HITACHI
Inspire the Next

Global-Active Device
User Guide

Hitachi Virtual Storage Platform G200, G400, G600, G800
Hitachi Virtual Storage Platform G1000

Hitachi Data Systems

MK-92RD8072-08
November 2015

Contents

Preface

This document describes and provides instructions for using the global-active device (GAD) feature of the Hitachi Virtual Storage Platform G1000 (VSP G1000) and Hitachi Virtual Storage Platform G200, G400, G600, G800 (VSP G200, G400, G600, G800) storage systems.

Please read this document carefully to understand how to use this product, and maintain a copy for reference purposes.

Intended audience

This document is intended for system administrators, Hitachi Data Systems representatives, and authorized service providers who install, configure, and operate the Hitachi storage systems.

Readers of this document should be familiar with the following:

- Data processing and RAID storage systems and their basic functions.
- The Hitachi storage system and the *Hardware Guide* for the storage system.
- The Command Control Interface software.
- The Hitachi Command Suite (HCS) software.

Product version

This document revision applies to:

- Hitachi Virtual Storage Platform G200, G400, G600, G800: firmware 83-02-0*x* or later
- Hitachi Virtual Storage Platform G1000: microcode 80-03-3*x* or later

Release notes

The release notes for this product are available on Hitachi Data Systems Support Connect: https://support.hds.com/en_us/contact-us.html. Read the release notes before installing and using this product. They may contain requirements or restrictions that are not fully described in this document or updates or corrections to this document.

Changes in this revision

- Added support for iSCSI remote copy connections on the VSP G1000.
- Added interoperability information for DP-VOLs with the Data Direct Mapping attribute (Requirements and restrictions on page 2-2).
- Added interoperability information for volumes with the T10 PI attribute (Requirements and restrictions on page 2-2).
- Added information about setting the ALUA mode (new section Setting the ALUA mode (VSP G1000 only) on page 4-45) and enabling the ALUA mode of existing pair volumes (new section Enabling the ALUA mode of an existing GAD pair (VSP G1000 only) on page 6-4).
- Added information about setting a nonpreferred cross path (new section Setting a nonpreferred cross path (VSP G1000 only) on page 4-49).
- Added information about setting preferred and nonpreferred paths using ALUA and HDLM (new section Setting preferred and nonpreferred paths (VSP G1000 only) on page 3-9).

- Added information about changing preferred path settings (new section Changing preferred path settings (VSP G1000 only) on page 6-6).
- Added information about completing SIMs using Hitachi Device Manager - Storage Navigator (see the new topic *Completing SIMs for GAD* in the online help).
- Updated the list of system option modes (SOMs) (System option modes on page 3-2):
  - Added SOM 1080.
  - Updated the description of SOMs 689 and 784.
- Updated the list of CCI error codes (Table 12-1 Error codes and details when operating CCI (when SSB1 is 2E31, B901, B90A, B90B, B912 or D004) on page 12-8).

Referenced documents

Hitachi Virtual Storage Platform G1000 documents:

- *Command Control Interface Installation and Configuration Guide*, MK-90RD7008
- *Command Control Interface User and Reference Guide*, MK-90RD7010
- *Hitachi Virtual Storage Platform G1000 Hardware Guide*, MK-92RD8007
- *Hitachi Virtual Storage Platform G1000 Provisioning Guide for Open Systems*, MK-92RD8014
- *Hitachi Device Manager - Storage Navigator Messages*, MK-92RD8017
- *Hitachi Universal Replicator User Guide*, MK-92RD8023
- *Hitachi Universal Volume Manager User Guide*, MK-92RD8024
- *Hitachi Virtual Storage Platform G1000 Product Overview*, MK-92RD8051

Hitachi Virtual Storage Platform G200, G400, G600, G800 documents:

- *Provisioning Guide for Hitachi Virtual Storage Platform Gx00 and Fx00 Models*, MK-94HM8014
- *System Administrator Guide for Hitachi Virtual Storage Platform Gx00 and Fx00 Models* , MK-94HM8016
- *Hitachi Device Manager - Storage Navigator Messages*, MK-94HM8017
- *Hitachi Virtual Storage Platform G200, G400, G600, G800 Hitachi Universal Replicator User Guide*, MK-94HM8023
- *Hitachi Virtual Storage Platform G200, G400, G600, G800 Hitachi Universal Volume Manager User Guide*, MK-94HM8024

Hitachi Command Suite documents:

- *Hitachi Command Suite User Guide*, MK-90HC172
- *Hitachi Command Suite Messages*, MK-90HC178

Document conventions

This document uses the following terminology conventions:

Convention	Description
Hitachi Virtual Storage Platform G1000	Refers to all models of the Hitachi Virtual Storage Platform G1000 storage systems, unless otherwise noted.
Hitachi Virtual Storage Platform G200, G400, G600, G800 VSP Gx00 models	Refers to all models of the Hitachi Virtual Storage Platform G200, G400, G600, G800 storage systems, unless otherwise noted.
VSP G series	Refers to all models of the Hitachi Virtual Storage Platform G200, G400, G600, G800 and Hitachi Virtual Storage Platform G1000 storage systems, unless otherwise noted.

This document uses the following typographic conventions:

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

Convention	Description
↑*value*↑ ceiling ceiling(*value*)	Ceiling function (round up *value* to the next integer)
_ (underlined text)	Default value

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.
	Important	Calls attention to important or additional information.
	Caution	Warns the user of adverse conditions or consequences (for example, disruptive operations).
	WARNING	Warns the user of severe conditions or consequences (for example, destructive operations).

Convention 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 KB	1,000 (10^3) bytes
1 MB	1,000 KB or $1{,}000^2$ bytes
1 GB	1,000 MB or $1{,}000^3$ bytes
1 TB	1,000 GB or $1{,}000^4$ bytes
1 PB	1,000 TB or $1{,}000^5$ bytes
1 EB	1,000 PB or $1{,}000^6$ bytes

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

Logical capacity unit	Value
1 block	512 bytes

Logical capacity unit	Value
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 documentation is available on Hitachi Data Systems Support Connect: https://support.hds.com/en_us/documents.html. Check this site for the most current documentation, including important updates that may have been made after the release of the product.

Getting help

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

Hitachi Data Systems Community is a new global online community for HDS customers, partners, independent software vendors, employees, and prospects. It is the destination to get answers, discover insights, and make connections. **Join the conversation today!** Go to community.hds.com, register, and complete your profile.

Comments

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

Thank you!

1

Overview of global-active device

This chapter provides an overview of the global-active device feature of the Hitachi Virtual Storage Platform G200, G400, G600, G800 and Hitachi Virtual Storage Platform G1000 storage systems.

About global-active device

Global-active device (GAD) enables you to create and maintain synchronous, remote copies of data volumes on the Hitachi Virtual Storage Platform G1000 and Hitachi Virtual Storage Platform G200, G400, G600, G800 storage systems. 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 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 Hitachi Dynamic Link Manager (HDLM) 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
- Load balancing through migration of virtual storage machines without storage impact

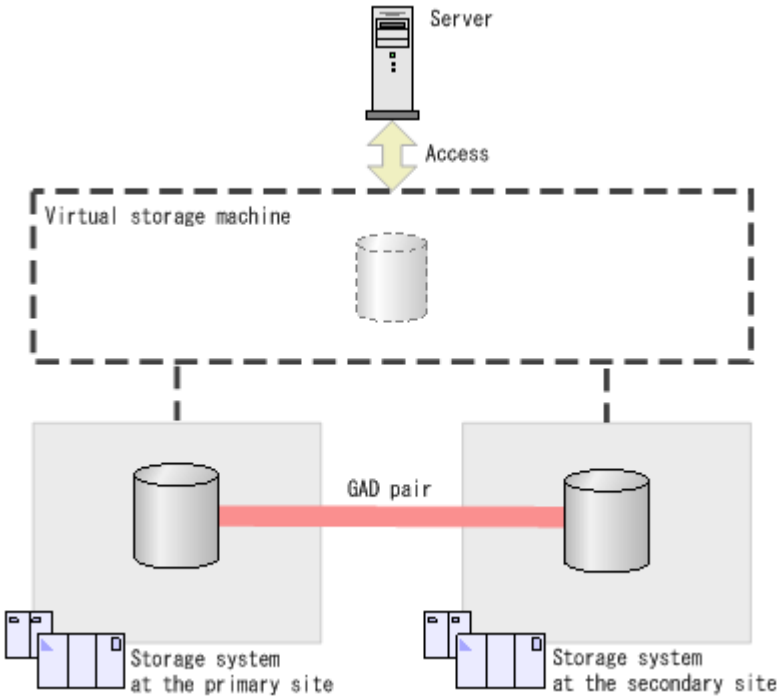

Related topics

- Global-active device solutions on page 1-3

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, as shown in the following illustration, to provide continuous server I/O to the data volume.

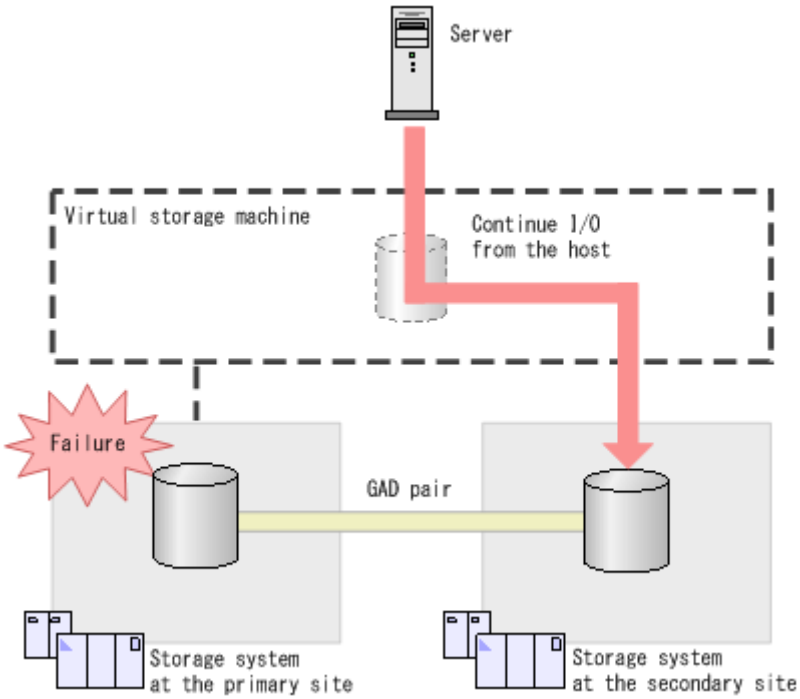

Failover clustering without storage impact

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.

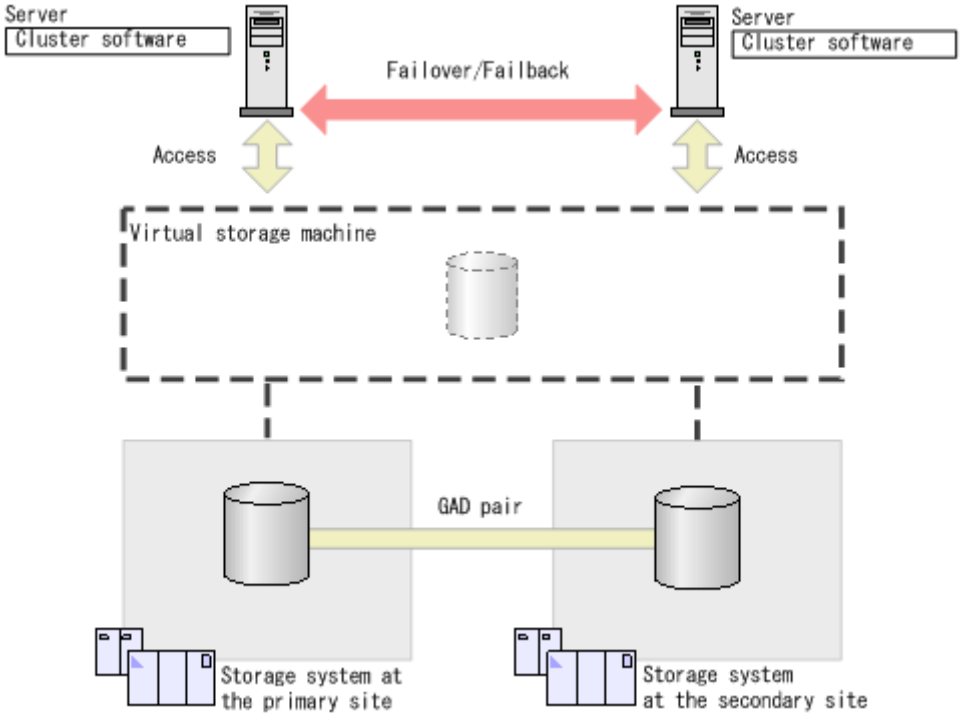

Server load balancing without storage impact

When the I/O load on a virtual storage machine at the primary site increases, as shown in the following figure, global-active device enables you to migrate the virtual machine to the paired server without performing any operations on the storage systems.

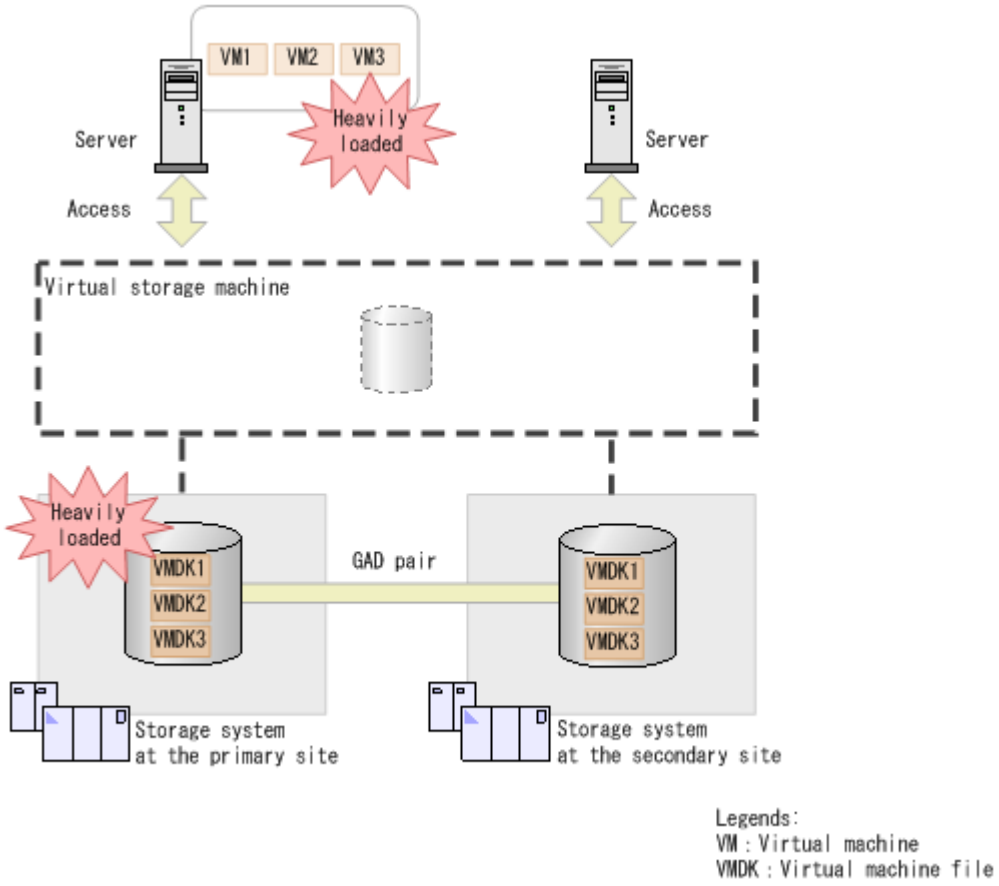

As shown in this example, the server virtualization function is used to migrate virtual machine VM3 from the primary-site server to the secondary-site server. Because the GAD primary and secondary volumes contain the same data, you do not need to migrate any data between the storage systems.

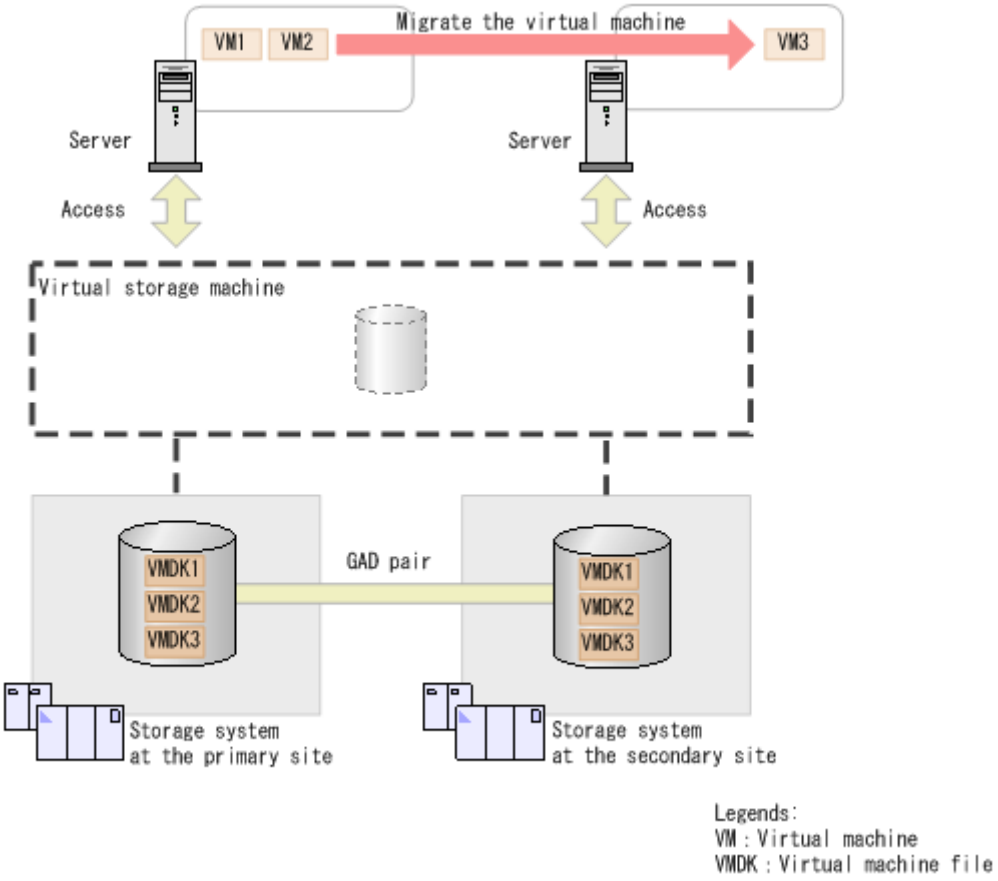

System configurations for GAD solutions

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: • Continuous server I/O (if a failure occurs in a storage system) • Migration of a virtual storage machine of a server 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.

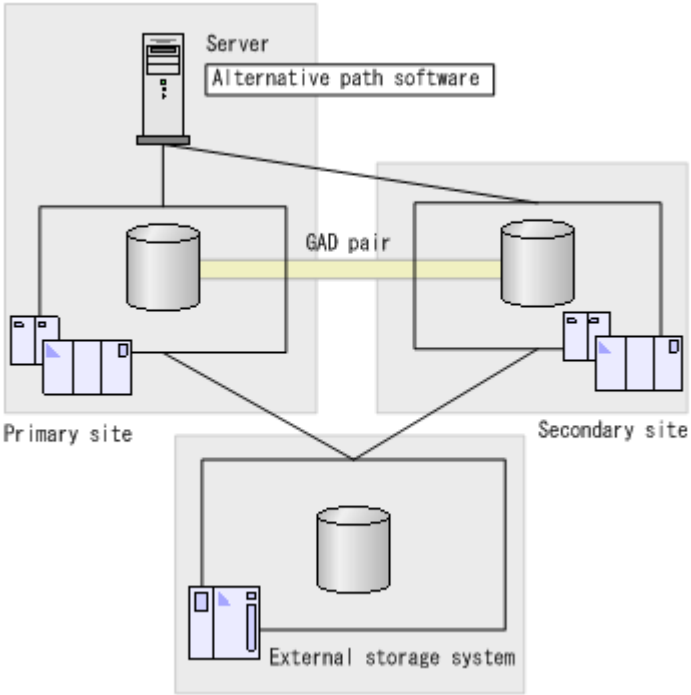

Server-cluster configuration

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.

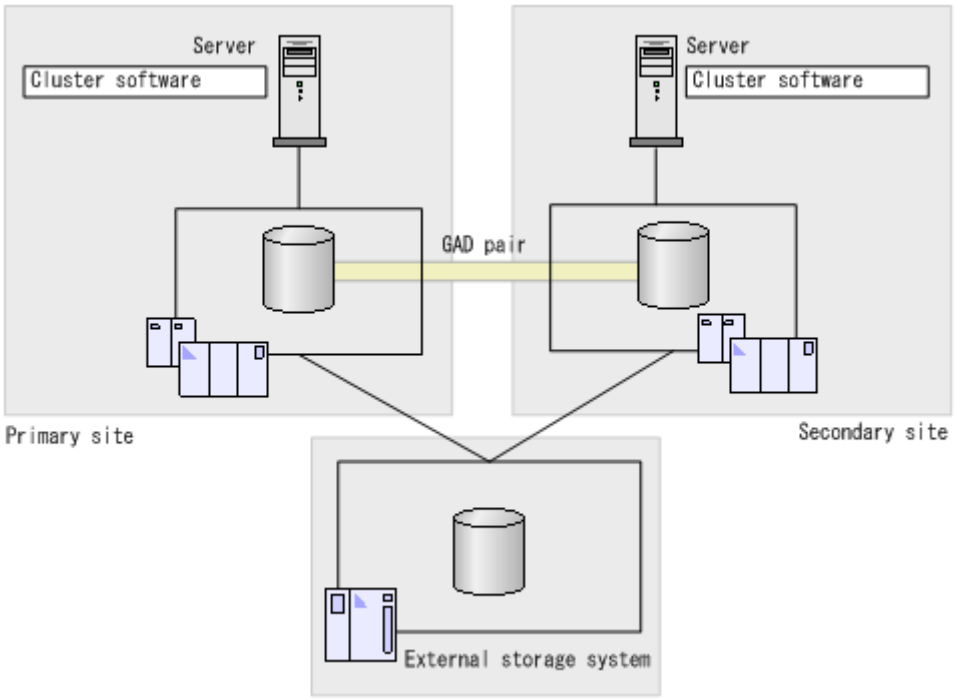

Cross-path configuration

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.

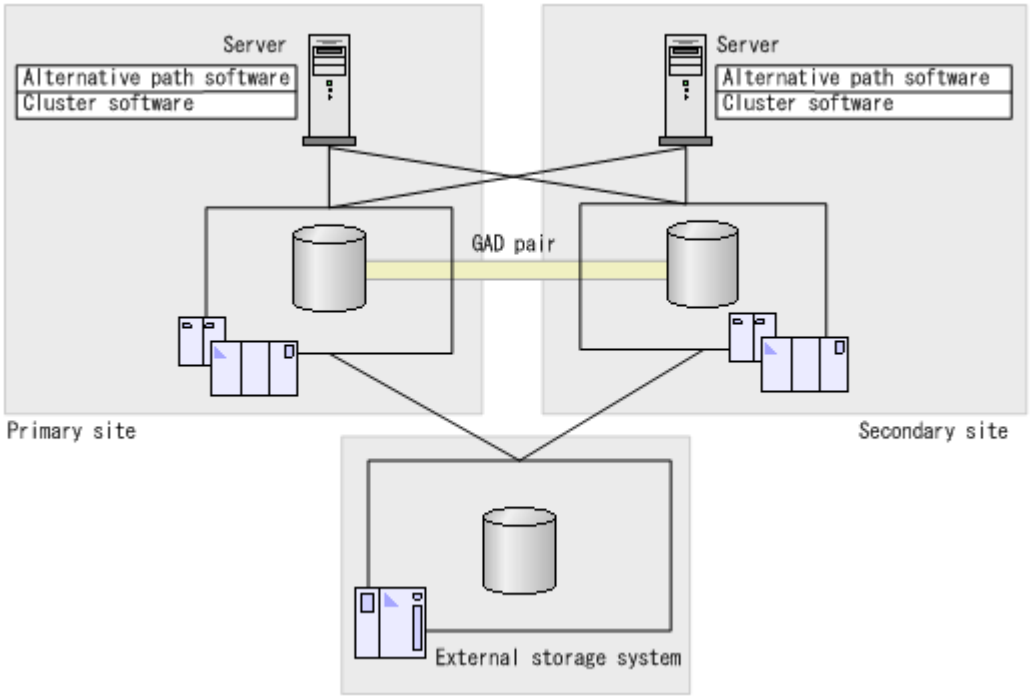

Global-active device and global storage virtualization

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.

Related topics

- About the virtual ID on page 1-9

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

Virtual information required by the server	Resource on which virtual IDs are set
Product	Resource group
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*.

Related topics

- Global-active device and global storage virtualization on page 1-9

GAD status monitoring

GAD operations are managed based on the following information:

- Pair status
- I/O mode of the P-VOL and S-VOL
- GAD status, which is a combination of pair status and I/O mode

Related topics

- GAD status on page 1-11
- GAD status transitions on page 1-12
- I/O modes on page 1-14
- Relationship between GAD status, pair status, and I/O mode on page 1-15

- Global-active device and server I/O on page 1-16

GAD status

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. This status is issued when you do the following: • Prepare a quorum disk. • Copy data from the P-VOL to the S-VOL.	No	P-VOL and S-VOL	P-VOL
Mirrored	The pair is operating normally.	Yes	P-VOL and S-VOL	P-VOL and S-VOL
Quorum disk blocked (VSP G1000 only)	Quorum disk is blocked, but the data is mirrored. I/O from the server to the S-VOL cannot be accepted.	Yes	P-VOL and S-VOL	P-VOL and S-VOL
Suspended	The pair is suspended. I/O from the server is sent to the volume with the latest data. When a failure occurs or the pair is suspended, the status changes to Suspended.	No	P-VOL or S-VOL	P-VOL or S-VOL
Blocked	I/O is not accepted by either pair volume. This status occurs when: • 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. If more than one failure occurs at the same time, the GAD status changes to Blocked.	No	None	P-VOL and S-VOL
* For details on how to determine which volume has the latest data, see Relationship between GAD status, pair status, and I/O mode on page 1-15.				

Related topics

- GAD status monitoring on page 1-10

GAD status transitions

The GAD status changes depending on the pair operation and failure. The following illustration shows the GAD pair status transitions.

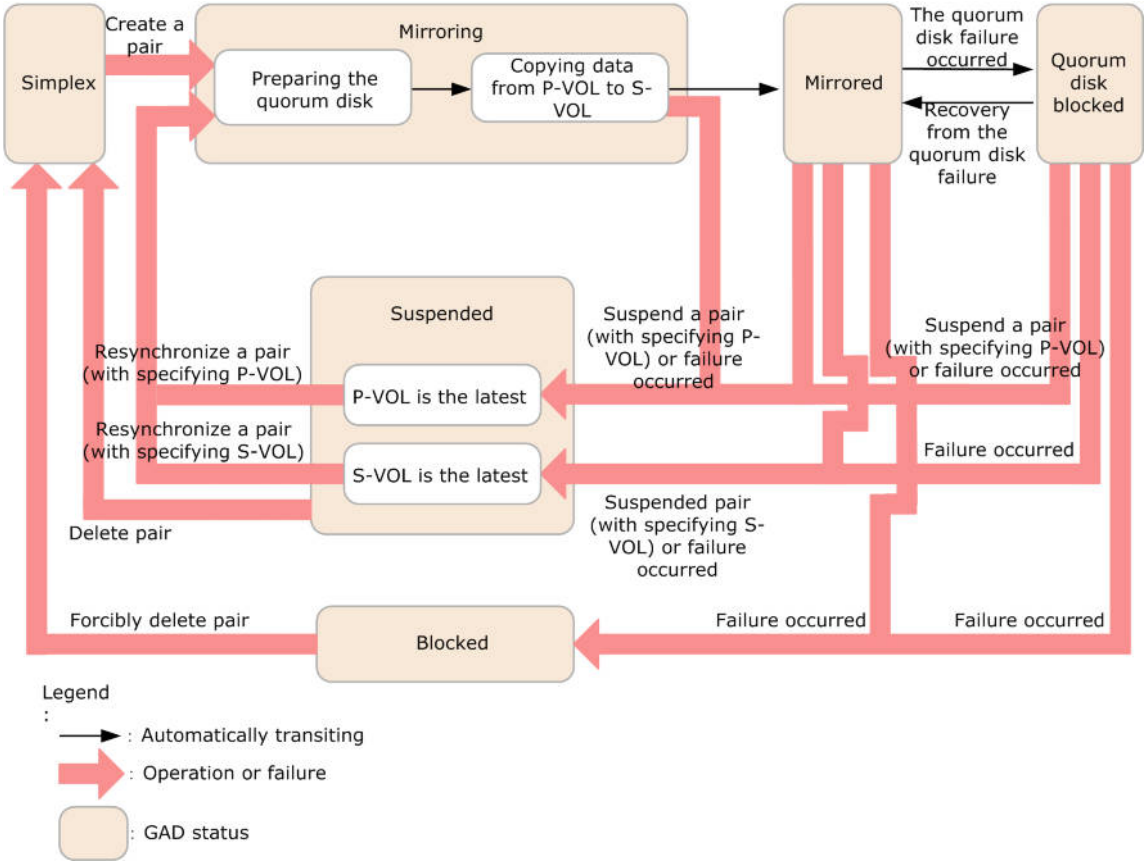

If you resynchronize a pair specifying the P-VOL, I/O continues on the P-VOL. If you resynchronize a pair specifying the S-VOL, data flow switches from the S-VOL to the P-VOL, and then I/O continues on the new P-VOL.

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 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 statuses displayed by CCI and Hitachi Command Suite are slightly different.

Pair status		Description
CCI	HCS	
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.

Pair status		Description
CCI	HCS	
	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. For details about the suspend types, see GAD suspend types on page 1-13.		

Related topics

- GAD status monitoring on page 1-10
- GAD suspend types on page 1-13

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".
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.

Suspend type	Volume	Description
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.

Related topics

- GAD status monitoring on page 1-10

I/O modes

The following table lists and describes the GAD I/O modes, which represent the I/O actions on the P-VOL and the S-VOL of a GAD pair. As shown in the following table, the I/O modes displayed by CCI and Hitachi Command Suite are slightly different.

I/O mode			Read processing	Write processing
I/O mode	CCI[1]	HCS		
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:

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).

I/O mode			Read processing	Write processing
I/O mode	CCI[1]	HCS		
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.				

Related topics

- GAD status monitoring on page 1-10

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

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
	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 (VSP G1000 only)	Not applicable	PAIR	Mirror (RL)	PAIR	Block	P-VOL and S-VOL
Suspended	Pair operation	PSUS	Local	SSUS	Block	P-VOL
	Failure	PSUE*	Local	PSUE	Block	P-VOL
		PSUE*	Local	SMPL	--	P-VOL
		PSUE*	Local	N	N	P-VOL
	Pair operation	PSUS	Block	SSWS	Local	S-VOL
	Failure	PSUE	Block	SSWS*	Local	S-VOL
		SMPL	--	SSWS*	Local	S-VOL
		N	N	SSWS*	Local	S-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	
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
* If the server does not issue the write I/O, the pair status might be PAIR, depending on the failure location.						

Related topics

- GAD status monitoring on page 1-10
- Global-active device and server I/O on page 1-16

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. This topic provides a detailed description of read and write I/O processing for each GAD status.

Related topics

- Server I/O (GAD status: Mirrored) on page 1-16
- Server I/O (GAD status: Mirroring or Quorum disk blocked) on page 1-17
- Server I/O when the GAD status is Suspended on page 1-18
- Server I/O when the GAD status is Blocked on page 1-19
- GAD status monitoring on page 1-10

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.

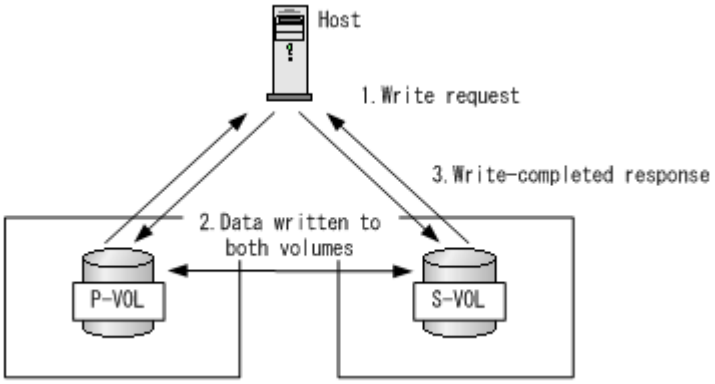

Read requests are read from the volume connected to the server and then sent to the server. There is no communication between the primary and secondary storage systems.

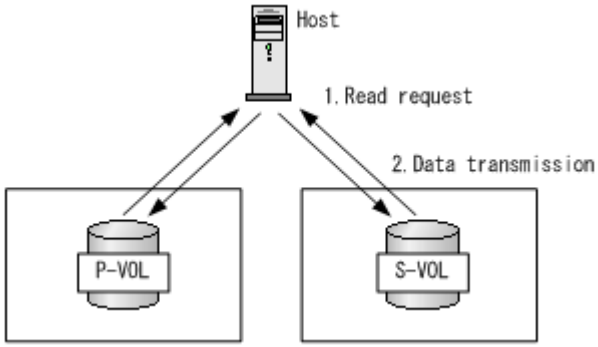

Related topics

- GAD status monitoring on page 1-10
- Global-active device and server I/O on page 1-16

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

When the GAD status is Mirroring or Quorum disk blocked (VSP G1000 only), the I/O mode for the P-VOL is Mirror (RL), and the I/O mode for the S-VOL is Block.

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.

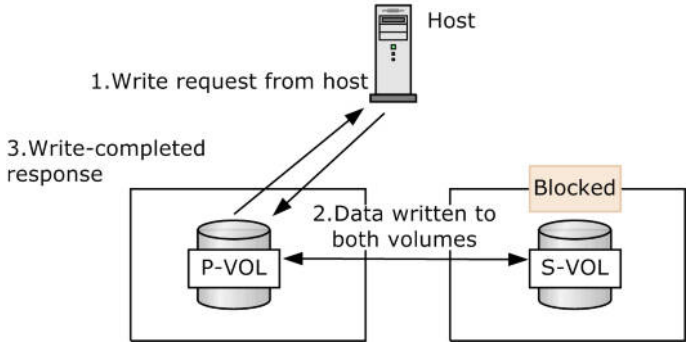

Read requests are read by the P-VOL and then sent to the host. There is no communication between the primary and secondary storage systems.

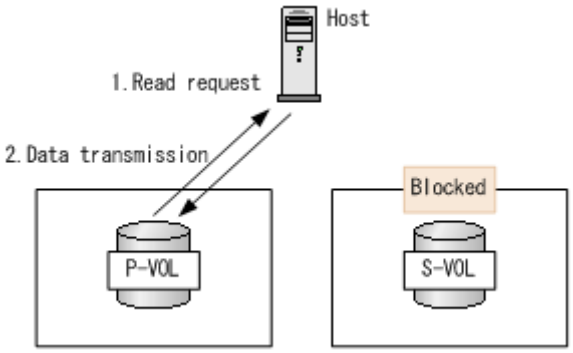

Related topics

- GAD status monitoring on page 1-10
- Global-active device and server I/O on page 1-16

Server I/O when the GAD status is Suspended

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.

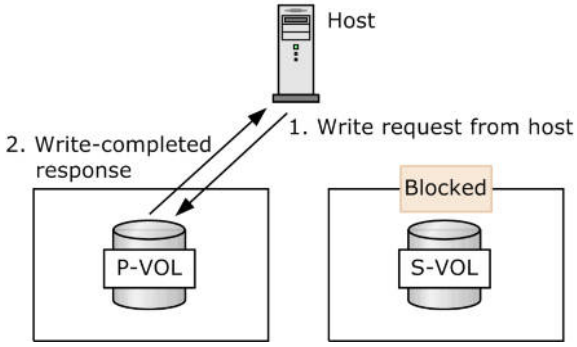

Read requests are read by the P-VOL and then sent to the host. There is no communication between the primary and secondary storage systems.

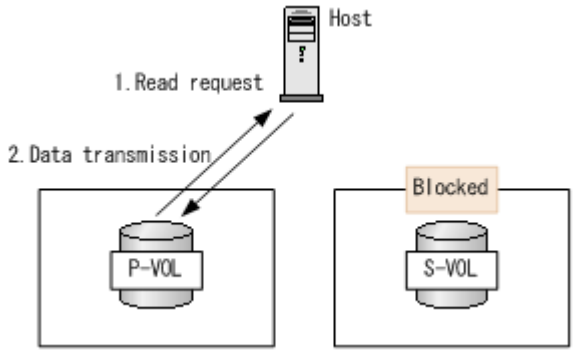

Related topics

- GAD status monitoring on page 1-10
- Global-active device and server I/O on page 1-16

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.

Related topics

- GAD status monitoring on page 1-10
- Global-active device and server I/O on page 1-16

Quorum disk and server I/O

The quorum disk is a volume virtualized from an external storage system. 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 primary and secondary storage systems check the quorum disk every 500 ms for the physical path statuses.

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

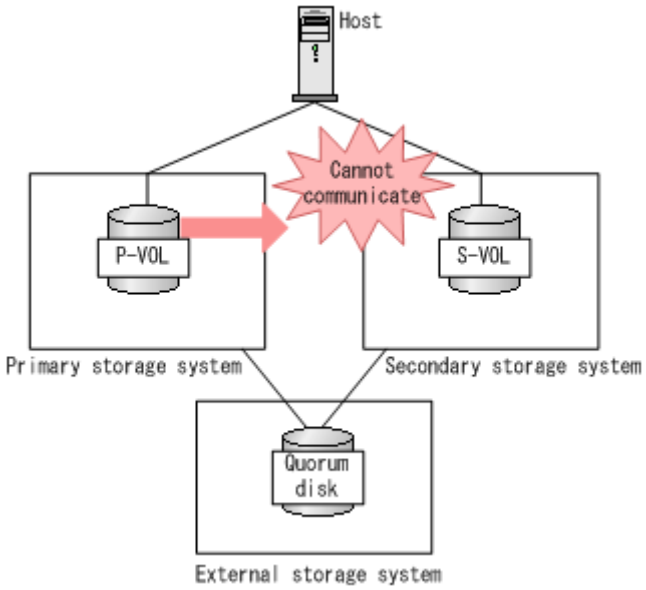

1. The primary storage system cannot communicate over the data path and writes this status to the quorum disk.
2. When the secondary storage system detects from the quorum disk that a path failure has occurred, it stops accepting read/write.
3. The secondary storage system communicates to the quorum 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.

Related topics

- I/O stoppage detected in the counterpart system on page 1-20
- I/O stoppage not detected in the counterpart system on page 1-21
- I/O stoppage and data duplication with blocked quorum disk (VSP G1000) on page 1-21

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.

Related topics

- Quorum disk and server I/O on page 1-19

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.

- 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.

Related topics

- Quorum disk and server I/O on page 1-19

I/O stoppage and data duplication with blocked quorum disk (VSP G1000)

This section describes the I/O stoppage and data duplication when a failure occurs on the quorum disk. GAD pairs that meet both of the following requirements can keep the same data in P-VOL and S-VOL if the quorum disk is blocked when the pair status is PAIR. If a failure occurs on the P-VOL when the quorum disk is blocked, you must delete the pair, and you can restart operations using the data in the S-VOL.

- The microcode version of the primary and secondary storage systems must be 80-02-4*x* or later.

- The pair must be created, resynchronized, or swap resynchronized after the microcode of the primary and secondary storage systems has been upgraded to the specified level (80-02-4*x* or later).

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

For a pair created, resynchronized, or swap resynchronized with microcode 80-02-4*x* or later

When the microcode version of both the primary and secondary storage systems is 80-02-4*x* or later, and when the pair is created, resynchronized, or swap resynchronized after the microcode of both storage systems has been upgraded to 80-02-4*x* or later, I/O stoppage and data duplication are as follows:

- **When the quorum disk is blocked when the pair status is PAIR:** The primary and secondary storage systems communicate by 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 duplication can be maintained by using the remote paths between the primary and secondary storage systems.
- **When the quorum disk is blocked when 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 when 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:** 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 duplication 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 with microcode 80-02-3*x* or earlier

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

- **When the quorum disk is blocked when 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 when 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.

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:** 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.

