtual LDEV range(s) from the list.

Control	Description
 Add	Adds a new Virtual LDEV range to the list. The Configure Virtual LDEV Range page of the wizard (see below) is displayed to enable a Virtual LDEV Range to be added.
Virtual Serial(s)	The Virtual Serial(s) and associated LDEV Ranges that will be used.

Create Node - Hitachi Block Device

Configure Virtual LDEV Range

VSM Serial Number

Start of Virtual LDEV range

End of Virtual LDEV range

Cancel Discard Previous **Apply**

Figure 12 Hitachi Block Wizard - Configure Virtual LDEV Range



Caution: These ranges are used to control the virtual LDEV IDs used for non-GAD replications and snapshots. They must be defined to exclude the IDs of any GAD volumes. Failure to provide such a range (or providing an incorrect range) may result in ID clashes when attempting to set up GAD replications.

Control	Description
VSM Serial Number	Enter the serial number of the VSM you intend to use within HDID.
Start of LDEV range	Enter the lower limit of the LDEV range to use for allocation.
End of LDEV range	Enter the lower limit of the LDEV range to use for allocation.







Figure 13 Hitachi Block Device Wizard - Specify ports used for provisioning

Control	Description
 Edit	Enabled only if one provisioning port is selected. The Specify Port page of the wizard (see below) is displayed to enable a port to be added.
 Delete	Enabled only if one or more provisioning ports are selected. Deletes the port(s) from the list.
 Add	Adds a new provisioning port to the list. The Specify Port page of the wizard (see below) is displayed to enable a port to be added.
Provisioning Port(s)	Lists the ports that will be used for provisioning.

Create Node - Hitachi Block Device

Specify Port

Port

CL 1 - A

Cancel Discard Previous Apply

Figure 14 Hitachi Block Device Wizard - Specify Port



Note: If more than one provisioning port is selected, then the port with the least amount of LUNs will be used.

Control	Description
Port	<p>Enter the port identifier in the following format:</p> <p>CL_c-s</p> <p>where:</p> <ul style="list-style-type: none"> c is the physical channel number in the range 1...n s is the physical slot number in the range A...Z

Create Node - Hitachi Block Device

Summary of 'myBlock'

Proxy Node
Client1

Storage Device Serial Number
410297

Username
lanH

LDEV Provisioning Range
All Available

Configured Command Devices

Type	LDEV ID	IP Address	Port
Fibre	Any Available	-	-

Ports used for Provisioning

Port
CL1-A

VSM Virtual LDEV Ranges

VSM Serial Number	Start	End
123456	0x00	0xff

Cancel Previous **Finish**

Figure 15 Hitachi Block Wizard - Summary

Control	Description
Summary	Summary of the settings entered.

Adding a storage system in Storage Advisor

Onboarding a storage system is the process of associating it with Storage Advisor. After the storage system is onboarded, manage it from the Storage Advisor dashboard.

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

Before you begin

Storage Advisor requires 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 Storage Advisor has access to all custom resource groups and meta resource groups.

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

Procedure

1. On the Storage Advisor dashboard, click **Storage Systems** on the left pane.

2. Click the plus sign (+) to add a storage system.
3. Enter values for the following parameters on the **Onboard Storage System** page.
 - **IP Address:** For a storage system with an SVP, enter the IP address of the external service processor for the storage system you want to discover.
 - **User name and password:** Log in as a user that has administrator privileges on this storage system. For example, you can log in as the user `maintenance`.
4. Click **Submit**.

Result

The Jobs tab is updated with the job called `Create Storage System`. If multiple storage systems are being added, there will be a job for each one.

Wait a while for the storage system to be added. Refresh the Jobs tab to verify that storage system is onboarded.

The dashboard shows 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 where you can see the newly added storage system.

When a storage system is onboarded, Storage Advisor goes through an initialization process where it gathers the information about the current configuration of the storage system. During this time you will see that the ports, volumes, pools, and parity groups in the storage system are "Not accessible". Once the initialization is complete, you can see the port, pool, volume, and parity group information in the storage system details.



Note: If operations are performed outside of Storage Advisor, it takes time to update in Storage Advisor (approximately 20 minutes, depending on cache refresh).

Next steps

1. In the parity group inventory 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.

Adding servers

Use Storage Advisor to add servers so you can attach volumes.

You can add multiple server parameters from a file, or add one server at a time.

There are two methods of adding servers:

- 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, in the order specified: Name, Description, IPAddress, OSType, WWNs (comma separated list of WWNs). All fields are required except Description and IPAddress. Valid OSType values are as follows:

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

Procedure

1. On the Storage Advisor dashboard, click **Servers**. Then click the plus sign (+) to open the **Add Server** page.



Note: iSCSI is not currently supported for high availability.

Add Servers

CSV Import

Fibre Servers

SERVER NAME	DESCRIPTION	IP ADDRESS	OS TYPE	WWN
Host Name	Description	IP Address	HP_UX	
50:00:00:00:00:00, 50:00:00:00:00:00:01				

iSCSI Servers

SERVER NAME	DESCRIPTION	IP ADDRESS	OS TYPE	CHAP USER	CHAP SECRET	iSCSI NAMES
Host Name	Description	IP Address	HP_UX			
				sysadmin	*****	iqn.1992-01.com.company:db, iqn.1992-01.com.c...

Cancel Reset Submit

2. On the **Add Server** page, 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 will populate the page. Example:

```
Name,Description,IPAddress,OSType,WWNS
Esxi,ESXI HOST,10.30.90.200,VMWARE_EX,10:00:00:05:33:26:f7:21
Win,WINDOWS
HOST,10.30.91.80,WIN_EX,"10:00:00:05:33:26:f7:37,10:00:00:05:33:26:f7:36"
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:e0:fd"
ESXi_Cisco_2,ESXi HOST connected to Cisco
Fabric,,VMWARE_EX,"100000053326df1a,100000053326df1b"
```

- Click the plus sign (+) in the table to add a row and enter the required information for Fibre Channel. You can add more servers by clicking the plus sign again.

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

Next steps

Create volumes and attach them to the server.

Registering Data Instance Director

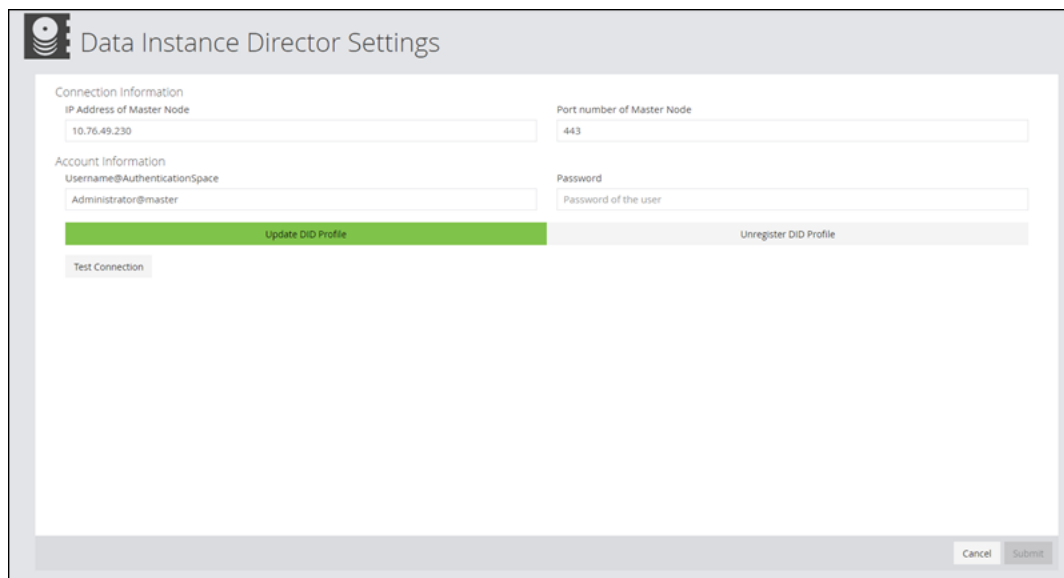
You can use Storage Advisor to register Data Instance Director settings.

Before you begin

Data Instance Director is installed. The user with permissions to perform pair management operations is identified.

Procedure

1. From the **Settings** menu, click **Data Instance Director Settings**.
2. Under **Connection Information**, enter the IP address of the Master node. The port number of the Master node displays automatically.
3. Under **Account Information**, enter the user name and password of the Data Instance Director user who can perform pair management operations for high availability.
4. Click **Test Connection** to verify that you connected to the Master node successfully.
5. Click **Submit**.



The screenshot shows the 'Data Instance Director Settings' page. It is divided into two main sections: 'Connection Information' and 'Account Information'. In the 'Connection Information' section, the 'IP Address of Master Node' is set to '10.76.49.230' and the 'Port number of Master Node' is '443'. The 'Account Information' section shows the 'Username@AuthenticationSpace' as 'Administrator@master' and the 'Password' field is empty. Below these fields are two buttons: 'Update DID Profile' (highlighted in green) and 'Unregister DID Profile'. At the bottom left of the form is a 'Test Connection' button. At the bottom right are 'Cancel' and 'Submit' buttons.

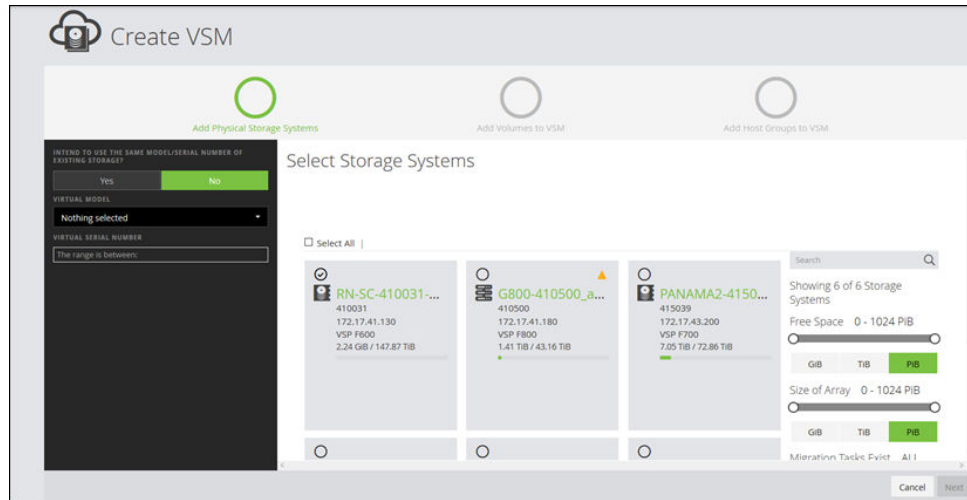
Creating virtual storage machines

You can use Storage Advisor 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** page.



2. On the **Create VSM** page, select the storage systems you want to use.
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.
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 step and add volumes later.
5. Click **Next** again to add host groups. Add host groups from one storage system at a time and click the plus sign (+) to add them to the list.
6. Click **Submit** to create a job to add the virtual storage machine.

Create, attach, and protect volumes with high availability

Storage Advisor allows you to create, attach, and protect volumes in a single page using global-active device technology.

Before you begin

Register Data Instance Director in Storage Advisor.

Procedure

1. Click **Servers** on the dashboard to open the **Servers** page or navigate to the detail page for a server.
2. Select a server, then select **Create, Attach and Protect Volumes with High Availability**.

STORAGE SYSTEM
410031

SUBSCRIBED CAPACITY
147% (208.61 TiB)

VIRTUAL STORAGE MACHINE
Nothing selected

Capacity

141.65 TiB Total

VOLUME LABEL	LABEL SUFFIX	NUMBER OF VOLUMES	SIZE	POOL TYPE	POOL TIER	POOL	CAPACITY SAVING
Volume	0	1	1 GiB	Thin	Silver	Auto Selected	No

3. Configure volumes for the specified storage system.

You can switch to another storage system by using the drop-down **Storage System** list. If you want to add the volume to a virtual storage machine, use the **Virtual Storage Machine** list. If you don't choose a VSM the meta-resource group will be used.

- Select the number of volumes.
- Enter the volume label and select a suffix for it.
- Select the size.
- Select the volume unit: **GiB**, **TiB**, or **PiB**.
- Select the pool type: **Thin** or **Tiered**.
- For a Thin pool, select the tier: **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 Storage Advisor selects the best pool for provisioning the volume based on utilization and tier requirements.

4. If desired, select a type of **Capacity Saving: Compression** or **Deduplication and Compression**.



Note: Capacity saving can be set for volumes based on tiered pools only for VSP F1500, VSP G1000, and VSP G1500 models 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**.

- When you have made your choices, click the plus sign (+) to add the volume row to the list of volumes that will be created. Add more rows as needed.
- Click **Next** to choose attachment settings.