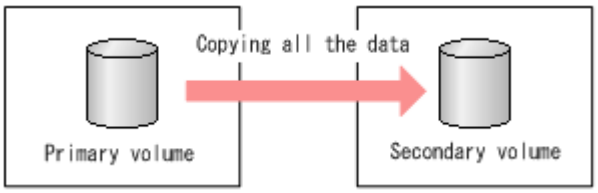

- **Differential copy:** Only the differential data between the P-VOL and the S-VOL is copied. Differential copy is used when the GAD status changes from Suspended to Mirrored.

  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.

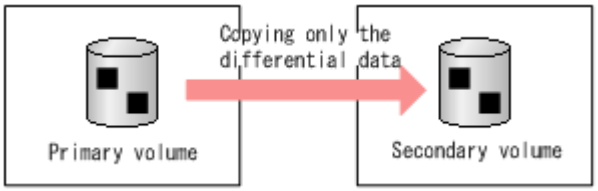

GAD consistency groups

By using consistency groups, you can manage multiple GAD pairs as a group. The GAD pairs in a GAD 3DC delta resync (GAD+UR) configuration must be registered to a consistency group.

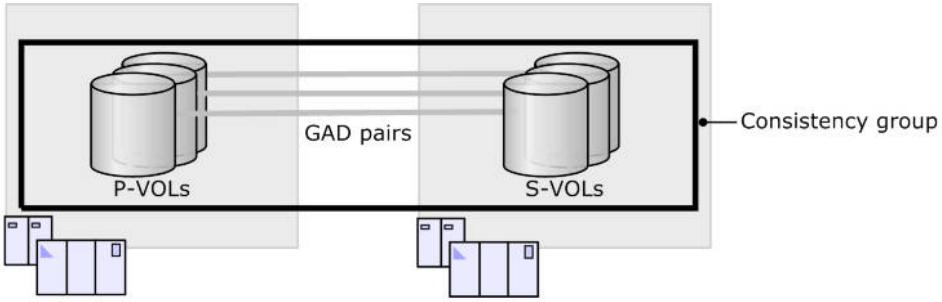

Registering GAD pairs to consistency groups enables you to perform operations on all GAD pairs in a consistency group at the same time. In addition, when a failure occurs, the GAD pairs are suspended by consistency group (concurrent suspension).

For details about storage system support (models, microcode) for consistency groups, see Requirements and restrictions on page 2-2.

Note: When you register GAD pairs to a consistency group, you should use the cross-path configuration. For details, see System configurations for GAD solutions on page 1-6.

Related topics

- Operations on GAD pairs by consistency group on page 1-24
- Suspension of GAD pairs by consistency group on page 1-25
- Use cases for consistency groups on page 1-27
- GAD consistency group statuses on page 1-31

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.

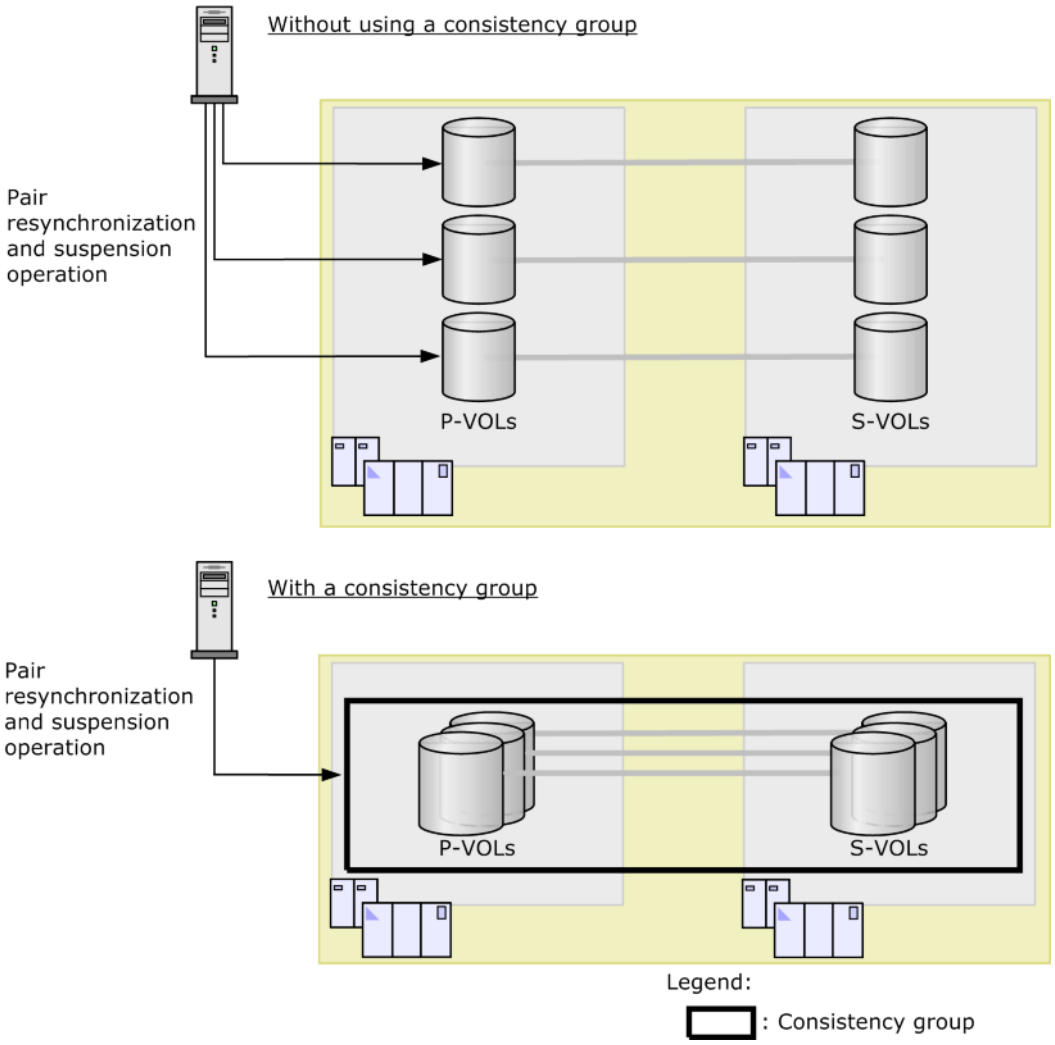

For details about storage system support (models, microcode) for consistency groups, see Requirements and restrictions on page 2-2.

Related topics

- GAD consistency groups on page 1-23

Suspension of GAD pairs by consistency group

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. When a failure occurs, data consistency is guaranteed 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.

- 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.

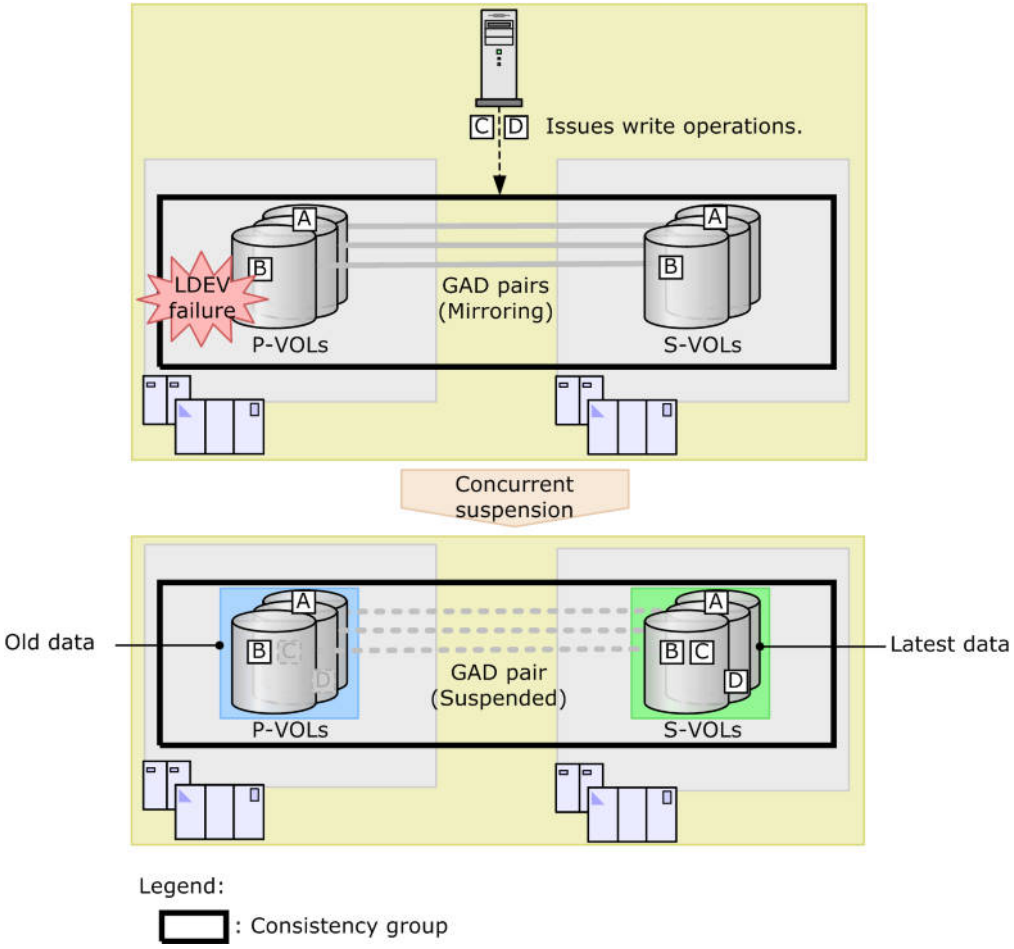

For example, a server issues write operations A to D. After the storage system receives write operation B, all GAD pairs registered to the consistency group change to the Suspended state because of an LDEV failure in the primary volume. In such a case, write operations A and B received before the GAD 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 (models, microcode) for consistency groups, see Requirements and restrictions on page 2-2.

Related topics

- GAD consistency groups on page 1-23

Use cases for consistency groups

The following sections describe specific use cases for GAD consistency groups.

For details about storage system support (models, microcode) for consistency groups, see Requirements and restrictions on page 2-2.

Batch failover

By using consistency groups, you can perform a remote site batch failover operation for GAD pairs by consistency group.

When consistency groups are not used, a remote site failover operation is performed only for the applications that access the volume where the failure occurred.

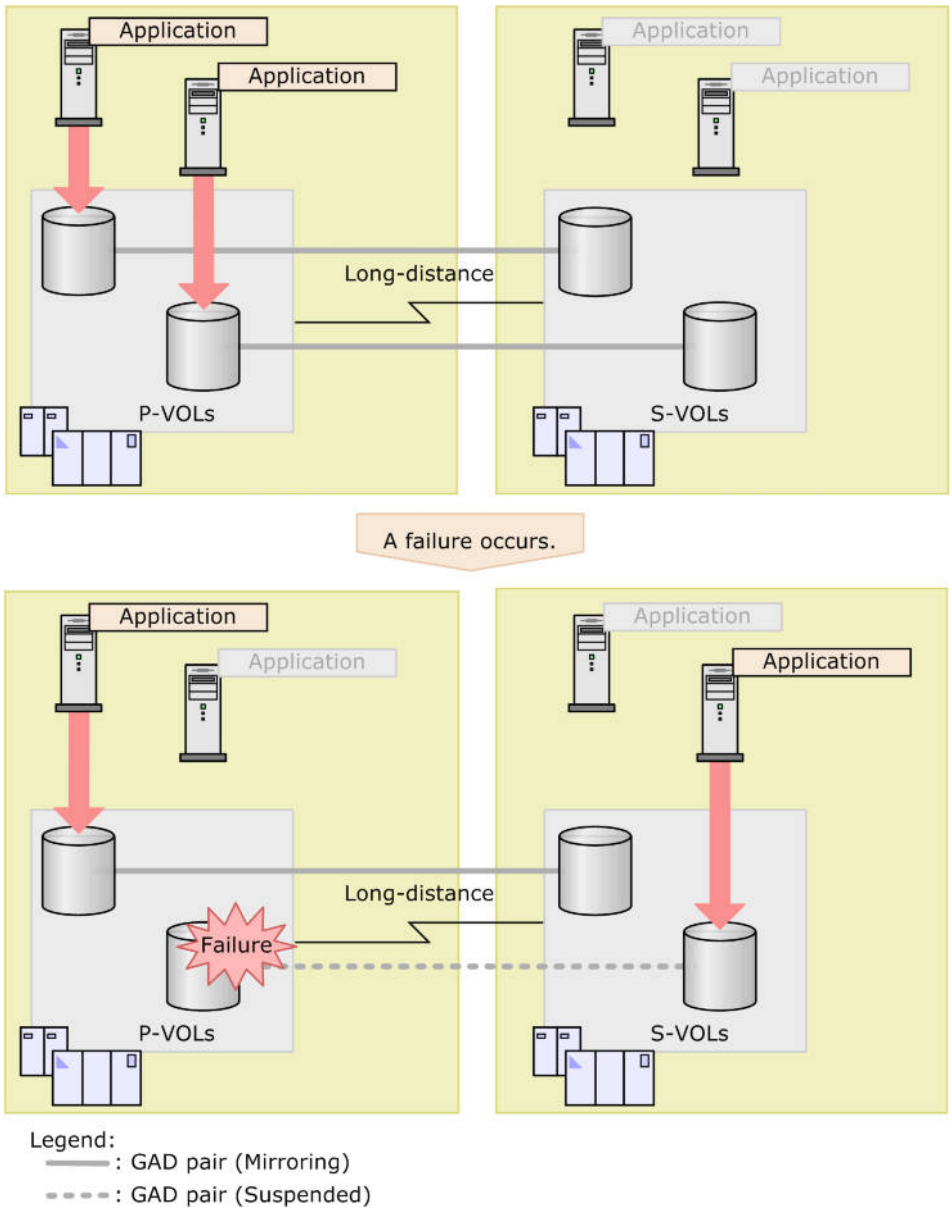

When using consistency groups, if a failure occurs, you can perform a remote site failover operation for all applications that access the volume, together with the GAD pairs in the consistency group.

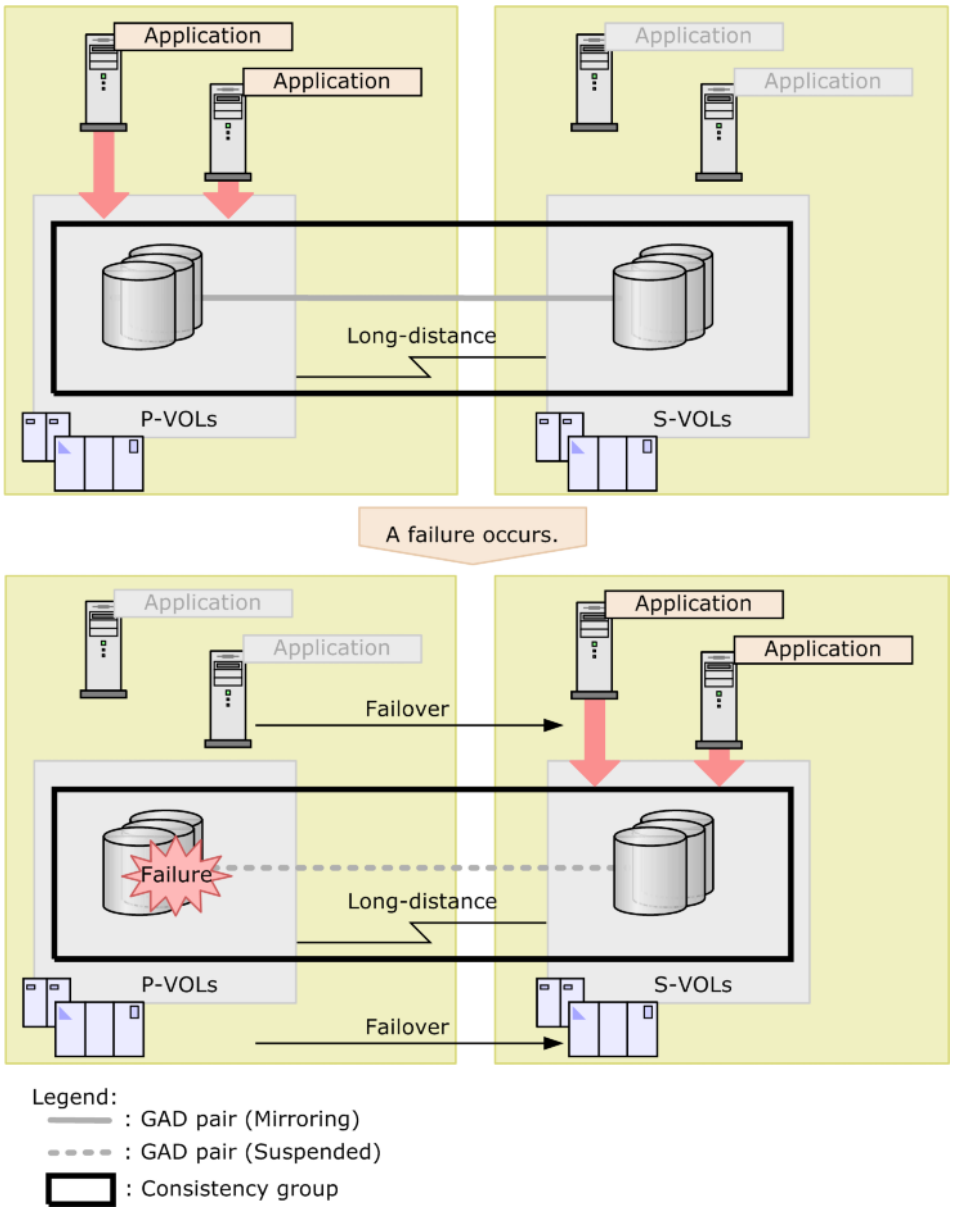

Related topics

- GAD consistency groups on page 1-23

Resuming operations by using consistent backup data

If GAD pairs change to the Suspended state, I/O from servers continues only for the volume that has the most recent data. While GAD pairs are in the Suspended state, if a failure occurs in the storage system that has the most recent data, thus making it impossible to access the most recent data, you can resume operations from the point when GAD pair suspension started by using the consistent data (old data).

For example, assume that GAD pairs changed to the Suspended state due to a path failure of the primary volume. At this point, the primary volume has finished performing the write operation data up to data B.

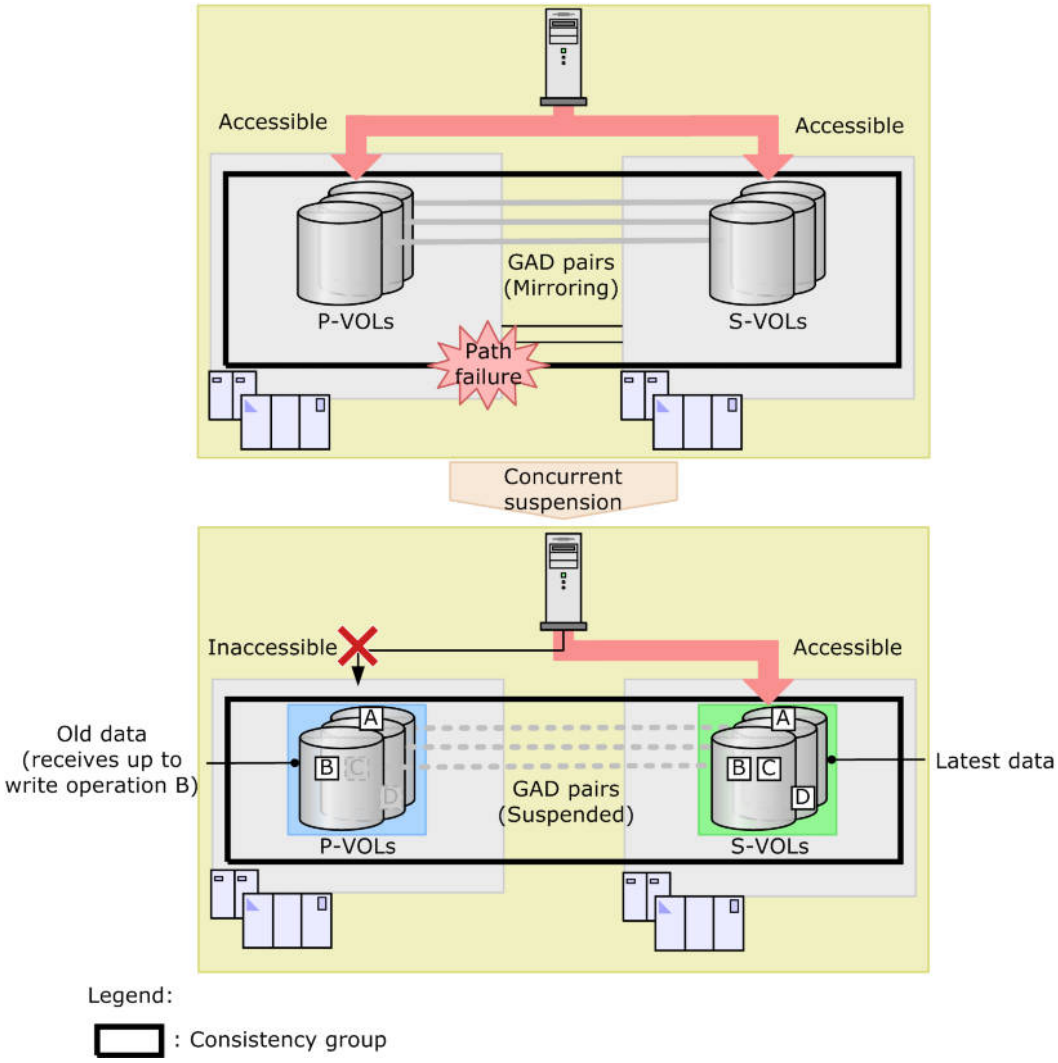

Then, a failure occurred in the storage system of the secondary volume, making it impossible to access the most recent data in that volume. In such a case, after deleting the GAD pairs, you can resume the write processing for data C by using the primary volume.

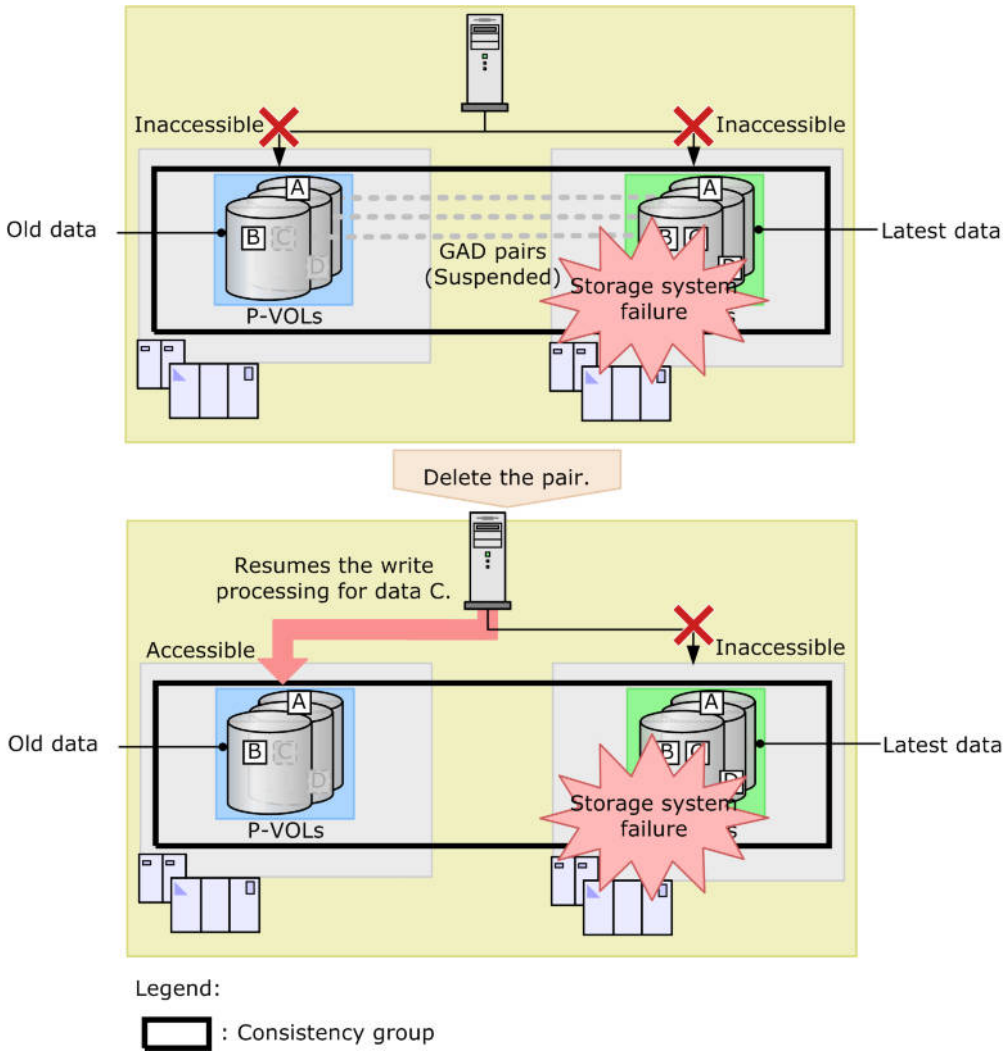

Related topics

- GAD consistency groups on page 1-23

GAD consistency group statuses

The following table describes the statuses of GAD consistency groups. You can view the status of a consistency group by using Device Manager - Storage Navigator.

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.

Status	Description
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.
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.

Related topics

- GAD consistency groups on page 1-23

Global-active device components

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

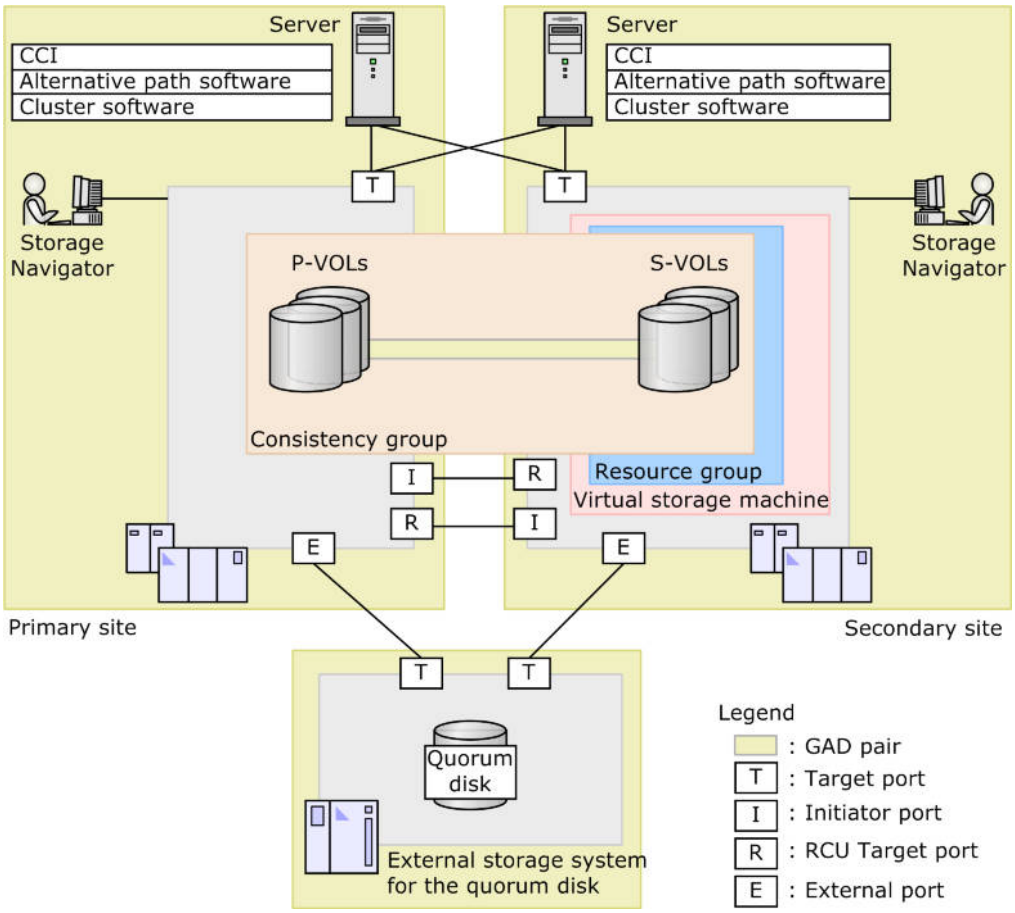

Storage systems

The primary and secondary storage systems must be VSP G1000 or VSP G200, G400, G600, G800 storage systems.

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.

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 (models, microcode) for consistency groups, see Requirements and restrictions on page 2-2.

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 of 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. Both Fibre Channel and iSCSI remote copy connections are supported. 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) on the VSP G200, G400, G600, G800.

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 the CLI software for the storage system:

- Hitachi Command Suite
- Command Control Interface

Hitachi Command Suite

The Hitachi Command Suite (HCS) software enables you to configure and manage GAD pairs and monitor and manage your global-active device environment.

- When one Device Manager server manages both global-active device storage systems, you can access all required functions for your GAD setup from the HCS **Set up Global-Active Device** window.
- When performing operations on GAD pairs during a failure, or when adding a ShadowImage or Thin Image pair volume to a GAD pair volume for additional data protection, perform the operation from Replication Manager.
- Hitachi Command Suite does not provide access to all global-active device operations. For example, the operation to forcibly delete a GAD pair can only be performed using Device Manager - Storage Navigator.

Command Control Interface

The Command Control Interface (CCI) command-line interface (CLI) software can be used to configure the global-active device environment and create and manage GAD pairs. CCI is also used to perform disaster recovery procedures.

Configuration workflow for global-active device

The following table lists the global-active device configuration tasks and indicates the location of the instructions for the tasks.

Configuration task	Operation target	CCI	HCS
Installing global-active device	Primary and secondary storage systems	Not available.	VSP G1000: Section on installing a software application in the Hitachi Command Suite User Guide VSP G200, G400, G600, G800: *System Administrator Guide for Hitachi Virtual Storage*

Configuration task		Operation target	CCI	HCS
				Platform Gx00 and Fx00 Models
Creating command devices		Primary and secondary storage systems	Creating the command devices on page 4-11	Configuring pair management servers on page 7-9
Creating and executing CCI configuration definition files		Server. (With HCS, this is the pair management server.)	Creating the configuration definition files on page 4-11	Monitoring and managing global-active device pairs on page 7-11
Connecting primary and secondary storage systems	Changing port attributes (VSP G1000 only)	Primary and secondary storage systems	Setting the port attributes (VSP G1000 only) on page 4-13	Setting up a global-active device environment on page 7-2
	Adding remote connections	Primary and secondary storage systems	Adding remote connections on page 4-15	
Creating the quorum disk	Changing the port attribute to External (VSP G1000 only)	Primary and secondary storage systems	Creating the quorum disk on page 4-17, Setting the port attributes for connecting the external storage system on page 4-17	
	Mapping the external volume	Primary and secondary storage systems	Creating external volume groups on page 4-19	
	Setting the quorum disk	Primary and secondary storage systems	Setting external volumes as quorum disks on page 4-24	
Setting up the secondary storage system	Creating a VSM	Secondary storage system	Setting up the secondary storage system on page 4-26	
	Setting the GAD reserve attribute	Secondary storage system	Setting the GAD reserve attribute on the S-VOL on page 4-33	Section on allocating volumes in the *Hitachi Command Suite User Guide*
	Adding an LU path to the S-VOL	Secondary storage system	Adding an LU path to the S-VOL on page 4-40	
Updating CCI configuration definition files		Server	Editing CCI configuration definition files on page 4-42	Monitoring and managing global-active device pairs on page 7-11
Creating GAD pair		Primary storage system	Creating the GAD pair on page 4-43	Allocating global-active device pairs on page 7-12

Configuration task	Operation target	CCI	HCS
Adding alternate path to the S-VOL	Server	Adding an alternate path to the S-VOL on page 4-48	Section on optimizing HBA configurations in the *Hitachi Command Suite User Guide*

2

System requirements

This chapter provides the system requirements for global-active device (GAD) 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	• **Model:**
	  ◦ Between VSP G1000 and VSP G1000.
	  ◦ Between VSP G200, G400, G600, G800 and VSP G200, G400, G600, G800.
	Note:
	  ◦ GAD is not supported between VSP G1000 and VSP G200, G400, G600, G800.
	  ◦ GAD is not supported on the VSP Fx00 models.
	• **Microcode:**
	  ◦ VSP G1000: DKCMAIN 80-01-4*x* or later
	  ◦ VSP G200, G400, G600, G800: DKCMAIN 83-01-2*x* or later
	• **Global-active device license:** The GAD feature must be installed and enabled on the primary and secondary storage systems.
	• **Controller emulation type:** The controller emulation type of the primary and secondary storage systems must be same.
	• **Shared memory/control memory:**
	  ◦ VSP G200: You can use GAD only with control memory in the basic part. Adding control memory expands the capacity of the pairs being created.
	  ◦ VSP G400, VSP G600, VSP G800: Additional control memory is required in the primary and secondary storage systems. For information on adding control memory and setting the GAD dedicated area, contact Hitachi Data Systems customer support.
	  ◦ VSP G1000: Additional shared memory is required in the primary and secondary storage systems.
	Note: The terms "shared memory" and "control memory" refer to the same hardware on the storage systems. The term "shared memory" is used for VSP G1000, and the term "control memory" is used for VSP G200, G400, G600, G800.

GAD 3DC delta resync (GAD+UR)	• All storage systems in a GAD+UR configuration must be VSP G1000 systems. • Microcode: DKCMAIN version 80-02-4x or later • The GAD pairs in a GAD+UR configuration must be registered to a consistency group.
External storage systems (for quorum disk)	• The storage system must be supported for attachment using Universal Volume Manager. For details, see the *Hitachi Universal Volume Manager User Guide* for the storage system.

Item	Requirements and restrictions
	• The maximum distance between the external storage system and the primary site and secondary site is 1,500 km.
Licensed capacity	• The page size assigned to the virtual volume is counted as a licensed capacity for GAD. • 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	Refer to the Hitachi Data Systems interoperability matrix: http://www.hds.com/products/interoperability
SCSI commands	• 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. • The ALUA command is not supported (VSP G1000 only).
Physical paths connecting the primary and secondary storage systems	• Maximum number of physical paths: 8 • Maximum distance between the primary and secondary storage systems: 100 km • Port type: Fibre Channel, iSCSI (VSP G200, G400, G600, G800 83-01-2*x* and later, VSP G1000 80-03-3*x* and later) • Port attribute (VSP G1000 only): 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	• 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. • Protocol: All remote paths in a path group must be the same protocol, either Fibre Channel or iSCSI. Remote paths for Fibre Channel and iSCSI 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. You can also use multiple path groups with the same combination of the storage systems at the primary and the secondary sites. • When using the System connection type and not the CU connection type (specified on the Add Remote Connection window), specify different paths and path

Item	Requirements and restrictions
	groups for TrueCopy, Universal Replicator, and Universal Replicator for Mainframe secondary storage systems.

Virtual storage machines (VSMs)	• Maximum number of VSMs per storage system: 8
	• Maximum number of GAD volumes per VSM:
	  ◦ VSP G200: 2,048
	  ◦ VSP G400, VSP G600: 4,096
	  ◦ VSP G800: 16,384
	  ◦ VSP G1000: 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.
	The VSP G1000 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 G1000 only)	• DKCMAIN 80-02-*xx* or later: A volume in a resource group that was migrated from a VSP or USP V/VM storage system to the VSP G1000 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 to the VSP G1000 cannot be used as a GAD volume.

Maximum number of GAD pairs	• VSP G200: 2,046
	If CCI is used in the In-band method and one volume (virtual, internal, or external) is used as a command device, the maximum number of GAD pairs is 2,045.
	• VSP G400, VSP G600: 4,094
	If CCI is used in the In-band method and one volume (virtual, internal, or external) is used as a command device, the maximum number of GAD pairs is 4,093.
	• VSP G800:
	  ◦ If CCI is used in the In-band method and one internal volume is used as a command device, the maximum number of GAD pairs is 14,079.
	  ◦ If CCI is used in the In-band method and one virtual volume or one external volume is used as a command device, the maximum number of GAD pairs is 14,078.
	• VSP G1000:
	  ◦ 65,279 when you create all pairs with internal volumes

Item	Requirements and restrictions
	If CCI is used in the In-band method and one internal volume is used as a command device, the maximum number of GAD pairs is 65,278. If CCI is used in the In-band method and one virtual volume or one external volume is used as a command device, the maximum number of GAD pairs is 65,277. ◦ 63,231 when you create all pairs with DP-VOLs and external volumes If CCI is used in the In-band method and one internal volume is used as a command device, the maximum number of GAD pairs is 63,231. If CCI is used in the In-band method and one virtual volume or one external volume is used as a command device, the maximum number of GAD pairs is 63,230. • 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 on page 3-30.
Pair volumes	• 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. ◦ 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 (VSP G1000 only) ◦ External volumes (VSP G1000 only) • Emulation type (VSP G1000 only): OPEN-V. • Volume size: The P-VOL and S-VOL must be equal in size. • Maximum volume size: DP-VOL: same as the maximum size of a DP-VOL. For details, see the *Provisioning Guide* for the storage system. Internal volume: 3,145,663 MB (6,442,317,824 blocks) External volume: 4,194,304 MB (8,589,934,592 blocks) • SAN boot: You can use GAD pair volumes for SAN boot. • 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

Item	Requirements and restrictions
	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 (VSP G200, G400, G600, G800): The same value must be set for the T10 PI attribute of the P-VOL and the S-VOL.

Quorum disks	• Maximum number of quorum disks: 32 per storage system in the primary storage system and secondary storage system.
	• Maximum number of pairs per quorum disk:
	  ◦ VSP G200: 2,046
	  ◦ VSP G400, VSP G600: 4,094
	  ◦ VSP G800: 14,079
	  ◦ VSP G1000:
	  63,231 when you create all pairs with DP-VOLs and external volumes.
	  65,279 when you create all pairs with internal volumes.
	• Emulation type (VSP G1000 only): 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.
	• Interoperability: A GAD quorum disk cannot also be used as a High Availability Manager quorum disk for a different storage system model or as a GAD quorum disk for a different storage system model. A GAD quorum disk can be shared only by the same storage system models.
	• Requirements for the external storage system volume:
	  ◦ The T10 PI attribute must not be enabled (VSP G200, G400, G600, G800 only).
	  ◦ The Data Direct Mapping attribute must not be set.
	  ◦ The volume must not belong to the system resource group (VSP G200, G400, G600, G800 only).

Consistency groups	• Storage system model: Consistency groups are currently supported only on VSP G1000 storage systems. • Maximum number of consistency groups per storage system: 256 (CTG ID 0-255)

Item	Requirements and restrictions
	• Maximum number of pairs per consistency group: 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	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, VSP G1000 DKCMAIN version 80-03-31-00/00 or later is required. Refer to the Hitachi Data Systems interoperability matrix: http://www.hds.com/products/interoperability
Cluster software	Cluster software is required for the server-cluster and cross-path GAD configurations. Refer to the Hitachi Data Systems interoperability matrix: http://www.hds.com/products/interoperability
User interfaces	• Hitachi Command Suite: 8.0.1 or later HCS must be connected to the primary and secondary storage systems. • Command Control Interface: ◦ VSP G200, G400, G600, G800: 01-34-03/00 or later ◦ VSP G1000: 01-32-03/00 or later CCI must be installed on the host connected to the primary and secondary storage systems. The CCI command device is required on the primary and secondary storage systems.

Interoperability requirements

This section describes the interoperability of global-active device (GAD) with the other features of the VSP G1000 and VSP G200, G400, G600, G800 storage systems.

Volume types that can be used for GAD

The following table lists the volume types and specifies whether the volume can be used for GAD operations.

Volume type	Used as GAD P-VOL?	Used as GAD S-VOL?	Used as quorum disk?
Dynamic Provisioning / Dynamic Tiering / Active flash			

Volume type	Used as GAD P-VOL?	Used as GAD S-VOL?	Used as quorum disk?
Virtual volume	Yes[1]	Yes[1]	No
Pool volume	No	No	No
ShadowImage / Thin Image			
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 for VSP G1000 No for VSP G200, G400, G600, G800	Yes[5] for VSP G1000 No for VSP G200, G400, G600, G800	No
S-VOL	No	No	No
Journal volume	No	No	No
Universal Volume Manager			
External volume	Yes[1] for VSP G1000 No for VSP G200, G400, G600, G800	Yes[1] for VSP G1000 No for VSP G200, G400, G600, G800	Yes
Data Retention Utility			
Volume with access attribute	Yes	Yes[2]	No
Volume Migration			
Source volume	No	No	No
Target volume	No	No	No
Cache Residency Manager (VSP G1000 only)			
The volume on which Cache Residency Manager is set	No	No	No
Hitachi Virtual LUN			
Virtual LUN volume	Yes for VSP G1000 No for VSP G200, G400, G600, G800	Yes for VSP G1000 No for VSP G200, G400, G600, G800	Yes[3]
LUN Manager			
The volume on which paths are defined	Yes	Yes	No

Volume type	Used as GAD P-VOL?	Used as GAD S-VOL?	Used as quorum disk?
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.[4]

Notes:

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. 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.
3. 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.
4. You cannot encrypt a nonencrypted quorum disk in the external storage system from the primary or secondary storage system.
5. GAD S-VOL is used as UR delta resync pair P-VOL.

Related topics

- Dynamic Provisioning / Dynamic Tiering / Active flash on page 2-10
- ShadowImage on page 2-10
- Thin Image on page 2-16
- Use cases for pairing GAD volumes with SI or HTI on page 2-21
- Universal Replicator (VSP G1000 only) on page 2-21
- Data Retention Utility on page 2-23
- LUN Manager on page 2-24
- Virtual Partition Manager on page 2-24
- Volume Shredder on page 2-25
- Performance Monitor on page 2-25

Dynamic Provisioning / Dynamic Tiering / Active flash

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

Note:

- Only allocated page capacity is counted as GAD license capacity. Page capacity or license capacity counted toward the P-VOL and for the S-VOL might differ, because page capacity for the volumes changes according to the operation, for example, tier relocation or reclaiming zero pages.
- You cannot add capacity to a DP-VOL that is used as a GAD pair volume. To do so, delete the pair, add the capacity to the DP-VOL, and then re-create the pair.

Related topics

- Volume types that can be used for GAD on page 2-7

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.

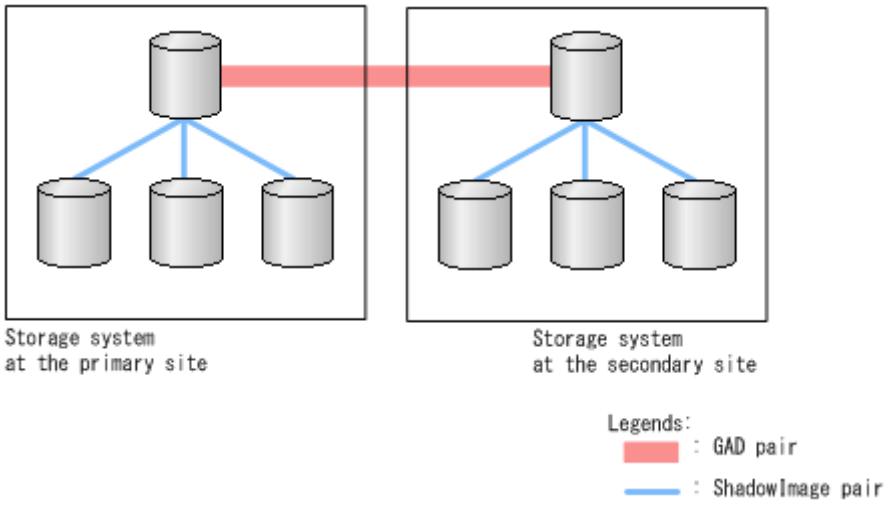

Because the server recognizes a GAD pair as one volume, it sees the volume as paired with six ShadowImage volumes.

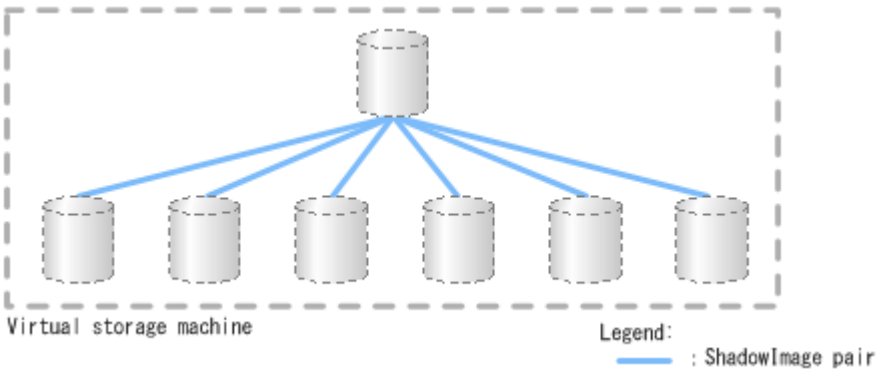

You can create three additional, cascaded SI pairs using the SI S-VOLs. This means that up to nine SI pairs can be created with the GAD P-VOL, and nine SI pairs can be created with the GAD S-VOL.

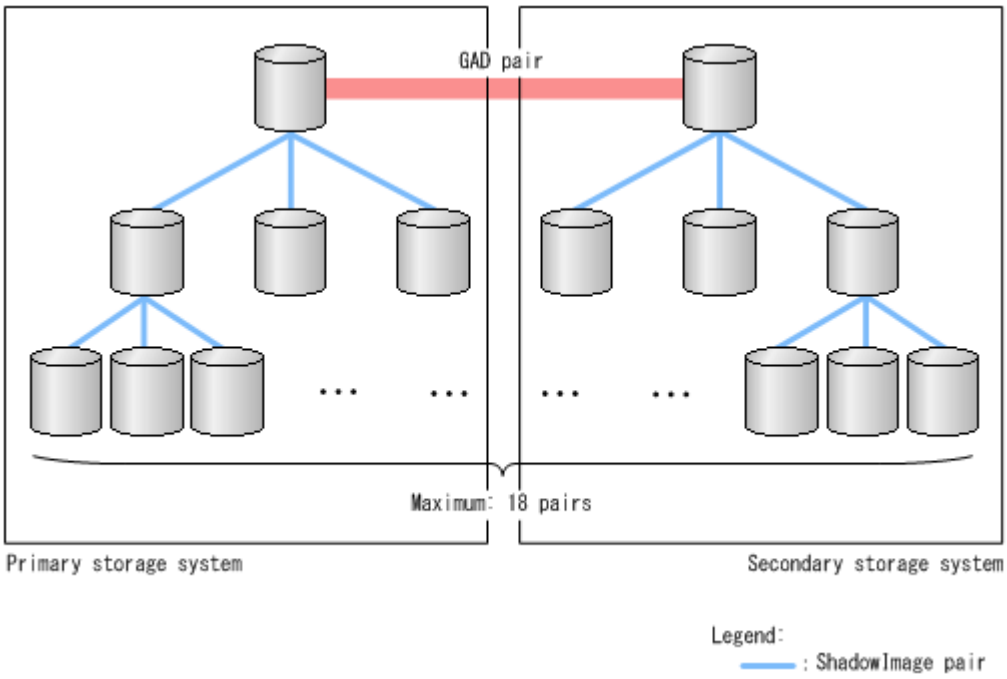

Note:

- 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.

Related topics

- Limitations when sharing GAD and ShadowImage volumes on page 2-12
- SI operations and GAD pair status on page 2-12
- GAD operations and SI pair status on page 2-14

Limitations when sharing GAD and ShadowImage volumes

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.

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

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 2-1 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 2-2 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.

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
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 2-3 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 2-4 SI operations when GAD status is Quorum disk blocked (VSP G1000 only)

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
	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 2-5 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
	Block	Primary	No	No	Yes	Yes	Yes	No	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
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.									

Table 2-6 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 2-7 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

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
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:

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 2-8 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
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.

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

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.

Related topics

- ShadowImage on page 2-10

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.

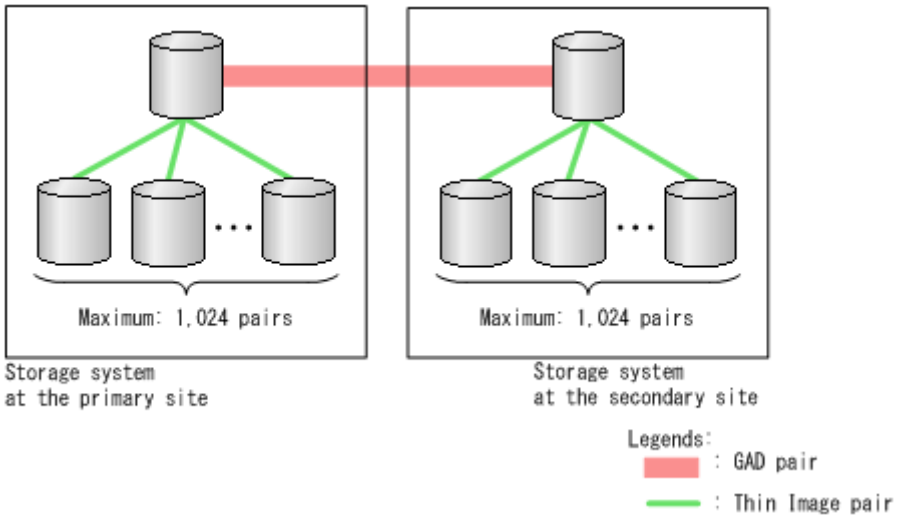

Because the server recognizes the GAD pair as one volume, it sees the volume as paired with 2,048 HTI volumes.

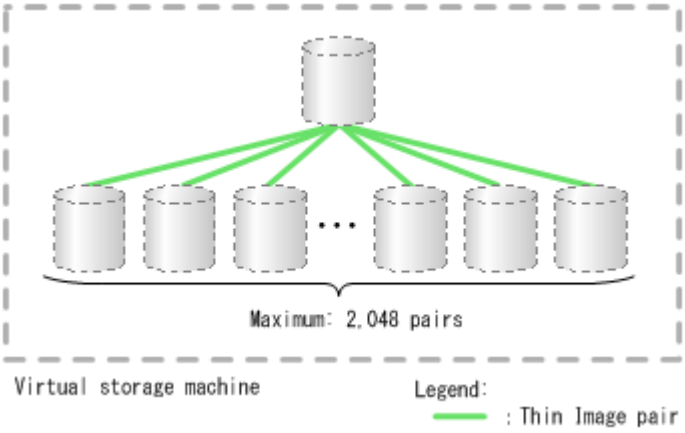

Note:

- 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.

Related topics

- Limitations for using both GAD and Thin Image on page 2-17
- Thin Image operations and GAD status on page 2-17
- GAD operations and Thin Image pair status on page 2-19

Limitations for using both GAD and Thin Image

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.

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

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 2-9 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 2-10 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 2-11 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 2-12 Thin Image operations when GAD status is Quorum disk blocked (VSP G1000 only)

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/ Suspend copy
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 Thin Image copy destination.									

Table 2-13 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 2-14 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

Related topics

- Thin Image on page 2-16

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 2-15 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 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
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:

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 2-16 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 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
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]
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.

TI pair status	GAD pair operations							
	Create GAD 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

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. 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 HTI pair volume.
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.

Related topics

- Thin Image on page 2-16

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, 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 (VSP G1000 only)

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: The primary storage system, the GAD secondary storage system, and the UR storage system must all be VSP G1000 systems with DKCMAIN microcode version 80-02-4*x* 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.

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

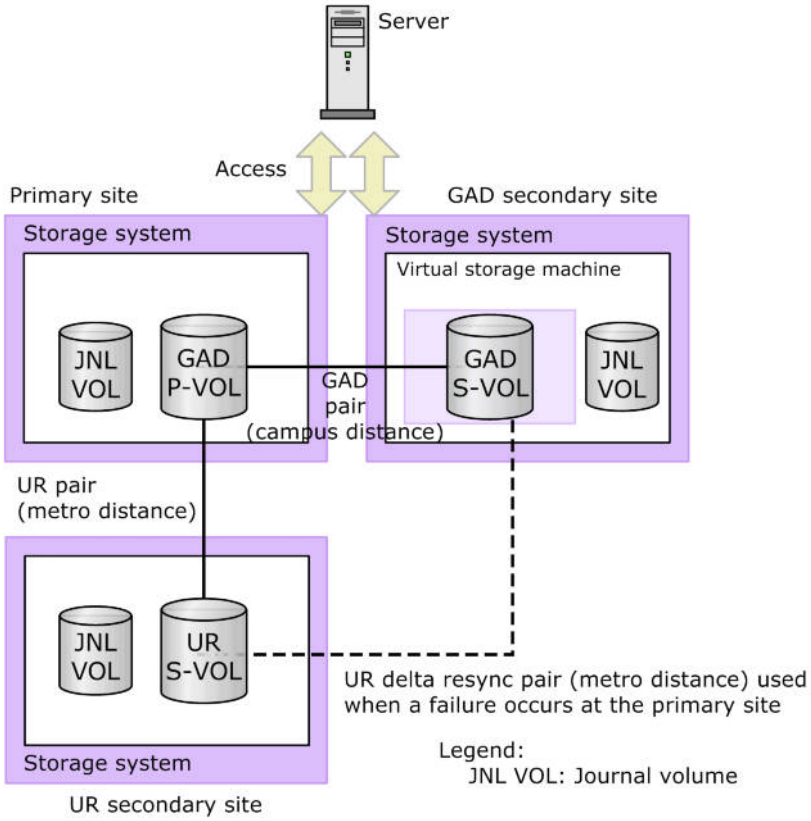

In a GAD+UR configuration:

- 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.

For details about the GAD 3DC delta resync (GAD+UR) configuration, see Chapter 5, GAD 3DC delta resync (GAD+UR) operations on page 5-1.

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.

Related topics

- GAD status and I/O allowance by access attribute on page 2-23

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 (VSP G1000 only)	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
Mirroring	Read/Write	Read/Write	Ends normally	Rejected
Suspended (when the I/O mode of the primary volume is Local and the I/O	Read Only or Protect	Read/Write	Depends on the attribute*	Rejected

GAD statuses	Access attribute		I/O	
	P-VOL	S-VOL	P-VOL	S-VOL
mode of the secondary volume is Block)	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.				

Related topics

- Data Retention Utility on page 2-23

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. Keep the number of LDEV from which you remove the path at one time to 100 or fewer.

Virtual Partition Manager

GAD pair volumes and quorum disks can migrate CLPRs.

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.

However, 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 server I/O data for GAD volumes that is added to Performance Monitor depends on the GAD status, as shown the following tables.

Table 2-17 Writes to GAD volumes to be added to Performance Monitor

GAD status	P-VOL	S-VOL
Mirrored	The sum of the following values: • Number of writes to the P-VOL • Number of RIOs to the P-VOL from the S-VOL	The sum of the following values: • Number of reads from the S-VOL • Number of RIOs to the S-VOL from the P-VOL
Quorum disk blocked (VSP G1000 only)	Number of writes to the P-VOL	Number of RIOs to the S-VOL from the P-VOL
Mirroring	Number of writes to the P-VOL	Number of RIOs to the S-VOL from the P-VOL
Suspended (when the P-VOL has the latest information)	Number of writes to the P-VOL	Not counted*

GAD status	P-VOL	S-VOL
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 2-18 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 (VSP G1000 only)	Number of reads from the P-VOL	Not counted*
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 2-19 Relation between amount of I/O added to Performance Monitor and amount of server I/O

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 writes to the P-VOL and S-VOL
Quorum disk blocked (VSP G1000 only)	The same as the number of writes to the P-VOL	The same as the number of reads from the P-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.							

Related topics

- Performance Monitor on page 2-25

3

Planning for global-active device

This chapter provides planning information for configuring the primary and secondary storage systems, data paths, pair volumes, and quorum disk for global-active device operations.

Storage system preparation

To prepare the storage systems for global-active device operations:

- Make sure that the primary, secondary, and external storage systems meet the global-active device system requirements described in chapter 2.
- Make sure that the primary storage system is configured to report sense information to the host. The secondary storage system should also be attached to a host server to report sense information in the event of a problem with an S-VOL or the secondary storage system.
- If power sequence control cables are used, set the power source selection switch for the cluster to "Local" to prevent the server from cutting the power supply to the primary storage system. In addition, make sure that the secondary storage system is not powered off during GAD operations.
- Establish the physical paths between the primary and secondary storage systems. Switches and channel extenders can be used. For details, see Planning physical paths on page 3-12.
- Review the shared memory/control memory requirements for the primary and secondary storage systems in Requirements and restrictions on page 2-2. Make sure that the cache in both storage systems works normally. Pairs cannot be created if cache requires maintenance.

  Configure the cache on the secondary storage system so that it can adequately support the remote copy workload and all local workload activity. When the cache memory and the shared memory/control memory in the storage system become redundant, you can remove them. For instructions on adding and removing cache and shared memory/control memory, see Adding and removing cache and shared memory/control memory on page 3-7.

  When determining the amount of cache required for GAD, consider the amount of Cache Residency Manager data (VSP G1000 only) that will also be stored in cache.
- Make sure that the appropriate host modes and host mode options (HMOs) are set. For details, see the *Provisioning Guide* for the storage system.
  - HMO 78, the nonpreferred path option, must be configured to specify nonpreferred paths for HDLM operations.
  - HMOs 49, 50, and 51 can be used to improve response time of host I/O for distance direct connections (up to 10 km Long Wave).
- Make sure that the appropriate system option modes (SOMs) are set on your storage systems. For details about SOMs that apply to remote copy operations, see System option modes on page 3-2.

System option modes

The system option modes (SOMs) enable you to tailor the storage system to unique operating requirements and are set on the storage system by the Hitachi Data Systems representative.

The following table lists and describes the SOMs that apply to remote copy operations, including global-active device and TrueCopy. The settings of these SOMs affects both global-active device and TrueCopy operations. For a complete list of the SOMs, see the *Hardware Guide* for the storage system.

Mode	Description
689	Allows you to slow the initial copy and resync operations when the write-pending rate on the RCU exceeds 60%. • ON: The initial copy and resync copy operations are slowed down when the write-pending rate on the RCU exceeds 60% and when the write-pending rate of the MP PCB to which the S-VOL belongs exceeds 60%. • OFF (default): The initial copy and resync copy operations are not slowed down when the write-pending rate on the RCU or MP PCB exceeds 60%. **Notes:** 1. SOM 689 can be set online. 2. The microcode on both MCU and RCU must support this mode. 3. This mode should be set when requested by the user. 4. Set this mode to ON when GAD is installed, as the performance degradation is more likely to occur due to active-active I/Os. 5. If the write-pending status remains at 60% or higher on the RCU for a long time, it takes extra time for the initial copy and resync copy to be completed due to the slower copy operations. 6. Do not set this mode if the primary or secondary system is connected to USP V/VM with microcode earlier than 60-02-*xx*-*xx*/*xx*. If this mode is applied and the write-pending rate reaches 60%, pair suspend might occur. 7. As this mode is enabled per storage system, in an environment where TC and GAD are used, the mode is applied to both program products. When GAD is installed in a storage system that already uses TC, TC initial copy might take a longer time.
784	SOM 784 can reduce the MIH watch time of RIO for a TC or GAD pair internally so that update I/Os can continue by using an alternate path without MIH or time-out occurrence in the environment where open host time-out time is short (15 seconds or less). VSP G1000: This mode is applied to TC and GAD (80-01-42-00/00 and later). VSP G200, G400, G600, G800: This mode is applied to TC (83-01-01-*x*0/00 and later) and GAD (83-01-21-*x*0/00 and later). • ON: The MIH time of RIO is internally reduced so that, even though a path failure occurs between storage systems in the environment where host MIH time is set to 15 seconds, update I/Os can be processed by using an alternate path promptly, lowering the possibility of host MIH occurrence. • OFF (default): The operation is processed in accordance with the TC or GAD specification. **Notes:** 1. This mode also applies to TCz operations. For details, see the *Hitachi TrueCopy® for Mainframe User Guide*. 2. The mode is applied to the environment where open-systems host time-out time is set to 15 seconds or less. 3. The mode is applied to reduce RIO MIH time to 5 seconds.

Mode	Description
	4. This function is available for all the TC and GAD pairs on the storage system, unable to specify the pairs that are using this function or not. 5. For a TC or GAD pair with the mode effective (RIO MIH time is 5 seconds), the setting of RIO MIH time made at RCU registration (default is 15 seconds, which can be changed within range from 10 to 100 seconds) is invalid. However, RIO MIH time displayed on Device Manager - Storage Navigator and CCI is not 5 seconds but is what was set at RIO registration. 6. If a failure occurs on the switched path between storage systems, open server time-out might occur. 7. If an MP to which the path between storage systems belongs is overloaded, switching to an alternate path delays, and host time-out might occur. 8. If an RIO retry occurs due to other factors than RIO MIH (5 seconds), such as a check condition report issued from RCU to MCU, the RIO retry is performed on the same path instead of an alternate path. If a response delay to the RIO occurs constantly on this path due to path failure or link delay, host time-out might occur due to response time accumulation for each RIO retried within 5 seconds. 9. Even though the mode is set to ON, if open host time-out time is set to 10 seconds or less, host time-out might occur due to a path failure between DKCs. 10. Operation commands are not available for promptly switching to an alternate path. 11. The mode works for the pair for which initial pair creation or resync operation is executed. 12. Microcode downgrade to an unsupported version cannot be executed unless all the TC and GAD pairs are suspended or deleted. 13. For GAD pairs, the mode is effective if the microcode is in a version that supports GAD. 14. SOM 784 is supported only when the port type of all remote paths in a single remote connection is Fibre Channel. If the port type of all remote paths in a single remote connection is iSCSI or if iSCSI and Fibre Channel are mixed, this SOM is not supported.
1050* (VSP G1000 only)	Enables creation of pairs using user capacity in excess of 1.8 PB per system by managing differential BMP in hierarchical memory for pair volumes whose capacity is 4 TB (open) or 262,668 Cyl (mainframe) or less. • ON: For pair volumes of 4 TB (open)/262,668 Cyl (mainframe) or less, differential BMP is managed in hierarchical memory that performs caching to CM/PM using HDD as a master and enables creation of pairs using user capacity in excess of 1.8 PB per system. • OFF (default): For pair volumes of 4 TB (open)/262,668 Cyl (mainframe) or less, differential BMP is managed in shared memory (SM) as usual so that the user capacity to create pairs is limited to 1.8 PB per system. Also, differential MPB management can be switched from hierarchical memory to SM by performing a resync operation for pairs whose volume capacity is 4 TB (open)/262,668 Cyl (mainframe) or less. **Notes:** 1. This mode is related to SOM 1058. 2. Hierarchical memory is used for volumes over 4 TB (open)/262,668 Cyl (mainframe) even when the mode is set to OFF.

Mode	Description
	3. When the mode is turned OFF after operating with the mode set to ON, if resynchronizing pairs whose volume capacity is 4 TB (open)/262,668 Cyl (mainframe) or less managed in the hierarchical memory, the pair resync operation might fail due to lack of the SM capacity. 4. If a pair resync operation fails, pairs exceeding 1.8 PB of user capacity must be deleted. 5. OPEN non-DP-VOLs do not support hierarchical memory differential management so that the mode cannot be applied. 6. To downgrade the microcode, the procedure described in Maintenance Manual (Function ID 01:03 described in 03 MICRO-FC) is required. Refer to the manual before applying the mode. 7. Apply this mode when you want to create pairs using user capacity exceeding 1.8 PB per system. 8. Depending on the I/O pattern, I/O response might be affected. 9. Depending on the condition, initial copy time might be affected. 10. The differential BMP managed in hierarchical memory does not use the shared memory, but the differential data is managed in a pool with which the volumes are associated. For the pool capacity necessary for the differential data management, refer to "Differential Data" in the manuals for TrueCopy, Universal Replicator, and global-active device. 11. Before using this SOM to increase pair volumes on the storage system, check the resources on the storage system to make sure all performance requirements are met.
1058* (VSP G1000 only)	Changes differential BMP management from SM to hierarchical memory so that the number of pairs to be created on a system and user capacity used for pairs increase. • For mainframe systems, all pairs can be managed in hierarchical memory so that pairs can be created by all LDEVs. • For open systems, pairs that can only be managed in SM use SM so that the number of pairs that can be created using non-DP-VOLs increases. **Mode 1058 = ON**: • **When SOM 1050 is set to ON:** o By resynchronizing mainframe VOLs of 262,668 Cyl or less, the differential BMP management is switched from SM to hierarchical memory. Hierarchical memory management remains as is. o By resynchronizing open VOLs (DP-VOLs only) of 4 TB or less, the differential BMP management is switched from SM to hierarchical memory. Hierarchical memory management remains as is. • **When SOM 1050 is set to OFF:** o By resynchronizing mainframe VOLs of 262,668 Cyl or less, the differential BMP management is switched from hierarchical memory to SM. SM management remains as is. o By resynchronizing open VOLs (DP-VOLs only) of 4 TB or less, the differential BMP management is switched from hierarchical memory to SM. SM management remains as is. **Mode 1058 = OFF (default)**: • **When SOM 1050 is set to ON:** o The differential BMP management does not change by resynchronizing pairs.

Mode	Description

	• **When SOM 1050 is set to OFF:**
	  ◦ By resynchronizing mainframe VOLs of 262,668 Cyl or less, the differential BMP management is switched from hierarchical memory to SM. SM management remains as is.
	  ◦ By resynchronizing open VOLs (DP-VOLs only) of 4 TB or less, the differential BMP management is switched from hierarchical memory to SM. SM management remains as is.
	Notes:
	1. Depending on the I/O pattern, I/O response might be affected.
	2. Depending on the condition, initial copy time might be affected.
	3. This mode is related to SOM 1050.
	4. OPEN non-DP-VOLs do not support hierarchical memory differential management so that the mode cannot be applied.
	5. To downgrade the microcode, the procedure described in Maintenance Manual (Function ID 01:03 described in 03 MICRO-FC) is required. Refer to the manual before applying the mode.
	6. Apply this mode when you want to increase the number of pairs that can be created on a system or the user capacity for pairs.
	7. The differential BMP managed in hierarchical memory does not use shared memory, but the differential data is managed in a pool with which the volumes are associated. For the pool capacity necessary for the differential data management, refer to "Differential Data" in the manuals for TrueCopy, Universal Replicator, and global-active device.
	8. Before using this SOM to increase pair volumes on the storage system, check the resources on the storage system to make sure all performance requirements are met.

1070	Allows you to reduce the time required to change the status of all pairs in a consistency group when you suspend or resynchronize a GAD pair by consistency group. When you suspend or resynchronize a GAD pair by consistency group, the command responds to the pair operation, and then the status of the pairs in the consistency group changes in sequence at a constant frequency. • ON: More pairs are processed per cycle than when SOM 1070 is OFF, which reduces the time required to change the status of all pairs in a consistency group. The MP usage rate might increase due to the increased pair status processing. • OFF (default): Fewer pairs are processed per cycle than when SOM 1070 is ON. The MP usage rate increases even when SOM 1070 is OFF, but the increase is less than when SOM 1070 is ON. For details about storage system support (models, microcode) for consistency groups, see Requirements and restrictions on page 2-2.
1080 (VSP G1000 only)	The quorum disk is blocked if the health check on the quorum disk is delayed and does not finish within 15 seconds. If the health check is delayed, you can change the timeout time until a retry occurs on the alternate path. • ON: The timeout time for the health check command is set to three seconds. The timeout time for ABTS is set to 15 seconds. If ABTS finishes within 12 seconds and a retry attempt is made to perform a health check on the alternate path, the quorum disk will not be blocked. If ABTS does not finish within 12 seconds, the quorum disk will

Mode	Description
	be blocked. If ABTS finishes within 15 seconds, the other commands running on the same path will not be reset. If ABTS does not finish within 15 seconds, the other commands running on the same path will be reset. • OFF: The timeout time for the health check command is set to three seconds. The timeout time for ABTS is set to three seconds. If ABTS does not finish within three seconds and a retry attempt is made to perform a health check on the alternate path, the quorum disk will not be blocked. If ABTS finishes within three seconds, the other commands running on the same path will not be reset. If ABTS does not finish within three seconds, the other commands running on the same path will be reset.
* For microcode 80-02-4*x* or later, you should use hierarchical memory for all remote copy pairs. Using SOMs 1050 and 1058 enables the use of hierarchical memory for volumes of 4 TB or less, which increases the total amount of replicated capacity above 1.8 PB. **WARNING:** If you allocate more than 1.8 PB of replicated capacity for volumes of 4 TB or less and then later want to return to using shared memory/control memory, you will have to delete some pairs to reduce the total replicated capacity for volumes of 4 TB or less to below 1.8 PB.	

Adding and removing cache and shared memory/control memory

This section provides workflows for adding and removing cache and shared memory/control memory in a storage system in which GAD pairs already exist.

Adding and removing cache memory

Use the following workflow to add or remove cache memory in a storage system in which GAD pairs already exist:

1. Identify the status of the GAD volumes in the storage system.
2. If a GAD volume is in the COPY status, wait until the status changes to PAIR, or suspend the GAD pair.

   Do not add or remove cache memory when any volumes are in the COPY status.
3. When the status of all volumes has been confirmed, cache memory can be added to or removed from the storage system by your service representative. Contact Hitachi Data Systems customer support for adding or removing cache memory.
4. After the addition or removal of cache memory is complete, resynchronize the pairs that you suspended in step 2.

Adding shared memory/control memory

Use the following workflow to add shared memory/control memory to a storage system in which GAD pairs already exist:

1. Identify the status of the GAD volumes in the storage system.

2. If a GAD volume is in the COPY status, wait until the status changes to PAIR, or suspend the GAD pair.

   Do not add shared memory/control memory when any volumes are in the COPY status.
3. When the status of all volumes has been confirmed, shared memory/control memory can be added to the storage system by your service representative. Contact Hitachi Data Systems customer support for adding shared memory/control memory.
4. After the addition of shared memory/control memory is complete, resynchronize the pairs that you suspended in step 2.

Removing shared memory used as 64KLDEV (VSP G1000)

Use the following workflow to remove shared memory used as 64KLDEV (VSP G1000):

1. Identify the status of all volumes with an LDEV ID of 0x4000 or higher.
2. If a volume with an LDEV ID of 0x4000 or higher is used by a GAD pair, delete the GAD pair.

   Do not remove shared memory used in Extension1 when any volume with an LDEV ID of 0x4000 or higher is used by a GAD pair.
3. When the status of all volumes with an LDEV ID of 0x4000 or higher has been confirmed, shared memory can be removed from the storage system by your service representative. Contact Hitachi Data Systems customer support for removing shared memory.

Removing shared memory used in TC/UR/GAD (VSP G1000)

Use the following workflow to remove shared memory used in TC/UR/GAD (VSP G1000):

1. Identify the status of all volumes.
2. If a volume is used by a TC/UR/GAD pair, delete the TC/UR/GAD pair.

   Do not remove shared memory used in TC/UR/GAD when any volume is used by a GAD pair.
3. When the status of all volumes has been confirmed, shared memory can be removed from the storage system by your service representative. Contact Hitachi Data Systems customer support for adding or removing cache memory.

Removing control memory used in Extension1 (VSP Gx00 models)

Use the following workflow to remove control memory used in Extension1 (VSP G200, G400, G600, G800):

1. Identify the status of all volumes with an LDEV ID of 0x4000 or higher.
2. If a volume with an LDEV ID of 0x4000 or higher is used by a GAD pair, delete the GAD pair.

   Do not remove control memory used in Extension1 when any volume with an LDEV ID of 0x4000 or higher is used by a GAD pair.

3. When the status of all volumes with an LDEV ID of 0x4000 or higher has been confirmed, control memory can be removed from the storage system by your service representative. Contact Hitachi Data Systems customer support for removing control memory.

Removing control memory used in Extension2 (VSP Gx00 models)

Use the following workflow to remove control memory used in Extension2:

1. Identify the status of all volumes.
2. If a volume is used by a GAD pair, delete the GAD pair.

   Do not remove control memory used in Extension2 when any volume is used by a GAD pair.
3. When the status of all volumes has been confirmed, control memory can be removed from the storage system by your service representative. Contact Hitachi Data Systems customer support for adding or removing control memory.

Planning system performance

Remote copy operations can affect the I/O performance of host servers and the primary and secondary storage systems. You can minimize the effects of remote copy operations and maximize efficiency and speed by changing your remote connection options and remote replica options.

Your Hitachi Data Systems account team can help you analyze your workload and optimize copy operations. Using workload data (MB/s and IOPS), you can determine the appropriate amount of bandwidth, number of physical paths, and number of ports for your global-active device system. When these are properly determined and sized, the data path should operate free of bottlenecks under all workload levels.

Related topics

- Setting preferred and nonpreferred paths using HDLM (VSP G1000 only) on page 3-10

Setting preferred and nonpreferred paths (VSP G1000 only)

If an alternate path that connects a server and a storage system in a GAD configuration contains a short-distance straight path and a long-distance cross path, I/O performance varies depending on the path. By setting the short-distance straight path as a preferred I/O path, you can suppress I/Os to and from the inefficient long-distance cross path. As a result, overall system performance can be improved.

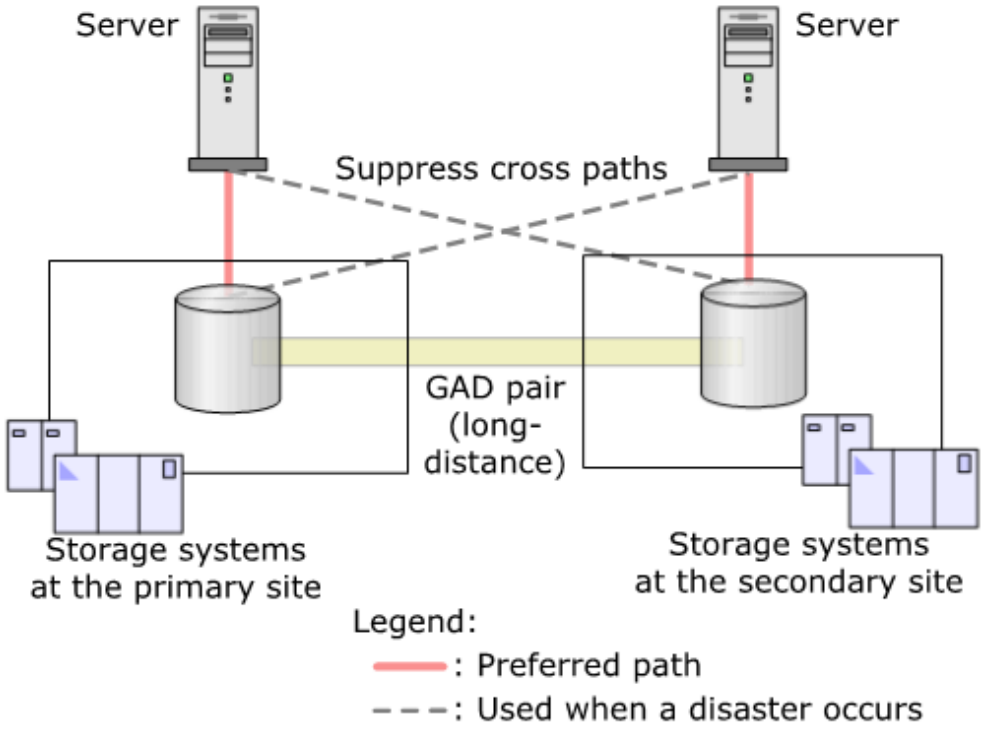

Related topics

- Setting preferred and nonpreferred paths using ALUA on page 3-10
- Setting preferred and nonpreferred paths using HDLM (VSP G1000 only) on page 3-10

Setting preferred and nonpreferred paths using ALUA

When you perform Asymmetric Logical Unit Access (ALUA) in a cross-path configuration, you can specify the preferred path to use for issuing an I/O request from a server to a storage system. To specify the preferred path, you must enable the ALUA mode in the storage system, and set the path to use as the preferred path.

Related topics

- System configurations for GAD solutions on page 1-6
- Planning system performance on page 3-9
- Setting preferred and nonpreferred paths using HDLM (VSP G1000 only) on page 3-10

Setting preferred and nonpreferred paths using HDLM (VSP G1000 only)

Hitachi Dynamic Link Manager (HDLM) allows you to specify alternate paths to be used for normal global-active device operations by using host mode options. Other paths are used when failures occur in the paths (including

alternate paths) that should be used for normal operations. Host mode option (HMO) 78, the nonpreferred path option, must be configured to specify nonpreferred paths, which are used when failures occur.

When VSP G1000 settings are applied to HDLM, the attribute of the HDLM path to which HMO 78 was set changes to non-owner path. The attribute of the HDLM path to which HMO 78 was not set changes to owner path. For details, see the documents for HDLM version 8.0.1 or later.

Hitachi Dynamic Link Manager (VSP G200, G400, G600, G800 only)

Hitachi Dynamic Link Manager (HDLM) allows you to specify alternate paths to be used for normal global-active device operations. Other paths are used when failures occur in the paths (including alternate paths) that should be used for normal operations. Host mode option (HMO) 78, the nonpreferred path option, must be configured to specify nonpreferred paths, which are used when failures occur.

For example, if servers and storage systems are connected in a cross-path configuration, I/O response is prolonged because the primary-site server is distant from the secondary storage system, and the secondary-site server is distant from the primary storage system. Normally in this case you use paths between the primary server and primary storage system and paths between the secondary server and secondary storage system. If a failure occurs in a path used in normal circumstances, you will use the paths between the primary server and secondary storage system, and paths between the secondary server and primary storage system.

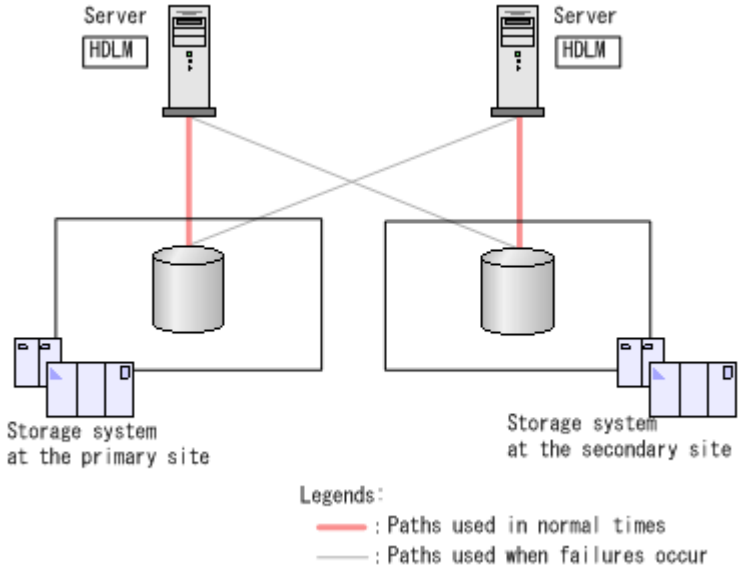

After you incorporate the storage system settings to HDLM, the attribute of the HDLM path to which HMO 78 was set changes to non-owner path. If HMO 78 is not set to the path, the HDLM path attribute changes to owner path.

Related topics

- Planning system performance on page 3-9

- Setting preferred and nonpreferred paths using ALUA on page 3-10

Planning physical paths

When configuring physical paths to connect the storage systems at the primary and secondary sites, make sure that the paths can handle all of the data that could be transferred to the primary and secondary volumes under all circumstances. Review the information in the following sections when planning physical paths:

- Determining the required bandwidth on page 3-12
- Fibre Channel connections on page 3-12
- iSCSI data path requirements on page 3-14
- Connection types on page 3-15

Note: Use the same protocol for data paths between a host and a storage system and between primary and secondary storage systems. When different protocols are used in the data paths (for example, Fibre Channel data paths between the host and storage system and iSCSI data paths between the storage systems), make sure the timeout period for commands between the host and the storage system is equal to or greater than the timeout period for commands between the storage systems.

Determining the required bandwidth

You must have sufficient bandwidth to handle all data transfers at all workload levels. The amount of required bandwidth depends on the amount of server I/O to the primary volumes. To identify the required bandwidth, you must first collect the write workload data under all workload conditions, including peak write workload, and then measure and analyze the data. You can use performance-monitoring software such as Hitachi Tuning Manager or Hitachi Performance Monitor to collect the workload data.

Related topics

- Planning physical paths on page 3-12

Fibre Channel connections

For Fibre Channel connections, direct, switch, and extender connections are supported. Use short-wave (optical multi-mode) or long-wave (optical single-mode) optical fiber cables to connect the storage systems at the primary and secondary sites. The required cables and network relay devices differ depending on the distance between the primary and secondary storage systems, as described in the following table.