Attach volumes and select secondary servers

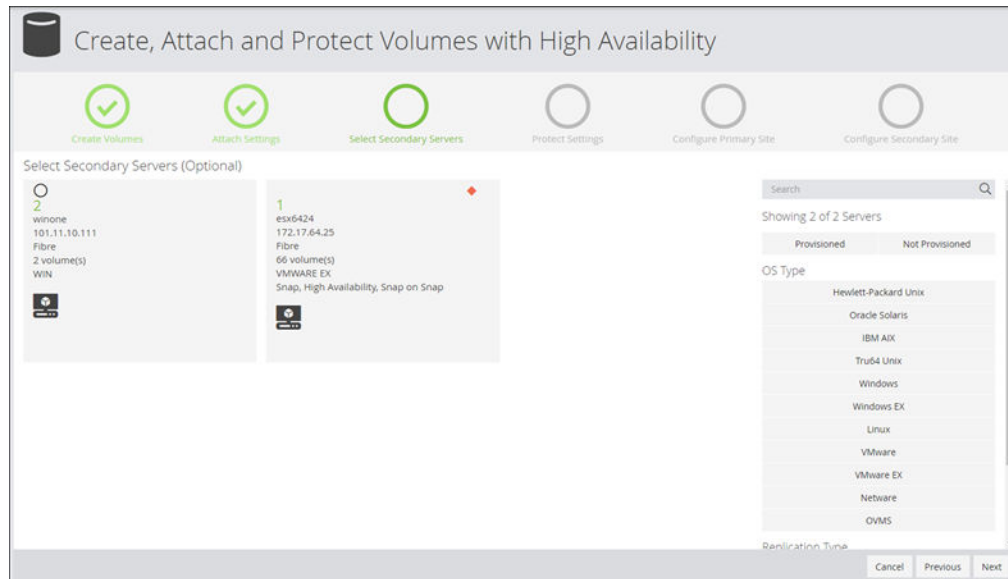
Before you begin

- In the Settings menu, access the Data Instance Director Settings page and register Data Instance Director .
- Make sure that a quorum disk is available.
- Make sure that a virtual storage machine is available unless you want use the meta-resource group of the primary storage system. The virtual storage machine you select must have one or more undefined resources available.

# OF VOL.	LABEL	LABEL S...	SIZE	POOL TYPE	TIER	POOL
1			1 GB	HDP	Silver	Auto Selected

Procedure

1. 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 Storage Advisor.
2. The prepopulated **Host Mode Option** will depend 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.
Storage Advisor 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.
3. Select the **LUN Alignment**.
By default, Storage Advisor uses the LUN number that is common to the servers. If attachment is to only one server, this setting has no effect.
4. The **Auto Create Zone** is set to **No** by default. You can set it to **Yes** to automatically create zones.
5. Leave **ALUA** set to **Enabled** if you want to set preferred paths.
6. Click **Next** to proceed to the **Select Secondary Servers (Optional)** panel.



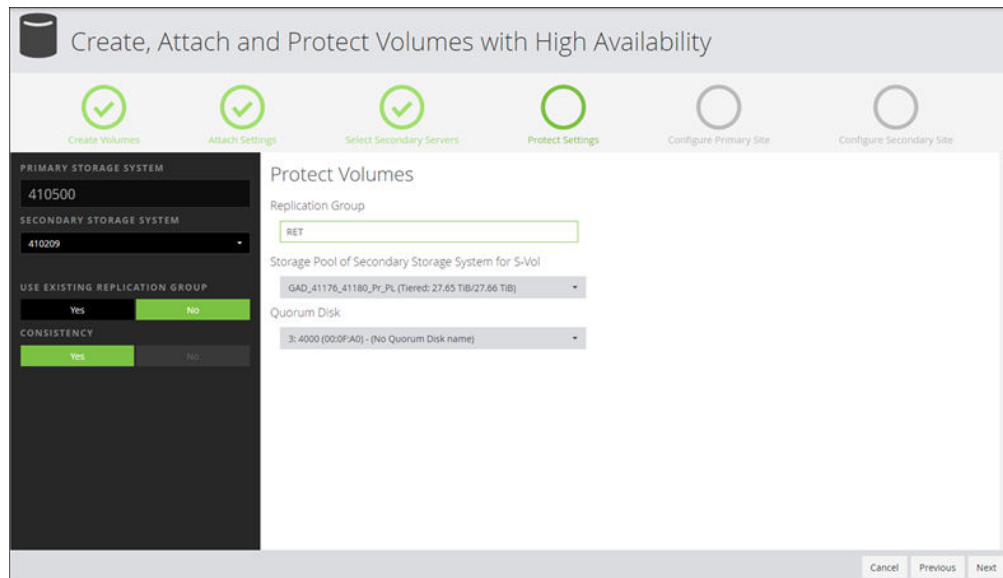
7. Optionally, select a secondary server and click **Next** to proceed to the **Protect Volumes** panel.

Select volume protection options

Select a secondary storage system and other options.

Procedure

1. Select a secondary storage system.
2. Choose whether to use an existing replication group.



3. **Consistency** is set to **Yes** and cannot be edited. Using consistency means that copy operations will run on all pairs in the group simultaneously.

4. Select a replication group from the options in the **Replication Group** list or enter a new name. Up to 26 alphanumeric characters, and also 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.
5. Select the pool to use for S-Vols and select a quorum disk.
6. Click **Next** to configure the primary site.

Configuring connections to the primary and secondary sites

Configure the primary and secondary sites to complete high availability provisioning.

Procedure

1. In the **Configure Primary Site** panel, connect ports in the primary storage system to the server in the primary site and to the server in the secondary site, if there is one. You can configure preferred and unpreferred paths.



2. 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.
3. Click **Submit** to create a job to create volumes, attach to servers and set up data protection.
4. You can monitor the job in the **Jobs** page.



Note: If the job does not complete successfully, access Data Instance Director 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 Storage Advisor in Data Instance Director, specify LDEV IDs in decimal format per line and not in hex format or range format.

Dashboard / Nodes + / fk_rg12112018_qd0 / Edit

Edit Node - Hitachi Block Host 'fk_rg12112018_qd0'

Specify Logical Devices

Enter Logical Devices using any of the following formats:

- LDEV_ID - for a single logical device, e.g., 100, 0x10
- LDEV_ID-LDEV_ID - for a logical device range, e.g., 200-299, 0x01-0x0F
- Host Group ID - for all logical devices within the host group, e.g., CL1-A-0, CL10-A-0

Logical Devices

169
177
139

One entry per line.

Cancel Previous Next

Protect existing volumes with high availability

You can protect existing volumes with high availability.

Procedure

1. Access the **Volumes** page and select one or more attached volumes. They cannot belong to different VSMs. Select **Protect Volumes with High Availability**.

Protect Volumes with High Availability

Attach Settings | Select Secondary Servers | Protect Settings | Configure Primary Site | Configure Secondary Site

STORAGE SYSTEM
410209

VIRTUAL STORAGE MACHINE ID
425207-VSPF400-F600andVSPG...

HOST MODE
AutoSelect

HOST MODE OPTION
40 - VVol expansion, 73 - Support Option for WS...

LUN ALIGNMENT
Yes No

AUTO CREATE ZONE
Yes No

ALUA ENABLED
Yes No

Primary Servers related to selected Volumes

SERVER NAME	SERVER ID	OS TYPE	ATTACHED VOLUME IDS
win9121	1	Windows EX	804 (00:03:24)

Cancel 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 Storage Advisor.

- The prepopulated **Host Mode Option** will depend 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.

Storage Advisor 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.

- Select the **LUN Alignment**.

By default, Storage Advisor uses the LUN number that is common to the servers. If attachment is to only one server, this setting has no effect.

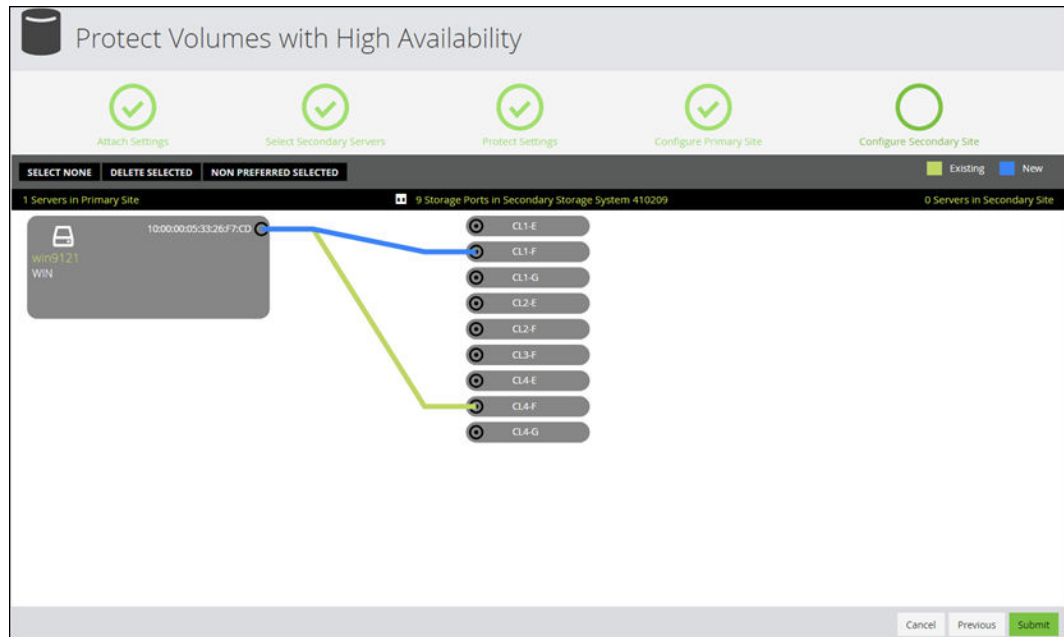
- The **Auto Create Zone** is set to **No** by default. You can set it to **Yes** to automatically create zones.

- ALUA Enabled** is set to **Yes**. If it is set to **No**, you must rediscover volume or restart the OS after job completion.

- Click **Next** to select secondary servers, if applicable.

- Click **Next** to choose a secondary storage system and replication group, and quorum disk.

9. In the **Protect Volumes** panel, select a secondary storage system.
10. Choose whether to use an existing replication group.
11. **Consistency** is set to **Yes** and cannot be edited. Using consistency means that copy operations will 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 also 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.
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** panel, you can view the path of the primary server. If you need to change it, you will need to edit the LUN path before applying data protection.




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** page.



Note: If the job does not complete successfully, access Data Instance Director 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 Storage Advisor in Data Instance Director, specify LDEV IDs in decimal format per line and not in hex format or range format.



Chapter 4: High availability provisioning for VSP G1x00, F1500, VSP Gx00 and VSP Fx00

Prerequisites to provisioning for high availability

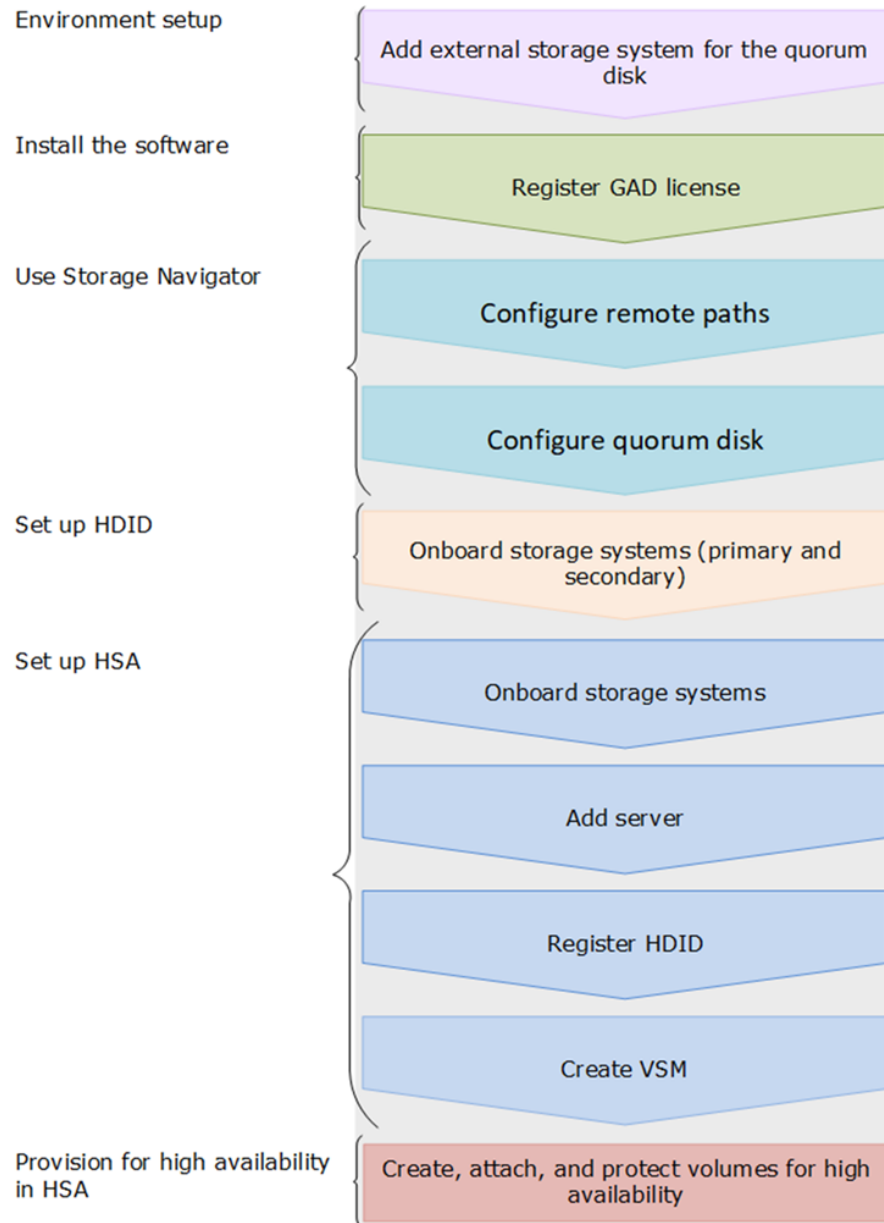
Review the conditions required to provision for high availability.