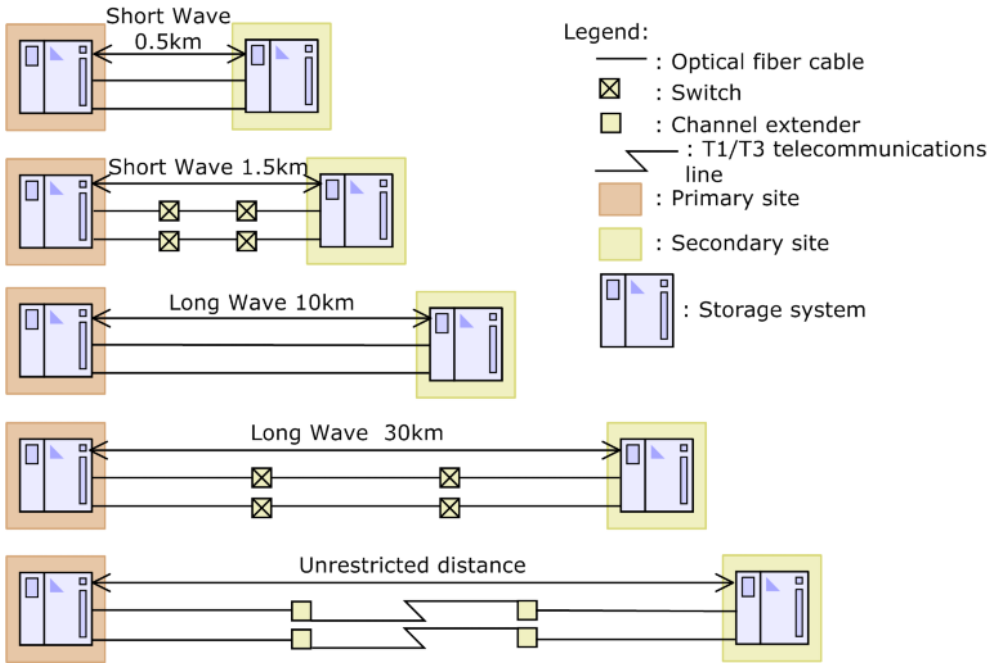

Distance between storage systems	Cable type	Network relay device
Up to 1.5 km	Short wave (optical multi-mode)	Switches are required if the distance is 0.5 to 1.5 km.
1.5 to 10 km	Long wave (optical single-mode)*	Not required.
10 to 30 km	Long wave (optical single-mode)*	Switches must be used.
30 km or longer	Communication line	An authorized third-party channel extender is required.
* Long wave cannot be used for FCoE (VSP G1000 only).		

No special settings are required for the storage system if switches are used in a Fibre Channel environment.

Long-wave (optical single-mode) cables can be used for direct connection at a maximum distance of 10 km. The maximum distance that might result in the best performance differs depending on the link speed, as shown in the following table. For details about the availability of serial-channel GAD connections, contact Hitachi Data Systems customer support.

Link speed	Maximum distance for best performance
1 Gbps	10 km
2 Gbps	6 km
4 Gbps	3 km
8 Gbps	2 km
16 Gbps	1 km

Related topics

- Planning physical paths on page 3-12

iSCSI data path requirements

For iSCSI data paths, direct, switch, and extender connections are supported. The following table describes the considerations for ports and network settings when you configure a system using iSCSI data paths.

Item	Considerations
Ports	• When the parameter settings for an iSCSI port are changed, the iSCSI connection is disconnected temporarily and then reconnected. Change parameter settings only when the I/O load is low to avoid impact to the system. • Although a log might be output in the host if you change the settings of an iSCSI port connected to a host, this does not indicate a failure. In a system that monitors system logs, an alert might be issued. If an alert is issued, change the iSCSI port setting, and then check if the host is reconnected. • **Delayed ACK:** When you use iSCSI for the connection between storage systems, and the same port for the connection to the host, disable the Delayed ACK setting for the port (in the Edit Ports window set Delayed ACK to Disable, default = Enable). If Delayed ACK is set to Enable, it might take time for the host to recognize a volume used by GAD pairs. For example, it takes up to 8 minutes for the host to recognize 2,048 volumes. • **Selective ACK:** Make sure Selective ACK is enabled for iSCSI ports. Selective ACK is set to Enable by default (Edit Ports window). Do not change this setting. • In an environment in which a delay occurs in a line between storage systems, such as long-distance connections, you must set an optimal window size of iSCSI ports in storage systems at the primary and secondary sites after verifying various sizes. The maximum value you can set is 1024 KB. The default window size is 64 KB, so you must change this setting. • iSCSI ports do not support the fragmentation (splitting packets) functionality. When the value for the maximum transfer unit (MTU) of a switch is smaller than the MTU value of the iSCSI port, packets are lost, and communication might not be performed correctly. Set the same MTU value for the iSCSI port and the switch. The MTU value for the iSCSI port must be greater than 1500. In a WAN environment in which the MTU value is smaller than 1500, fragmented data cannot be sent or received. In this environment, set a smaller value for the maximum segment size (MSS) of the WAN router according to the WAN environment, and then connect the iSCSI port. Alternatively, use iSCSI in an environment in which the MTU value is 1500 or greater. • A port can be used for connections to the host (target attribute) and to a storage system (initiator attribute). However, to minimize the impact on the system if a failure occurs either in the host or in a storage system, you should connect the port for

Item	Considerations
	the host and for the storage system to separate front end modules.
Network settings	• Disable the spanning tree setting on the port of a switch connecting to an iSCSI port. If you enable the spanning tree functionality of a switch, packets might not be looped in the network when the link is up or down. If this occurs, the packets might be blocked for approximately 30 seconds. If you must enable the spanning tree settings, enable the port fast functionality of the switch. • In a network path between storage systems, if you use a line that has a slower transfer speed than the iSCSI port, packets are lost, and the line quality is degraded. Configure the system so that the transfer speed for the iSCSI ports and the lines is the same. • Delay in a line between storage systems varies depending on the system environment. Therefore, validate the system first, and then check the optimum window size settings of iSCSI ports. If the impact of the line delay is determined to be unacceptable, consider using devices for optimizing or accelerating the WAN speed. • When iSCSI is used, packets are sent or received using TCP/IP. Because of this, the amount of packets might exceed the capacity of a communication line, or packets might be resent. As a result, performance might be greatly affected. Use Fibre Channel data paths for critical systems that require high performance.

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 on page 3-15
- Connection using switches on page 3-17
- Connection using channel extenders on page 3-18

Direct connection

With a direct connection, the two storage systems are connected directly to each other.

Note: You do not need to set the port attributes (Initiator, RCU Target, Target) on the VSP G200, G400, G600, G800.

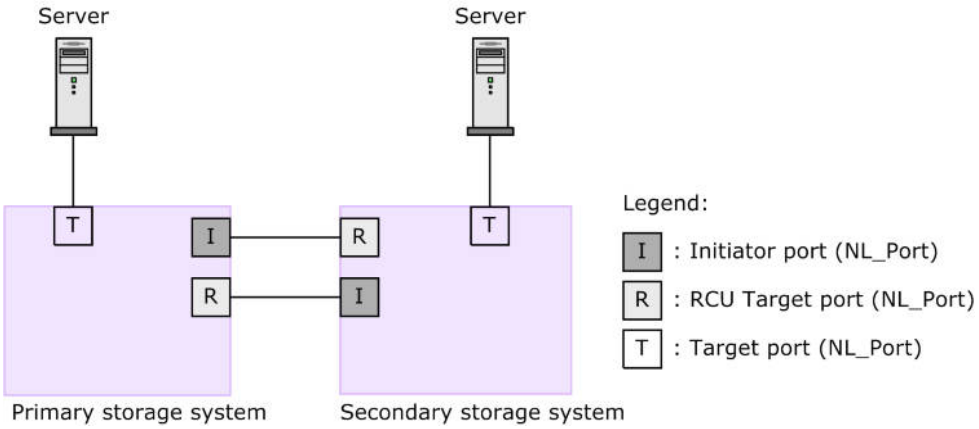

You can use the following host mode options (HMOs) to improve response time of host I/O by improving response time between the storage systems for distance direct connections (up to 10 km Long Wave) when the open package is used.

- HMO 49 (BB Credit Set Up Option1)
- HMO 50 (BB Credit Set Up Option2)
- 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, the protocol used for the connection between the storage systems, and the setting of HMO 51. The link speed that can be specified differs on each condition either.

Package name	Protocol	HMO 51 setting	Fabric setting	Topology: remote replication ports	Link speed that can be specified
16FC8 (VSP G1000) 4HF8 (VSP G200, G400, G600, G800)	8 Gbps FC	OFF	OFF	FC-AL	• 2 Gbps • 4 Gbps • 8 Gbps
	8 Gbps FC	ON	OFF	Point-to-Point	• 2 Gbps • 4 Gbps • 8 Gbps
16FC16 (VSP G1000)	16 Gbps FC	OFF	OFF	FC-AL	• 4 Gbps • 8 Gbps
	16 Gbps FC	ON	OFF	Point-to-Point	• 4 Gbps • 8 Gbps • 16 Gbps
8FC16 (VSP G1000)	16 Gbps FC	OFF	OFF	FC-AL	• 4 Gbps • 8 Gbps
	16 Gbps FC	ON	OFF	Point-to-Point	• 4 Gbps

Package name	Protocol	HMO 51 setting	Fabric setting	Topology: remote replication ports	Link speed that can be specified
2HF16 (VSP G200, G400, G600, G800)					• 8 Gbps • 16 Gbps

Related topics

- Connection types on page 3-15

Connection using switches

With a switch connection, optical fiber cables can be connected.

Note: You do not need to set the port attributes (Initiator, RCU Target, Target) on the VSP G200, G400, G600, G800.

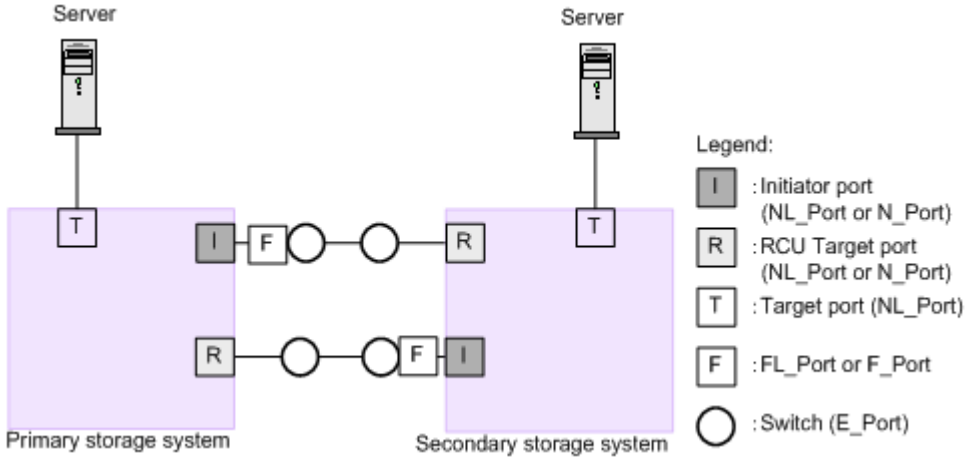

Switches from some vendors (for example, McData ED5000) require F_port.

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 100 km) and the open package is used.

- HMO 49 (BB Credit Set Up Option1)
- HMO 50 (BB Credit Set Up Option2)
- 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 either.

Package name	Protocol	HMO 51 setting	Fabric setting	Topology: Initiator and RCU Target	Link speed that can be specified
16FC8 (VSP G1000) 4HF8 (VSP G200, G400, G600, G800)	8 Gbps FC	OFF	ON	Point-to-Point	• 2 Gbps • 4 Gbps • 8 Gbps
	8 Gbps FC	ON	ON	Point-to-Point	• 2 Gbps • 4 Gbps • 8 Gbps
8FC16 (VSP G1000) 2HF16 (VSP G200, G400, G600, G800)	16 Gbps FC	OFF	ON	Point-to-Point	• 4 Gbps • 8 Gbps • 16 Gbps
	16 Gbps FC	ON	ON	Point-to-Point	• 4 Gbps • 8 Gbps • 16 Gbps
16FC16 (VSP G1000)	16 Gbps FC	OFF	ON	Point-to-Point	• 4 Gbps • 8 Gbps • 16 Gbps
	16 Gbps FC	ON	ON	Point-to-Point	• 4 Gbps • 8 Gbps • 16 Gbps
16FE10 (VSP G1000)	10 Gbps FCoE	OFF	ON	Point-to-Point	10 Gbps
	10 Gbps FCoE	ON	ON	Point-to-Point	10 Gbps

Related topics

- Connection types on page 3-15

Connection using channel extenders

Channel extenders and switches should be used for long-distance connections.

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

Note: You do not need to set the port attributes (Initiator, RCU Target, Target) on the VSP G200, G400, G600, G800.

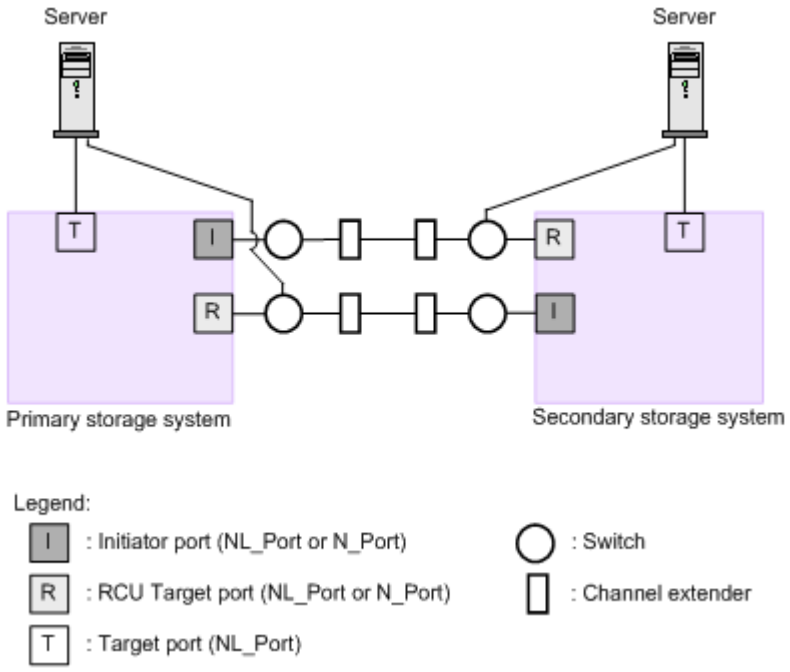

Note:

- 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 Hitachi Data Systems customer support for more information.
- Make sure that your channel extenders can support remote I/O. For details, contact Hitachi Data Systems 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.

Related topics

- Connection types on page 3-15

Planning ports (VSP G1000 only)

Data is transferred from Initiator ports in one storage system to RCU Target ports in the other system. Once you have identified the peak write workload, which is the amount of data transferred during peak periods, you can determine the amount of bandwidth and the number of Initiator and RCU Target ports required for your system.

The following describes the port attributes that you must set on the VSP G1000. For details on setting the port attributes, see the *Hitachi Virtual Storage Platform G1000 Provisioning Guide for Open Systems*.

- **Initiator ports:** Send remote copy commands and data to the RCU Target ports on a connected storage system. One Initiator port can be connected to a maximum of 64 RCU Target ports.

Caution: Do not add or delete a remote connection or add a remote path at the same time that the SCSI path definition function is in use.

- **RCU Target ports:** Receive remote copy commands and data from the Initiator ports on a connected storage system. One RCU Target port can be connected to a maximum of 16 Initiator ports.

  The number of remote paths that can be specified does not depend on the number of ports. The number of remote paths can be specified for each remote connection.
- **Target ports:** Connect the storage system to the host servers. When a server issues a write request, the request is sent from a Target port on the storage system to a VSP G1000 volume.
- **External ports:** Connect the storage system to external storage systems configured using Universal Volume Manager. The external storage system for the GAD quorum disk is connected to an external port on the primary and secondary storage systems.

Planning the quorum disk

An external storage system must be prepared for the GAD quorum disk. Review the information in the following sections when planning the quorum disk:

- Installation of the external storage system on page 3-20
- Relationship between the quorum disk and number of remote connections on page 3-22
- Relationship between quorum disks and consistency groups on page 3-26
- Response time from the external storage system on page 3-28
- Cache pending rate of the CLPR to which the quorum disk is assigned on page 3-28

Installation of the external storage system

The external storage system can be installed in the following locations, depending on the number of sites in your configuration:

- In a three-site configuration, the external storage system is installed in a third site away from the primary and secondary sites. I/O from servers continues if any failure occurs at the primary site, the secondary site, or the site where the external storage system is installed.

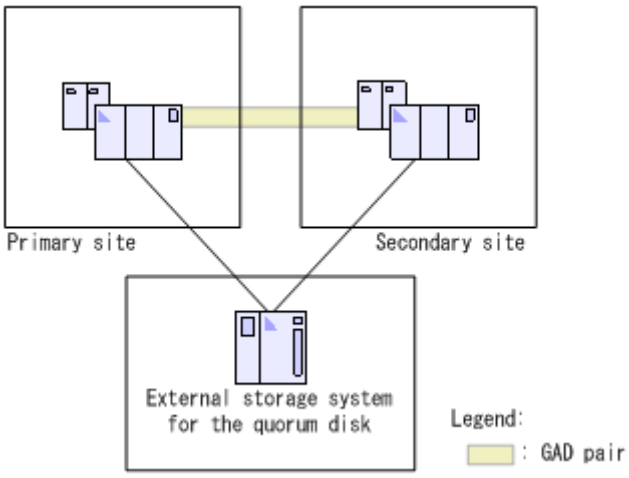

- In a two-site configuration, the external storage system is installed at the primary site. If failure occurs at the secondary site, I/O from servers will continue. However, if a failure occurs at the primary site, I/O from servers will stop.

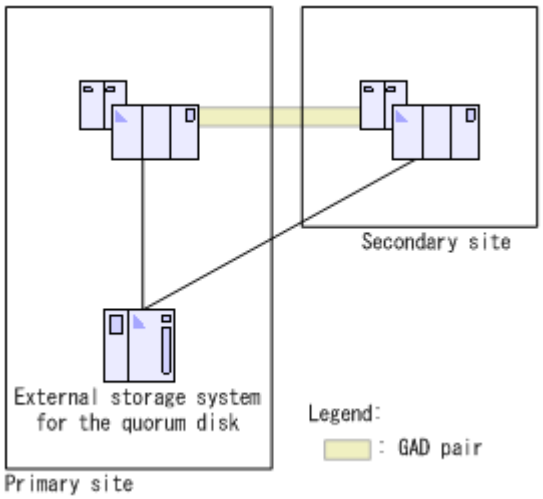

At the secondary site, you cannot install any external storage system for quorum disks.

Note: When you use iSCSI in the remote paths between the primary storage system and the external storage system for the quorum disk or between the secondary storage system and the external storage system for the quorum disk, the quorum disk blockade might occur due to one remote path failure.

Related topics

- Planning the quorum disk on page 3-20

Relationship between the quorum disk and number of remote connections

When you use multiple remote connections, you should prepare as many quorum disks as remote connections to avoid the possibility of a single remote connection failure causing the suspension of the GAD pairs that are using the other normal remote connections. Simultaneously, you must make a combination of one quorum disk, one remote connection from the primary storage system to the secondary storage system, and one remote connection from the secondary storage system to the primary storage system.

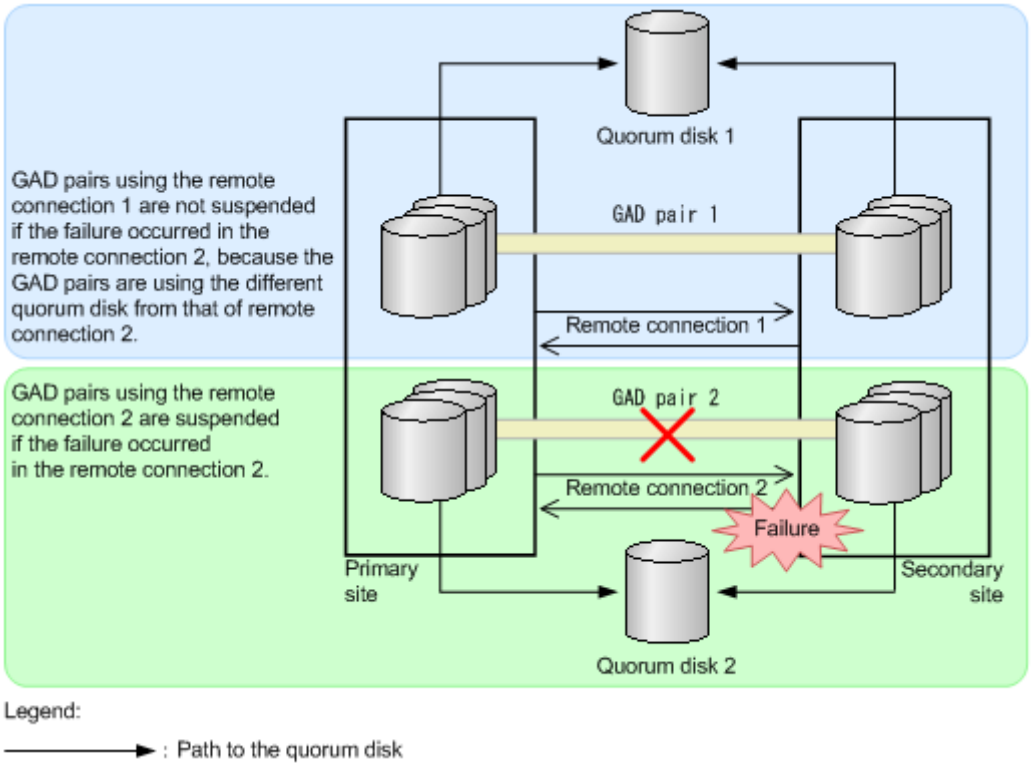

Tip: If you are planning to manage many GAD pairs using one quorum disk, if more than 8 physical paths are necessary for the remote connection, you can configure the system with one quorum disk for two or more remote connections.

When all paths used in the remote connection are blocked, the GAD pairs will be suspended in units of quorum disks. In the configuration shown below, the GAD pairs that are using remote connection 1 will be suspended even if the failure occurred at remote connection 2. Also, when a failure occurs at the path from the volume at the primary site or the secondary site to the quorum disk, the GAD pairs that are using the same quorum disk will be suspended.

- Suspended pairs depending on failure location (quorum disk not shared) on page 3-23
- Suspended pairs depending on failure location (quorum disk shared) on page 3-25

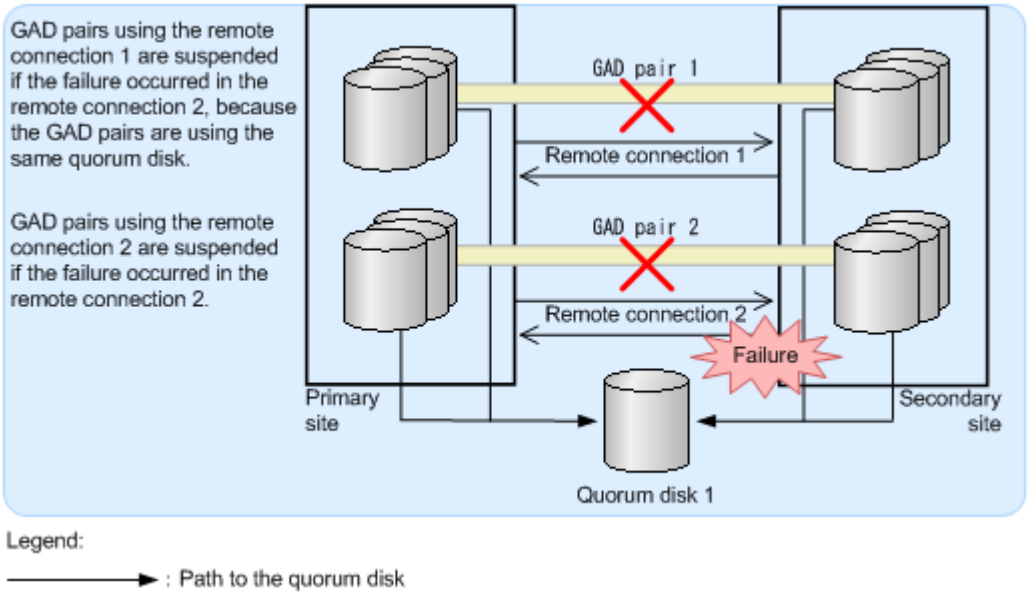

Related topics

- Planning the quorum disk on page 3-20

Suspended pairs depending on failure location (quorum disk not shared)

When the same number of quorum disks as the remote connections are used, only GAD pair that uses the failed remote connection, a quorum disk or a path to the quorum disk is suspended. The GAD pair that uses the normal remote connection, quorum disk and path to the quorum disk can keep the status being mirrored. The following figure shows the relationship between the failure locations and the GAD pair suspended by the failure.

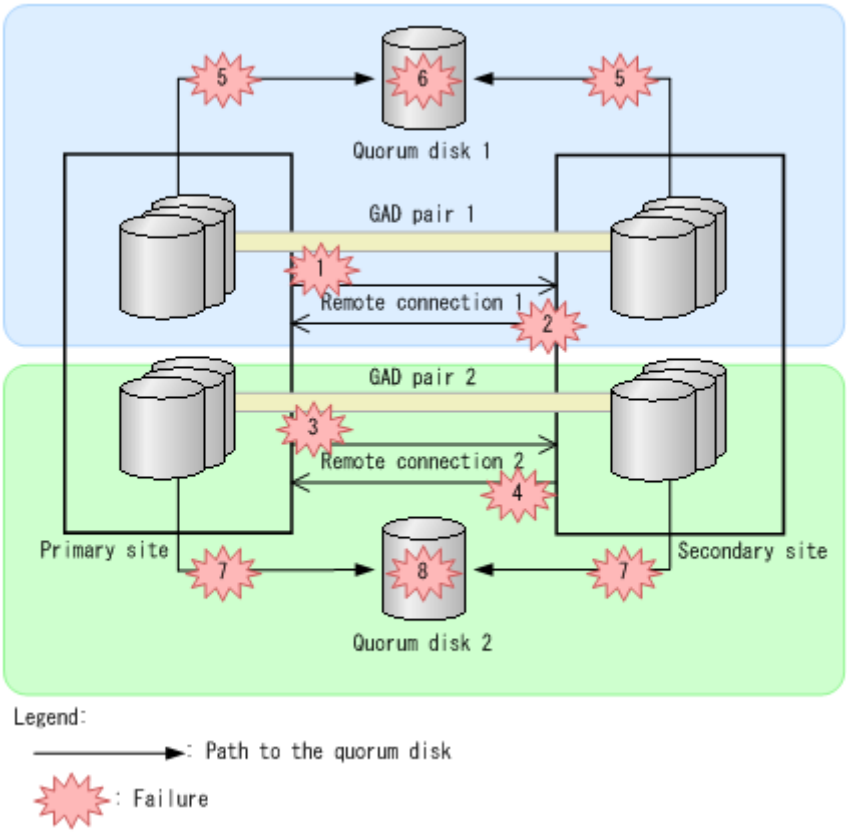

#	Failure locations	GAD pair 1	GAD pair 2
1	Remote connection 1 from the primary site to the secondary site	Suspended	Not suspended
2	Remote connection 1 from the secondary site to the primary site	Suspended	Not suspended
3	Remote connection 2 from the primary site to the secondary site	Not suspended	Suspended
4	Remote connection 2 from the secondary site to the primary site	Not suspended	Suspended
5	Path to the quorum disk 1	VSP G1000: Not suspended* VSP G200, G400, G600, G800: Suspended	Not suspended
6	Quorum disk 1	VSP G1000: Not suspended* VSP G200, G400, G600, G800: Suspended	Not suspended
7	Path to the quorum disk 2	Not suspended	VSP G1000: Not suspended*

#	Failure locations	GAD pair 1	GAD pair 2
			VSP G200, G400, G600, G800: Suspended
8	Quorum disk 2	Not suspended	VSP G1000: Not suspended* VSP G200, G400, G600, G800: Suspended
* For VSP G1000 the GAD pair is not suspended, but I/O mode of the S-VOL changes to Block.			

Related topics

- Relationship between the quorum disk and number of remote connections on page 3-22

Suspended pairs depending on failure location (quorum disk shared)

When a quorum disk is shared by more than one connections, all GAD pairs which share a quorum disk are suspended, regardless of the failure locations, as shown below.

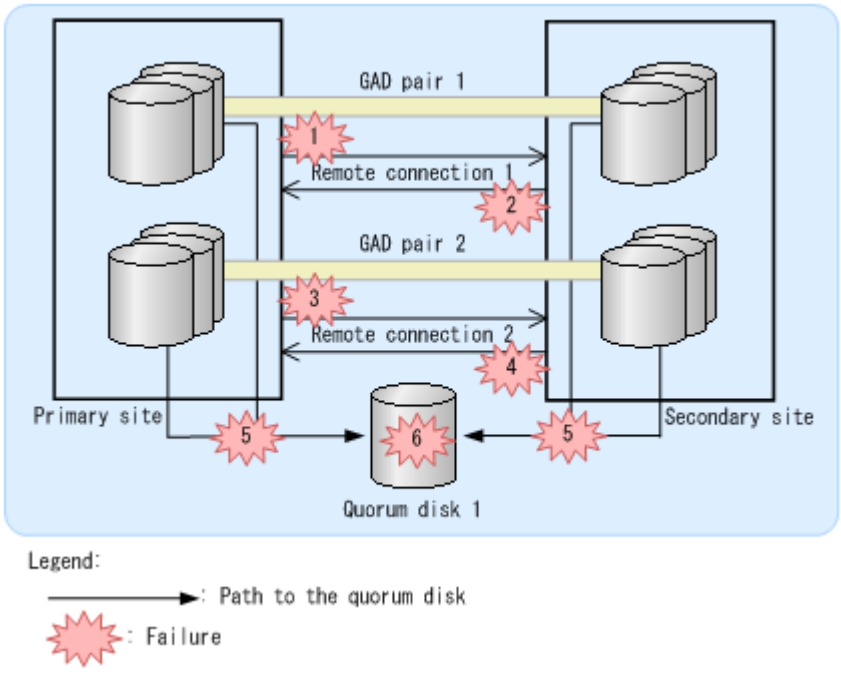

#	Failure locations	GAD pair 1	GAD pair 2
1	Remote connection 1 from the primary site to the secondary site	Suspended	Suspended
2	Remote connection 1 from the secondary site to the primary site	Suspended	Suspended

#	Failure locations	GAD pair 1	GAD pair 2
3	Remote connection 2 from the primary site to the secondary site	Suspended	Suspended
4	Remote connection 2 from the secondary site to the primary site	Suspended	Suspended
5	Path to the quorum disk 1	VSP G1000: Not suspended* VSP G200, G400, G600, G800: Suspended	VSP G1000: Not suspended* VSP G200, G400, G600, G800: Suspended
6	Quorum disk 1	VSP G1000: Not suspended* VSP G200, G400, G600, G800: Suspended	VSP G1000: Not suspended* VSP G200, G400, G600, G800: Suspended
* For VSP G1000 the GAD pair is not suspended, but I/O mode of the S-VOL changes to Block.			

Related topics

- Relationship between the quorum disk and number of remote connections on page 3-22

Relationship between quorum disks and consistency groups

A single quorum disk can be shared by multiple consistency groups. When creating GAD pairs to be registered to different consistency groups, you can specify the same quorum disk ID.

For details about storage system support (models, microcode) for consistency groups, see Requirements and restrictions on page 2-2.

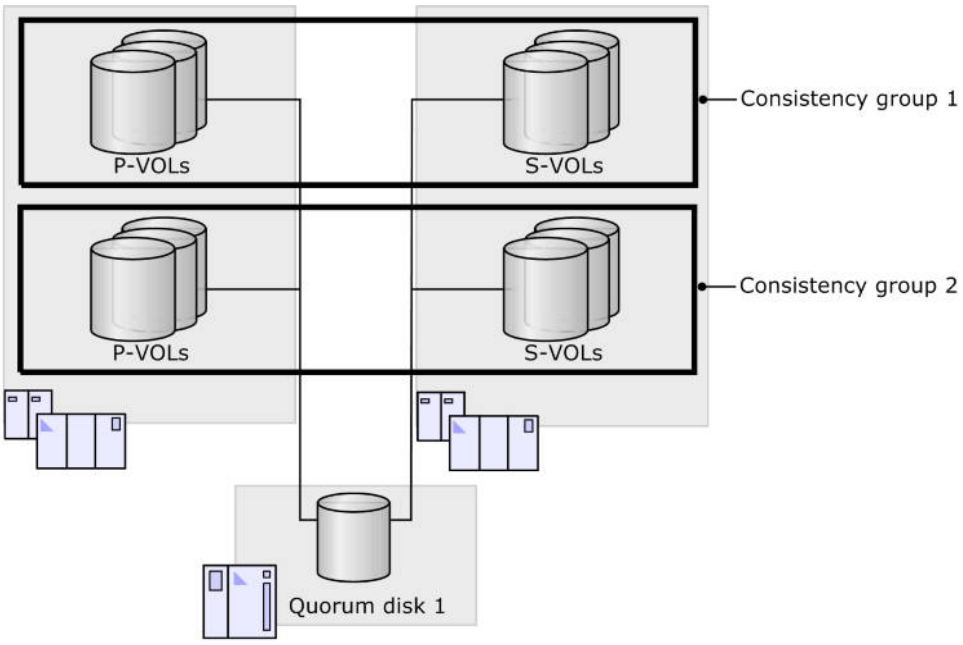

Pairs registered to the same consistency group must use the same quorum disk. When creating pairs in a single consistency group, you cannot specify multiple quorum disk IDs.

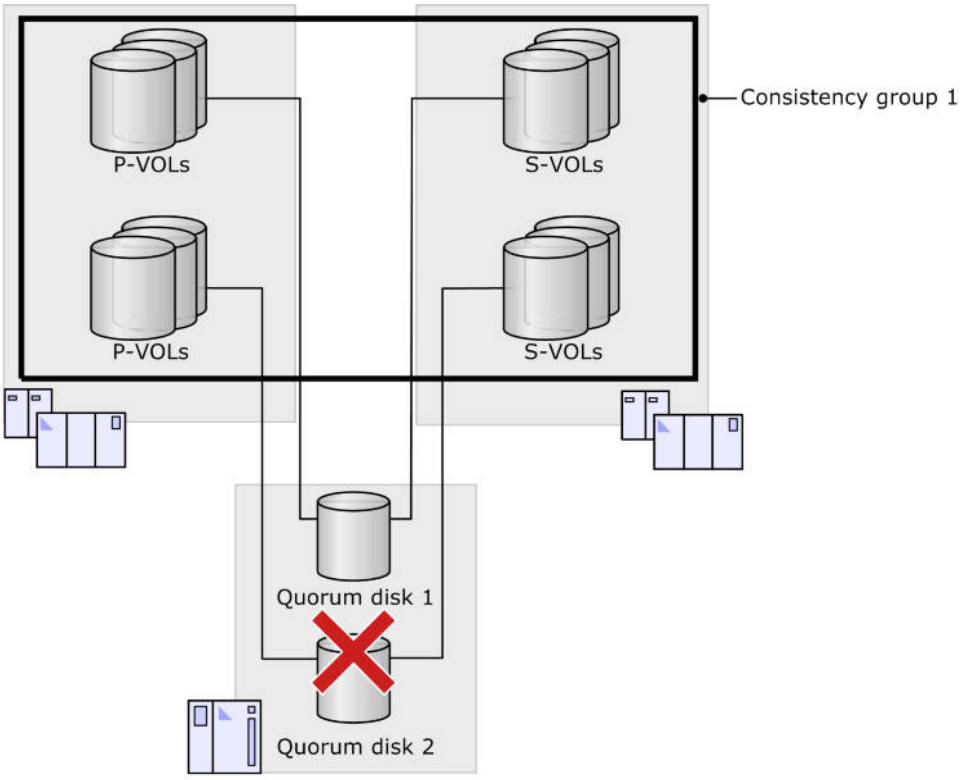

For details about storage system support (models, microcode) for consistency groups, see Requirements and restrictions on page 2-2.

Related topics

- Planning the quorum disk on page 3-20

Response time from the external storage system

If the response time from the external storage system for quorum disks is delayed for more than one second, GAD pairs might be suspended by some failures. You should monitor the response time of the quorum disks regularly using Performance Monitor on the primary or secondary storage system. Specify **External storage** > **Logical device** > **Response time (ms)** on the monitoring objects. If the response time exceeds 100 ms, review the configuration and consider the following actions:

- Lower the I/O load, if the I/O load of volumes other than the quorum disk is high in the external storage system.
- Remove the causes of the high cache load, if the cache load is high in the external storage system.
- Lower the I/O load of the entire external storage system, when you perform maintenance of the external storage system. Alternatively, perform maintenance on the external storage system with settings that will minimize the impact to the I/O, referring to the documentation for the external storage system.

Related topics

- Planning the quorum disk on page 3-20

Cache pending rate of the CLPR to which the quorum disk is assigned

If the write-pending rate of the CLPR to which the quorum disk (external volume) on the primary or secondary storage systems is assigned is high, the I/O performance of the GAD pair volumes might decrease or the GAD pairs might be suspended by some failure. To address this situation:

1. Use Performance Monitor on the primary or secondary storage system to perform regular monitoring of the write-pending rate of the CLPR to which the quorum disks are assigned (specify **Cache** > **Write Pending Rate (%)** on the monitoring objects). For details, see the *Performance Guide* for the storage system.
2. If the write-pending rate exceeds 70%, review your configuration and consider the following actions:
    - Lower the I/O load in the storage system.
    - If the cache load is high:

      - Lower the I/O load.

      - Migrate the quorum disk to a CLPR for which the cache load is low.

      - Add cache memory to increase the cache capacity of the storage system.

- The cache pending rate might exceed 70% temporarily due to failures on the primary and secondary storage systems. To prevent the I/O performance of the GAD pair volumes from decreasing or the GAD pairs from being suspended by failures related to this situation, the write-pending rate should be below 35% under normal conditions.

Planning GAD pairs and pair volumes

This section describes planning for differential data management, calculating the maximum number of GAD pairs, and the requirements for primary and secondary volumes related to the GAD configuration.

Related topics

- Differential data on page 3-29
- Maximum number of GAD pairs on page 3-30
- S-VOL resource group and storage system: same serial number and model on page 3-32

Differential data

Differential data is managed by the bitmap in units of tracks. A track that receives a write command while the pair is split is managed as differential data in the bitmap. When the pair is resynchronized, the differential data is copied to the S-VOL in units of tracks.

VSP G1000: When a GAD pair contains a DP-VOL that is larger than 4,194,304 MB (8,589,934,592 blocks), the differential data is managed by the pool to which the GAD pair volume is related.

VSP G200, G400, G600, G800: When you use a DP-VOL under either of the following conditions, the differential data is managed by the pool to which the GAD pair volume is related:

- You create a GAD pair by using a DP-VOL that is larger than 4,194,304 MB (8,589,934,592 blocks).
- You create a GAD pair by using a DP-VOL with the Advanced System Settings No. 5 or No. 6 enabled in Device Manager - Storage Navigator.

  Advanced System Settings No. 5: Manage differential bitmaps in DP pool at pair create and resync operations for 4TB or less TC/UR/GAD pairs.

  Advanced System Settings No. 6: Manage differential bitmaps in DP pool at pair create operations for 4TB or less TC/UR/GAD pairs.
- You resynchronize a GAD pair by using a DP-VOL with the Advanced System Settings No. 5 or No. 6 enabled in Device Manager - Storage Navigator.

In this case, additional pool capacity (up to 4 pages, depending on the software configuration) is required for each increase of user data size by 4,123,168,604,160 bytes (~4 TB). For a GAD pair with a DP-VOL that is larger than 4,194,304 MB (8,589,934,592 blocks), data management might

fail due to insufficient pool capacity. If this occurs, all of the P-VOL data (all tracks) is copied to the S-VOL when the pair is resynchronized.

Note: VSP G200, G400, G600, G800: If you enable Advanced System Settings No. 5 or No. 6 in Device Manager - Storage Navigator, the differential data is managed by the pool to which the GAD pair volume is related even if the volume is smaller than 4,194,304 MB (8,589,934,592 blocks). This enables the total capacity of the pair volumes to be increased over 1.8 PB. If you want the differential data to be managed by the control memory again after you create pairs whose total capacity is larger than 1.8 PB, you must reduce the total capacity of the pairs that are smaller than 4,194,304 MB(8,589,934,592 blocks) to 1.8 PB by deleting some pairs. When the differential data is managed by the control memory, the total capacity of the TC, UR, and GAD pairs is 1.8 PB. For example, if the total capacity of the GAD pairs is already 1.8 PB, you cannot create any TC or UR pairs.

For instructions on releasing the differential data (pages) managed in a pool, see Releasing the differential data managed in a pool on page 4-49.

Maximum number of GAD pairs

The maximum number of GAD pairs per storage system is specified in Requirements and restrictions on page 2-2. The maximum number of pairs per storage system is subject to restrictions, such as the number of cylinders used in volumes or the number of bitmap areas used in volumes.

The maximum number of pairs is 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 (total number of DP-VOLs plus external volumes: 63,232 for VSP G1000).

In the calculation formulas below, "ceiling" is the function that rounds up the value inside the parentheses to the next integer. "Floor" is the function that rounds down the value inside the parentheses to the next integer.

Note: If the volume size is larger than 4,194,304 MB (8,589,934,592 blocks), bitmap area is not used. Therefore, the calculation for the bitmap areas is not necessary when creating GAD pairs with DP-VOLs that are larger than 4,194,304 MB (8,589,934,592 blocks).

Related topics

- Planning GAD pairs and pair volumes on page 3-29
- Calculating the number of cylinders on page 3-30
- Calculating the number of bitmap areas on page 3-31
- Calculating the number of available bitmap areas on page 3-31
- Calculating the maximum number of pairs on page 3-32

Calculating the number of cylinders

To calculate the number of cylinders, start by calculating the number of logical blocks, which indicates volume capacity measured in blocks.

```
number-of-logical-blocks = volume-capacity (in bytes) / 512
```

Then use the following formula to calculate the number of cylinders:

```
number-of-cylinders = ceiling(ceiling(number-of-logical-blocks / 512) / 15)
```

Related topics

- Maximum number of GAD pairs on page 3-30
- Calculating the number of bitmap areas on page 3-31

Calculating the number of bitmap areas

Calculate the number of bitmap areas using the number of cylinders.

```
number-of-bitmap-areas = ceiling((number-of-cylinders × 15) / 122,752)
```

122,752 is the differential quantity per bitmap area. The unit is bits.

Note: You must calculate the number of required bitmap areas for each volume. If you calculate the total number of cylinders in multiple volumes and then use this number to calculate the number of required bitmap areas, the calculation results might be incorrect.
The following are examples of correct and incorrect calculations, assuming that one volume has 10,017 cylinders and another volume has 32,760 cylinders.

- Correct:
  ```
  ceiling((10,017 × 15) / 122,752) = 2
  ceiling((32,760 × 15) / 122,752) = 5
  ```
  The calculation result is seven bitmap areas in total.
- Incorrect:
  ```
  10,017 + 32,760 = 42,777 cylinders
  ceiling((42,777 × 15) / 122,752) = 6
  ```
  The calculation result is six bitmap areas in total.

Related topics

- Maximum number of GAD pairs on page 3-30
- Calculating the number of cylinders on page 3-30
- Calculating the number of available bitmap areas on page 3-31

Calculating the number of available bitmap areas

The total number of bitmap areas available in the storage system is:

- VSP G200: 36,000
- VSP G400, VSP G600, VSP G800, VSP G1000: 65,536

The number of bitmap areas is shared by TrueCopy, TrueCopy for Mainframe, Universal Replicator, Hitachi Universal Replicator for Mainframe, and GAD. If you use these software products, subtract the number of bitmap areas required for these products from the total number of bitmap areas in the storage system, and then use the formula in the next section to calculate the maximum number of GAD pairs. For details about calculating the number of bitmap areas required for the other software products, see the appropriate user guide.

Calculating the maximum number of pairs

Use the following values to calculate the maximum number of pairs:

- The number of bitmap areas required for pair creation.
- The total number of available bitmap areas in the storage system, or the number of available bitmap areas calculated in Calculating the number of available bitmap areas on page 3-31.

  Calculate the maximum number of pairs using the following formula with the total number of bitmap areas in the storage system (or the number of available bitmap areas) and the number of required bitmap areas, as follows:

  ```
  maximum-number-of-pairs-that-can-be-created = floor(total-number-of-bitmap-areas-in-storage-system / number-of-required-bitmap-areas)
  ```

Related topics

- Maximum number of GAD pairs on page 3-30
- Calculating the number of available bitmap areas on page 3-31

S-VOL resource group and storage system: same serial number and model

You can create GAD pairs specifying a volume in a resource group that has the same serial number and model as the storage system for the S-VOL. In this case, you must specify a volume in the resource group (virtual storage machine) whose serial number and model are same as the secondary storage system for the P-VOL.

When you create GAD pairs, the virtual LDEV ID of the P-VOL is copied to the virtual LDEV ID of the S-VOL. In the following figure, the copied virtual LDEV ID of the P-VOL is equal to the original virtual LDEV ID of the S-VOL. The volume in a resource group that has the same serial number and the same model as the storage system and whose original LDEV ID is equal to the virtual LDEV ID will be treated as a normal volume but as a virtualized volume by the global storage virtualization function.

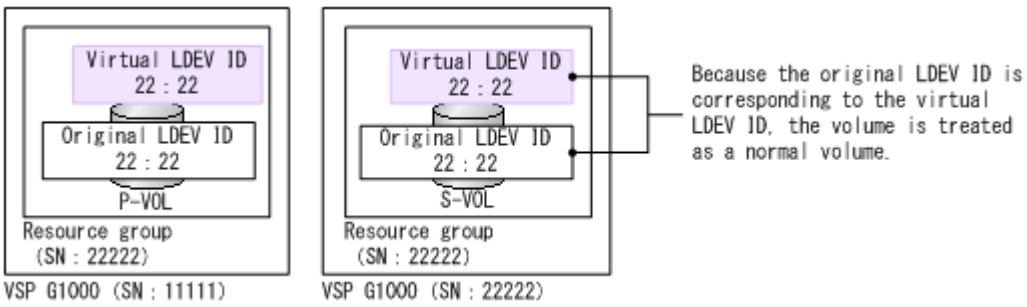

When virtual information is copied from the P-VOL to the S-VOL and a normal volume requirement is not met, as in the following examples, you cannot create GAD pairs:

- The copied virtual SSID of the P-VOL is not the same as the original SSID of the S-VOL.
- The copied virtual emulation type of the P-VOL is not the same as the original emulation type of the S-VOL.

  The virtual emulation type includes the virtual CVS attribute (-CVS). The storage system does not support LUSE, so LUSE configuration (*n) volumes are not supported as P-VOLs.

4

Configuration and pair management using CCI

This chapter describes and provides instructions for using CCI commands to configure a global-active device system and manage GAD pairs.

- Adding an alternate path to the S-VOL
- Setting a nonpreferred cross path (VSP G1000 only)
- Releasing the differential data managed in a pool

Global-active device system configuration

The following illustration shows a completed GAD system configuration, which includes the following key components:

- Host server connected to the primary and secondary storage systems with management software (alternate path or cluster or both, depending on the system configuration).
- Primary and secondary storage systems at the primary and secondary sites. The following system components are configured on the storage systems:
  - The volumes that will become the primary and secondary volumes of GAD pairs
  - The GAD feature installed on the primary and secondary storage systems
  - A virtual storage machine on the secondary storage system. Virtual storage machines can be configured on both the primary and secondary storage systems, but in this sample system configuration only the secondary storage system has a virtual storage machine.
  - A resource group on the secondary storage system. Resource groups can be configured on both the primary and secondary storage systems, but in this sample configuration only the secondary storage system has a resource group.
  - An external volume group on each storage system for the quorum disk
  - Remote paths between the storage systems
  - A CCI command device on each storage system
- External storage system with the quorum disk connected to both storage systems using Universal Volume Manager

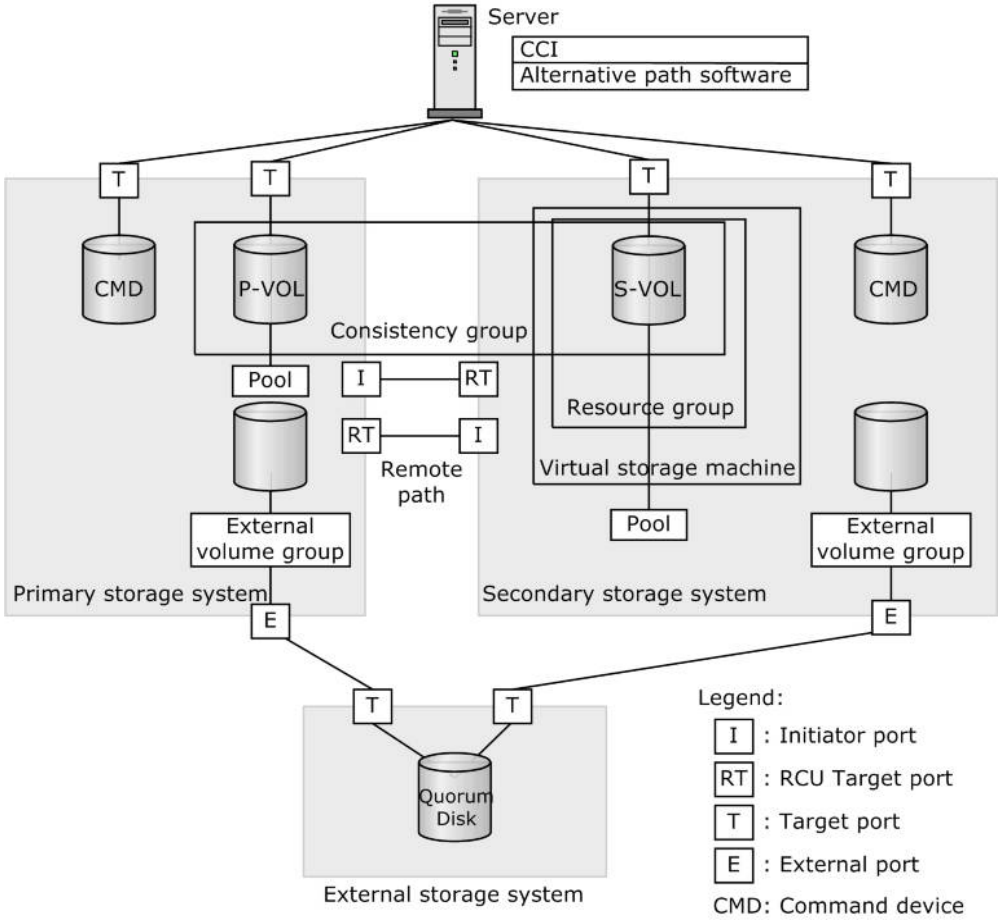

Note:

- Consistency groups are currently supported only on the VSP G1000.
- You do not need to set the port attributes (Initiator, RCU Target, Target) on the VSP G200, G400, G600, G800.

Primary storage system settings

The primary storage system components used in the procedures and examples in this chapter have the following settings.

Primary storage system

Model	Serial number
VSP G1000	11111*
* The serial number for the VSP G200, G400, G600, G800 is a six-digit number that begins with "4" (for example, 411111).	

Primary volume

Actual LDEV ID	Capacity	Port attribute	Port name	LU number
22:22	1,024,000 blocks	Target	CL1-A	0

Ports for remote connections

Port name	Port attribute
CL3-A	Initiator
CL4-A	RCU Target

External volume for the quorum disk

Actual LDEV ID	Port attribute	Port name	External volume group number	Path group ID	LU number	Quorum disk ID
99:99	External	CL5-A	1-1	1	0	0

Pair

Pair type	Pair group name	Device name (LDEV ID)	Mirror ID	CTG ID
GAD	oraHA	dev1(22:22)	0	2

Secondary storage system settings

The secondary storage system components used in the procedures and examples in this chapter have the following settings.

Secondary storage system

Model	Serial number
VSP G1000	22222*
* The serial number for the VSP G200, G400, G600, G800 is a six-digit number that begins with "4" (for example, 422222).	

Secondary volume

Actual LDEV ID	Capacity	Port attribute	Port name	LU number
44:44	1,024,000 blocks	Target	CL1-C	0

Ports for remote connections

Port name	Port attribute
CL3-C	RCU Target
CL4-C	Initiator

External volume for the quorum disk

Actual LDEV ID	Port attribute	Port name	External volume group number	Path group ID	LU number	Quorum disk ID
88:88	External	CL5-C	1-2	1	0	0

Resource group

Resource group name	Virtual storage machine	
	Model	Serial number
HAGroup1	VSP G1000	11111*
* The serial number for the VSP G200, G400, G600, G800 is a six-digit number that begins with "4" (for example, 411111).		

Host group

Host group ID	Host group name	Usage
CL1-C-1	1C-G00	For the S-VOL

Pool

Pool ID	Pool name	Pool volume
0	HA_POOL	77:77

Pair

Pair type	Pair group name	Device name (LDEV ID)	Mirror ID	CTG ID
GAD	oraHA	dev1(44:44)	0	2

CCI server configuration

The CCI server configuration used in the procedures and examples in this chapter has the following settings.

CCI instances and configuration definition files

Instance number	Configuration definition files	Usage
0	horcm0.conf	For the operation of the primary storage system
1	horcm1.conf	For the operation of the secondary storage system
100	horcm100.conf	For the operation of the primary storage system from the viewpoint of the virtual storage machine (serial number: 11111)
101	horcm101.conf	For the operation of the secondary storage system from the viewpoint of the virtual storage machine (serial number: 11111)

For operations involving virtual storage machines, the parameters specified in the `raidcom` command and the objects displayed by the `raidcom` command are based on the virtual ID. In the procedures and examples in this chapter, there is no virtual storage machine defined in the primary storage system, but you can operate the primary storage system as if there is a virtual storage machine with the same serial number and model as the primary storage system.

External storage system settings

The external storage system used in the procedures and examples in this chapter has the following settings.

External storage system

Model	Serial number
VSP G1000	33333*

Model	Serial number
* The serial number for the VSP G200, G400, G600, G800 is a six-digit number that begins with "4" (for example, 433333).	

WWN

Storage system at destination	WWN
Storage system at the primary site	50060e8007823520
Storage system at the secondary site	50060e8007823521

Workflow for creating a GAD environment

1. Initial state on page 4-8
2. Adding the external storage system for the quorum disk on page 4-9
3. Verifying the physical data paths on page 4-9
4. Creating the command devices on page 4-11
5. Creating the configuration definition files on page 4-11
6. Starting CCI on page 4-13
7. Connecting the primary and secondary storage systems on page 4-13
8. Creating the quorum disk on page 4-17
9. Setting up the secondary storage system on page 4-26
10. Updating the CCI configuration definition files on page 4-41
11. Creating the GAD pair on page 4-43
12. Adding an alternate path to the S-VOL on page 4-48

Note: This chapter provides CCI examples and instructions using the in-band method of issuing CCI commands. You can also issue GAD commands using the out-of-band method. For details about the in-band and out-of-band methods, see the *Command Control Interface User and Reference Guide*.

Initial state

The initial state before GAD configuration consists of one host, one primary storage system, and one secondary storage system.

- Primary and secondary storage systems:
  - Additional shared memory/control memory for GAD is installed in both storage systems.
  - The GAD feature is installed on both storage systems.
  - Resource group 0 exists by default in both storage systems.

- Dynamic Provisioning virtual volumes (DP-VOLs) are configured and have LU paths defined. These volumes will become the primary volumes of GAD pairs.

- Host:
  - The required management software for your configuration, alternate path and /or cluster software, is installed.
  - The CCI software is installed.

Note: The creation of GAD pairs is not affected by the presence or absence of server I/O to the DP-VOLs.

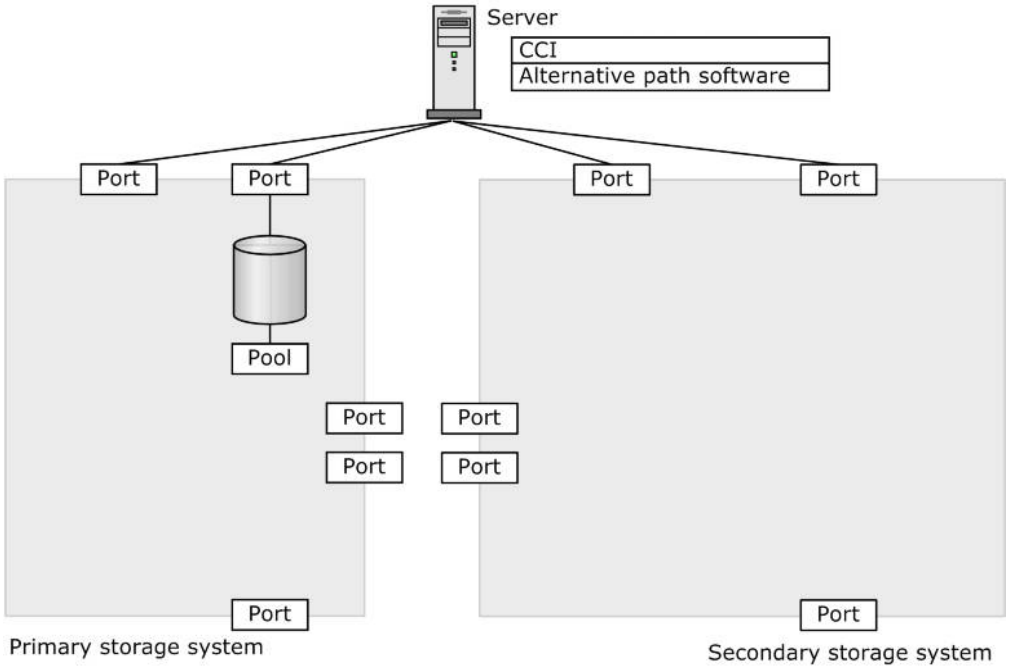

Adding the external storage system for the quorum disk

Install an external storage system for the quorum disk. The storage system must be supported by Universal Volume Manager for connection as external storage.

Related topics

- Requirements and restrictions on page 2-2
- Planning the quorum disk on page 3-20

Verifying the physical data paths

Make sure that the following physical data paths are connected and configured:

- From the primary storage system to the secondary storage system: two or more paths
- From the secondary storage system to the primary storage system: two or more paths
- From the primary storage system to the external storage system: two or more paths
- From the secondary storage system to the external storage system: two or more paths
- From the host to the primary storage system: two or more paths
- From the host to the secondary storage system: two or more paths

The following figure shows the physical data paths (redundant paths not shown). Although only one path is required for each location, it is important that you connect the storage systems using at least two physical paths. If you connect nodes using only one physical path, an unexpected failover might occur in the server, or the GAD pairs might be suspended, even though only one path or hardware failure has occurred. When maintenance is performed on the physical paths between storage systems, the GAD pairs must be suspended.

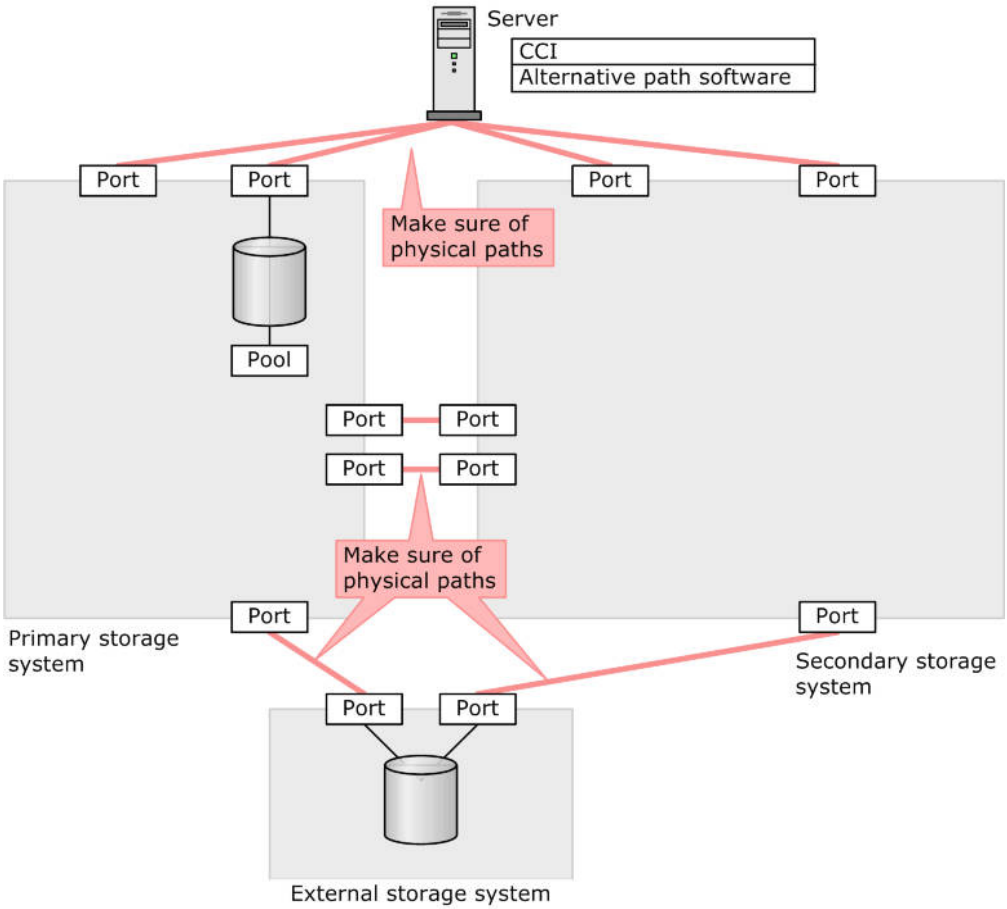

Creating the command devices

A command device (CMD) is required on each storage system for communication between CCI and the storage system. The command device must be created in resource group 0 in the primary storage system and in the secondary storage system. After the command devices have been created, host recognition must be set to the command devices.

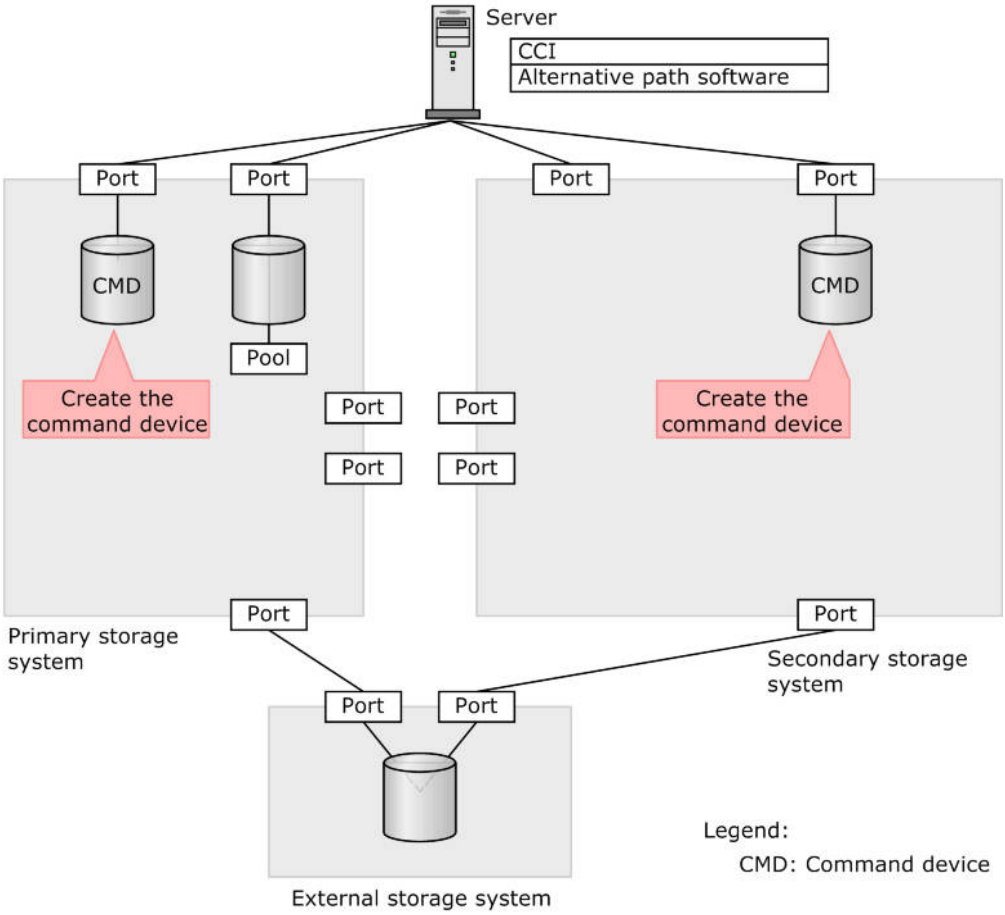

1. Using Hitachi Command Suite, allocate a command device in Resource Group 0 in both storage systems and enable user authentication. For details about creating a command device, see the *Provisioning Guide* for the storage system.
2. If necessary, change the topology and fabric settings for the ports defined to the command devices.
3. Define the volume to the port connected to the host.

Creating the configuration definition files

You must create four HORCM configuration definition files on the host for your GAD environment:

- One that describes the primary storage system and P-VOLs
- One that describes the secondary storage system and S-VOLs

- One for operating the virtual storage machine (SN: 11111) on the primary storage system
- One for operating the virtual storage machine (SN: 11111) on the secondary storage system

The configuration definition files for the examples in this chapter are shown below. For details about creating CCI configuration definition files, see the *Command Control Interface Installation and Configuration Guide*. The examples below show files on a Windows host.

Note: When specifying the serial number for VSP G1000 using CCI, add a "3" at the beginning of the serial number. For example, for serial number 11111, enter 311111.

HORCM file for the primary storage system: horcm0.conf

```
HORCM_MON
#ip_address        service         poll(10ms)     timeout(10ms)
localhost          31000           -1             3000

HORCM_CMD
\\.\PhysicalDrive0
```

HORCM file for the secondary storage system: horcm1.conf

```
HORCM_MON
#ip_address        service         poll(10ms)     timeout(10ms)
localhost          31001           -1             3000

HORCM_CMD
\\.\PhysicalDrive1
```

HORCM file for the virtual storage machine (SN: 11111) on the primary storage system: horcm100.conf

```
HORCM_MON
#ip_address        service        poll(10ms)      timeout(10ms)
localhost         31100          -1              3000
HORCM_CMD
\\.\PhysicalDrive0

HORCM_VCMD
# redefine Virtual DKC Serial# as unitIDs
311111
```

HORCM file for the virtual storage machine (SN: 11111) on the secondary storage system: horcm101.conf

```
HORCM_MON
#ip_address service poll(10ms) timeout(10ms)
localhost 31101 -1 3000
HORCM_CMD
\\.\PhysicalDrive1

HORCM_VCMD
# redefine Virtual DKC Serial# as unitIDs
311111
```

Starting CCI

After creating the CCI configuration definition files, you can start the CCI software. Because you are not yet operating the virtual storage machine, you only need to start instances (0 and 1). You do not yet need to start instances (100 and 101) for the virtual storage machine.

Procedure (Windows shown)

1. Start CCI instances 0 and 1.

```
horcmstart 0 1
starting HORCM inst 0
HORCM inst 0 starts successfully.
starting HORCM inst 1
HORCM inst 1 starts successfully.
```

2. Enter the user name and password, and perform user authentication.

```
raidcom -login <username> <password> -IH0
raidcom -login <username> <password> -IH1
```

The `-IH` option in this example is used for each command to specify an instance. You can also perform the operation using a shell for each instance. To start the shell, specify an instance number to the environment variable HORCMINST, and then execute the command.

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.

- Setting the port attributes (VSP G1000 only) on page 4-13
- Adding remote connections on page 4-15

Setting the port attributes (VSP G1000 only)

The Initiator and RCU Target port attributes must be set on the primary and secondary storage system ports for GAD command and data transfer. Initiator ports and RCU Target ports are required on both the primary and secondary storage systems.

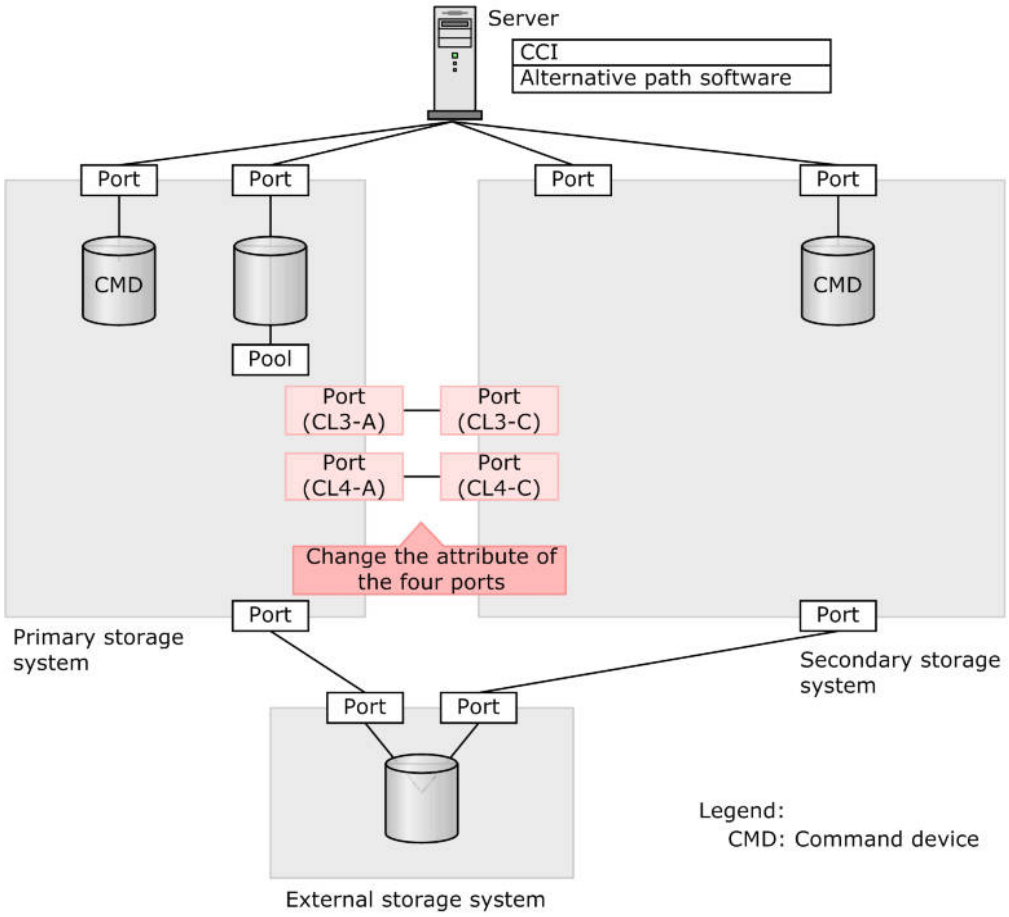

Procedure

1. Change the attribute of port (CL3-A) on the primary storage system to Initiator.
   ```
   raidcom modify port -port CL3-A -port_attribute MCU -IH0
   ```
2. Change the attribute of port (CL4-A) on the primary storage system to RCU Target.
   ```
   raidcom modify port -port CL4-A -port_attribute RCU -IH0
   ```
3. Change the attribute of port (CL3-C) on the secondary storage system to RCU Target.
   ```
   raidcom modify port -port CL3-C -port_attribute RCU -IH1
   ```
4. Change the attribute of port (CL4-C) on the secondary storage system to Initiator.
   ```
   raidcom modify port -port CL4-C -port_attribute MCU -IH1
   ```

Use the same procedure to change the port attributes for the alternate paths. The alternate paths are not shown in the illustration.

Check command and output examples

1. Display the port information for the primary storage system.
   ```
   raidcom get port -IH0
   PORT   TYPE  ATTR  SPD  LPID  FAB  CONN  SSW   SL   Serial#
   ```

```
WWN                PHY_PORT
(snip)
CL3-A  FIBRE MCU   AUT    E8  N    FCAL  N      0    311111
50060e80072b6720   -
(snip)
CL4-A  FIBRE RCU   AUT    97  N    FCAL  N      0    311111
50060e80072b6730   -
(snip)
```

2. Display the port information for the secondary storage system. Confirm that the port attributes have been changed as intended.

```
raidcom get port -IH1
PORT   TYPE  ATTR  SPD  LPID  FAB  CONN  SSW   SL   Serial#
WWN                PHY_PORT
(snip)
CL3-C  FIBRE RCU   AUT    D6  N    FCAL  N      0    322222
50060e800756ce22   -
(snip)
CL4-C  FIBRE MCU   AUT    7C  N    FCAL  N      0    322222
50060e800756ce32   -
(snip)
```

Adding remote connections

Add bidirectional remote connections between the primary and secondary storage systems. Specify the same path group ID to the bidirectional remote connections.

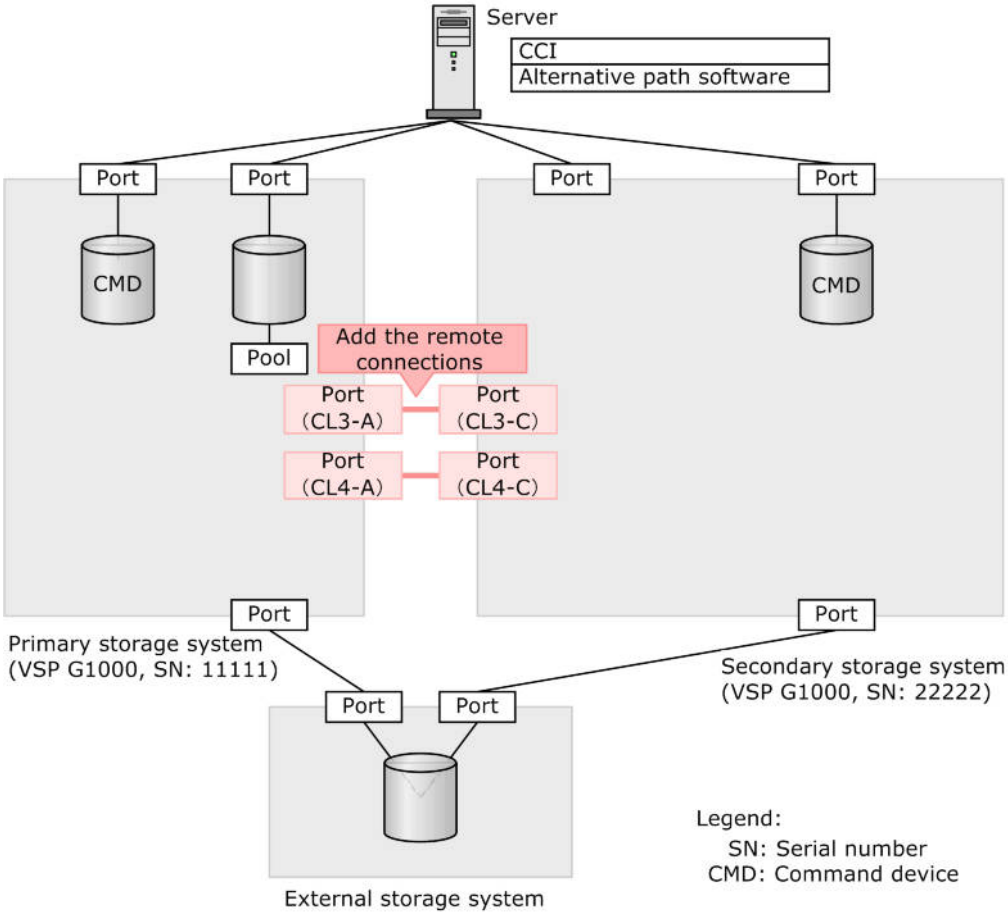

Note:

- When specifying the serial number for VSP G1000 using CCI, add a "3" at the beginning of the serial number. For example, for serial number 11111, enter 311111.
- To specify the VSP G1000, enter `R800`.
- The VSP G1000 is displayed as `R8` in command output.
- To specify the VSP G200, G400, G600, G800, enter `M800`.
- The VSP G200, G400, G600, G800 is displayed as `M8` in command output.

Procedure

1. Add a remote connection with path group ID 0 from primary storage system port (CL3-A) to secondary storage system port (CL3-C).

```
raidcom add rcu -cu_free 322222 R800 0 -mcu_port CL3-A -rcu_port
CL3-C -IH0
```

2. Confirm that asynchronous command processing has completed.

```
raidcom get command_status -IH0
HANDLE   SSB1    SSB2    ERR_CNT        Serial#     Description
00c3        -       -          0         311111     -
```

3. Add a remote connection with path group ID 0 from secondary storage system port (CL4-C) to primary storage system port (CL4-A).

```
raidcom add rcu -cu_free 311111 R800 0 -mcu_port CL4-C -rcu_port
CL4-A -IH1
```

4. Confirm that asynchronous command processing has completed.

```
raidcom get command_status -IH1
HANDLE   SSB1    SSB2    ERR_CNT        Serial#     Description
00c3        -       -          0         322222     -
```

Add the alternate paths between the storage systems using the `raidcom add rcu_path` command. These alternate paths are not shown in the illustration.

Check command and output examples

1. On the primary storage system, display remote connection information.

```
raidcom get rcu -cu_free 322222 R800 0 -IH0
Serial# ID  PID  MCU RCU  M/R T  PNO  MPORT  RPORT  STS_CD
SSIDs ...
 322222 R8    0    -   -  RCU F    0  CL3-A  CL3-C  NML_01  -
```

2. On the secondary storage system, display the remote connection information, and confirm that the serial number, model, and port name of the storage system are correct and that the path status is normal.

```
raidcom get rcu -cu_free 311111 R800 0 -IH1
Serial# ID  PID  MCU RCU  M/R T  PNO  MPORT  RPORT  STS_CD
SSIDs ...
 311111 R8    0    -   -  RCU F    0  CL4-C  CL4-A  NML_01  -
```

Creating the quorum disk

When a failure occurs, the quorum disk is used by the primary and secondary storage systems to determine which pair volume contained the latest data when the failure occurred.

This section provides instructions for setting up the quorum disk. You will map the disk on the external storage system to the primary and secondary storage systems. Make sure the external volume is formatted before proceeding. You should be familiar with Universal Volume Manager to set up the quorum disk.

An external volume for a quorum disk must be mapped to one external volume group.

- Setting the port attributes for connecting the external storage system on page 4-17
- Creating external volume groups on page 4-19
- Creating external volumes on page 4-21
- Setting external volumes as quorum disks on page 4-24

Setting the port attributes for connecting the external storage system

This section provides instructions for setting the ports on the primary and secondary storage systems (VSP G1000 only) to the "External" attribute in preparation for connecting to the external storage system.

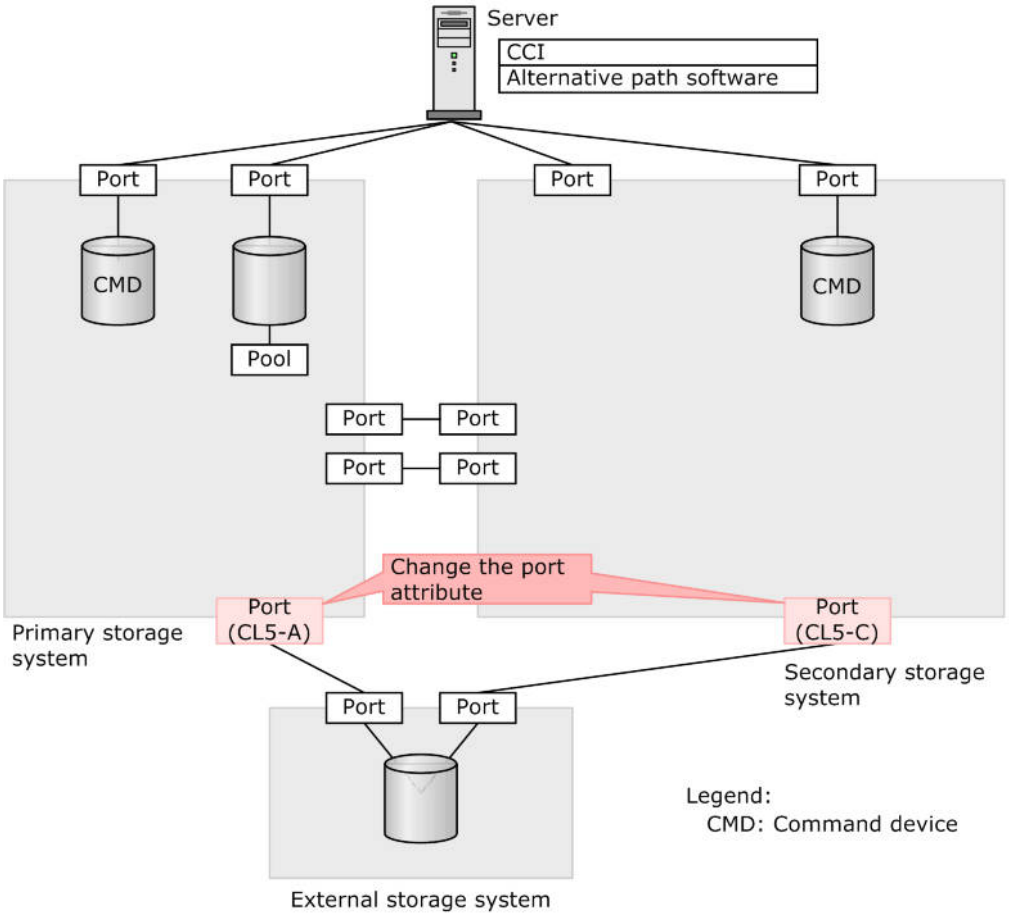

Procedure

1. Change the attribute of the port (CL5-A) on the primary storage system to External.

   ```
   raidcom modify port -port CL5-A -port_attribute ELUN -IH0
   ```

2. Change the attribute of the port (CL5-C) on the secondary storage system to External.

   ```
   raidcom modify port -port CL5-C -port_attribute ELUN -IH1
   ```

Check command and output examples

1. Display port information for the primary storage system.

   ```
   raidcom get port -IH0
   PORT   TYPE  ATTR   SPD  LPID  FAB  CONN  SSW   SL   Serial#
   WWN              PHY_PORT
   (snip)
   CL5-A  FIBRE ELUN   AUT    E4  N    FCAL  N      0    311111
   50060e80072b6740  -
   (snip)
   ```

2. Display the port information for the secondary storage system. Confirm that the port attributes have been changed as intended.