Make sure that the following conditions are met before provisioning for high availability:

- Data Instance Director v6.6.1 or later is registered in Storage Advisor.
- The registered Storage Advisor satisfies the following conditions
 - Both primary and secondary storage systems are registered with Data Instance Director.
 - Global-active device licenses for both primary and secondary storage systems are installed in Data Instance Director.
- At least one Fibre Channel server must be registered.
- The primary and secondary storage systems must have the following configured:
 - A quorum disk. (The same disk and ID should be assigned in each storage system.)
 - Remote paths , which should be configured bidirectionally.
 - - If the VSM in the primary storage system is using the meta_resource, a VSM with at least one undefined host group must be configured in Storage Advisor for the secondary storage system.
 - If the VSM in the primary storage system is not using the meta_resource, a VSM with at least one undefined volume and one undefined host group must be configured for the primary storage system and, a VSM with at least one undefined host group must be configured for the secondary storage system in Storage Advisor.
- Make sure that both primary and secondary storage systems:
 - Are onboarded in Storage Advisor.
 - Have at least one Fibre port.
 - Have global-active device licenses.
 - Have enough room to create volumes and remote pairs including CTG.
 - Have a Thin or Tiered pool with enough capacity.

Configuration workflow for high availability

The following figure shows the workflow for configuring and provisioning for high availability.



Register global-active device license

Use license keys to install software.

Before you begin

- You must have the Storage Administrator (Initial Configuration) role to perform this task.
- You must install a license key for each software application before you use it.

Procedure

For VSP Fx00 models and VSP Gx00 models, perform the following steps:

1. From the **Maintenance Utility** menu, click **Licenses**.
2. In the **Licenses** window, click **Install**.
3. Select whether to enter a key code or specify a license key file:
 - **License Key Code:** enter the license key code for the software.
 - **License Key File:** Specify a license key file to install the software. Click **Browse** and specify the license key file. You can use a file name of up to 200 alphanumeric characters, excluding several symbols (" \ ; : * ? < > | / ,). The file extension is ".plk".
4. Click **Apply**.
5. Click **OK** in the confirm window.

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

For Virtual Storage Platform F1500 and VSP G1x00 versions, perform the following steps:

6. From the **Settings** menu, click **Environmental Settings**.
7. In the **Environmental Settings** window, click **Install Licenses**.
8. Select whether to enter a key code or specify a license key file:
 - **License Key Code:** enter the license key code for the software.
 - **License Key File:** Specify a license key file to install the software. Click **Browse** and specify the license key file. You can use a file name of up to 200 alphanumeric characters, excluding several symbols (" \ ; : * ? < > | / ,). The file extension is ".plk".
9. Click **Add** and then click **Finish**.
10. Click **Apply**.

Adding the external storage system for the quorum disk

Install an external storage system for the quorum disk .

The storage system must be supported by Universal Volume Manager for connection as external storage.

Connecting the primary and secondary storage systems

To connect the primary and secondary storage systems, you will first set the port attributes on both storage systems, physically connect the storage systems, and then add the remote paths between the storage systems.

Defining port attributes for the remote connection

Set the port attributes for the ports that connect the primary and secondary storage systems.

Before you begin

- The Storage Administrator (System Resource Management) role is required.
- The number of hosts connected to a target port must be limited to 128 or fewer to avoid disconnection.

Procedure

1. Click **Storage Systems**, and then expand the Storage Systems tree.
2. In the Storage Systems tree, click **Ports/Host Groups/iSCSI Targets**.
3. In the **Ports/Host Groups/iSCSI Targets** window, click the **Ports** tab, select the port to be changed, and click **Edit Ports**.
4. In the **Edit Ports** window, select the appropriate **Port Attribute: RCU Target** or **Initiator**. Set the primary storage system port as the RCU Target and the secondary storage system port as the Initiator.
5. Click **Finish** and then **Apply**.
6. Repeat step 4, but set the secondary storage system port as the RCU Target and the primary storage system port as the Initiator.

Adding a remote connection

Add a remote connection to register a storage system of a secondary site to the storage system of a primary site. Also, add a remote connection from the storage system of a secondary site to the storage system of a primary site. When a remote connection is added, both storage systems are ready to perform operations on GAD. You can also set a remote path between storage systems when the remote connection is added.



Note:

- Remote path operations cannot be performed during microcode/firmware exchange processing. Before performing remote path operations, make sure that microcode/firmware exchange processing is complete.
- Remote path operations cannot be performed when microcode/firmware exchange processing has been interrupted (for example, due to user cancellation or error). Before performing remote path operations, make sure that microcode/firmware exchange processing has completed normally.

Before you begin

- Storage Administrator (Remote Copy) role is required.
- Physical paths are set.
- VSP 5000 series: The port attributes of local and remote storage systems are defined for GAD.
- You know the remote storage system model, serial number, and path group ID.

Procedure

1. In the **Storage Systems** tree, select **Replication**, and then **Remote Connections**.
2. Select the **Connections (To)** tab.
3. Click **Add Remote Connection** to open the **Add Remote Connection** window.
4. (VSP 5000 series) In **Connection Type**, select **System**.
5. Set each items in **Remote Storage System**.
 - **Model:** Select the remote storage system model.
 Select 7 for VSP F1500 and VSP G1x00, 18 for VSP Fx00 models and VSP Gx00 models.
 VSP Gx00 models and VSP Fx00 models (18) can be selected only when **System** is selected for **Connection Type**.
 - **Serial Number:** Enter the five or six digit serial number for the remote storage system.
 Specify the serial number for the storage system when using a volume in a virtual storage machine. Do not specify the serial number for the virtual storage machine.
6. Set the items in **Remote Paths**.
 - **Path Group ID:** Select the ID for the path group.
 - **Minimum Number of Paths:** Specify the minimum number or paths that are required for each remote storage systems that are connected to the current local storage system. When the number of normal paths become fewer than the value specified in **Minimum Number of Paths**, the local storage system suspends all GAD pairs that will be affected, and prevents the server performance from being harmed due to insufficient number of paths.
 - For **Select Type**, select the port type. Then select the port to be used for both the local storage system and the remote storage system. To add more paths, click **Add Path**. You can also add and delete paths later as needed using the **Add Remote Paths** window.
7. Enter the **RIO MIH Time**, if necessary.
 You can enter a value between 10 and 100. The default setting is 15.
 The RIO MIH (Remote I/O Missing Interrupt Handler) is the waiting time from when copy starts until when it ends. This value applies to the slots which received the request of copying data between storage systems.

8. Enter **Round Trip Time** in milliseconds, if necessary.

The round trip time is a time limit for data to travel from the P-VOL to the S-VOL. This value is the reference value to control the copy pace of the initial copy automatically when the initial copy is performed, and to lessen the impact to the response time of the remote I/O for the update I/O.

9. Click **Finish**.

10. In the **Confirm** window, check the settings you made, and then enter the task name in **Task Name**.

11. Click **Apply**.

The task is registered. The **Tasks** window appears if the **Go to tasks window for status** check box is selected.

12. On the **Tasks & Alerts** tab, you can click the task name to view progress and other details.

Creating the quorum disk

Defining port attributes on VSP G1000,VSP G1500, and VSP F1500 for connecting the primary and secondary storage systems.

Set the ports on the primary and secondary storage systems to the "External" attribute in preparation for connecting to the external storage system.

Before you begin

- The Storage Administrator (System Resource Management) role is required.
- The number of hosts connected to a target port must be limited to 128 or fewer to avoid disconnection.

Procedure

1. Click **Storage Systems**, and then expand the Storage Systems tree.
2. In the Storage Systems tree, click **Ports/Host Groups/iSCSI Targets**.
3. In the **Ports/Host Groups/iSCSI Targets** window, click the **Ports** tab, select the port to be changed, and click **Edit Ports**.
4. In the **Edit Ports** window, select the **Port Attribute: External**.
5. Click **Finish** and then **Apply**.

Discover and add external volumes

Scan for and add external volumes.

Procedure

1. In Storage Navigator, expand the Storage Systems tree and click **External Storage**.

2. Click **Add External Volumes**.
3. In the **Add External Volumes** window, click **Create External Path Group** and enter an **Initial External Path Group ID**. Click **Discover External Target Ports**.
4. Select an **External Storage System** and then click **Discover External Target Ports**.
5. Select the ports you want to scan and click **Add** to move them to **Selected External Ports**.
6. Click **OK** and then **Next**.
If there are no external volumes discovered, you can add one to the host group.
7. In the **Add External Volumes** window click **Next**.
8. Enter an LDEV name and then select the volumes you want to use in **Discovered External Volumes**.
9. Click **Finish** and then **Apply**.

Adding the quorum disk

Add the quorum disk on the primary and secondary storage systems.

Before you begin

- Storage Administrator (Provisioning) role is required.
- The mapping of volumes of the external storage system for the quorum disk has been completed.

Procedure

1. In the Storage Systems tree, select Replication, and then Remote Connections.
2. Select the Quorum Disks tab.
3. Click Add Quorum Disks to open the Add Quorum Disks window.
4. Select Quorum Disk ID.
5. In the Available LDEVs table, select the volume you want to set to the quorum disk.
6. Select Remote Storage Systems.
7. Click Add. To remove the selected quorum disks from the Selected Quorum Disks table, select the quorum disk, and then click Remove.
8. Click Finish.
9. In the Confirm window, check the settings you made, and then enter the task name in Task Name.
10. Click Apply. The task is registered. The Tasks window appears if the Go to tasks window for status check box is selected.

Adding block storage to Data Instance Director


Onboard the primary and secondary storage systems to Data Instance Director

Hitachi Block Storage Node Wizard

This wizard is launched when a new Hitachi Block Node is added to the Nodes Inventory.

To add a block node to the Nodes Inventory, click Nodes in the navigation sidebar. Click the plus sign (+) in the Nodes Inventory and select Storage in the Create Node page. Select Hitachi Block Device as the Storage Type to launch the Hitachi Block Storage Node Wizard.

Figure 16 Hitachi Block Device Wizard - Specify Node Name

Control	Description
Node Name	Enter a name for the Hitachi Block storage node.
I confirm that ...	<p>This checkbox must be checked to proceed with the node configuration.</p> <div>  Caution: Resources or replication relationships created or adopted by HDID must only be managed, modified and deleted via HDID. Failure to do so will cause unpredictable consequences and is not supported unless specifically advised to do so by the HDID documentation. </div>

Create Node - Hitachi Block Device

Allocate node to Access Control Resource Group

This node will be added to the 'default' resource group. Select an additional resource group as required.

Name	Description
<input type="radio"/> myResourceGroup	A user defined resource group

Cancel Previous **Next**

Figure 17 Hitachi Block Device Wizard - Allocate node to Access Control Resource Group

Control	Description
Resource Groups	Select the resource group(s) to which this node will be allocated for the purposes of RBAC. All nodes are automatically allocated to the 'default' resource group.

Create Node - Hitachi Block Device

Select proxy node



Proxy Node

Select a Node ▼

Proxy node requires a block command device configured and CCI / RAIDCOM management tools installed.

Cancel Previous Next

Figure 18 Hitachi Block Device Wizard - Select proxy node

Control	Description
Proxy Node	<p>Select an HDID node to act as a proxy.</p> <div>  Caution: ISM nodes and their associated CMDs used to control storage devices must not be shared with other applications. </div> <div>  Note: The proxy node is responsible for interfacing with the Block storage device. It can be a Windows or Linux machine with the HDID Client software installed and must be connected via a command device to the Block storage device. The command device must <i>only</i> have user authentication enabled. The proxy node must have supported version of CCI installed. Refer to the Command Control Interface documentation available at https://knowledge.hitachivantara.com. </div>





Figure 19 Hitachi Block Device Wizard - Select Metadata Directory

Control	Description
Common Metadata Directory	<div>Enter a directory on the Proxy node where HDID can place metadata files related to Block snapshots and replications. Click the Browse button to open the Path Dialog if required.</div> <div> Note: The metadata directory is defined once for this proxy and all storage nodes on this proxy will use this setting. It cannot be changed after initial configuration.</div>

Create Node - Hitachi Block Device

Specify Device

☒ Select from detected storage devices

Select a Storage Serial Number

If the desired serial does not appear in this list then it may not have an available fibre command device on the selected proxy.