   ```
   raidcom get port -IH1
   PORT   TYPE  ATTR   SPD  LPID  FAB  CONN  SSW   SL   Serial#
   ```

```
WWN                PHY_PORT
(snip)
CL5-C  FIBRE ELUN   AUT    D5  N    FCAL  N      0    322222
50060e800756ce42  -
(snip)
```

Creating external volume groups

Create external volume groups for the quorum disk to map the disk on the external storage system to the primary and secondary storage systems.

- Verify that the volumes in the external storage system are formatted.
- Use the `raidcom discover lun -port` command to verify that the same E_VOL_ID_C value (volume identifier included in the SCSI Inquiry command of the external volume) is displayed for the primary and secondary storage systems.

For details about creating external volume groups, see the *Hitachi Universal Volume Manager User Guide*.

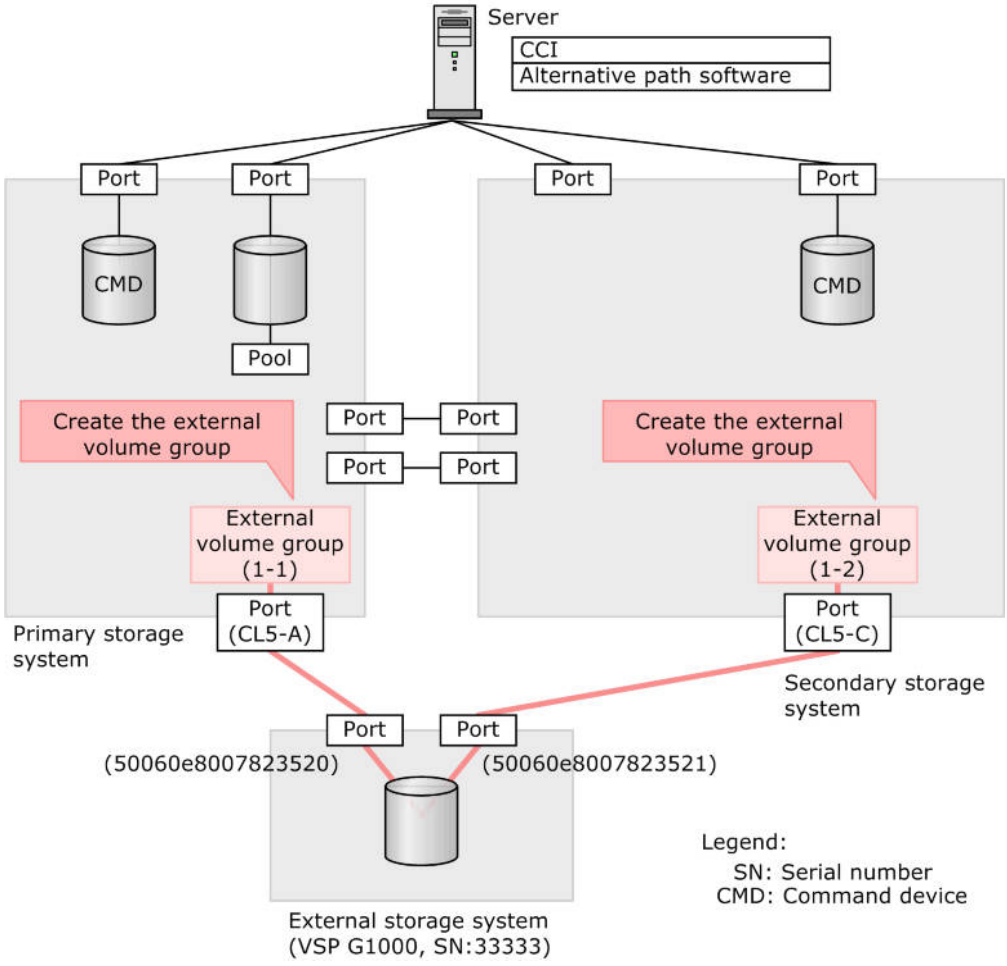

Procedure

1. Search for information about the external storage system port that is connected to port (CL5-A) on the primary storage system.

```
raidcom discover external_storage -port CL5-A -IH0
PORT   WWN               PM  USED  Serial#  VENDOR_ID PRODUCT_ID
CL5-A  50060e8007823520  M   NO    33333    HITACHI   VSP G1000
```

Note: The PRODUCT_ID for the VSP G200, G400, G600, G800 is VSP Gx00.

2. Display the LU that is defined to external storage system port (50060e8007823520) that is connected to primary storage system port (CL5-A). Check the LU number, and note the value shown in the E_VOL_ID_C field.

```
raidcom discover lun -port CL5-A -external_wwn 50060e8007823520 -
IH0
PORT   WWN              LUN  VOL_Cap(BLK)  PRODUCT_ID  E_VOL_ID_C
CL5-A  50060e8007823520   0      61440000  OPEN-V      HITACHI
500308235AAAA
```

3. Map the LU (0) that is defined to the external storage system port (50060e8007823520) that is connected to the primary storage system port (CL5-A). Specify 1 for the path group ID, and specify 1-1 for the external volume group number.

```
raidcom add external_grp -path_grp 1 -external_grp_id 1-1 -port
CL5-A -external_wwn 50060e8007823520 -lun_id 0 -IH0
```

4. Confirm that asynchronous command processing has completed.

```
raidcom get command_status -IH0
HANDLE   SSB1    SSB2    ERR_CNT        Serial#     Description
00c3        -       -          0         311111     -
```

5. Search for information about the external storage system port that is connected to port (CL5-C) on the secondary storage system.

```
raidcom discover external_storage -port CL5-C -IH1
PORT     WWN                PM  USED   Serial#  VENDOR_ID
PRODUCT_ID
CL5-C    50060e8007823521   M   NO     33333    HITACHI   VSP
G1000
```

Note: The PRODUCT_ID for the VSP G200, G400, G600, G800 is VSP Gx00.

6. Display the LU that is defined to external storage system port (50060e8007823521) that is connected to secondary storage system port (CL5-C). Check the LU number, and confirm that the E_VOL_ID_C field displays the same value as in step 2.

```
raidcom discover lun -port CL5-C -external_wwn 50060e8007823521 -
IH1
PORT     WWN              LUN  VOL_Cap(BLK)  PRODUCT_ID
E_VOL_ID_C
CL5-C    50060e8007823521   0      61440000  OPEN-V      HITACHI
500308235AAAA
```

7. Map the LU (0) that is defined to external storage system port (50060e8007823521) that is connected to secondary storage system port

(CL5-C). Specify 1 for the path group ID, and specify 1-2 for the external volume group number.

```
raidcom add external_grp -path_grp 1 -external_grp_id 1-2 -port
CL5-C -external_wwn 50060e8007823521 -lun_id 0 -IH1
```

8. Confirm that asynchronous command processing has completed.

```
raidcom get command_status -IH1
HANDLE   SSB1    SSB2    ERR_CNT        Serial#     Description
00c3        -       -          0         322222     -
```

Check command and output examples

1. On the primary storage system, display information about the external path to the volume in the external storage system.

```
raidcom get path -path_grp 1 -IH0
PHG GROUP STS CM IF MP# PORT   WWN                 PR LUN PHS
Serial# PRODUCT_ID LB PM
  1 1-1   NML E  D    0 CL5-A  50060e8007823520     1   0 NML
33333 VSP G1000  N  M
```

Note: The `PRODUCT_ID` for the VSP G200, G400, G600, G800 is `VSP Gx00`.

2. On the secondary storage system, display information about the external path to the volume in the external storage system. Confirm that the external storage system information is correct, including serial number, model, and WWN, and confirm that the path status and volume status are normal.

```
raidcom get path -path_grp 1 -IH1
PHG GROUP STS CM IF MP# PORT   WWN                 PR LUN PHS
Serial# PRODUCT_ID LB PM
  1 1-2   NML E  D    0 CL5-C  50060e8007823521     1   0 NML
33333 VSP G1000  N  M
```

Note: The `PRODUCT_ID` for the VSP G200, G400, G600, G800 is `VSP Gx00`.

Creating external volumes

Using capacity in the external storage system, you will create virtual external volumes on the primary and secondary storage systems that will be mapped to the quorum disk.

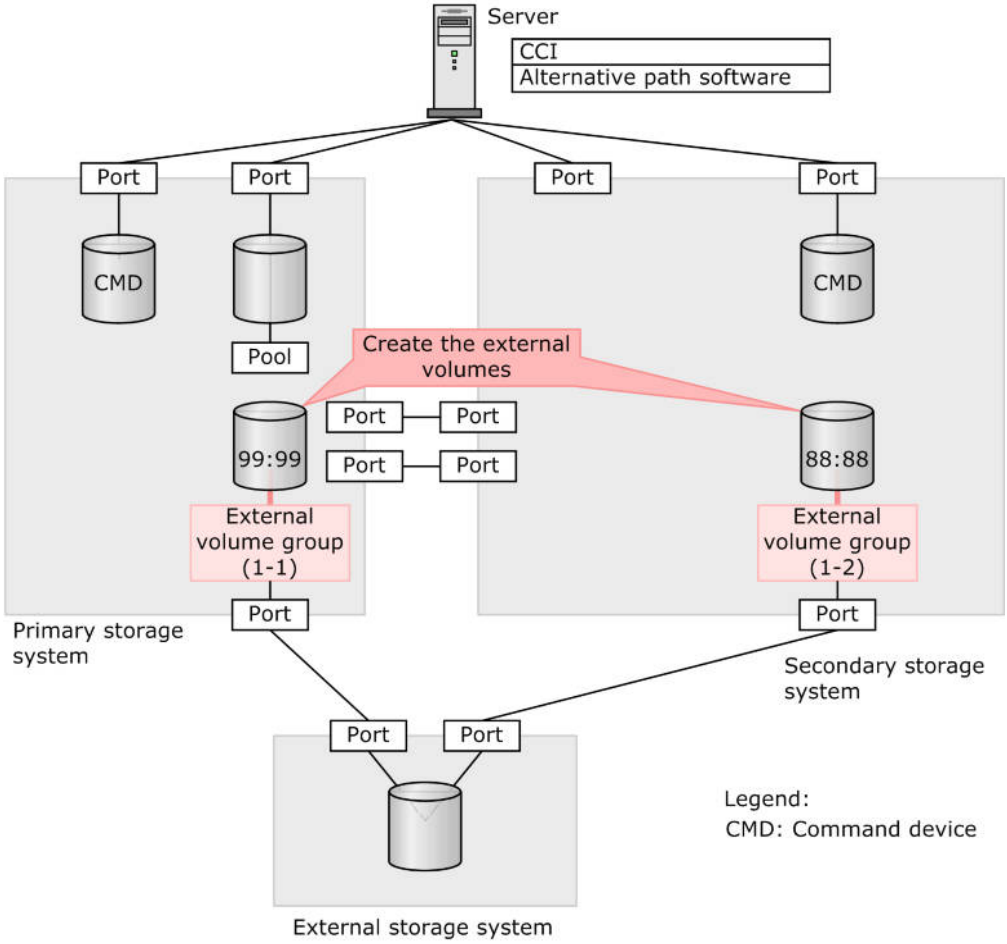

Procedure

1. Specify external volume group (1-1) assigned to the primary storage system to create an external volume whose LDEV ID is 0x9999. Allocate all capacity in the external volume group.

   ```
   raidcom add ldev -external_grp_id 1-1 -ldev_id 0x9999 -capacity
   all -IH0
   ```

2. Confirm that asynchronous command processing has completed.

   ```
   raidcom get command_status -IH0
   HANDLE   SSB1    SSB2    ERR_CNT        Serial#     Description
   00c3        -       -          0         311111     -
   ```

3. Specify external volume group (1-2) assigned to the secondary storage system to create an external volume whose LDEV ID is 0x8888. Allocate all free space in the external volume group.

   ```
   raidcom add ldev -external_grp_id 1-2 -ldev_id 0x8888 -capacity
   all -IH1
   ```

4. Confirm that asynchronous command processing has completed.

   ```
   raidcom get command_status -IH1
   HANDLE   SSB1    SSB2    ERR_CNT        Serial#     Description
   00c3        -       -          0         322222     -
   ```

Check command and output examples

1. Display information about the volume (LDEV ID: 0x9999).

```
raidcom get ldev -ldev_id 0x9999 -fx -IH0
Serial#  : 311111
LDEV : 9999
SL : 0
CL : 0
VOL_TYPE : OPEN-V-CVS
VOL_Capacity(BLK) : 61440000
NUM_PORT : 0
PORTs :
F_POOLID : NONE
VOL_ATTR : CVS : ELUN
E_VendorID : HITACHI
E_ProductID : OPEN-V
E_VOLID :
484954414348492035303330324236374141414100000000000000000000000000000000
E_VOLID_C : HITACHI 500308235AAAA................
NUM_E_PORT : 1
E_PORTs : CL5-A-0 0 50060e8007823520
LDEV_NAMING :
STS : NML
OPE_TYPE : NONE
OPE_RATE : 100
MP# : 0
SSID : 0007
RSGID : 0
```

2. Display the information about the volume (LDEV ID: 0x8888). Confirm that the information about the external volume is correct.

```
raidcom get ldev -ldev_id 0x8888 -fx -IH1
Serial#  : 322222
LDEV : 8888
SL : 0
CL : 0
VOL_TYPE : OPEN-V-CVS
VOL_Capacity(BLK) : 61440000
NUM_PORT : 0
PORTs :
F_POOLID : NONE
VOL_ATTR : CVS : ELUN
E_VendorID : HITACHI
E_ProductID : OPEN-V
E_VOLID :
484954414348492035303330324236374141414100000000000000000000000000000000
E_VOLID_C : HITACHI 500308235AAAA................
NUM_E_PORT : 1
E_PORTs : CL5-C-0 0 50060e8007823521
LDEV_NAMING :
STS : NML
OPE_TYPE : NONE
OPE_RATE : 100
MP# : 0
SSID : 0005
RSGID : 0
```

Setting external volumes as quorum disks

This section provides instructions for setting the virtualized external volumes in the primary and secondary storage systems as quorum disks. The same quorum disk ID must be set to the primary and secondary storage systems.

The serial number and model of the paired storage system is specified for the `-quorum_enable` option of the `raidcom modify ldev` command.

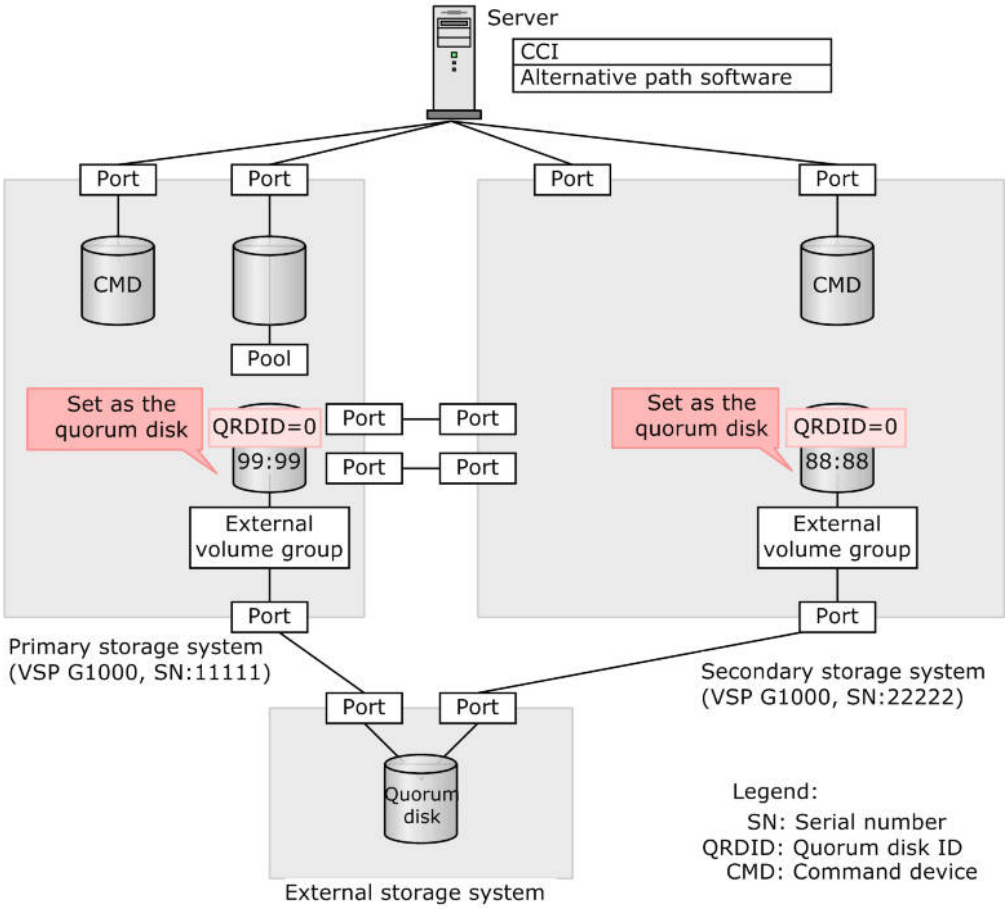

Note:

- When specifying the serial number for VSP G1000 using CCI, add a "3" at the beginning of the serial number. For example, for serial number 11111, enter 311111.
- To specify the VSP G1000, enter `R800`.
- The VSP G1000 is displayed as `R8` in command output.
- To specify the VSP G200, G400, G600, G800, enter `M800`.
- The VSP G200, G400, G600, G800 is displayed as `M8` in command output.

Procedure

1. Specify 0 to the quorum disk ID to set the volume (0x9999) in the primary storage system (serial number 22222, entered as 322222) as a quorum disk. Specify the storage system's serial number and model.

```
raidcom modify ldev -ldev_id 0x9999 -quorum_enable 322222 R800 -
quorum_id 0 -IH0
```

2. Confirm that asynchronous command processing has completed.

```
raidcom get command_status -IH0
HANDLE   SSB1    SSB2    ERR_CNT        Serial#     Description
00c3        -       -          0         311111     -
```

3. Specify 0 to the quorum disk ID to set the volume (0x8888) in the secondary storage system (serial number 11111, entered as 311111) as a quorum disk. Specify the storage system's serial number and model.

```
raidcom modify ldev -ldev_id 0x8888 -quorum_enable 311111 R800 -
quorum_id 0 -IH1
```

4. Confirm that asynchronous command processing has completed.

```
raidcom get command_status -IH1
HANDLE   SSB1    SSB2    ERR_CNT        Serial#     Description
00c3        -       -          0         322222     -
```

Check command and output examples

1. Display the information about the volume (LDEV ID: 0x9999).

```
raidcom get ldev -ldev_id 0x9999 -fx -IH0
Serial#  : 311111
LDEV : 9999
SL : 0
CL : 0
VOL_TYPE : OPEN-V-CVS
VOL_Capacity(BLK) : 61440000
NUM_PORT : 0
PORTs :
F_POOLID : NONE
VOL_ATTR : CVS : ELUN : QRD
E_VendorID : HITACHI
E_ProductID : OPEN-V
E_VOLID :
484954414348492035303330324236374141414100000000000000000000000000
0000000
E_VOLID_C : HITACHI 500308235AAAA.................
NUM_E_PORT : 1
E_PORTs : CL5-A-0 0 50060e80072b6750
LDEV_NAMING :
STS : NML
OPE_TYPE : NONE
OPE_RATE : 100
MP# : 0
SSID : 0007
QRDID : 0
QRP_Serial# : 322222
QRP_ID : R8
RSGID : 0
```

2. Display the information about volume (LDEV ID: 0x8888). Confirm that the following values are correct:
   - QRDID (quorum disk ID)
   - QRP_Serial# (serial number of the storage system that forms a GAD pair)

- QRP_ID (model of the storage system that forms a GAD pair)

```
raidcom get ldev -ldev_id 0x8888 -fx -IH1
Serial#  : 322222
LDEV : 8888
SL : 0
CL : 0
VOL_TYPE : OPEN-V-CVS
VOL_Capacity(BLK) : 61440000
NUM_PORT : 0
PORTs :
F_POOLID : NONE
VOL_ATTR : CVS : ELUN : QRD
E_VendorID : HITACHI
E_ProductID : OPEN-V
E_VOLID :
484954414348492035303330324236374141414100000000000000000000000000
0000000
E_VOLID_C : HITACHI 500308235AAAA.................
NUM_E_PORT : 1
E_PORTs : CL5-C-0 0 50060e80072b6760
LDEV_NAMING :
STS : NML
OPE_TYPE : NONE
OPE_RATE : 100
MP# : 0
SSID : 0005
QRDID : 0
QRP_Serial# : 311111
QRP_ID : R8
RSGID : 0
```

Setting up the secondary storage system

This section provides instructions for creating a virtual storage machine (VSM) in the secondary storage system and configuring it for GAD pair operations.

To create a virtual storage machine, you add resources such as host group IDs and LDEV IDs to a resource group that is created for the virtual storage machine. You can also reserve the host group and the volume ID by only adding them to the resource group. Create a host group and a volume (actual volume) by specifying the reserved IDs so that the host group and the volume can be used on a GAD pair.

The following procedures describe how to create a GAD environment. If appropriate, you can use existing storage system resources, for example, Dynamic Provisioning pools and DP-VOLs that have already been created.

Related topics

- Creating a resource group (P-VOL not registered to a VSM) on page 4-27
- Creating a resource group (P-VOL registered to a VSM) on page 4-28
- Reserving a host group ID on page 4-30
- Deleting the virtual LDEV ID of the S-VOL on page 4-31

- Reserving an LDEV ID for the S-VOL on page 4-32
- Setting the GAD reserve attribute on the S-VOL on page 4-33
- Creating a host group on page 4-35
- Creating a pool on page 4-37
- Creating the S-VOL on page 4-39
- Adding an LU path to the S-VOL on page 4-40

Creating a resource group (P-VOL not registered to a VSM)

When GAD setup is complete, the host sees the P-VOL and S-VOL of each pair as a single volume in a single storage system. Resource groups are created in the secondary storage system using the primary storage system's serial number and model as virtual information, so that the P-VOL and S-VOL of each pair share the same virtual storage machine information.

A virtual storage machine is composed of multiple resource groups that have the same virtual information. When you create a resource group and specify the virtual serial number and model, the resource group is registered in the virtual storage machine. If the virtual storage machine does not already exist in the storage system, it is created automatically when the resource group is created. The following illustration shows the creation of a resource group when the P-VOL is not already registered to a virtual storage machine.

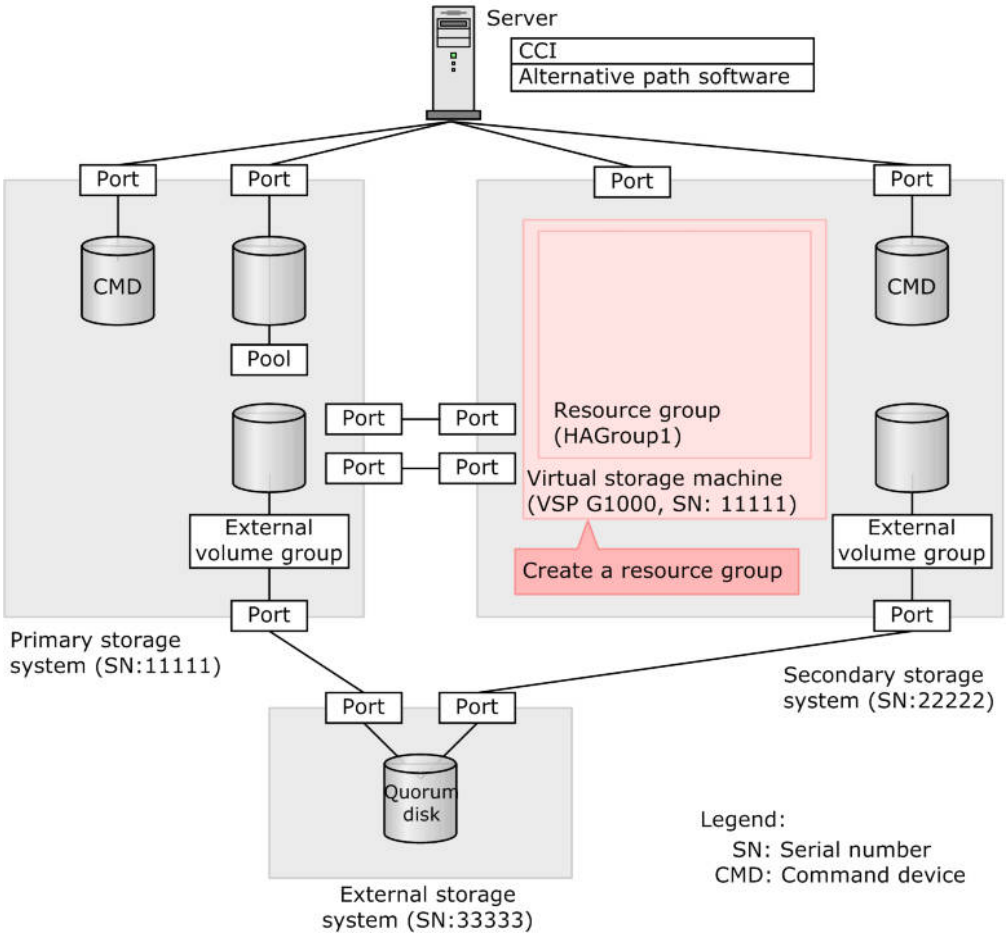

Note:

- When specifying the serial number for VSP G1000 using CCI, add a "3" at the beginning of the serial number. For example, for serial number 11111, enter 311111.
- To specify the VSP G1000, enter `R800`.
- The VSP G1000 is displayed as `R8` in command output.
- To specify the VSP G200, enter `M800S`.
- To specify the VSP G400 or VSP G600, enter `M800M`.
- To specify the VSP G800, enter `M800H`.
- The VSP G200, G400, G600, G800 is displayed as `M8` in command output.

Procedure

Specify the primary storage system's serial number and model for the virtual storage machine you are creating on the secondary storage system.

```
raidcom add resource -resource_name HAGroup1 -virtual_type 311111
R800 -IH1
```

Check command and output examples

Display the information about the resource groups of the secondary storage system. Information about all resource groups is displayed. Confirm the resource group name, resource group number, virtual serial number, and virtual model.

```
raidcom get resource -key opt -IH1
RS_GROUP            RGID   V_Serial#  V_ID   V_IF    Serial#
meta_resource          0      322222  R8     Y        322222
HAGroup1               1      311111  R8     Y        322222
```

Note: If you need to delete the virtual information set to the resource group, you must delete the resource group:

```
raidcom delete resource -resource_name HAGroup1 -IH1
```

Creating a resource group (P-VOL registered to a VSM)

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.

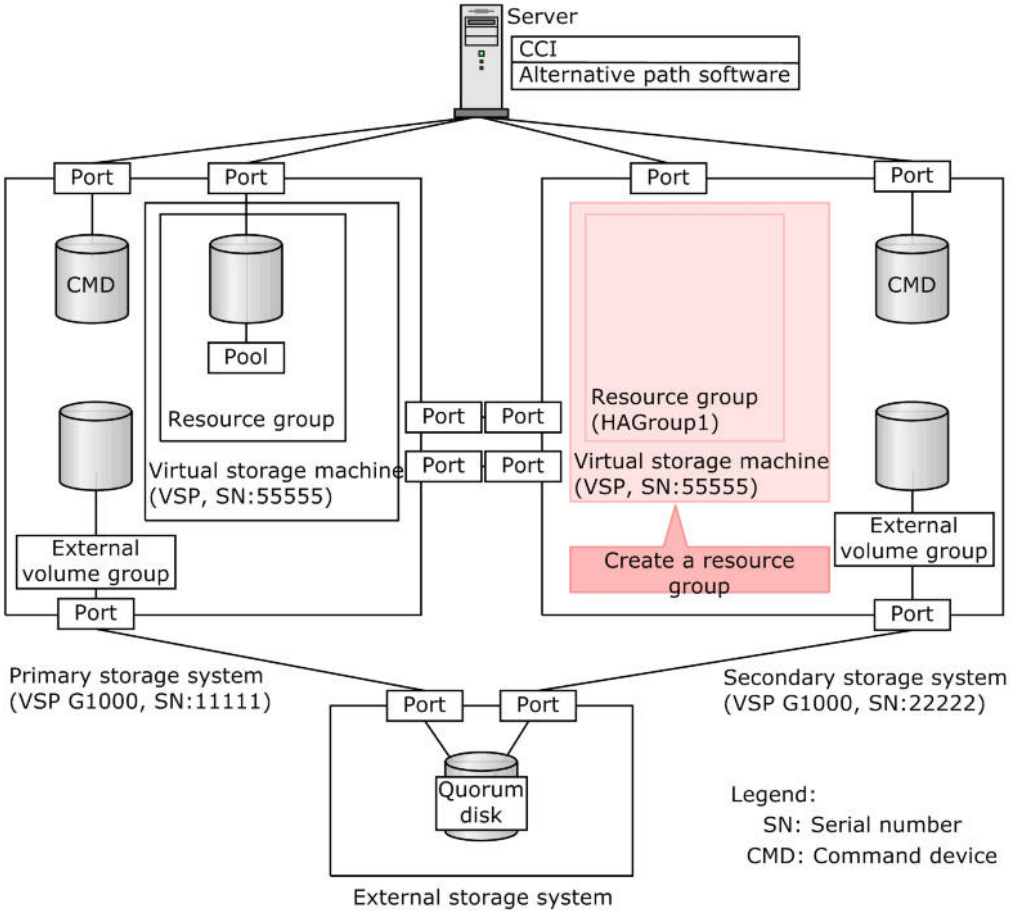

Note:

- When specifying the serial number for VSP G1000 using CCI, add a "3" at the beginning of the serial number. For example, for serial number 11111, enter 311111.
- To specify the VSP G1000 enter `R800`.
- To specify the VSP enter `R700`.
- To specify the VSP G200 enter `M800S`.
- To specify the VSP G400 or VSP G600 enter `M800M`.
- To specify the VSP G800 enter `M800H`.
- To specify the HUS VM enter `M700`.

Procedure

Specify the serial number and model of the virtual storage machine to which the P-VOL is registered for the virtual storage machine you are creating on the secondary storage system. Command for the sample configuration shown above:

```
raidcom add resource -resource_name HAGroup1 -virtual_type 55555
R700 -IH1
```

Reserving a host group ID

In the secondary storage system's resource group, you must reserve a host group ID to be used by the S-VOL.

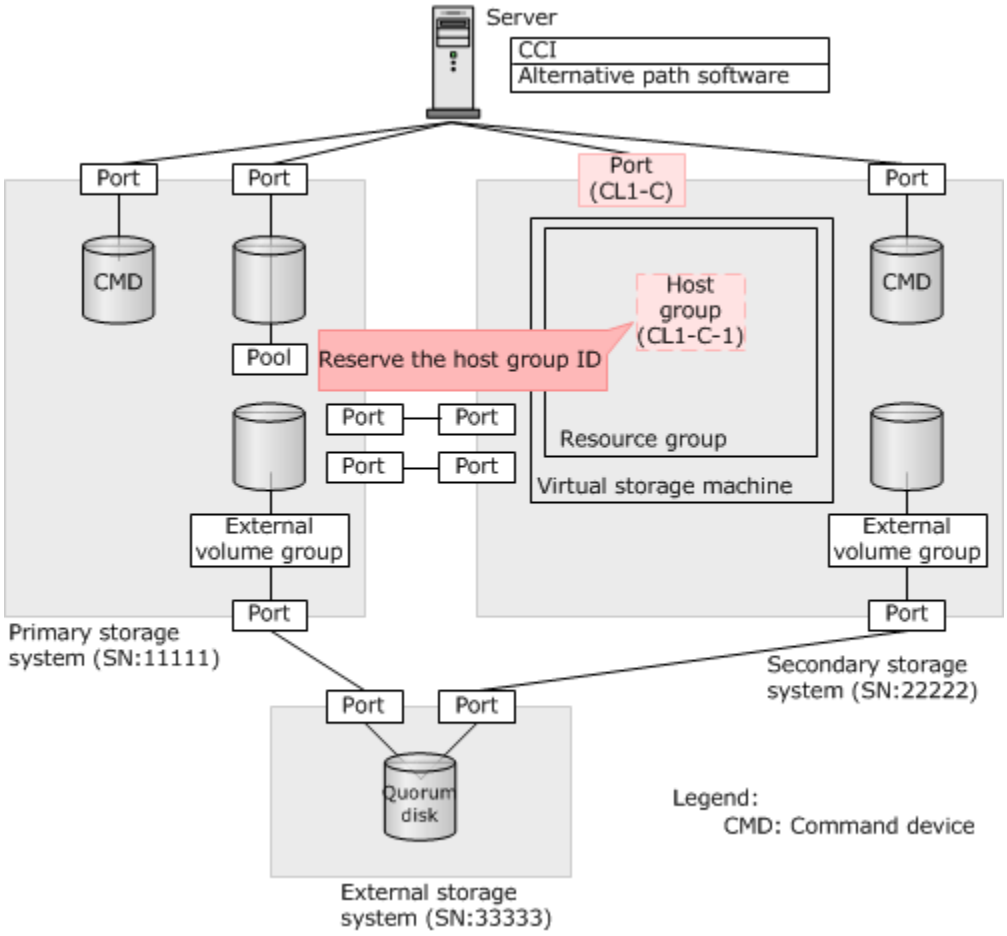

Procedure

Reserve a host group ID (CL1-C-1) in resource group (HAGroup1).

```
raidcom add resource -resource_name HAGroup1 -port CL1-C-1 -IH1
```

Check command and output examples

Display information about the host group that is set to port (CL1-C). Confirm that the port name, host group ID, and host group name are correct.

```
raidcom get host_grp -port CL1-C -resource 1 -IH1
PORT   GID  GROUP_NAME                        Serial# HMD
HMO_BITs
CL1-C    0  1C-G00                            322222 WIN
```

Note: If you reserve a host group for which no actual volume is defined in the resource group, specifying the `-key host_grp` option for the check command allows you to display the reserved host group. The following example shows the result of executing the check command.

```
raidcom get host_grp -port CL1-C -key host_grp -resource 1 -IH1
PORT   GID  GROUP_NAME                        Serial# HMD
HMO_BITs
CL1-C    0  1C-G00                            322222 WIN
CL1-C    1  HAVol                             322222 WIN
CL1-C    2  -                                 322222 -
CL1-C    3  -                                 322222 -
CL1-C    4  -                                 322222 -
CL1-C    5  -                                 322222 -
```

As shown in this example, the host groups with host group ID 0 to 5 are reserved in resource group 1. Actual volumes are defined for the host groups with host group ID 0 and 1. The host groups with host group ID 2 to 5 are reserved in the resource group, but actual volumes are not defined for them. The host groups with host group ID 6 to 254 are not displayed, because they are not reserved in resource group 1.

Deleting the virtual LDEV ID of the S-VOL

Delete temporarily the virtual LDEV ID of the volume to be added to the virtual storage machine.

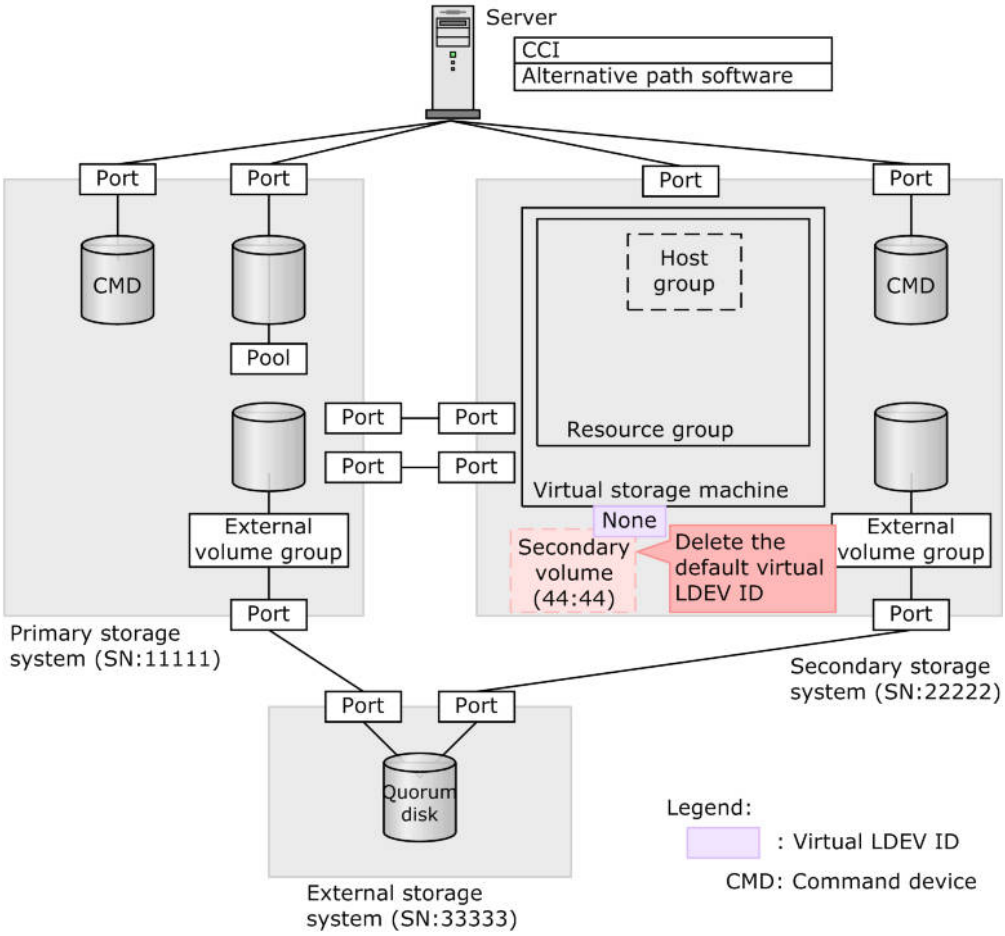

Procedure

Delete the virtual LDEV ID of the volume (0x4444).

```
raidcom unmap resource -ldev_id 0x4444 -virtual_ldev_id 0x4444 -IH1
```

Check command and output examples

Display information about the volume (LDEV ID: 0x4444). For the volume whose virtual LDEV ID was deleted, `fffe` is displayed for `VIR_LDEV` (virtual LDEV ID).

```
raidcom get ldev -ldev_id 0x4444 -fx -IH1
Serial#  : 322222
LDEV : 4444 VIR_LDEV : fffe
SL : -
CL : -
VOL_TYPE : NOT DEFINED
SSID : -
RSGID : 0
```

Note: If you need to reconfigure a deleted virtual LDEV ID, use the `raidcom map resource` command (example: `raidcom map resource -ldev_id 0x4444 -virtual_ldev_id 0x4444 -IH1`). The default virtual LDEV ID is the same as the actual LDEV ID. After reconfiguring the virtual LDEV ID, use the check command to confirm that the virtual LDEV ID is the same as the actual LDEV ID.

Reserving an LDEV ID for the S-VOL

In the newly created resource group, you will reserve an LDEV ID so that the volume is available to become the target volume of a global-active device pair.