☐ Specify by IP

IP Address

Port Number

Cancel Previous Next

Figure 20 Hitachi Block Device Wizard - Specify Device

Control	Description
Select from detected storage devices	Select this option to specify the hardware storage device by serial number. A list of storage device serial numbers, available to the proxy node selected in the previous step, is displayed in the dropdown menu below.
Specify by IP	<p>Select this option to specify the hardware storage device using the IP address and port number of an IP command device on the storage device. Additional fibre and IP command devices can be added at a step later in the wizard.</p> <div> <p> Note:</p> <p>For HUS VM storage devices, use the IP address of the SVP.</p> <p>For VSP storage devices, use the IP address of CTL1 or CTL2. DO NOT use the IP address of the SVP.</p> </div>

Create Node - Hitachi Block Device

Specify credentials for device

Storage Device Serial Number

410297



Username

Password

The device account requires the following roles: Storage Administrator (Provisioning, Local Copy, Remote Copy) and Security Administrator (View Only, View and Modify).

Cancel Previous Next

Figure 21 Hitachi Block Device Wizard - Specify credentials for device

Control	Description
Storage Device Serial Number	Displays the serial number of the Block device specified in the previous step.
Username	<p>Enter the username for the Block device.</p> <p> Note: The username specified must be a member of the <i>Storage Administrator (Provisioning, Local Copy, Remote Copy)</i> and <i>Security Administrator (View Only, View and Modify)</i> on the Block device. If the Block device cannot be accessed or its credentials are invalid then the node will fail authorization. The configuration wizard can be reopened to correct any errors.</p>
Password	<p>Enter the password for the Block device.</p> <p> Note: The password for authorizing a Block device must contain only useable CCI command characters: A-Za-z0-9'-./: @ \ _</p>

Create Node - Hitachi Block Device

Specify LDEV Provisioning Range

LDEV Range

☒ All
☐ User defined

Start

0x00

End

0x00

Cancel Previous **Next**




Figure 22 Hitachi Block Device Wizard - Specify LDEV Provisioning Range

Control	Description
All	Select this option if you want HDID to automatically detect the LDEV range from which snapshots and replications should be allocated.
User defined	Select this option if you want to manually specify the LDEV range from which snapshots and replications should be allocated.
Start	Enabled only if User defined is selected. Enter the lower limit of the LDEV range to use for allocation.
End	Enabled only if User defined is selected. Enter the upper limit of the LDEV range to use for allocation.

Create Node - Hitachi Block Device

Command Device Specification and Priority

If any command devices are specified HDID will only use these command devices. If no command devices are specified HDID will attempt to use any fibre based command device available to the proxy. IP based command devices must be specified for HDID to use them. HDID will attempt commands on command devices in the order they are specified, failing over to the next if there is an issue with a command device.

☐ Select All (0)   

Priority	Type	LDEV ID	IP Address	Port	Status
+					
<input type="radio"/> 1	Fibre	Any Available	-	-	

Cancel Previous **Next**



Figure 23 Hitachi Block Device Wizard - Command Device Specification and Priority




The user able to specify zero or more fibre or IP command devices in priority order.

If no command devices are specified, then HDID will attempt to control the hardware storage device via any fibre connected command device, available to the Proxy Node specified, in an order specified by HORCM.

If one or more command devices are specified, then HDID will attempt to control the hardware storage device via a command device in the order specified by the user. If the first command device fails, HDID will progress to the next. If all specified command devices fail then the operation fails. HDID will not attempt to use any command devices that are not specified, even if they are available.

for example, it is possible to specify a specific fibre command device, followed by any fibre command device, followed by a specific IP command device.

Control	Description
 Edit	Enabled when only one command device is selected. Launches the Configure Command Device wizard below to allow the settings to be edited.
 Increase Priority	Enabled when only one command device is selected. Increases the priority of the selected command device.

Control	Description
 Decrease Priority	Enabled when only one command device is selected. Decreases the priority of the selected command device.
 Delete	Enabled when one or more command devices are selected. Deletes selected command device.
 Add	Launches the Configure Command Device wizard below to guide you through setting up a fibre or IP command device.

Create Node - Hitachi Block Device

Configure Command Device

Fibre

IP

Fibre Command Device Options

☒ Use any available fibre command device

☐ Select from detected fibre command devices

Select a Fibre Command Device ▼

Cancel Discard Previous Apply

Figure 24 Hitachi Block Device Wizard - Configure Command Device - Fibre

Control	Description
Use any available fibre command device	Select this option to insert an entry in the command device list that allows HDID to use any available fibre command device.

Control	Description
Select from detected fibre command devices	Select this option to insert a specific fibre command device in the list. The detected fibre command devices are displayed in the dropdown menu below using their decimal LDEV ID.

Create Node - Hitachi Block Device

Configure Command Device

Fibre

IP

IP Command Device Options

IP Address

Port Number

Cancel Discard Previous Apply

Figure 25 Hitachi Block Device Wizard - Configure Command Device - IP



Note: When configuring IP command devices for VSP storage devices, we recommend adding one for CTL1 and one for CTL2, to maintain dual redundancy.



Control	Description
IP Address	<p>Enter the IP address of the command device to add to the list.</p> <p>Note:</p> <p>For HUS VM storage devices, use the IP address of the SVP.</p> <p>For VSP storage devices, use the IP address of CTL1 or CTL2. DO NOT use the IP address of the SVP.</p>

Control	Description
Port Number	Enter the port number of the command device.

Create Node - Hitachi Block Device

Specify LDEV Ranges for each VSM

Define a virtual LDEV range for each VSM serial you intend to use within HDID. The ranges are used to control the virtual LDEV IDs used for replications and snapshots (excluding GAD), and should be defined to not include the IDs of any GAD volumes. Failure to provide such a range (or providing an incorrect range) may result in ID clashes when attempting to set up GAD replications.

☐ Select All (0)  



Virtual Serial	Start	End
+		


Cancel Previous Next

Figure 26 Hitachi Block Wizard - Specify LDEV Ranges for each VSM



Note: GAD replications require P-VOLs and S-VOLs to have matching virtual serial numbers and virtual LDEV IDs. To avoid virtual LDEV ID collisions between GAD volumes and non-GAD S-VOLs (created by HDID for other types of replications and snapshots), it is possible to define virtual LDEV ID ranges to be used by those non-GAD operations. Virtual LDEV ranges can be specified for each VSM (Virtual Storage Machine) to be used.

Control	Description
 Edit	Enabled only if one Virtual LDEV range is selected. The Configure Virtual LDEV Range page of the wizard (see below) is displayed to enable a port to be added.
 Delete	Enabled only if one or more Virtual LDEV ranges are selected. Deletes the Virtual LDEV range(s) from the list.

Control	Description
 Add	Adds a new Virtual LDEV range to the list. The Configure Virtual LDEV Range page of the wizard (see below) is displayed to enable a Virtual LDEV Range to be added.
Virtual Serial(s)	The Virtual Serial(s) and associated LDEV Ranges that will be used.

Create Node - Hitachi Block Device

Configure Virtual LDEV Range

VSM Serial Number

Start of Virtual LDEV range

End of Virtual LDEV range

Figure 27 Hitachi Block Wizard - Configure Virtual LDEV Range



Caution: These ranges are used to control the virtual LDEV IDs used for non-GAD replications and snapshots. They must be defined to exclude the IDs of any GAD volumes. Failure to provide such a range (or providing an incorrect range) may result in ID clashes when attempting to set up GAD replications.

Control	Description
VSM Serial Number	Enter the serial number of the VSM you intend to use within HDID.
Start of LDEV range	Enter the lower limit of the LDEV range to use for allocation.
End of LDEV range	Enter the lower limit of the LDEV range to use for allocation.







Figure 28 Hitachi Block Device Wizard - Specify ports used for provisioning

Control	Description
 Edit	Enabled only if one provisioning port is selected. The Specify Port page of the wizard (see below) is displayed to enable a port to be added.
 Delete	Enabled only if one or more provisioning ports are selected. Deletes the port(s) from the list.
 Add	Adds a new provisioning port to the list. The Specify Port page of the wizard (see below) is displayed to enable a port to be added.
Provisioning Port(s)	Lists the ports that will be used for provisioning.

Create Node - Hitachi Block Device

Specify Port

Port

CL 1 - A

Cancel Discard Previous Apply

Figure 29 Hitachi Block Device Wizard - Specify Port



Note: If more than one provisioning port is selected, then the port with the least amount of LUNs will be used.

Control	Description
Port	<p>Enter the port identifier in the following format:</p> <p>CL_c-s</p> <p>where:</p> <ul style="list-style-type: none"> c is the physical channel number in the range 1...n s is the physical slot number in the range A...Z

Create Node - Hitachi Block Device

Summary of 'myBlock'

Proxy Node
Client1

Storage Device Serial Number
410297

Username
lanH

LDEV Provisioning Range
All Available

Configured Command Devices

Type	LDEV ID	IP Address	Port
Fibre	Any Available	-	-

Ports used for Provisioning

Port
CL1-A

VSM Virtual LDEV Ranges

VSM Serial Number	Start	End
123456	0x00	0xff

Cancel Previous **Finish**

Figure 30 Hitachi Block Wizard - Summary

Control	Description
Summary	Summary of the settings entered.

Adding a storage system in Storage Advisor

Onboarding a storage system is the process of associating it with Storage Advisor. After the storage system is onboarded, manage it from the Storage Advisor dashboard.

Before you begin

Storage Advisor requires 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 Storage Advisor has access to all custom resource groups and meta resource groups.

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

Procedure

1. On the Storage Advisor dashboard, click **Storage Systems** on the left pane.
2. Click the plus sign (+) to add a storage system.

3. Enter values for the following parameters on the **Onboard Storage System** page.
 - **IP Address:** For a storage system with an SVP, enter the IP address of the external service processor for the storage system you want to discover.
 - **User name and password:** Log in as a user that has administrator privileges on this storage system. For example, you can log in as the user `maintenance`.
4. Click **Submit**.

Result

The Jobs tab is updated with the job called `Create Storage System`. If multiple storage systems are being added, there will be a job for each one.

Wait a while for the storage system to be added. Refresh the Jobs tab to verify that storage system is onboarded.

The dashboard shows 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 where you can see the newly added storage system.

When a storage system is onboarded, Storage Advisor goes through an initialization process where it gathers the information about the current configuration of the storage system. During this time you will see that the ports, volumes, pools, and parity groups in the storage system are "Not accessible". Once the initialization is complete, you can see the port, pool, volume, and parity group information in the storage system details.



Note: If operations are performed outside of Storage Advisor, it takes time to update in Storage Advisor (approximately 20 minutes, depending on cache refresh).

Next steps

1. In the parity group inventory 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.

Adding servers

Use Storage Advisor to add servers so you can attach volumes.

You can add multiple server parameters from a file, or add one server at a time.

There are two methods of adding servers:

- 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, in the order specified: Name, Description, IPAddress, OSType, WWNs (comma separated list of WWNs). All fields are required except Description and IPAddress. Valid OSType values are as follows:

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

Procedure

1. On the Storage Advisor dashboard, click **Servers**. Then click the plus sign (+) to open the **Add Server** page.



Note: iSCSI is not currently supported for high availability.

2. On the **Add Server** page, 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 will populate the page. Example:

```
Name,Description,IPAddress,OSType,WWNS
Esxi,ESXI HOST,10.30.90.200,VMWARE_EX,10:00:00:05:33:26:f7:21
Win,WINDOWS
HOST,10.30.91.80,WIN_EX,"10:00:00:05:33:26:f7:37,10:00:00:05:33:26:f7:36"
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:e0:fd"
ESXi_Cisco_2,ESXi HOST connected to Cisco
Fabric,,VMWARE_EX,"100000053326df1a,100000053326df1b"
```

- Click the plus sign (+) in the table to add a row and enter the required information for Fibre Channel. You can add more servers by clicking the plus sign again.

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

Next steps

Create volumes and attach them to the server.

Registering Data Instance Director

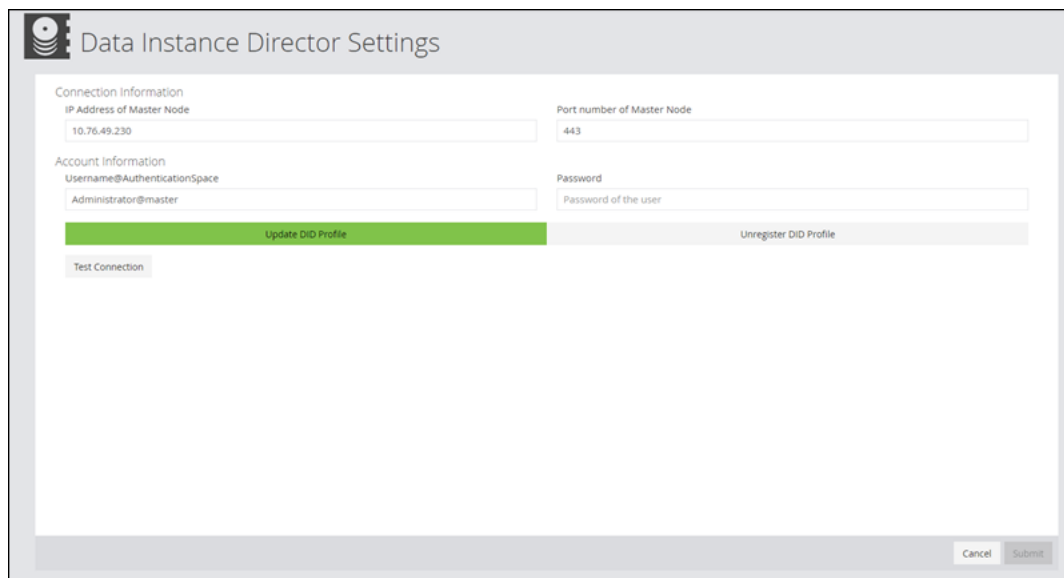
You can use Storage Advisor to register Data Instance Director settings.

Before you begin

Data Instance Director is installed. The user with permissions to perform pair management operations is identified.

Procedure

1. From the **Settings** menu, click **Data Instance Director Settings**.
2. Under **Connection Information**, enter the IP address of the Master node. The port number of the Master node displays automatically.
3. Under **Account Information**, enter the user name and password of the Data Instance Director user who can perform pair management operations for high availability.
4. Click **Test Connection** to verify that you connected to the Master node successfully.
5. Click **Submit**.



The screenshot shows the 'Data Instance Director Settings' web page. It has a header with a logo and the title. Below the header, there are two main sections: 'Connection Information' and 'Account Information'. In the 'Connection Information' section, there is a text box for 'IP Address of Master Node' containing '10.76.49.230' and a text box for 'Port number of Master Node' containing '443'. In the 'Account Information' section, there is a text box for 'Username@AuthenticationSpace' containing 'Administrator@master' and a text box for 'Password' with the placeholder 'Password of the user'. Below these text boxes are two buttons: 'Update DID Profile' (highlighted in green) and 'Unregister DID Profile' (disabled). At the bottom left of the form area is a 'Test Connection' button. At the bottom right of the page are 'Cancel' and 'Submit' buttons.


Creating virtual storage machines

You can use Storage Advisor 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** page.



2. On the **Create VSM** page, select the storage systems you want to use.
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.
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 step and add volumes later.
5. Click **Next** again to add host groups. Add host groups from one storage system at a time and click the plus sign (+) to add them to the list.
6. Click **Submit** to create a job to add the virtual storage machine.

Create, attach, and protect volumes with high availability

Storage Advisor allows you to create, attach, and protect volumes in a single page using global-active device technology.

Before you begin

Register Data Instance Director in Storage Advisor.

Procedure

1. Click **Servers** on the dashboard to open the **Servers** page or navigate to the detail page for a server.
2. Select a server, then select **Create, Attach and Protect Volumes with High Availability**.

Create, Attach and Protect Volumes with High Availability

STORAGE SYSTEM: 410031
 SUBSCRIBED CAPACITY: 147% (208.61 TiB)
 VIRTUAL STORAGE MACHINE: Nothing selected

VOLUME LABEL: Volume LABEL SUFFIX: 0 NUMBER OF VOLUMES: 1 SIZE: 1 GiB
 POOL TYPE: Thin POOL TIER: Silver POOL: Auto Selected CAPACITY SAVING: No

NUM...	VOLUME LABEL	LABEL ...	SIZE	UNIT	POOL ...	POOL ...	POOL ...	CAPAC...
1	Volume		1	GiB	Thin	Silver	Auto Selected	No

Cancel Next

3. Configure volumes for the specified storage system.

You can switch to another storage system by using the drop-down **Storage System** list. If you want to add the volume to a virtual storage machine, use the **Virtual Storage Machine** list. If you don't choose a VSM the meta-resource group will be used.

- Select the number of volumes.
- Enter the volume label and select a suffix for it.
- Select the size.
- Select the volume unit: **GiB**, **TiB**, or **PiB**.
- Select the pool type: **Thin** or **Tiered**.
- For a Thin pool, select the tier: **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 Storage Advisor selects the best pool for provisioning the volume based on utilization and tier requirements.

4. If desired, select a type of **Capacity Saving: Compression** or **Deduplication and Compression**.



Note: Capacity saving can be set for volumes based on tiered pools only for VSP F1500, VSP G1000, and VSP G1500 models 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**.

- When you have made your choices, click the plus sign (+) to add the volume row to the list of volumes that will be created. Add more rows as needed.
- Click **Next** to choose attachment settings.

Attach volumes and select secondary servers


Before you begin

- In the Settings menu, access the Data Instance Director Settings page and register Data Instance Director .
- Make sure that a quorum disk is available.
- Make sure that a virtual storage machine is available unless you want use the meta-resource group of the primary storage system. The virtual storage machine you select must have one or more undefined resources available.

# OF VOL.	LABEL	LABEL S...	SIZE	POOL TYPE	TIER	POOL
1			1 GB	HDP	Silver	Auto Selected

Procedure

1. 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 Storage Advisor.
2. The prepopulated **Host Mode Option** will depend 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.
Storage Advisor 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.
3. Select the **LUN Alignment**.
By default, Storage Advisor uses the LUN number that is common to the servers. If attachment is to only one server, this setting has no effect.
4. The **Auto Create Zone** is set to **No** by default. You can set it to **Yes** to automatically create zones.
5. Leave **ALUA** set to **Enabled** if you want to set preferred paths.
6. Click **Next** to proceed to the **Select Secondary Servers (Optional)** panel.



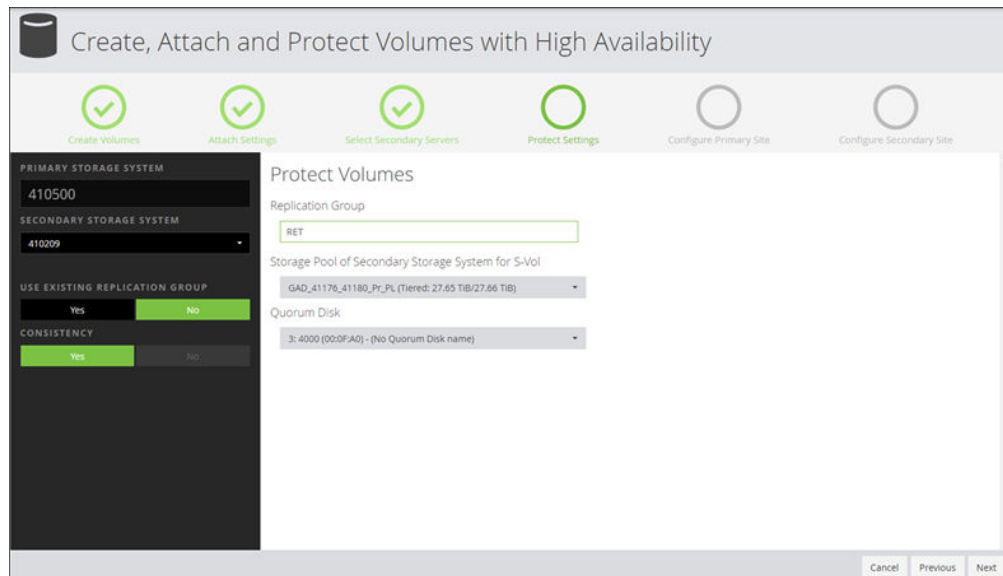
7. Optionally, select a secondary server and click **Next** to proceed to the **Protect Volumes** panel.

Select volume protection options

Select a secondary storage system and other options.

Procedure

1. Select a secondary storage system.
2. Choose whether to use an existing replication group.



3. **Consistency** is set to **Yes** and cannot be edited. Using consistency means that copy operations will run on all pairs in the group simultaneously.


4. Select a replication group from the options in the **Replication Group** list or enter a new name. Up to 26 alphanumeric characters, and also 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.
5. Select the pool to use for S-Vols and select a quorum disk.
6. Click **Next** to configure the primary site.

Configuring connections to the primary and secondary sites

Configure the primary and secondary sites to complete high availability provisioning.

Procedure

1. In the **Configure Primary Site** panel, connect ports in the primary storage system to the server in the primary site and to the server in the secondary site, if there is one. You can configure preferred and unpreferred paths.



2. 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.
3. Click **Submit** to create a job to create volumes, attach to servers and set up data protection.
4. You can monitor the job in the **Jobs** page.



Note: If the job does not complete successfully, access Data Instance Director 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 Storage Advisor in Data Instance Director, specify LDEV IDs in decimal format per line and not in hex format or range format.

Dashboard / Nodes + / fk_rg12112018_qd0 / Edit

Edit Node - Hitachi Block Host 'fk_rg12112018_qd0'

Specify Logical Devices

Enter Logical Devices using any of the following formats:

- LDEV_ID - for a single logical device, e.g., 100, 0x10
- LDEV_ID-LDEV_ID - for a logical device range, e.g., 200-299, 0x01-0x0F
- Host Group ID - for all logical devices within the host group, e.g., CL1-A-0, CL10-A-0

Logical Devices

169
177
139

One entry per line.

Cancel Previous Next

Protect existing volumes with high availability

You can protect existing volumes with high availability.

Procedure

1. Access the **Volumes** page and select one or more attached volumes. They cannot belong to different VSMs. Select **Protect Volumes with High Availability**.

Protect Volumes with High Availability

Attach Settings | Select Secondary Servers | Protect Settings | Configure Primary Site | Configure Secondary Site

STORAGE SYSTEM
410209

VIRTUAL STORAGE MACHINE ID
425207-VSPF400-F600andVSPG...

HOST MODE
AutoSelect

HOST MODE OPTION
40 - V-Vol expansion, 73 - Support Option for WS...

LUN ALIGNMENT
Yes No

AUTO CREATE ZONE
Yes No

ALUA ENABLED
Yes No

Primary Servers related to selected Volumes

SERVER NAME	SERVER ID	OS TYPE	ATTACHED VOLUME IDS
win9121	1	Windows EX	804 (00:03:24)

Cancel 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 Storage Advisor.

- The prepopulated **Host Mode Option** will depend 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.

Storage Advisor 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.

- Select the **LUN Alignment**.

By default, Storage Advisor uses the LUN number that is common to the servers. If attachment is to only one server, this setting has no effect.

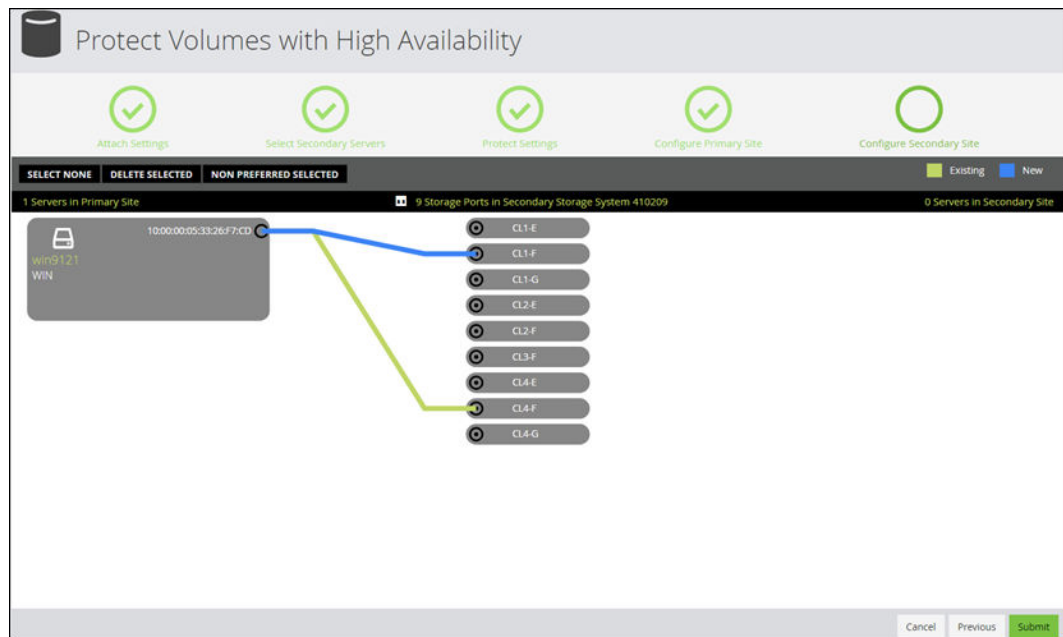
- The **Auto Create Zone** is set to **No** by default. You can set it to **Yes** to automatically create zones.

- ALUA Enabled** is set to **Yes**. If it is set to **No**, you must rediscover volume or restart the OS after job completion.