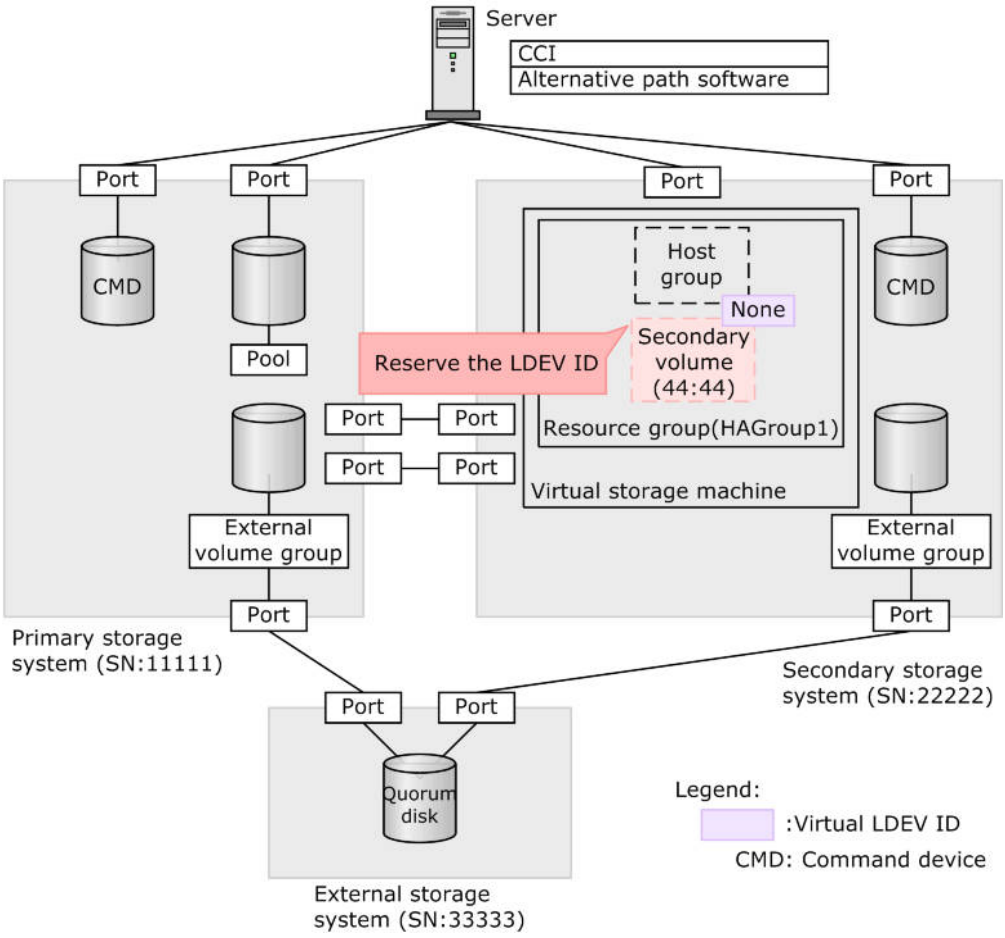

Procedure

Reserve the LDEV ID (0x4444) in the resource group (HAGroup1).

```
raidcom add resource -resource_name HAGroup1 -ldev_id 0x4444 -IH1
```

Check command and output examples

Display the information about volume (LDEV ID: 0x4444). Confirm that the number of the resource group in which the LDEV ID was reserved is displayed for `RSGID`.

```
raidcom get ldev -ldev_id 0x4444 -fx -IH1
Serial#  : 322222
LDEV : 4444 VIR_LDEV : fffe
SL : -
CL : -
VOL_TYPE : NOT DEFINED
SSID : -
RSGID : 1
```

Setting the GAD reserve attribute on the S-VOL

When you create a GAD pair, the P-VOL's LDEV ID is set as the virtual LDEV ID of the S-VOL. Before the pair can be created, the GAD reserve attribute

must be set on the volume that will become the S-VOL. Setting the reserve attribute deletes the virtual LDEV ID of the volume, so that the virtual LDEV ID of the P-VOL can be set to the volume when the GAD pair is created.

Note: Setting the GAD reserve attribute (deleting the virtual LDEV ID) makes the volume unusable by some software products (for example, ShadowImage). For details, see the interoperability requirements for GAD.

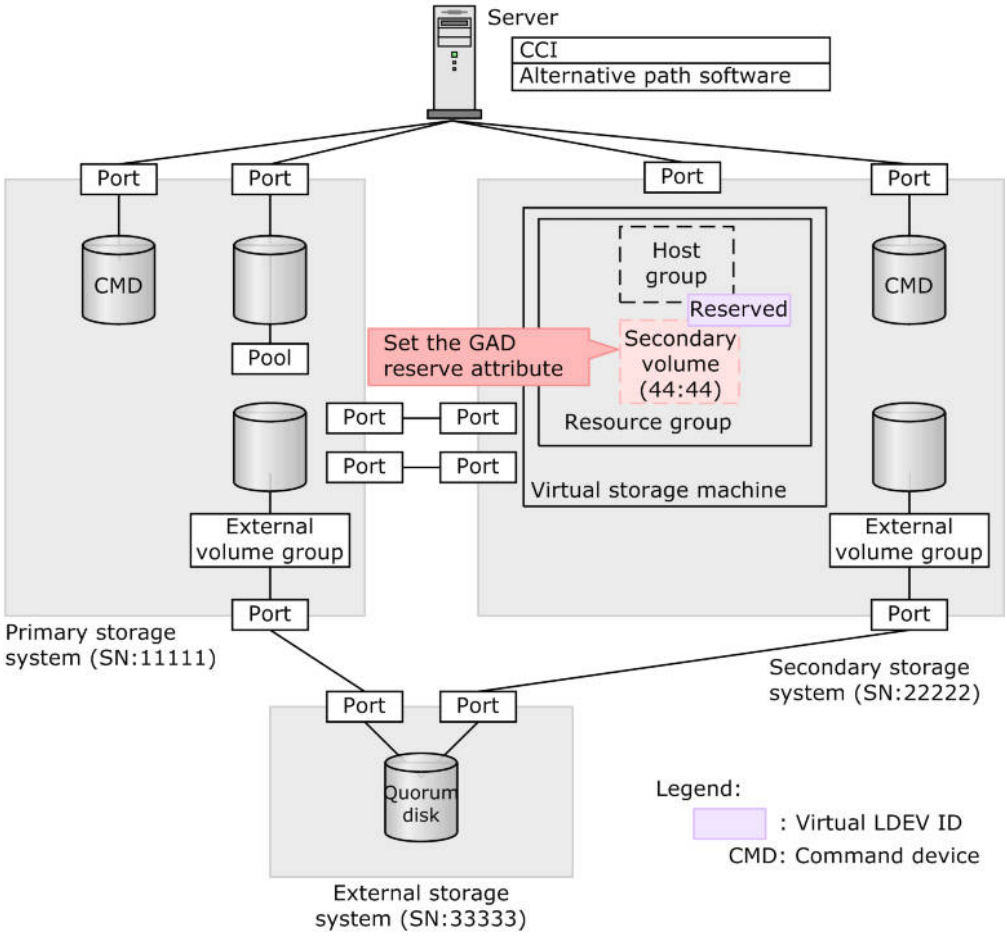

Procedure

Set the GAD reserve attribute to the LDEV ID (0x4444).

```
raidcom map resource -ldev_id 0x4444 -virtual_ldev_id reserve -IH1
```

Check command and output examples

Display the information about the volume (LDEV ID: 0x4444). For the LDEV ID to which the reserve attribute was set, `ffff` is displayed for `VIR_LDEV` (virtual LDEV ID).

```
raidcom get ldev -ldev_id 0x4444 -fx -IH1
Serial#  : 322222
LDEV : 4444 VIR_LDEV : ffff
SL : -
CL : -
VOL_TYPE : NOT DEFINED
```

```
SSID : -
RSGID : 1
```

Note: If you need to release the GAD reserve attribute, use the `raidcom unmap resource` command (example: `raidcom unmap resource -ldev_id 0x4444 -virtual_ldev_id reserve -IH1`). After releasing the GAD reserve attribute, use the check command to confirm that `fffe` is displayed for `VIR_LDEV` (virtual LDEV ID).

Creating a host group

Create a host group to be used by the S-VOL of the GAD pair. The requirement to create a host group depends on the host group ID used by the S-VOL:

- When the host group ID used by the S-VOL is 0

  The host group whose host group ID is 0 exists by default. Set a host mode to the host group (Windows host mode shown).

  Procedure

  ```
  raidcom modify host_grp -port CL1-C-0 -host_mode WIN -IH1
  ```

  Check command and output examples

  Display the information about the host group that is set for the port (CL1-C), and confirm that the host mode is correct.

  ```
  raidcom get host_grp -port CL1-C -IH1
  PORT GID GROUP_NAME Serial# HMD HMO_BITs
  CL1-C 0 1C-G00 322222 WIN
  ```

- When the host group ID used by the S-VOL is 1 or higher

  A host group whose host group ID is 1 or higher does not exist by default. Create a host group used by the S-VOL of the GAD pair, and set the host mode.

Note: If you create a new host group but do not reserve the new host group ID in the resource group, add the new host group ID to the resource group as described in Reserving a host group ID on page 4-30.

If you connect multiple server hosts of different platforms to a port, create a host group for each platform.

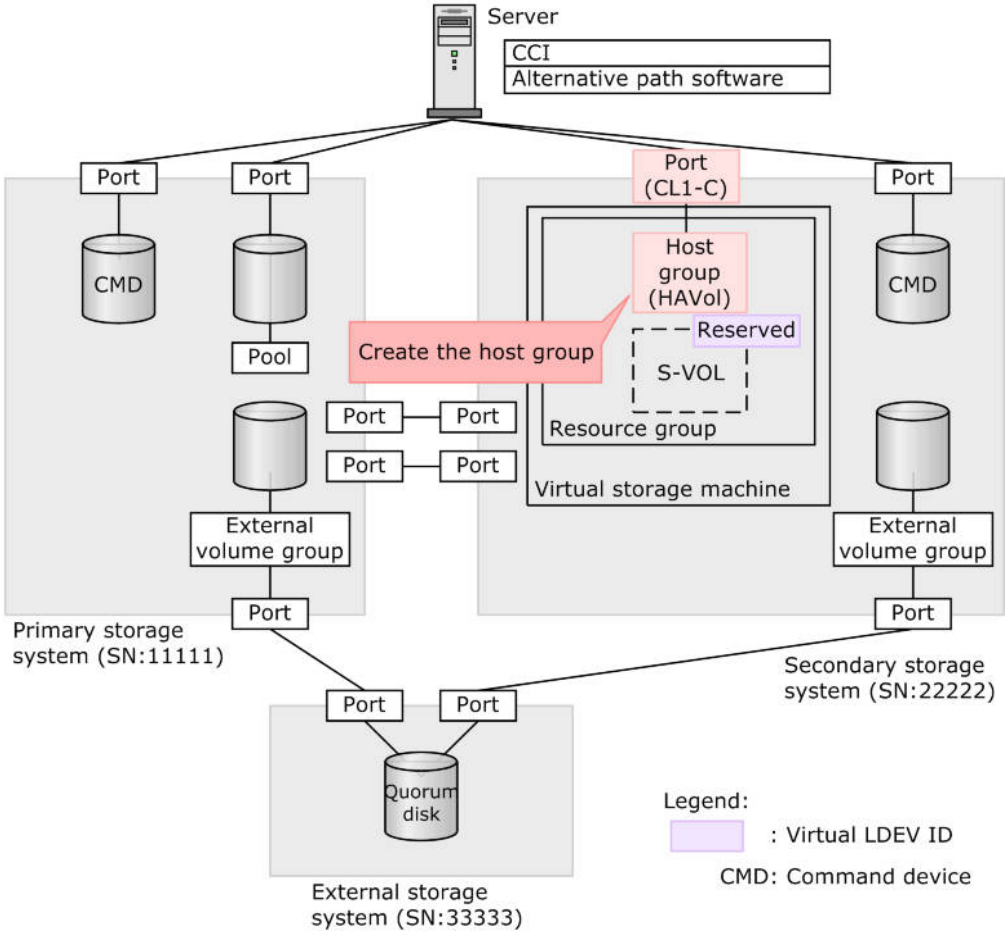

Procedure (for CL1-C-1)

1. Create a host group (HAVol) with host group ID 1 on port (CL1-C).
   ```
   raidcom add host_grp -port CL1-C-1 -host_grp_name HAVol -IH1
   ```
2. Set the host mode for the new host group (Windows host mode shown).
   ```
   raidcom modify host_grp -port CL1-C-1 -host_mode WIN -IH1
   ```
3. Reserve host group (CL1-C-1) to resource group 1.
   ```
   raidcom add resource -resource_name HAGroup1 -port CL1-C-1 -IH1
   ```

Check command and output examples

Display the information about the host group that is set for port (CL1-C). Confirm that the port name, host group ID, and host group name are correct.

```
raidcom get host_grp -port CL1-C -IH1
PORT   GID  GROUP_NAME                        Serial# HMD
HMO_BITs
CL1-C    0  1C-G00                            322222 WIN
CL1-C    1  HAVol                             322222 WIN
```

Related topics

- Reserving a host group ID on page 4-30

Creating a pool

After creating host groups, you must create a pool volume, format the volume, and create a Dynamic Provisioning pool.

Note: This example uses DP-VOLs. If you are using normal volumes (VSP G1000 only), you do not need to create a pool.

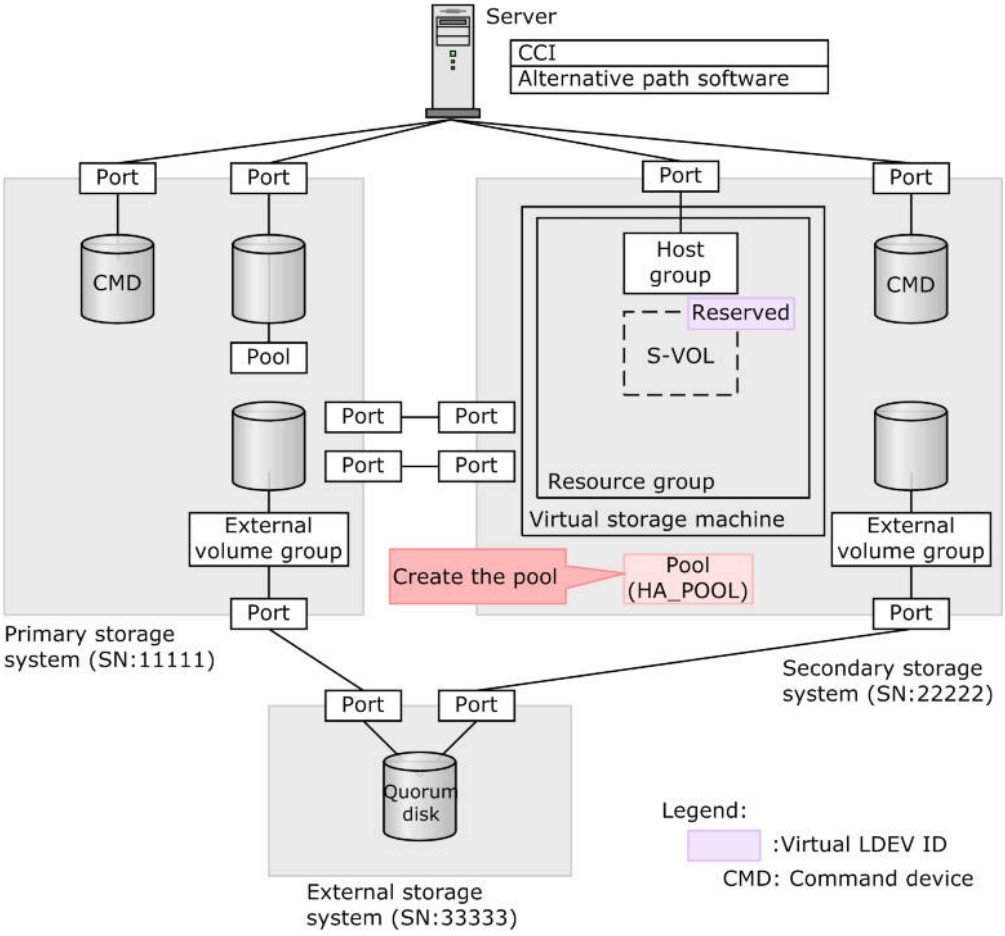

Procedure

1. Specify a parity group (13-4) to create a volume (pool volume) whose LDEV ID is 0x7777. The capacity is 100 GB.

   ```
   raidcom add ldev -ldev_id 0x7777 -parity_grp_id 13-4 -capacity
   100G -IH1
   ```

2. Confirm that asynchronous command processing has completed.

   ```
   raidcom get command_status -IH1
   HANDLE   SSB1    SSB2    ERR_CNT        Serial#     Description
   00c3        -       -          0         322222     -
   ```

3. Perform Quick Format to the volume (0x7777).

   ```
   raidcom initialize ldev -operation qfmt -ldev_id 0x7777 -IH1
   ```

4. Confirm that asynchronous command processing has completed.

```
raidcom get command_status -IH1
HANDLE   SSB1    SSB2    ERR_CNT        Serial#     Description
00c3       -       -          0         322222     -
```

5. Specify 0x7777 to the pool volume to create a pool for Dynamic Provisioning whose pool ID is 0 and whose pool name is HA_POOL.

```
raidcom add dp_pool -pool_id 0 -pool_name HA_POOL -ldev_id
0x7777 -IH1
```

6. Confirm that asynchronous command processing has completed.

```
raidcom get command_status -IH1
HANDLE   SSB1    SSB2    ERR_CNT        Serial#     Description
00c3       -       -          0         322222     -
```

Check command and output examples

1. Confirm that volume (LDEV ID: 0x7777) is set for the pool volume of pool (pool ID: 0).

```
raidcom get ldev -ldev_id 0x7777 -fx -IH1
Serial#  : 322222
LDEV : 7777
SL : 0
CL : 0
VOL_TYPE : OPEN-V-CVS
VOL_Capacity(BLK) : 209715200
NUM_LDEV : 1
LDEVs : 7777
NUM_PORT : 0
PORTs :
F_POOLID : 0
VOL_ATTR : CVS : POOL
RAID_LEVEL  : RAID1
RAID_TYPE   : 2D+2D
NUM_GROUP : 1
RAID_GROUPs : 13-04
DRIVE_TYPE  : DKR5E-J1R2SS
DRIVE_Capa : 2286717474
LDEV_NAMING :
STS : NML
OPE_TYPE : NONE
OPE_RATE : 100
MP# : 0
SSID : 0007
RSGID : 0
```

2. Check the pool capacity.

```
raidcom get dp_pool -IH1
PID  POLS U(%)  AV_CAP(MB)  TP_CAP(MB) W(%) H(%) Num LDEV#
LCNT  TL_CAP(MB)
000  POLN   0       98196       98196   70   80    1 30583
0          0
```

3. Check the pool name.

```
raidcom get pool -key opt -IH1
PID  POLS U(%) POOL_NAME                              Seq# Num
LDEV# H(%) VCAP(%) TYPE PM
000  POLN   0  HA_POOL                              322222   1
30583  80       -  OPEN N
```

Creating the S-VOL

Specify the volume that will become the S-VOL using the GAD reserve attribute and LDEV ID mapped earlier. The S-VOL must be the same size as the P-VOL.

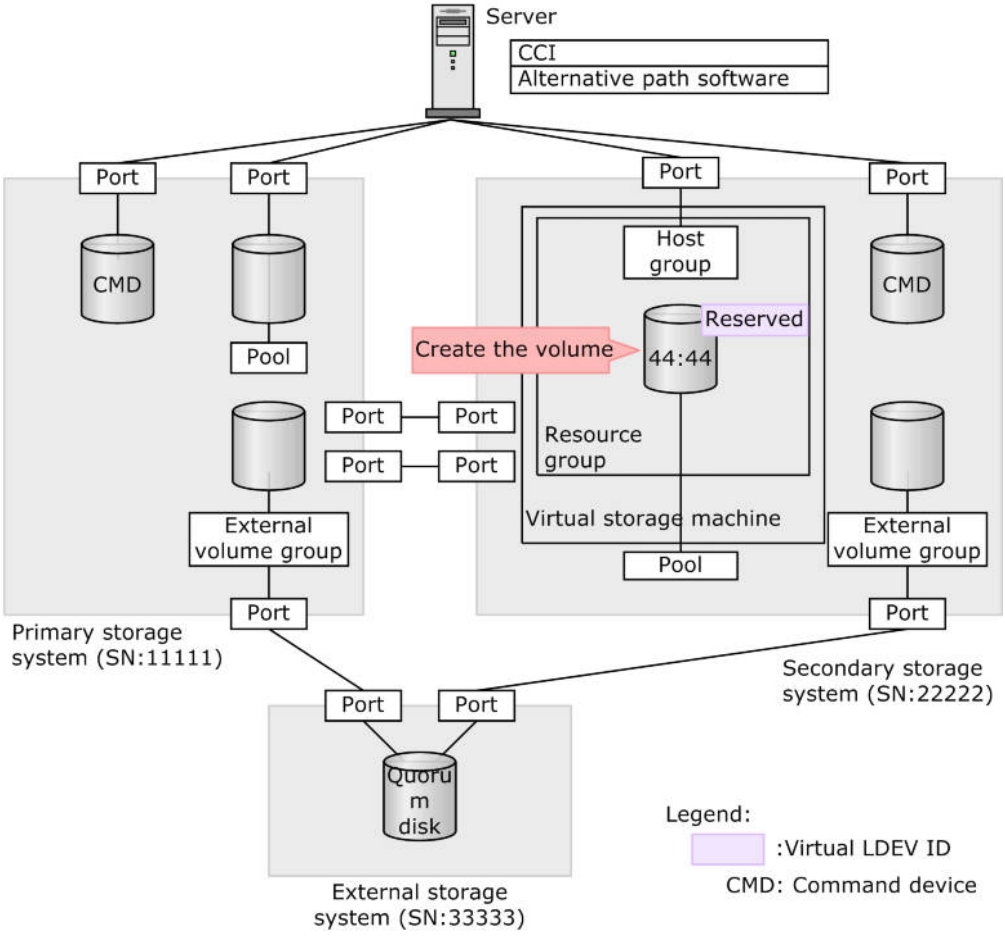

Procedure

1. Check the capacity of the P-VOL.

```
raidcom get ldev -ldev_id 0x2222 -fx -IH0
Serial#  : 311111
LDEV : 2222
SL : 0
CL : 0
VOL_TYPE : OPEN-V-CVS
VOL_Capacity(BLK) : 1024000
NUM_PORT : 0
PORTs :
F_POOLID : NONE
VOL_ATTR : CVS : HDP
B_POOLID : 0
LDEV_NAMING :
STS : NML
OPE_TYPE : NONE
OPE_RATE : 100
MP# : 0
```

```
SSID : 0005
Used_Block(BLK) : 0
RSGID : 0
```

2. In the Dynamic Provisioning pool with pool ID 0, create a virtual volume (DP-VOL) with a capacity of 1,024,000 blocks and LDEV ID = 0x4444.

```
raidcom add ldev -pool 0 -ldev_id 0x4444 -capacity 1024000 -IH1
```

3. Confirm that asynchronous command processing has completed.

```
raidcom get command_status -IH1
HANDLE   SSB1    SSB2    ERR_CNT        Serial#     Description
00c3        -       -          0         322222     -
```

Check command and output examples

Display the information for volume (LDEV ID: 0x4444). Confirm that the new volume satisfies the following requirements:

- The GAD reserve attribute is set (virtual LDEV ID (`VIR_LDEV`) is `ffff`).
- The volume has the same capacity as the P-VOL.
- The volume is a DP-VOL (volume attribute (`VOL_ATTR`) is `HDP`).

Note: This example uses DP-VOLs. This requirement does not apply if you are using normal volumes (VSP G1000 only).

```
raidcom get ldev -ldev_id 0x4444 -fx -IH1
Serial#  : 322222
LDEV : 4444 VIR_LDEV : ffff
SL : 0
CL : 0
VOL_TYPE : OPEN-V-CVS
VOL_Capacity(BLK) : 1024000
NUM_PORT : 0
PORTs :
F_POOLID : NONE
VOL_ATTR : CVS : HDP
B_POOLID : 0
LDEV_NAMING :
STS : NML
OPE_TYPE : NONE
OPE_RATE : 100
MP# : 0
SSID : 0009
Used_Block(BLK) : 0
RSGID : 1
```

Adding an LU path to the S-VOL

Add an LU path between the port connected to the server and the S-VOL. The host does not recognize the S-VOL, because the virtual LDEV ID has not yet been defined to the volume for the S-VOL.

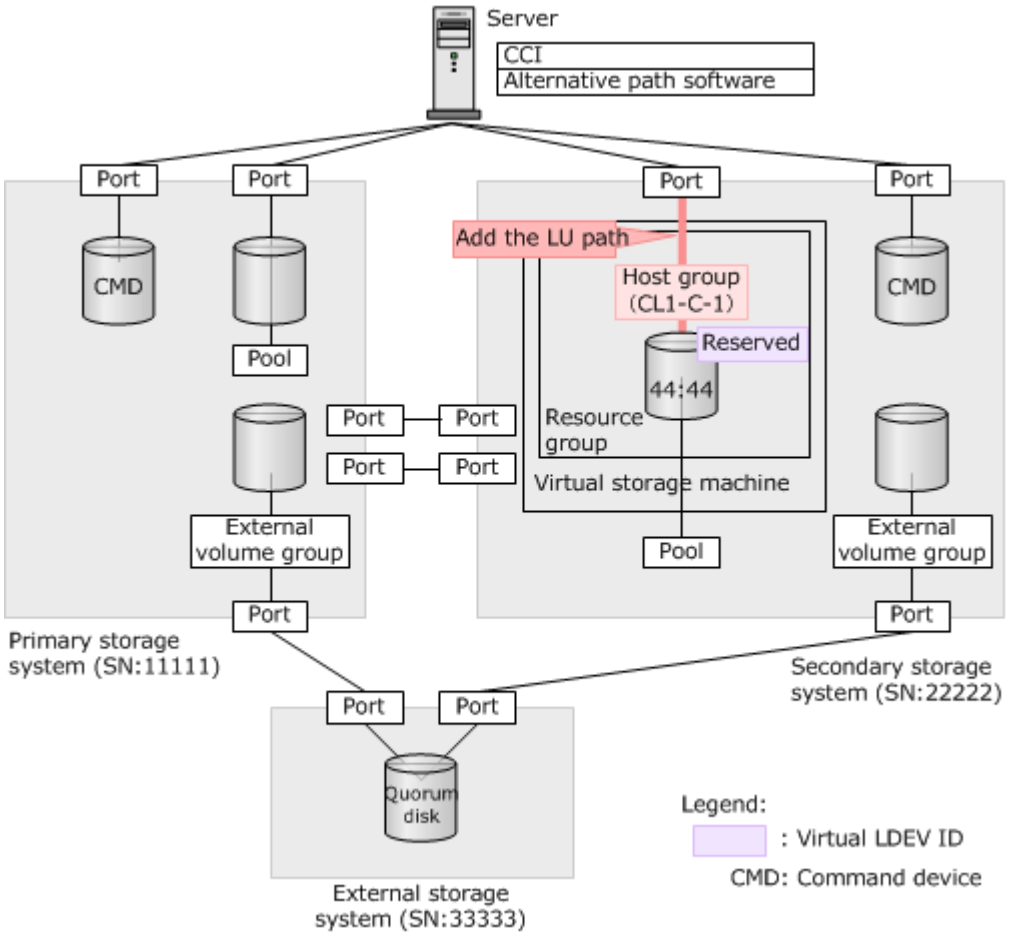

Procedure

Specify host group (CL1-C-1) and LU (0) to add an LU path to S-VOL (0x4444).

```
raidcom add lun -port CL1-C-1 -lun_id 0 -ldev_id 0x4444 -IH1
```

Check command and output examples

Display the information about the LU paths that are defined in host group (CL1-C-1).

```
raidcom get lun -port CL1-C-1 -fx -IH1
PORT   GID  HMD            LUN  NUM     LDEV  CM    Serial#  HMO_BITs
CL1-C    0  WIN            0    1       4444   -     322222
```

Updating the CCI configuration definition files

Before creating the GAD pair, you must update the CCI configuration definition files on the primary and secondary storage systems to add the information for the volumes that will become the P-VOL and S-VOL.

- Shutting down CCI on page 4-42
- Editing CCI configuration definition files on page 4-42

- Restarting CCI on page 4-43

Shutting down CCI

You must shut down both CCI instances before editing the configuration definition files.

Procedure (Windows shown)

Shut down instance 0 and instance 1.

```
horcmshutdown 0 1
inst 0:
HORCM Shutdown inst 0 !!!
inst 1:
HORCM Shutdown inst 1 !!!
```

Editing CCI configuration definition files

The following examples show the configuration definition files for a Windows host. Make sure to specify the actual LDEV IDs for the GAD pair volumes, not the virtual LDEV IDs.

Note: When specifying the serial number for VSP G1000 using CCI, add a "3" at the beginning of the serial number. For example, for serial number 11111, enter 311111.

Example of primary HORCM file, horcm0.conf

The underlined lines below show the updates for the volumes in the sample configuration in this chapter. Make sure to enter the information for your system in your configuration definition files.

```
HORCM_MON
#ip_address        service         poll(10ms)     timeout(10ms)
localhost          31000           -1             3000

HORCM_CMD
\\.\PhysicalDrive0

HORCM_LDEV
#GRP     DEV              SERIAL          LDEV#           MU#
oraHA    dev1             311111          22:22           h0

HORCM_INST
#GPR     IP_ADR              PORT#
oraHA    localhost           31001
```

Example of secondary HORCM file, horcm1.conf

The underlined lines below show the updates for the volumes in the sample configuration in this chapter. Make sure to enter the information for your system in your configuration definition files.

```
HORCM_MON
#ip_address        service         poll(10ms)     timeout(10ms)
```

```
localhost          31001          -1             3000

HORCM_CMD
\\.\PhysicalDrive1

HORCM_LDEV
#GRP     DEV              SERIAL          LDEV#           MU#
oraHA    dev1             322222          44:44           h0

HORCM_INST
#GPR     IP_ADR              PORT#
oraHA    localhost           31000
```

Restarting CCI

After editing the configuration definition files, restart both CCI instances.

Procedure (Windows shown)

Start instances 0 and 1.

```
horcmstart 0 1
starting HORCM inst 0
HORCM inst 0 starts successfully.
starting HORCM inst 1
HORCM inst 1 starts successfully.
```

Creating the GAD pair

After verifying that the virtual LDEV ID and enabling the ALUA mode, you can create the GAD pair.

- Verifying the virtual LDEV ID at the secondary site on page 4-43
- Setting the ALUA mode (VSP G1000 only) on page 4-45
- Creating a global-active device pair on page 4-46

Verifying the virtual LDEV ID at the secondary site

Before creating a GAD pair, check that the same virtual LDEV ID as that of the primary volume does not exist in the virtual storage machine of the secondary site, which has the same serial number and model as the primary storage system. If the same virtual LDEV ID as the primary volume exists, you cannot create the GAD pair.

Operate the virtual storage machine to check that the virtual LDEV ID does not exist. Specify the virtual storage machine for HORCM_VCMD of the configuration definition file, and then start CCI.

Procedure (Windows shown)

1. Start instances (100 and 101) for confirming the virtual LDEV IDs.

   ```
   horcmstart 100 101
   starting HORCM inst 100
   ```

```
HORCM inst 100 starts successfully.
starting HORCM inst 101
HORCM inst 101 starts successfully.
```

2. Confirm the P-VOL's virtual LDEV ID.

```
raidcom get ldev -ldev_id 0x2222 -key front_end -cnt 1 -fx -IH100
Serial# LDEV#  SL  CL VOL_TYPE          VOL_Cap(BLK) PID
ATTRIBUTE             Ports PORT_No:LU#:GRPNAME ...
311111  2222   0   0 OPEN-V-CVS             1024000   0
CVS|HDP                   0
```

3. Check that the same virtual LDEV ID as that of the primary volume does not exist in the virtual storage machine of the secondary site. After you execute this command, if virtual LDEV ID 0x2222 is not displayed, the same virtual LDEV ID (0x2222) as that of the primary volume does not exist in the virtual storage machine of the secondary site.

```
raidcom get ldev -ldev_id 0x2222 -key front_end -cnt 1 -fx -IH101
```

When you specify the virtual storage machine for HORCM_VCMD in the configuration definition file and execute the `raidcom get ldev` command by specifying the `-cnt` option, the virtual LDEV IDs in the range specified by the `-cnt` option are displayed.

Tip: To display the volume information as a list for each volume, use the `-key front_end` option for the `raidcom get ldev` command.

Revising the virtual LDEV ID at the secondary site

If the same virtual LDEV ID as the primary volume is displayed for the virtual storage machine of the secondary site, there might be errors in the GAD system implementation plan. Revise the system configuration.

The example when the same virtual LDEV ID as that of the P-VOL (0x2222) is assigned to the volume (LDEV ID: 0xfefe) in the virtual storage machine of the secondary machine is shown below.

Procedure

1. Check whether the same virtual LDEV ID as that of the primary volume is assigned to the virtual storage machine of the secondary site.

```
raidcom get ldev -ldev_id 0x2222 -key front_end -cnt 1 -fx -IH101
Serial# LDEV# SL CL VOL_TYPE VOL_Cap(BLK) PID ATTRIBUTE Ports
PORT_No:LU#:GRPNAME ...
311111 2222 - - NOT DEFINED
```

The virtual LDEV ID (0x2222) is assigned to the virtual storage machine of the secondary site.

2. Confirm the actual LDEV ID of the volume whose virtual LDEV ID is 0x2222.

```
raidcom get ldev -ldev_id 0x2222 -fx -IH101
Serial# : 311111 PHY_Serial# : 322222
LDEV : 2222 PHY_LDEV : fefe
SL : -
CL : -
VOL_TYPE : NOT DEFINED
```

```
SSID : -
RSGID : 1
```

In this example, the virtual LDEV ID (0x2222) is assigned to the volume whose actual LDEV ID is 0xfefe.

3. To use the virtual LDEV ID (0x2222) for a GAD pair volume, use the `raidcom unmap resource` command to remove assignment of the virtual LDEV ID (0x2222) from the volume whose LDEV ID is 0xfefe.

   ```
   raidcom unmap resource -ldev_id 0xfefe -virtual_ldev_id 0x2222 -IH1
   ```

4. Confirm that the assignment of the virtual LDEV ID (0x2222) is removed from the volume whose LDEV ID is 0xfefe.

   ```
   raidcom get ldev -ldev_id 0x2222 -key front_end -cnt 1 -fx -IH101
   ```

   When you specify the virtual storage machine for HORCM_VCMD in the configuration definition file, and execute the `raidcom get ldev` command by specifying the `-cnt` option, the virtual LDEV IDs existing in the specified range by the `-cnt` option are displayed.

   After you execute the above command, if the virtual LDEV ID 0x2222 is not displayed, the same virtual LDEV ID (0x2222) as that of the primary volume does not exist in the virtual storage machine of the secondary site.

Note: After releasing the virtual LDEV ID assignment, if you execute the `raidcom get ldev` command without specifying the `-cnt` option, the following error code and message are output:

```
raidcom: [EX_EGPERM] Permission denied with the Resource Group
```

In the example above, the virtual LDEV ID (0x2222) has not been defined after you released the virtual LDEV ID assignment. Therefore, the user of the virtual storage machine does not have access authority.
When a command is executed specifying a virtual storage machine (that is, using HORCM_VCMD), both the actual ID and the virtual ID of the specified resource must be assigned to the user. When the virtual storage machine is not specified (that is, using HORCM_CMD), the user can execute the command only if the actual ID of the specified resource is assigned to the user.

Setting the ALUA mode (VSP G1000 only)

To specify the preferred path in a cross-path configuration, you must enable the ALUA mode. Before creating a GAD pair, enable the ALUA mode of the P-VOL. If the ALUA mode of the P-VOL is enabled, the ALUA mode of the S-VOL is also enabled when a GAD pair is created.

Note:

- It is not necessary to set the ALUA mode if the cross-path configuration is not used.
- To enable the ALUA mode, VSP G1000 DKCMAIN version 80-03-31-00/00 or later is required.

Command example

```
raidcom modify ldev -ldev_id 0x2222 -alua enable
```

Check command and output example

Use the following command to verify that the ALUA mode is set correctly (`ALUA: Enable` is displayed).

```
raidcom get ldev -ldev_id 0x2222 -fx -IH0
Serial# : 311111
LDEV : 2222
SL : 0
CL : 0
VOL_TYPE : OPEN-V-CVS
VOL_Capacity(BLK) : 1024000
NUM_PORT : 0
PORTs :
F_POOLID : NONE
VOL_ATTR : CVS : HDP
B_POOLID : 0
LDEV_NAMING :
STS : NML
OPE_TYPE : NONE
OPE_RATE : 100
MP# : 0
SSID : 0005
Used_Block(BLK) : 0
ALUA : Enable
RSGID : 0
```

Related topics

- Suspending GAD pairs on page 6-3

Creating a global-active device pair

When GAD configuration is complete, you can start creating GAD pairs. When a pair is created, the P-VOL LDEV ID is set as the S-VOL's virtual LDEV ID. When the paircreate operation completes, the pair status becomes PAIR, and the P-VOL and S-VOL can accept I/O from the host. When a pair is deleted, the S-VOL's virtual LDEV ID is deleted, and the GAD reserve attribute remains set on the S-VOL.

Note: When you create a GAD pair, make sure that the available pool capacity for Dynamic Provisioning below the warning threshold is more than the capacity of the secondary volume. If you create a GAD pair at the secondary storage system when the available pool capacity below the warning threshold is less than the capacity of the secondary volume, SIM (SIM=620*xxx*, where *xxx* is the pool ID) is issued (the used capacity exceeds the warning threshold).

Note: You cannot create a GAD pair by using instances (100 and 101) for confirming the virtual LDEV IDs. To create a GAD pair, use instances (0 and 1) for operating storage systems.

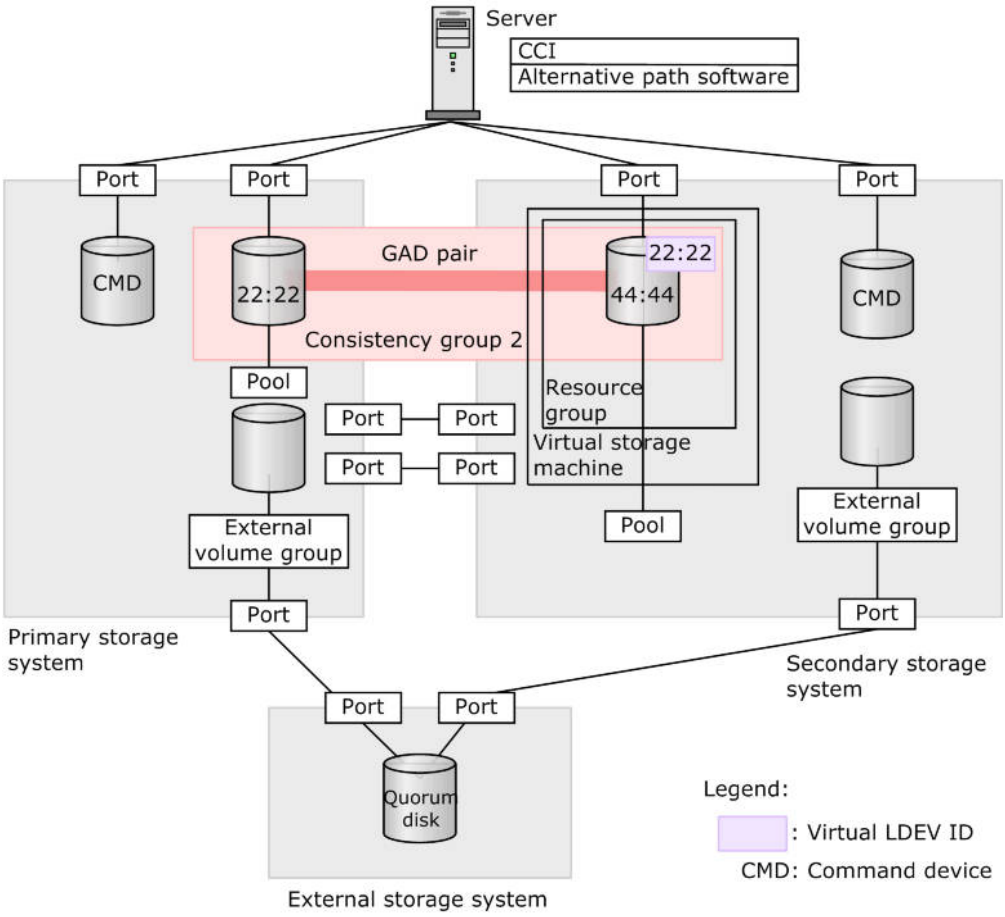

Note: Consistency groups are currently supported only on the VSP G1000.

Procedure

Specify 0 for the quorum disk ID and specify 2 for the consistency group ID (VSP G1000 only) to create a GAD pair.

```
paircreate -g oraHA -fg never 2 -vl -jq 0 -IH0
```

To register GAD pairs to a consistency group (VSP G1000 only), specify the -fg option when creating the pair. If you create a pair and do not specify a consistency group ID, the ID for an unused consistency group in the storage system will be automatically assigned.