- Click **Next** to select secondary servers, if applicable.

- Click **Next** to choose a secondary storage system and replication group, and quorum disk.

9. In the **Protect Volumes** panel, select a secondary storage system.
10. Choose whether to use an existing replication group.
11. **Consistency** is set to **Yes** and cannot be edited. Using consistency means that copy operations will 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 also 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.
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** panel, you can view the path of the primary server. If you need to change it, you will need to edit the LUN path before applying data protection.

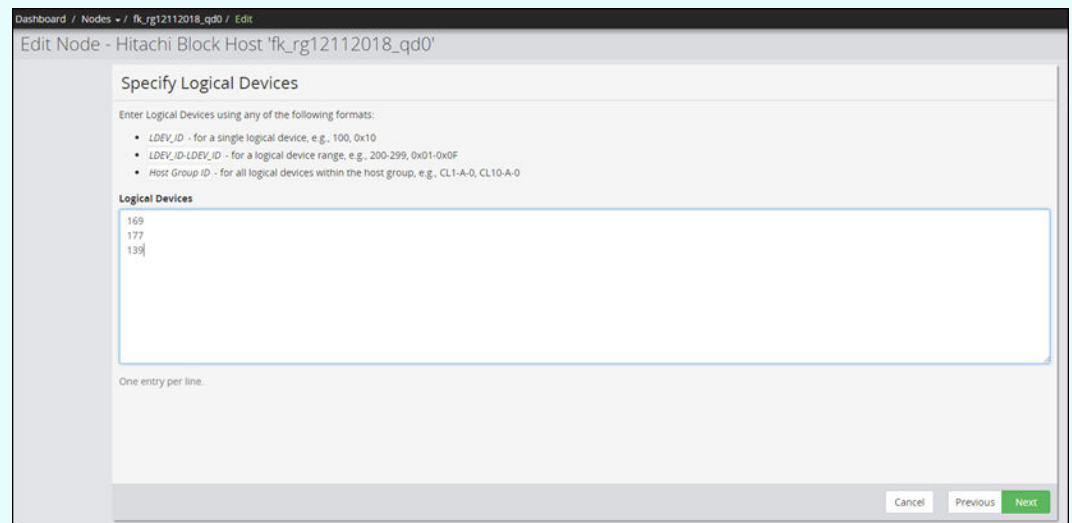


16. Click **Next** to configure ports in the secondary storage system to the server in the primary site, and to the server in the secondary site, if there is one.
17. Click **Submit** to create a job to protect volumes.
18. You can monitor the job in the **Jobs** page.



Note: If the job does not complete successfully, access Data Instance Director 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 Storage Advisor in Data Instance Director, specify LDEV IDs in decimal format per line and not in hex format or range format.



Monitoring and pair operations in Data Instance Director

Chapter 5: Monitoring and pair operations in Data Instance Director

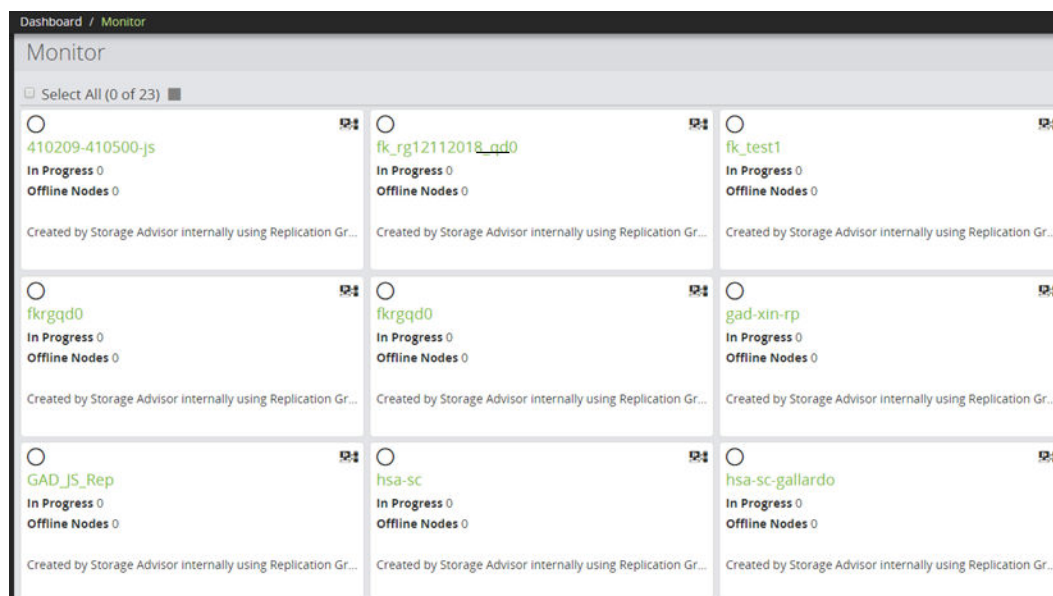
Use Data Instance Director to perform high availability pair operations and monitor data flows.

Before you begin

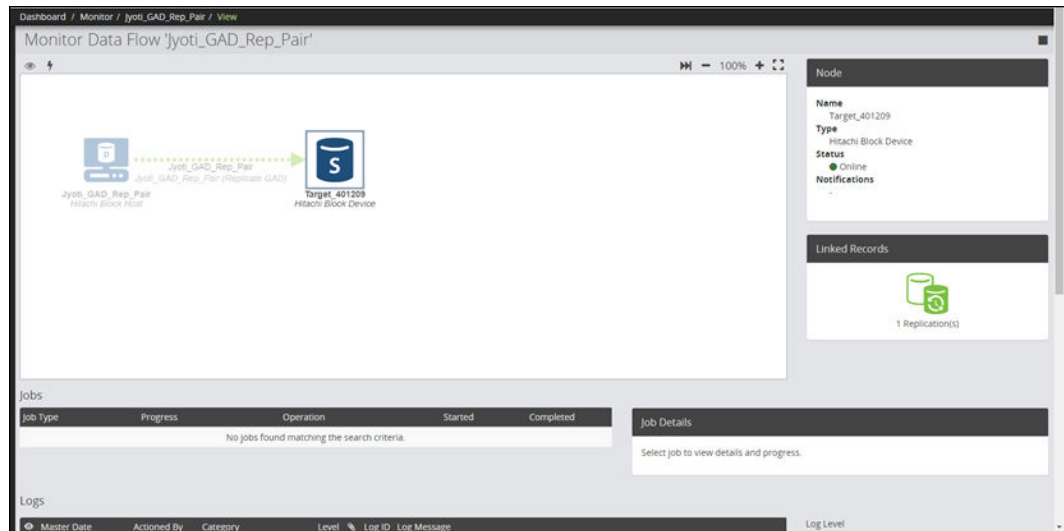
Data Instance Director must be registered in Storage Advisor, and the primary and secondary storage systems must be onboarded in Data Instance Director.

Procedure

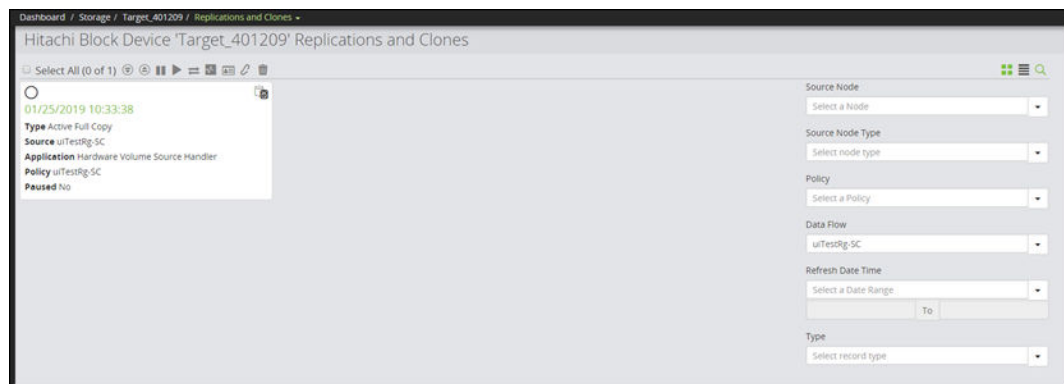
1. Log in to Data Instance Director and click **Monitoring** in the global navigation menu.
2. In the **Monitor** inventory page, you can see all data flows created by Storage Advisor.




3. Click a data flow to open the **Monitor** details page.




4. Click the target block device to expose any related replications in the **Linked Records** panel.
5. Click the **Replications** icon to view the details of related replications and perform operations on them.



6. The following operations are available in this page:
 - **Mount:** Enabled only if one or more Replications are selected. Opens the Hitachi Block Snapshot or Replication Mount Wizard to guide you through mounting the Replication.

 **Note:** The mount operation can take several minutes to complete.

 - **Unmount:** Enabled only if one mounted replication is selected. Unmounts the selected replication.

 **Note:** The unmount operation can take several minutes to complete.

 - **Pause:** Enabled only if one or more replications are selected. Pauses the replication.
 - **Resume:** Enabled only if one or more replications are selected. Resumes a paused replication.

- **Swap:** Enabled only if one or more replications are selected. Opens the Hitachi Block Replication Swap Wizard to guide you through swapping the replication direction.
- **Unsuspend:** If a Swap operation cannot be completed due to a P-VOL or data link fault between the primary and secondary device, the replication will enter the SSWS state (suspended for swapping) indicating that the swap is not yet complete. Unsuspend enables the replication process to be re-established once the cause has been rectified.
- **Transfer RBAC Permissions:** Allows RBAC ownership to be transferred from the current node to another node. Opens the Access Control Transfer Permissions Dialog.
- **Dissociate:** Enabled only if one or more replications are selected. Dissociates a replication that was previously adopted by Data Instance Director. Removes the selected replications from Data Instance Director, including state information such as direction and mount location. The replication remains active on the hardware devices. `DISSOCIATE` must be entered before the command is executed.



Caution: Dissociating a replication removes all knowledge of the replication from Data Instance Director, including state information such as direction and mount location.

Example

For more information about pair operations in Data Instance Director, refer to the *Hitachi Data Instance Director User Guide*.

Hitachi Block Snapshot or Replication Mount Wizard

This wizard is displayed when you mount a snapshot or replication from a Hitachi Block device.

This wizard allows you to expose volumes or mount the selected snapshot or replication.



Caution:

- When a snapshot or replication is mounted, it is made writable. Any changes made to the mounted original snapshot or replication will persist even after it has been unmounted.
- Any attempt to mount two or more copies of the same volume(s) simultaneously on the same machine will fail. The base OS will experience problems if there is already an instance of that file system mounted. This is because snapshots and replications share the same disk partition IDs as the original volume. In extreme conditions it can cause corruption to both disks.

**Note:**

- It is not possible to mount the SVOL of a GAD replication, paused or otherwise.
- For Oracle ASM the disks will be presented to the OS but will need to be manually mounted.
- When mounting a snapshot that contains a mounted sub directory, the subdirectory will be mounted as expected. However, the volume referenced by the subdirectory will also be mounted as a separate drive. Unmount will unmount both the expected and unexpected mounts.
- The mount operation can take several minutes to complete.
- If a mount operation needs to mount multiple disks and one of the mount operations fail, the snapshot/replication destination is shown as mounted rather than partially mounted. The logs will indicate the mount was only partially successful.

**Note:** Operating System Specific Behaviour

OS	Note
Linux	When mounting a Linux snapshot to a different Linux machine; in order for the user and group names to be displayed correctly the users and groups must have the same ID's as the source.
SUSE Linux	SUSE Linux is not able to perform automated mount operations if hosted on VMware. (RHEL and OEL Linux work as expected).
AIX & Solaris	The system command importvg is invoked by HDID to mount snapshots to the user specified location. importvg creates a directory for the user specified location plus an empty directory corresponding to the original mount point. Neither of these directories are removed by HDID when the snapshot is eventually unmounted, although neither will contain any data.

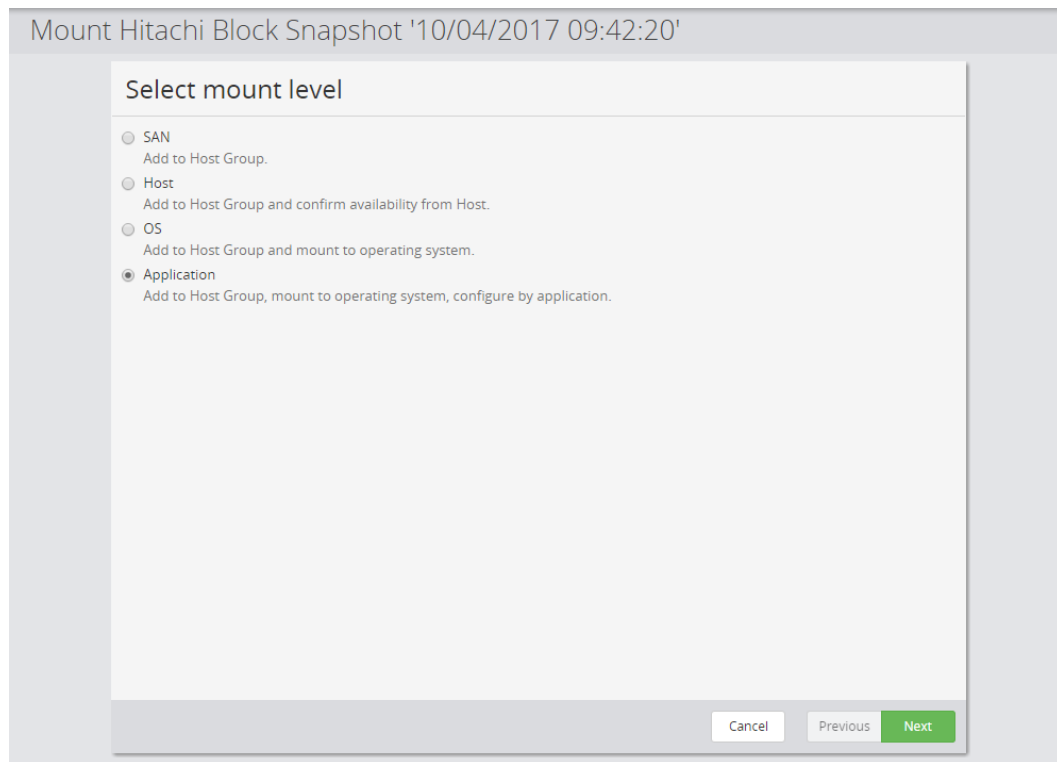


Figure 31 Mount Wizard - Select mount level

Control	Description
SAN	Adds the snapshot or replication to a Host Group. The wizard displays the Select Host Group page next.
Host	Adds the snapshot or replication to a Host Group and confirms that it is available from the specified Host. The wizard displays the Select Host Group and host page next.
OS	Adds the snapshot or replication to a Host Group and mounts it on the specified Host's operating system. The wizard displays the Select Host Group and host page next.
Application	<p>Displayed only for application snapshots. Adds the snapshot or replication to a Host Group and mounts it on the specified Host's operating system. The wizard displays the Select Host Group and host page next. The final step in the wizard provides one of the following application specific mount options:</p> <ul style="list-style-type: none"> Mount Wizard - Select Microsoft SQL Restore Options Mount Wizard - Select Oracle Restore Options

Mount Hitachi Block Snapshot '19/06/2017 11:35:15 ...


Select Host Group

Host Group

Select a Host Group

Cancel Previous Finish

Figure 32 Mount Wizard - Select Host Group (SAN mounts only)

Control	Description
Host Group	<p>Manually specify or select a host group to use to expose a snapshot or replication.</p> <p> Note: When exposing an LDEV, the host group specified must be in the same resource group as the secondary volumes.</p>

Mount Clone '16/08/2018 16:10:05'

Select Host Group and host

Host Group

☒ Automatically discover
☐ Selected

Select a Host Group ▼

OS Host

Select a Node ▼

VMware Node

Select a Node ▼



Only required for VMware based virtual machines.


Datastore

(Automatically Selected) ▼

Cancel Previous **Next**

Figure 33 Mount Wizard - Select Host Group and host (OS and Host mounts only)

Control	Description
Host Group	<p>Select from:</p> <ul style="list-style-type: none"> Automatically discover - Automatically choose a host group to use to expose a snapshot or replication. Selected - Manually specify or select a host group to use to expose a snapshot or replication. <p> Note: When mounting an LDEV to a specific host, the host must have a path to the storage via a host group in the same resource group as the secondary volumes.</p>
OS Host	<p>Specify an HDID <i>OS Node</i> to mount to or expose to.</p> <p> Note: Unless the user selects a host group, the machine where the volume is to be mounted must have an existing volume on the same storage device. If there is no connection between the mount host and the block storage device then HDID will fail the mount operation after a timeout of 30 minutes.</p>

Control	Description
VMware Node	Expose the volumes to the specified VMware host and mount them to the VM and RDM disks. <div>  Note: Exposing using a VMware host requires that a VMware Server node be configured in HDID and that the mount target VM has HDID Client and VMware Tools installed. </div>
Datastore	Specifies a destination datastore when mounting to a VMware VM which is part of a cluster, in which case the default datastore may not be a suitable place to save the RDM mount information. If the datastore field is left blank then mount information is saved alongside the VM.
Next/Finish	If Host was selected in the initial step then the Finish button is displayed. If OS was selected then the wizard displays the Select mount location page next.

Mount Hitachi Block Snapshot '01/02/2018 11:37:48'

Select mount mode

☐ Mount original
Modifications made while mounted are permanent

☒ Mount duplicate (cascaded snapshot)
Modifications made while mounted are lost when unmounted



Mount Pool

Select mount pool ▼

A mount pool is required because the snapshot is not fully provisioned and resides in a Thin Image pool.

Cancel Previous Next

Figure 34 Mount Wizard - Select mount mode

Control	Description																																				
Mount Original	<p>Mounts the original (Level 1) snapshot.</p> <div> Note: If the mounted snapshot is modified then those changes will remain when the snapshot is unmounted.</div>																																				
Mount duplicate (cascaded snapshot)	<p>Only enabled if the original (Level 1) snapshot was created in cascade mode. Mounts a copy of the original snapshot (i.e. a Level 2 snapshot).</p> <div> Note: The mounted snapshot will be discarded when it is unmounted.</div>																																				
Mount Pool	<p>Depending on the parameters specified for the snapshot operation on the data flow, a mount pool might or might not be required.</p> <p>The following table lists all the scenarios:</p> <table><tr><th>Snapshot Pool Type (specified on data flow)</th><th>Provisioning Type (specified on data flow)</th><th>Mount Mode</th><th>Mount Pool</th></tr><tr><td>Thin Image</td><td>Floating Device</td><td>Original</td><td>Required⁽¹⁾</td></tr><tr><td>Thin Image</td><td>Floating Device</td><td>Duplicate⁽⁴⁾</td><td>Required⁽¹⁾</td></tr><tr><td>Thin Image</td><td>Fully Provisioned</td><td>Original</td><td>N/A⁽³⁾</td></tr><tr><td>Thin Image</td><td>Fully Provisioned</td><td>Duplicate⁽⁴⁾</td><td>Optional⁽²⁾</td></tr><tr><td>Hybrid</td><td>Floating Device</td><td>Original</td><td>Optional⁽²⁾</td></tr><tr><td>Hybrid</td><td>Floating Device</td><td>Duplicate⁽⁴⁾</td><td>Optional⁽²⁾</td></tr><tr><td>Hybrid</td><td>Fully Provisioned</td><td>Original</td><td>N/A⁽³⁾</td></tr><tr><td>Hybrid</td><td>Fully Provisioned</td><td>Duplicate⁽⁴⁾</td><td>Optional⁽²⁾</td></tr></table>	Snapshot Pool Type (specified on data flow)	Provisioning Type (specified on data flow)	Mount Mode	Mount Pool	Thin Image	Floating Device	Original	Required ⁽¹⁾	Thin Image	Floating Device	Duplicate ⁽⁴⁾	Required ⁽¹⁾	Thin Image	Fully Provisioned	Original	N/A ⁽³⁾	Thin Image	Fully Provisioned	Duplicate ⁽⁴⁾	Optional ⁽²⁾	Hybrid	Floating Device	Original	Optional ⁽²⁾	Hybrid	Floating Device	Duplicate ⁽⁴⁾	Optional ⁽²⁾	Hybrid	Fully Provisioned	Original	N/A ⁽³⁾	Hybrid	Fully Provisioned	Duplicate ⁽⁴⁾	Optional ⁽²⁾
Snapshot Pool Type (specified on data flow)	Provisioning Type (specified on data flow)	Mount Mode	Mount Pool																																		
Thin Image	Floating Device	Original	Required ⁽¹⁾																																		
Thin Image	Floating Device	Duplicate ⁽⁴⁾	Required ⁽¹⁾																																		
Thin Image	Fully Provisioned	Original	N/A ⁽³⁾																																		
Thin Image	Fully Provisioned	Duplicate ⁽⁴⁾	Optional ⁽²⁾																																		
Hybrid	Floating Device	Original	Optional ⁽²⁾																																		
Hybrid	Floating Device	Duplicate ⁽⁴⁾	Optional ⁽²⁾																																		
Hybrid	Fully Provisioned	Original	N/A ⁽³⁾																																		
Hybrid	Fully Provisioned	Duplicate ⁽⁴⁾	Optional ⁽²⁾																																		

Control	Description
	<p>The following message is displayed in a blue rectangle to explain the options:</p> <ul style="list-style-type: none"> ▪ (1) "A mount pool is required because the snapshot is not fully provisioned and resides in a Thin Image pool." ▪ (2) "A mount pool is optional because the snapshot is fully provisioned and/or resides in a hybrid pool. If a mount pool is not specified the original or duplicate snapshot will use the cascade pool (if configured) or the snapshot pool of the original snapshot." ▪ (3) "A mount pool is not required because the secondary LDEVs are already fully provisioned." ▪ (4) "It is only possible to mount a duplicate of a snapshot created with cascade-mode enabled."

Mount Hitachi Block Snapshot '19/06/2017 12:35:16 ...

Specify mount location

Mount Location



☒ Original
☐ Drive starting at letter
☐ Directory

Select drive letter ▼

Browse

Cancel Previous Finish

Figure 35 Mount Wizard - Select mount location

Control	Description
Original	<p>The snapshot or replication is mounted at its original location.</p> <p> Note: Mounting at the original location will fail if there is already a volume mounted at that location.</p>
Drive starting at letter	<p>When mounting a snapshot or replication that contains multiple volumes, the first volume will mount at the specified drive and subsequent drives are used for each additional volume.</p>
Directory	<p>When mounting a snapshot or replication that contains multiple volumes, each volume will be assigned a separate subdirectory. Click Browse to view the drives and directories on the selected host. To create a new directory, type in the required path.</p> <p> Note: HDID does not check to make sure the directory selected as the mount point is empty. This means it is possible to mount a snapshot inside or even over the top of another mounted volume. This should be avoided.</p>

Mount Wizard - Select Microsoft SQL Restore Options

When mounting block snapshots or replications created by a policy containing an *MS SQL Server* classification, HDID will display the following additional wizard page that allows application specific options to be configured:



Note: Microsoft SQL server databases that are part of a multi-volume snapshot must be mounted using the Mount Location and Directory option in the **Select mount location** page of the **Mount Hitachi Block Snapshot or Replication** wizard, then recovered manually.

Mount Hitachi Block Snapshot '10/04/2017 09:42:20'

Select Microsoft SQL restore options

Restore Type

- ☒ Restore Only
- ☐ Add database to an instance and allow further manual recovery
- ☐ Add database to an instance and bring up database

Instance

Instance Name

Database Owner

Owner Name

Backup Database Name

Database Name

Cancel Previous Finish

Figure 36 Mount Wizard - Select Microsoft SQL Restore Options


Control	Description
Restore Only	Restores the backup without any further action.
Add database to an instance and allow further manual recovery	Puts the database in a <i>restoring</i> state, so that the additional backups can be restored. The restore sequence can restore other backups and replay transaction logs to support roll forward. When the database is in a <i>restoring</i> state, no users can access the database or the database contents.
Add database to an instance and bring up database	Recovers the database to a consistent state captured by the backup and brings the database online. No rolling-forward is performed to ensure that the transaction logs are not replayed.
Instance	Only enabled for Add database to an instance... restore types. The name of the instance.
Database Owner	Only enabled for Add database to an instance... restore types. The database owner's name.
Backup Database Name	Only enabled for Add database to an instance... restore types. The backup database's name.

Mount Wizard - Select Oracle Restore Options

When mounting block snapshots or replications created by a policy containing an *Oracle Database* classification, HDID will display the following additional wizard pages that allows application specific options to be configured:

Figure 37 Mount Wizard - Select Oracle Restore Options

Control	Description	Logs Reset Post Mount	Requires RMAN catalog	Requires control/ spfile in RMAN backup
Restore only	The database is simply mounted. It is left to the database administrator to recover manually.	No	No	No
Recover to last consistent state in backup	The database is recovered to the consistent state which was captured by the backup. The database is brought online. This type of mount can be performed with the data in the backup alone.	Depends (see note 1)	No	No

Control	Description	Logs Reset Post Mount	Requires RMAN catalog	Requires control/spfile in RMAN backup
Recover to point in time	A timestamp is entered which defines the point in time to recover. This requires a connection to the RMAN catalog and logs which are shared with/available to the host.	Depends (see note 1)	Yes	Yes (see note 2)
Recover to system change number (SCN)	A system change number is entered which defines the change point to recover. This requires a connection to the RMAN catalog and logs which are shared with/available to the host.	Depends (see note 1)	Yes	Yes (see note 2)
Recover to current position	<p>The database is recovered to the most current position possible. This requires a connection to the RMAN catalog and logs which are shared with/available to the host.</p> <div>  Note: When mounting, the current position is the latest point in time which is provided by the archive logs referenced in RMAN catalog. It does not include any archive logs or redo logs on the source machine which have not been backed up via RMAN. </div>	Depends (see note 1)	Yes	Yes (see note 2)



Note: (1) In the table above, logs will only be reset when Open Database is selected in the **Post Recovery Options** page of the wizard (see below).



Caution: (2) In the table above, for some recovery scenarios the RMAN catalog needs to hold a control file. RMAN can be configured to add a control/spfile backup every time an archive log backup is performed.

Mount Hitachi Block Snapshot '10/04/2017 12:57:22'

Select Post Recovery Options




☐ Change Oracle database ID (DBID)
☐ Change Oracle database unique name and SID

☐ Disable database schedule
☒ Open Database

Advanced Options

Cancel Previous Finish

Figure 38 Mount Wizard - Select Post Recovery Options

Control	Description
Change Oracle database ID (DBID)	<p>Only available when Recover to last consistent state in backup is selected. Creates a new DBID for the database.</p> <p> Tip: A DBID is a unique, Oracle generated number identifying each database. It is found in control files as well as datafile headers and is used to determine which database that file belongs to.</p>
Change Oracle database unique name and SID	<p>Only available when Recover to last consistent state in backup is selected. A new unique name and SID can be specified for the database.</p> <p> Tip: This changes the <code>unique_database_name</code> which is also used as the SID.</p>
Disable database schedule	<p>Only available when Recover to last consistent state in backup is selected. Disables database internal tasks scheduled for this database.</p> <p> Tip: The Oracle scheduler allows the administrator to schedule SQL commands as jobs. By selecting this option existing schedules will be disabled.</p>

Control	Description
Open Database	If selected, then after recovery the database will be placed in the <i>OPEN</i> state using the <i>RESETLOGS</i> or <i>NORESETLOGS</i> option, as per the requirements of the database. Otherwise the database will be left in the <i>MOUNT</i> state.
Advanced Options	Opens the Advanced Options page of the wizard.

Mount Hitachi Block Snapshot '10/04/2017 12:57:22'

Advanced Mount Options

Oracle Database Memory Target

☐

Database MEMORY_TARGET in GB. Will remove all other memory management related customization of the database.

Local Listener


☐

Network name of the Oracle Net local_listener.

Cancel Discard Previous **Apply**

Figure 39 Mount Wizard - Advanced Mount Options

This page of the wizard is displayed if the Advanced Options button is clicked on the **Select Post Recovery Options** page of the wizard.

Control	Description
Oracle database Memory Target	<p>Sets the database MEMORY_TARGET in GB. Entering a value here will remove all other memory management related customization of the database.</p> <div>  Tip: This allows Oracle databases from very powerful source systems to be deployed on less powerful systems. The PGA and SGA memory areas will be managed by Oracle withing the given memory target. </div>
Local Listener	Sets the network name of the Oracle Net local_listener.

Mount Hitachi Block Snapshot '10/04/2017 12:57:22'

Provide details for changing database ID or name

Password for sys user

Cancel Previous Finish

Figure 40 Mount Wizard - Provide details for changing database ID or name

This page of the wizard will only be displayed if either Change Oracle database ID (DBID) or Change Oracle database unique name and SID options are selected in a previous step.

Control	Description
Password for sys user	<p>Depending on the Oracle version:</p> <ul style="list-style-type: none"> Oracle 11g: The sys user password is required to change the Oracle database ID or database unique name. Oracle 12g: This field can be left empty.

Mount Hitachi Block Snapshot '10/04/2017 12:57:22'

Specify RMAN settings

The RMAN recovery catalog is used to store information about backups performed with the Oracle RMAN utility (e.g. transaction log backups). The catalog is used during database recovery.

RMAN Catalog Name

Name used in the SQL*Net connect string to connect to the RMAN catalog

Username

Password

Cancel Previous **Next**

Figure 41 Mount Wizard - Specify RMAN credentials

This page of the wizard is only displayed if Restore only or Recover to last consistent state in backup is selected in a previous step.

Control	Description
RMAN Catalog Name	For RMAN only. Enter the RMAN Catalog Name as it is entered in the SQL*Net connect string to connect to the RMAN catalog.
Username	For RMAN only. Enter the username for the RMAN catalog.
Password	For RMAN only. Enter the password for the RMAN catalog.

Hitachi Block Replication Swap Wizard

This wizard causes the direction of a replication to be reversed so that the secondary replicates to the primary.



Caution: When a replication is swapped, the secondary becomes writeable and the primary becomes read-only.

**Note:**

- If the swap cannot be completed then the replication pair enters the SSWS (suspended for swapping) state until the swap can be completed. The swap operation must be re-done by the user to complete a suspended swap.
- The flow direction of a replication pair should ONLY be determined by referring to the Storage Inventory. Primary and secondary volume information shown in the Log Attachments Dialog should not be used to infer the flow direction following a swap.

Hitachi Block Replication '19/06/2017 9:27:48 AM' Sw...

Confirm Swap

Confirm Swap

Swapping a replication can potentially cause the **loss or destruction of data**. If you are certain you want to perform this operation enter 'SWAP' in the field above.

Cancel Previous **Finish**

Figure 42 Replication Swap Wizard

Control	Description
Confirm Swap	The word SWAP must be explicitly typed in to confirm the action.

Access Control Transfer Permissions Dialog

This dialog is displayed when the permissions for a resource are being transferred from the current owner to new owner.

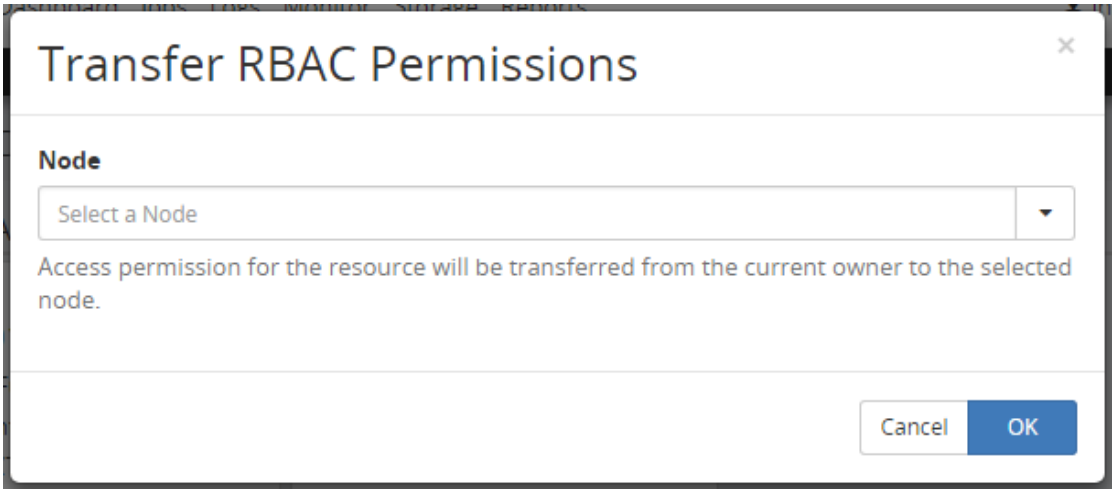


Figure 43 Transfer RBAC Permissions Dialog

Control	Description
Node	Select the node that will become the owner of this resource.

Chapter 6: Managing remote paths and quorum disks

You can use Storage Advisor Embedded to manage remote paths and quorum disks for Virtual Storage Platform G/F350, G/F370, G/F700, G/F900 storage systems.

Managing remote paths

Adding remote paths


Use Hitachi Storage Advisor Embedded to add remote paths to an already-configured remote path group to configure redundant remote paths between the two storage systems that make up a global-active device environment.

Before you begin

Identify the following items:

- The ID of the path group
- The port to be used on the connection-source storage system and the port of the connection-destination storage system

Procedure

1. In the navigation bar, click **Others > Remote Path Groups**.
2. Click the path group ID of the remote path group to which you want to add the remote path.
3. Click  (**Add Remote Paths**).

4. Specify the required items, and then add the remote path.

Make sure that the status of the remote path you added is **Normal**.



Note:

If the status of the remote path is not **Normal**, see the *Global-Active Device User Guide*.

Removing remote paths

When you change the configuration of a port or when the actual amount of data is less than the amount of data assumed at the time of design, remove remote paths from the remote path group. You cannot remove all of the remote paths from a remote path group. To remove all of the remote paths in a remote path group, remove the remote path group itself.

Before you begin

Identify the ID of the path group.

Procedure

1. In the navigation bar, click **Others > Remote Path Groups**.
2. Click the path group ID of the remote path group from which the remote path is to be removed.
3. From the list of remote paths, select one or more remote paths, and then click (**Remove Remote Paths**).


Removing remote path groups

When you no longer need to use a global-active device, remove the paths configured between the two storage systems for which global-active device is used.

Before you begin

Identify the ID of the path group of the remote path groups to be removed.

Procedure

1. In the navigation bar, click **Others > Remote Path Groups**.
2. Select one or more remote path groups and click  (**Delete Remote Path Groups**).

Managing quorum disks


Adding external paths to a quorum disk

To configure redundant paths to a quorum disk, add external paths to the quorum disk.

Before you begin

- Ensure that an external volume has been created, and that a quorum disk has been configured.
- Ensure that the two storage systems that make up the global-active device environment and the external storage system in which the quorum disk is configured are connected by physical paths.
- Ensure that the quorum disk is assigned a port on each of the two storage systems that make up the global-active device environment.
- Identify the following items:
 - The port to be used for external connection
 - The port information for the external storage system that is connected with a port for external connections:
 - When using Fibre Channel for the protocol you can specify WWN

Procedure

1. In the navigation bar, click **Others > External Volumes**.
2. Click  (**Add External Paths**) for the external volume to which you want to add a path.

EXTERNAL VOLUME NAME
ExtVol1
PARITY GROUP ID
E1-1000
PATH GROUP ID
50
MODEL
HITACHI VSP Gx00
SERIAL NUMBER
400012
EXTERNAL PATHS
CLS-A (50060e8012016040)
— 50060e8012000c60
CL6-A (50060e8012016050)
— 50060e8012000c70

Select Ports for External Path
PROTOCOL
FC | iSCSI
Type in keyword to search
0 selected
SELECT ID WWN
CL1-A 50060e8012016000
CL2-A 50060e8012016010
CL3-A 50060e8012016020
CL4-A 50060e8012016030
CL5-A 50060e8012016040
CL6-A 50060e8012016050
Cancel Next


3. Select the port for external connection, and then click **Next**.
4. Select the external path to be used, and then click **Submit**.
5. Click the volume name of the quorum disk to which you added an external path. On the volume details page, verify the external path you added.

If other external volumes are included in the same path group, external paths are also added to those volumes. To check information about volumes other than those for which you performed the operation, refresh the list of external volumes.

Changing the name of the external volume of a quorum disk

You can change the external volume name set for a quorum disk.

Procedure

1. In the navigation bar, click **Others** > **External Volumes**.
2. Click  for the external volume.
3. Enter a different external volume name.

Removing external paths to a quorum disk


Remove external paths set to a quorum disk.

Before you begin

Identify the following items:

- ID of the external volume
- Information about the external path to be removed

Procedure

1. In the navigation bar, click **Others > External Volumes**.
2. Click the volume name of the quorum disk whose external path you want to delete.
3. Select one or more external paths, and then click  (**Remove External Paths**).
If other external volumes are included in the same path group, external paths to those volumes are also removed. To check information about volumes other than those for which you performed the operation, refresh the list of external volumes.



Disabling the quorum disk settings and deleting the corresponding external volume

If you decide against using the global-active device environment, disable the settings of the quorum disk that is no longer required for the two storage systems for which global-active device is used. Delete the external volume used as the quorum disk if it is no longer required.

Before you begin

Identify the external volume name of the quorum disk for which settings are to be disabled.

Procedure


1. In the navigation bar, click **Others > External Volumes**.
2. Select one or more volumes used as the quorum disk whose settings you want to disable, and then click  (**Remove Quorum Disk Setting**).
3. Make sure that the target volumes are correct, and then click **Submit**.
4. To delete external volumes, select one or more external volumes, and click  (**Delete External Volume**).
5. Make sure that the target external volume is correct, select either **Delete with Destaging** or **Delete without Destaging**, and then click **Submit**.
If you select **Delete with Destaging**, the processing to write data stored in cache memory to the external volume (destage processing) is performed, and after the connection is disabled, the external volume is deleted. If you select **Delete without Destaging**, the destage processing is not performed, and the external volume is forcibly deleted.

Disabling the quorum disk settings

If the global-active device environment is no longer used, disable the quorum disk settings that are no longer required.

To delete an external volume that is configured as a quorum disk, click Others in the navigation bar and select External Volumes.

Procedure

1. In the navigation bar, click **Others > Quorum Disks**.
2. Select one or more quorum disks, and then click  (**Remove Quorum Disk Setting**).