The following example shows a command used to create a GAD pair without registering the pair to a consistency group:

```
paircreate -g oraHA -f never -vl -jq 0 -IH0
```

Check command and output examples

Confirm that a GAD pair is created.

```
pairdisplay -g oraHA -fxce -IH0
Group   PairVol(L/R) (Port#,TID, LU),Seq#,LDEV#.P/S,Status,Fence,
%,P-LDEV# M CTG JID AP EM       E-Seq# E-LDEV# R/W
```

```
oraHA    dev1(L)      (CL1-A-0, 0,    0)311111   2222.P-VOL PAIR
NEVER ,
100   4444 -    2    0   -   -             -        - L/M
oraHA    dev1(R)      (CL1-C-1, 0,    0)322222   4444.S-VOL PAIR
NEVER ,
100   2222 -    2    0   -   -             -        - L/M
```

Confirm that the copy progress reaches 100%.

Adding an alternate path to the S-VOL

Add an alternate path to the S-VOL on the host using the alternate path software. For some alternate path software the alternate path is added automatically. Make sure that the host has recognized the GAD secondary volume correctly.

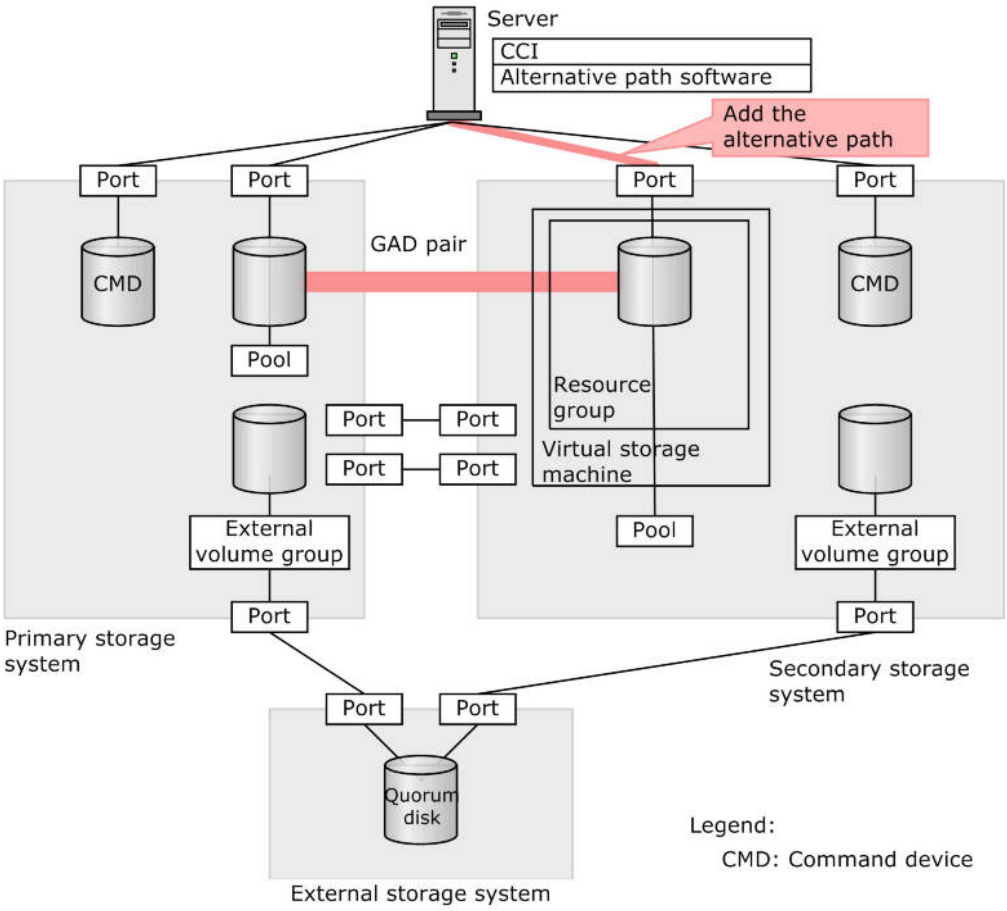

Caution: If Hitachi Dynamic Link Manager (HDLM) is installed on the server and host mode option 78 is set to the host group of the storage system, add the alternate path, and then execute the `dlnkmgr refresh -gad` command to incorporate the storage system settings to HDLM. For details about HDLM, see the HDLM user documentation.

Setting a nonpreferred cross path (VSP G1000 only)

After adding an alternate path to the S-VOL of a GAD pair from a server, set the asymmetric access status of the host group.

In a cross-path configuration, if the ALUA mode is enabled for the first time, all paths including cross paths between a server and a storage system are set as the preferred path. Therefore, you need to set the asymmetric access status of the cross path as a nonpreferred path. By doing so, paths between the server and the storage system in the primary site and between the server and the storage system in the secondary site can be preferably used.

Note:

- It is not necessary to set the asymmetric access status if the cross-path configuration is not used.
- Before setting the asymmetric access status, enable the ALUA mode. See Setting the ALUA mode (VSP G1000 only) on page 4-45.

Command example

```
raidcom modify lun -port CL1-A-0 -lun_id all -
asymmetric_access_state non_optimized
```

Check command and output example

Make sure that the ALUA mode (AL) is enabled (E) and that the nonpreferred access status (AAS) of the cross path is set to nonpreferred (AN).

```
raidcom get lun -port CL1-A-0 -key opt_page1 -fx
PORT  GID HMD        LUN NUM LDEV CM Serial# AL AAS
CL1-A   0 LINUX/IRIX   0   1 2222  -  311111 E  AN
```

Releasing the differential data managed in a pool

If differential data management fails due to insufficient pool capacity, you must release the differential data (pages) managed in the pool. You also need to release pages when you downgrade to a microcode version that does not support GAD pairs with DP-VOLs larger than 4,194,304 MB.

1. Delete all GAD pairs that use the V-VOL for which you want to release the pages using the **`pairsplit -S`** command.
2. Make sure that system option mode 755 is set to OFF. When system option mode 755 is set to ON, you cannot reclaim zero pages. For details about setting system option modes, contact your Hitachi Data Systems representative.
3. Restore the blocked pool. For details, see the *Provisioning Guide* for the storage system.
4. Release the V-VOL pages using the **`raidcom modify ldev`** command. Releasing pages might take some time. For details, see the *Provisioning Guide* for the storage system.

5

GAD 3DC delta resync (GAD+UR) operations

This chapter describes and provides instructions for using CCI commands to implement a GAD 3DC delta resync (GAD+UR) configuration.

For details about storage system support (models, microcode) for GAD+UR operations, see Requirements and restrictions on page 2-2.

GAD 3DC delta resync system configuration

The following figure shows a sample GAD 3DC delta resync (GAD+UR) configuration. The examples and procedures in this chapter describe this sample GAD+UR configuration.

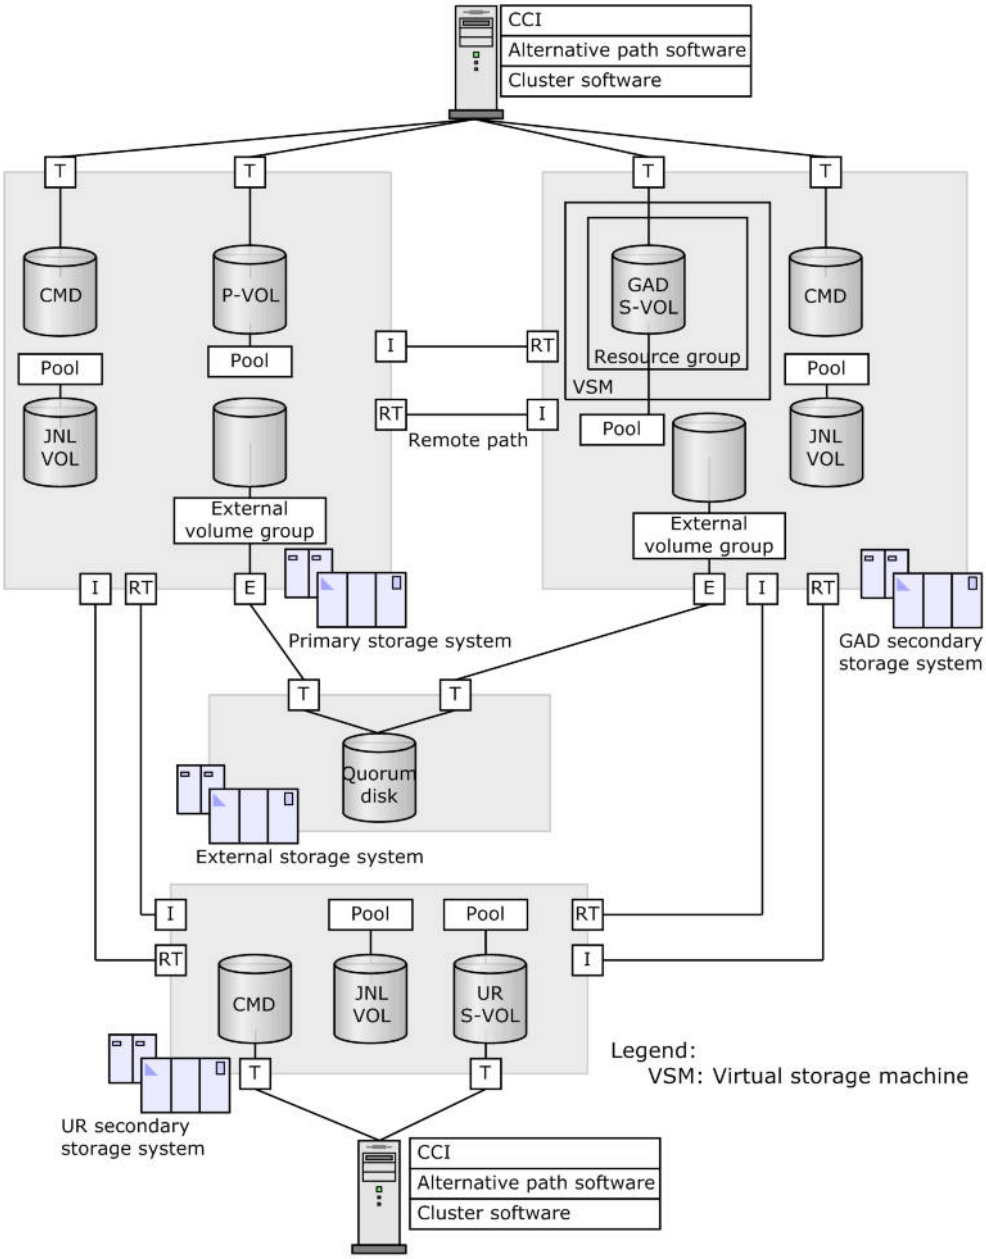

Figure 5-1 Sample GAD+UR configuration

Sites

The GAD+UR configuration spans the following four sites:

- Primary site: Site at which the primary storage system operates.

- GAD secondary site: Site at which the GAD secondary storage system operates.
- UR secondary storage system site: Site at which the UR secondary storage system operates.
- Quorum site: Site at which the external storage system containing the GAD quorum disk operates.

Storage systems

The GAD+UR configuration includes the following storage systems:

- Primary storage system: Contains the P-VOL, UR journal volume (JNL VOL), command device (CMD), and external volume for the quorum disk. The P-VOL and the JNL VOL are virtual volumes of HDP and are assigned capacities from different HDP pools.
- GAD secondary storage system: Contains the GAD S-VOL, delta UR journal volume, command device, external volume for the quorum disk, and the virtual storage machine (VSM). The GAD S-VOL must have the same virtual LDEV ID as the P-VOL. Therefore it is registered to the VSM.
- UR secondary storage system: Contains the UR S-VOL, UR/delta UR shared journal volume, command device, and external volume for the quorum disk. The UR secondary storage system must the same model as the primary and secondary storage systems.
- External storage system: Contains the GAD quorum disk.

Note: In this example the quorum disk resides in a separate external storage system. If desired, the quorum disk can reside in the UR secondary storage system.

For details about storage system support (models, microcode) for GAD+UR operations, see Requirements and restrictions on page 2-2.

Servers

There are two servers in this sample GAD+UR configuration. One server is connected to the primary storage system and the GAD secondary storage system. The other server is connected to the UR secondary storage system. The servers have the following components:

- CCI: Manages the configuration of the GAD/UR/delta UR pairs.
- Configuration definition file: Text file that contains the configuration information for the command device and the GAD/UR/delta UR pairs. It is used by CCI to manage the GAD/UR/delta UR pair configuration.
- Multi-path software: Controls the alternate paths from the host to the P-VOL or S-VOL.
- Cluster software at the primary site and GAD secondary site*: Provides clustering of the active-active configuration between the primary host and the GAD secondary host. In addition, the cluster software on the UR secondary host provides clustering of the active-passive configuration in

which the primary/GAD secondary hosts are active between the primary/GAD secondary hosts and the UR secondary host.

* This item does not apply to the sample configuration described in this chapter. In the sample in this chapter the primary host and GAD secondary host are combined.

- Cluster software at the UR secondary site: Provides clustering of the active-passive configuration in which the remote host becomes the standby host between the primary/secondary host and the remote host.
- Application (Active) at the primary site and GAD secondary site: Active application. Both primary and secondary hosts are active.
- Application (Standby) at the UR secondary site: Standby application.

Networks

The GAD+UR configuration has the following network configurations:

- Network between servers: All servers are connected to the same LAN.
- Networks between storage systems: There are two types of networks, both of which are connected using FC, between the storage systems:
  - Remote connection (primary-to-GAD secondary storage system, primary-to-UR secondary storage system, GAD secondary storage system-to-UR secondary storage system): Establish a bi-directional connection between the Initiator port and the RCU Target port in the primary, GAD secondary, and UR secondary storage systems.
  - External storage connection (primary-to-external storage, GAD secondary storage system-to-external storage system): Establish a connection between the external port in the primary/GAD secondary storage system and the Target port in the external storage system.
- Networks between servers and storage: The servers and storage systems are connected using FC. There are two types of networks between the servers and the storage systems:
  - Connection between the server and the primary/GAD secondary storage system: The server is connected to the primary storage system and GAD secondary storage system.
  - Connection between the server and the UR secondary storage system: The server is connected only to the UR secondary storage system.

GAD+UR pairs

When you create a GAD+UR configuration, the GAD pairs must already be registered to a consistency group.

With GAD+UR, when the statuses of the GAD pair (P-VOL and GAD S-VOL) and the UR pair (P-VOL and UR S-VOL) are PAIR and the UR delta resync pair (GAD S-VOL and UR S-VOL) is in the standby status, the operation is performed normally. When the UR delta resync pair status is HOLD, the pair status of the P-VOL is displayed as PSUS, and the journal mirror status is displayed as PJNS in CCI.

Primary storage system settings

The following tables list the settings for the sample primary storage system used in the examples and procedures in this chapter.

Primary storage system

Model	Serial number
VSP G1000	11111

Primary volume

Actual LDEV ID	HDP pool ID	Capacity	Port attribute	Port name	LU number
22:22	0	1,024,000 blocks	Target	CL1-A	1

HDP pool volume

Actual LDEV ID	Parity group ID	Capacity
aa:aa	1-1	100 GB
99:99	2-1	100 GB

HDP pool

Pool ID	Pool name	Pool volume	Use
0	PDP_POOL	aa:aa	P-VOL
1	PJNL_POOL	99:99	Journal

Host group

Host group ID	Host group name	Use
CL1-A-1	PVOL_HG1	P-VOL

Journal volume

Actual LDEV ID	HDP pool ID	Capacity
88:88	1	30 GB

Ports for connection between storage systems

Port name	Port attribute	Target storage system	Use
CL3-A	Initiator	GAD secondary storage system	GAD
CL4-A	RCU target	GAD secondary storage system	GAD
CL7-A	RCU target	UR secondary storage system	UR
CL6-A	Initiator	UR secondary storage system	UR
CL5-C	Target	External storage	Quorum disk

Pair

Pair type	Pair group name	Device name (LDEV ID)	Mirror ID	CTG ID
GAD	oraHA	dev1(22:22)	0	1
UR	oraREMOTE	dev2(22:22)	1	0

Quorum disk

Quorum disk ID
0

GAD secondary storage system settings

The following tables list the settings for the sample GAD secondary storage system used in the examples and procedures in this chapter.

Primary storage system

Model	Serial number
VSP G1000	22222

HDP pool volume

Actual LDEV ID	Parity group ID	Capacity
77:77	13-4	100 GB
99:99	2-1	100 GB

HDP pool

Pool ID	Pool name	Pool volume	Use
0	HA_POOL	77:77	GAD S-VOL
1	LJNL_POOL	99:99	Journal

GAD S-VOL

Actual LDEV ID	HDP pool ID	Capacity	Port attribute	Port name	LU number
44:44	0	1,024,000 block	Target	CL1-C	0

Host group

Host group ID	Host group name	Use
CL1-C-1	1C-G00	GAD S-VOL
CL1-B-1	LVOL_HG2	GAD S-VOL

Journal volume

Actual LDEV ID	HDP pool ID	Capacity
88:88	1	30 GB

Ports for connection between storage systems

Port name	Port attribute	Target storage system	Use
CL3-C	RCU target	Primary storage system	GAD
CL4-C	Initiator	Primary storage system	GAD
CL6-C	Initiator	UR secondary storage system	UR delta resync
CL7-C	RCU target	UR secondary storage system	UR delta resync
CL5-C	External	External storage	Quorum disk

Pair

Pair type	Pair group name	Device name (LDEV ID)	Mirror ID	CTG ID
GAD	oraHA	dev1(44:44)	0	1
UR delta resync	oraDELTA	dev3(44:44)	2	0

Quorum disk

Quorum disk ID
0

Resource group

Resource group ID	Resource group name	Virtual storage machine (VSM)	
		Model	Serial number
1	HAGroup1	VSP G1000	11111

UR secondary storage system settings

The following tables list the settings for the sample UR secondary storage system used in the examples and procedures in this chapter.

Primary storage system

Model	Serial number
VSP G1000	44444

HDP pool volume

Actual LDEV ID	Parity group ID	Capacity
77:77	13-4	100 GB
99:99	2-1	100 GB

HDP pool

Pool ID	Pool name	Pool volume	Use
0	VOL_POOL	77:77	UR S-VOL
1	RJNL_POOL	99:99	Journal

UR S-VOL

Actual LDEV ID	HDP pool ID	Capacity	Port attribute	Port name	LU number
66:66	0	1,024,000 blocks	Target	CL1-A	1

Host group

Host group ID	Host group name	Use
CL1-A-1	REMOTE	UR S-VOL

Journal volume

Actual LDEV ID	HDP pool ID	Capacity
88:88	1	30 GB

Ports for connection between storage systems

Port name	Port attribute	Target storage system	Use
CL6-A	RCU target	Primary storage system	UR
CL7-A	Initiator	Primary storage system	UR
CL6-C	RCU target	GAD secondary storage system	UR delta resync
CL7-C	Initiator	GAD secondary storage system	UR delta resync

Pair

Pair type	Pair group name	Device name (LDEV ID)	Mirror ID	CTG ID
UR	oraREMOTE	dev2(66:66)	1	1
UR delta resync	oraDELTA	dev3(66:66)	2	0

CCI server configuration

The following tables list the CCI configuration settings for the pair management server used in the examples and procedures in this chapter.

CCI settings for the pair management server for the primary and GAD secondary storage systems

Instance number	Configuration definition file	Use
0	horcm0.conf	GAD and UR
100	horcm100.conf	Operation in terms of VSM (serial number 11111) of primary storage system
1	horcm1.conf	GAD and UR delta resync
101	horcm101.conf	Operation in terms of VSM (serial number 11111) of GAD secondary storage system

CCI settings for the pair management server at the UR secondary site

Instance number	Configuration definition file	Use
2	horcm2.conf	UR and UR delta resync

Workflow for creating a GAD 3DC delta resync environment

Replication pairs are created in the following order when you set up a GAD 3DC delta resync (GAD+UR) environment:

1. GAD pair

Note: When you create a GAD+UR configuration, the GAD pairs must already be registered to a consistency group.

2. UR delta resync pair
3. UR pair

Workflow for creating a GAD+UR environment:

1. Preparing the UR secondary site on page 5-11
2. Installing the software products on page 5-12
3. Creating the command devices on page 5-12
4. Creating the UR delta resync pairs on page 5-12
5. Starting GAD+UR operations on page 5-30

Initial state

This chapter describes and provides instructions for using CCI to add a UR delta resync pair and a UR pair to a system that has a GAD pair. In the initial state before configuration, the necessary devices are installed at the primary site and GAD secondary site, the network required for the connections among the devices is physically connected, and the GAD pairs have already been created.

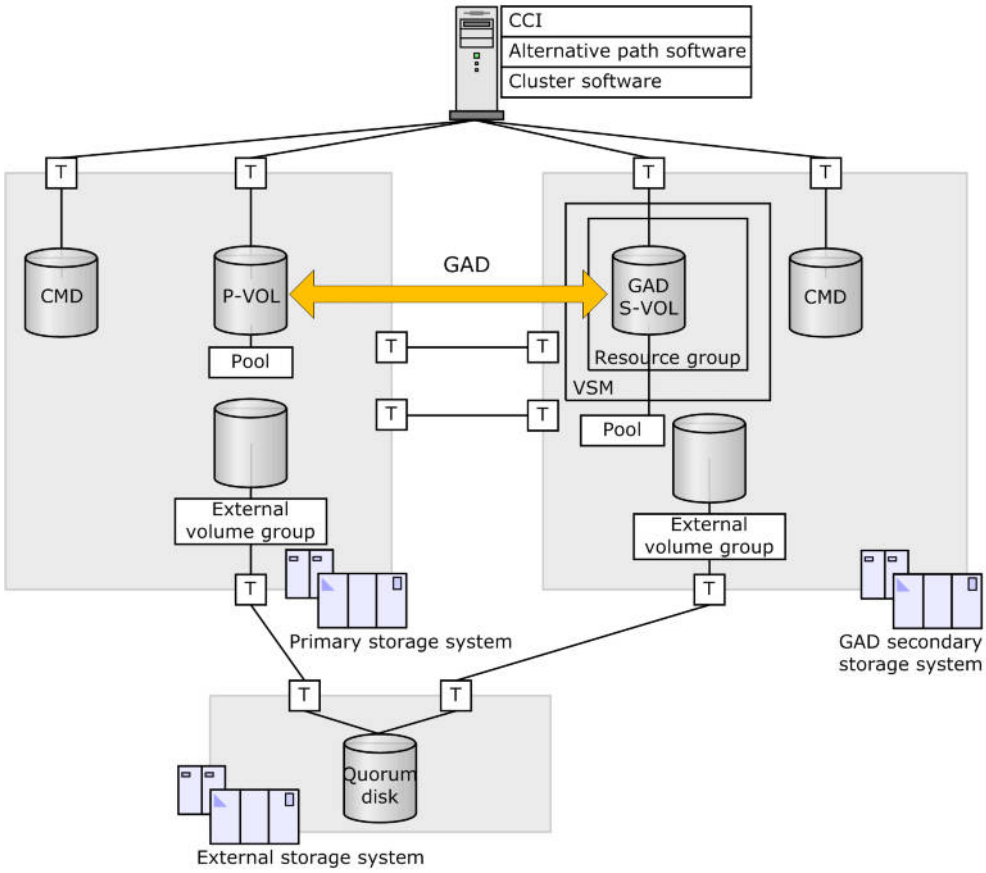

Preparing the UR secondary site

1. Install the storage system and server at the UR secondary site.
2. Install the following software on the server:
   - CCI
   - Multi-path software
   - Cluster software
3. Connect the storage system to the server.
4. Connect the storage system to the GAD storage systems with physical paths.

Installing the software products

Install the licenses for the following software products on the primary, GAD secondary, and UR secondary storage systems. For details, see the *Hitachi Command Suite User Guide*. For VSP G200, G400, G600, G800 see the *System Administrator Guide for Hitachi Virtual Storage Platform Gx00 and Fx00 Models*.

- Dynamic Provisioning
- TrueCopy
- Disaster Recovery Extended
- global-active device
- Universal Replicator

Creating the command devices

Create a CCI command device in Resource Group 0 in the UR secondary storage system, and make sure that the server recognizes it. Command devices should already exist in the primary and GAD secondary storage systems.

The following shows the flow of creating a command device. For details, see the *Command Control Interface Installation and Configuration Guide*.

1. Create the volume for the command device in the UR secondary storage system.
2. Format the created volume.
3. Define the created volume on the port connected to the server.
4. Enable the command device attribute of the created volume, and then enable user authentication of the new command device. For details about creating a command device, see the *Provisioning Guide* for the storage system.
5. If necessary, change the topology and fabric settings for the ports defined to the command devices.

Creating the UR delta resync pairs

This section describes the procedure for creating the UR delta resync pairs in your GAD 3DC delta resync (GAD+UR) configuration:

1. Creating the configuration definition file on page 5-13
2. Starting CCI on page 5-13
3. Adding the remote connections on page 5-14
4. Setting up the GAD secondary storage system on page 5-15
5. Setting up the UR secondary storage system on page 5-17
6. Updating the CCI configuration definition files on page 5-21
7. Creating the UR delta resync pairs on page 5-22

8. Creating the UR pairs on page 5-24

Creating the configuration definition file

Create the CCI configuration definition file in the server at the UR secondary site. The following shows the description example of the configuration definition file.

horcm2.conf

```
HORCM_MON
#ip_address    service    poll(10ms)     timeout(10ms)
localhost        31002      -1                   3000

HORCM_CMD
\\.\PhysicalDrive2
```

Starting CCI

Start the CCI instance in the server at the UR secondary site. The following shows the example of command execution.

1. Start CCI instances 1 and 2.
   - Server connected to the GAD secondary storage system
     ```
     horcmstart 1
     starting HORCM inst 1
     HORCM inst 1 starts successfully.
     ```
   - Server connected to the UR secondary storage system
     ```
     horcmstart 2
     starting HORCM inst 2
     HORCM inst 2 starts successfully.
     ```
2. Enter the user name and the password to authenticate the user.
   - Server connected to the GAD secondary storage system
     ```
     raidcom -login <username> <password> -IH1
     ```
   - Server connected to the UR secondary storage system
     ```
     raidcom -login <username> <password> -IH2
     ```

   The `-IH` option in this example is used for each command to specify an instance. You can also perform the operation using a shell for each instance. To start the shell, specify an instance number to the environment variable HORCMINST, and then execute the command.

Setting the port attributes

To establish the bi-directional remote connections between the GAD secondary storage system and the UR secondary storage system, first set the port attributes on each storage system. The following shows the example of command execution.

1. Set the attributes of Initiator port and RCU Target port in the GAD secondary storage system and the UR secondary storage system.

- Server connected to the GAD secondary storage system

```
raidcom modify port -port CL6-C -port_attribute MCU -IH1
raidcom modify port -port CL7-C -port_attribute RCU -IH1
```

- Server connected to the UR secondary storage system

```
raidcom modify port -port CL6-C -port_attribute RCU -IH2
raidcom modify port -port CL7-C -port_attribute MCU -IH2
```

Note: Use the same procedure to change the port attributes for the alternate paths.

2. Check the setting of the port attribute.

- Server connected to the GAD secondary storage system

```
raidcom get port -IH1
PORT   TYPE   ATTR  SPD  LPID  FAB  CONN  SSW  SL  Serial#
WWN               PHY_PORT
(Omitted)
CL6-C  FIBRE  MCU   AUT  E8    N    FCAL  N    0   322222
50060e80072b6720  -
(Omitted)
CL7-C  FIBRE  RCU   AUT  97    N    FCAL  N    0   322222
50060e80072b6730  -
(Omitted)
```

- Server connected to the UR secondary storage system

```
raidcom get port -IH2
PORT   TYPE   ATTR  SPD  LPID  FAB  CONN  SSW  SL  Serial#
WWN               PHY_PORT
(Omitted)
CL6-C  FIBRE  RCU   AUT  E8    N    FCAL  N    0   344444
50060e80072b6720  -
(Omitted)
CL7-C  FIBRE  MCU   AUT  97    N    FCAL  N    0   344444
50060e80072b6730  -
(Omitted)
```

Adding the remote connections

After setting the port attributes, add the remote connections between the GAD secondary storage system and UR secondary storage system. The following shows the example of command execution.

Note:

- When specifying the serial number for VSP G1000 using CCI, add a "3" at the beginning of the serial number. For example, for serial number 11111, enter 311111.
- To specify the VSP G1000, enter `R800`.
- The VSP G1000 is displayed as `R8` in command output.

1. Add a bi-directional remote connection between the GAD secondary storage system and UR secondary storage system. Add a remote connection whose path group ID is 0 from port CL6-C in the GAD secondary storage system to port CL6-C in the UR secondary storage

system, and add a remote connection whose path group ID is 0 from port CL6-C in the UR secondary storage system to port CL6-C in the GAD secondary storage system. Specify the same path group ID to the bidirectional remote connections. After executing the command, confirm that the asynchronous command processing ended normally.

- Server connected to the GAD secondary storage system

```
raidcom add rcu -cu_free 344444 R800 0 -mcu_port CL6-C -rcu_port
CL6-C -IH1
raidcom get command_status -IH1
HANDLE   SSB1    SSB2    ERR_CNT    Serial#    Description
00c3     -       -       0          322222     -
```

- Server connected to the UR secondary storage system

```
raidcom add rcu -cu_free 322222 R800 0 -mcu_port CL7-C -rcu_port
CL7-C -IH2
raidcom get command_status -IH2
HANDLE    SSB1    SSB2    ERR_CNT    Serial#    Description
00c3      -       -       0          344444     -
```

After adding the initial connections between the storage systems using the `raidcom add rcu` command, add the alternate paths between the storage systems using the `raidcom add rcu_path` command.

2. Check the settings of the remote connection.

- Server connected to the GAD secondary storage system

```
raidcom get rcu -cu_free 344444 R800 0 -IH1
Serial#    ID            PID
        MCU        RCU          M/R         T
PNO         MPORT
        RPORT    STS_CD    SSIDs
344444  R8           0             -
-               RCU
        F               0              CL6-C     CL6-C
        NML_01-
```

- Server connected to the UR secondary storage system

```
raidcom get rcu -cu_free 322222 R800 0 -IH2
Serial#    ID            PID
        MCU        RCU          M/R         T
PNO         MPORT
        RPORT    STS_CD    SSIDs
322222  R8           0             -
-               RCU
        F               0              CL7-C     CL7-C
        NML_01-
```

Setting up the GAD secondary storage system

To set up the GAD secondary storage system, you must create a pool volume of HDP pool for journal, create an HDP pool for journal, create a journal volume, and create a journal. The following shows the example of command execution.

1. Create a pool volume. Specify Parity group 2-1 and create volume 100 GB whose LDEV ID is 0x9999. After executing the command, confirm that the asynchronous command processing ended normally.

```
raidcom add ldev -ldev_id 0x9999 -parity_grp_id 2-1 -capacity
100G -IH1
raidcom get command_status -IH1
HANDLE     SSB1        SSB2        ERR_CNT
Serial#     Description
00c3       -               -               0
322222   -
```

2. Perform quick format of the created volume. After executing the command, confirm that the asynchronous command processing ended normally.

```
raidcom initialize ldev -operation qfmt -ldev_id 0x9999 -IH1
raidcom get command_status -IH1
HANDLE SSB1   SSB2        ERR_CNT                 Serial#
Description
00c3       -               -               0
322222   -
```

3. Create a pool. Specify the created volume 0x9999 in the pool volume and create an HDP pool whose pool ID is 1 and pool name is LJNL_POOL. After executing the command, confirm that the asynchronous command processing ended normally.

```
raidcom add dp_pool -pool_id 1 -pool_name LJNL_POOL -ldev_id
0x9999 -IH1
raidcom get command_status -IH1
HANDLE    SSB1    SSB2    ERR_CNT    Serial#    Description
00c3    -    -    0    322222    -
```

4. Confirm that the volume (LDEV ID: 0x9999) is set in the pool volume in the pool (Pool ID: 1).

```
raidcom get ldev -ldev_id 0x9999 -fx -IH1
Serial# : 322222
LDEV : 9999
(Omitted)
F_POOLID : 1
VOL_ATTR : CVS : POOL
(Omitted)
```

5. Check the pool capacity.

```
raidcom get dp_pool -IH1
PID    POLS    U(%)    AV_CAP(MB)    TP_CAP(MB)    W(%)
H(%)    Num    LDEV#    LCNT    TL_CAP(MB)
001    POLN    0    98196        98196        70    80    1
39321    0        0
```

6. Check the pool name.

```
raidcom get pool -key opt -IH1
PID    POLS    U(%)    POOL_NAME    Seq#    Num    LDEV#
H(%)    VCAP(%)    TYPE    PM
001    POLN    0    LJNL_POOL    322222    1    39321    80
-    OPEN    N
```

7. Create a journal volume. Create a virtual volume (DP-VOL) whose capacity is 30 GB and LDEV ID is 0x8888 in HDP pool ID 1. After executing the command, confirm that the asynchronous command processing ended normally.

```
raidcom add ldev -pool 1 -ldev_id 0x8888 -capacity 30G -IH1
raidcom get command_status -IH1
HANDLE    SSB1    SSB2    ERR_CNT    Serial#    Description
00c3      -       -       0          322222     -
```

8. Check the created volume.

```
raidcom get ldev -ldev_id 0x8888 -fx -IH1
Serial# : 322222
LDEV : 8888
(omitted)
VOL_ATTR : CVS : HDP
(omitted)
```

9. Create a journal. Specify the volume whose LDEV ID is 0x8888 to create a journal whose journal ID is 0. After executing the command, confirm that the asynchronous command processing ended normally.

```
raidcom add journal -journal_id 0 -ldev_id 0x8888 -IH1
raidcom get command_status -IH1
HANDLE    SSB1    SSB2    ERR_CNT    Serial#    Description
00c3      -       -       0          322222     -
```

10. Confirm that the volume (LDEV ID: 0x8888) is registered to the journal (journal ID: 0) and check the journal capacity.

```
raidcom get journal -IH1
JID   MU   CTG   JNLS   AP   U(%)   Q-Marker   Q-CNT   D-
SZ(BLK)   Seq#    Num   LDEV#
000   0    1     PJNN   4    21     43216fde   30
62914560    322222  1     34952
```

Setting up the UR secondary storage system

To set up the UR secondary storage system, you must create a pool volume of HDP pool for journal, create an HDP pool for journal, create a journal volume, create a journal, create a pool volume of HDP pool for the UR S-VOL, create an HDP pool for the UR S-VOL, create a UR S-VOL, create a host group, and add an LU path. The following shows the example of command execution.

Note: This example uses a pool volume of HDP pool for the UR S-VOL. This is not a requirement.

1. Use the following parameters and procedure to create pool volumes of HDP pool for journal, create HDP pool for journal, create journal volume, and create journal in the UR secondary storage system:

   Parameters:

   - RM instance ID: 2
   - Parity group : 2-1
   - Pool volume LDEV ID: 0x9999
   - Pool volume LDEV capacity : 100 GB
   - Pool ID: 1
   - Pool name : RJNL_POOL
   - Journal volume LDEV ID: 0x8888

- Journal volume LDEV capacity: 30 GB
- Journal ID: 0

Procedure:

a. Create a pool volume. Specify Parity group 2-1 and create volume 100 GB whose LDEV ID is 0x9999. After executing the command, confirm that the asynchronous command processing ended normally.

```
raidcom add ldev -ldev_id 0x9999 -parity_grp_id 2-1 -capacity
100G -IH2
raidcom get command_status -IH2
HANDLE     SSB1        SSB2        ERR_CNT
Serial#     Description
00c3       -               -               0
344444   -
```

b. Perform quick format of the created volume. After executing the command, confirm that the asynchronous command processing ended normally.

```
raidcom initialize ldev -operation qfmt -ldev_id 0x9999 -IH2
raidcom get command_status -IH2
HANDLE SSB1   SSB2        ERR_CNT                 Serial#
Description
00c3       -               -               0
344444   -
```

c. Create a pool. Specify the created volume 0x9999 in the pool volume and create an HDP pool whose pool ID is 1 and pool name is RJNL_POOL. After executing the command, confirm that the asynchronous command processing ended normally.

```
raidcom add dp_pool -pool_id 1 -pool_name RJNL_POOL -ldev_id
0x9999 -IH2
raidcom get command_status -IH2
HANDLE    SSB1    SSB2    ERR_CNT    Serial#    Description
00c3    -    -    0    344444    -
```

d. Confirm that the volume (LDEV ID: 0x9999) is set in the pool volume in the pool (Pool ID: 1).

```
raidcom get ldev -ldev_id 0x9999 -fx -IH2
Serial# : 344444
LDEV : 9999
(Omitted)
F_POOLID : 1
VOL_ATTR : CVS : POOL
(Omitted)
```

e. Check the pool capacity.

```
raidcom get dp_pool -IH2
PID    POLS    U(%)    AV_CAP(MB)    TP_CAP(MB)    W(%)
H(%)    Num    LDEV#    LCNT    TL_CAP(MB)
001    POLN    0    98196        98196        70    80
1    39321    0        0
```

f. Check the pool name.

```
raidcom get pool -key opt -IH2
PID    POLS    U(%)    POOL_NAME    Seq#    Num    LDEV#
H(%)    VCAP(%)    TYPE    PM
001    POLN    0    RJNL_POOL    344444    1    39321
80    -    OPEN    N
```

g. Create a journal volume. Create a virtual volume (DP-VOL) whose capacity is 30 GB and LDEV ID is 0x8888 in HDP pool ID 1. After executing the command, confirm that the asynchronous command processing ended normally.

```
raidcom add ldev -pool 1 -ldev_id 0x8888 -capacity 30G -IH2
raidcom get command_status -IH2
HANDLE    SSB1    SSB2    ERR_CNT    Serial#    Description
00c3    -    -    0    344444    -
```

h. Check the created journal volume.

```
raidcom get ldev -ldev_id 0x8888 -fx -IH2
Serial# : 344444
LDEV : 8888
(omitted)
VOL_ATTR : CVS : HDP
(omitted)
```

i. Create a journal. Specify the volume whose LDEV ID is 0x8888 to create a journal whose journal ID is 0. After executing the command, confirm that the asynchronous command processing ended normally.

```
raidcom add journal -journal_id 0 -ldev_id 0x8888 -IH2
raidcom get command_status -IH2
HANDLE    SSB1    SSB2    ERR_CNT    Serial#    Description
00c3      -       -       0          344444     -
```

j. Confirm that the volume (LDEV ID: 0x9999) is registered to the journal (journal ID: 0) and check the journal capacity.

```
raidcom get journal -IH2
JID   MU   CTG   JNLS   AP   U(%)   Q-Marker   Q-CNT   D-
SZ(BLK)   Seq#    Num   LDEV#
000   0    1     PJNN   4    21     43216fde   30
62914560    322222  1     34952
```

2. Use the following parameters and procedure to create pool volumes of HDP pool for the UR S-VOL and create an HDP pool for the UR S-VOL in the UR secondary storage system:

Parameters:

- RM instance ID: 2
- Parity group : 13-4
- Pool volume LDEV ID: 0x7777
- Pool volume LDEV capacity : 100 GB
- Pool ID: 0
- Pool name : VOL_POOL

Procedure:

a. Create a pool volume. Specify Parity group 13-4 and create volume 100 GB whose LDEV ID is 0x7777. After executing the command, confirm that the asynchronous command processing ended normally.

```
raidcom add ldev -ldev_id 0x7777 -parity_grp_id 13-4 -
capacity 100G -IH2
raidcom get command_status -IH2
HANDLE     SSB1        SSB2        ERR_CNT
Serial#     Description
```

```
00c3       -                -                0
344444   -
```

b. Perform quick format of the created volume. After executing the command, confirm that the asynchronous command processing ended normally.

```
raidcom initialize ldev -operation qfmt -ldev_id 0x7777 -IH2
raidcom get command_status -IH2
HANDLE SSB1   SSB2        ERR_CNT                 Serial#
Description
00c3       -                -                0
344444   -
```

c. Create a pool. Specify the created volume 0x7777 in the pool volume and create an HDP pool whose pool ID is 0 and pool name is VOL_POOL. After executing the command, confirm that the asynchronous command processing ended normally.

```
raidcom add dp_pool -pool_id 0 -pool_name VOL_POOL -ldev_id
0x7777 -IH2
raidcom get command_status -IH2
HANDLE    SSB1    SSB2    ERR_CNT    Serial#    Description
00c3    -    -    0    344444    -
```

d. Confirm that the volume (LDEV ID: 0x7777) is set in the pool volume in the pool (Pool ID: 0).

```
raidcom get ldev -ldev_id 0x7777 -fx -IH2
Serial# : 344444
LDEV : 7777
(Omitted)
F_POOLID : 0
VOL_ATTR : CVS : POOL
(Omitted)
```

e. Check the pool capacity.

```
raidcom get dp_pool -IH2
PID    POLS    U(%)    AV_CAP(MB)    TP_CAP(MB)    W(%)
H(%)    Num    LDEV#    LCNT    TL_CAP(MB)
000    POLN    0    98196        98196        70    80
0    30583    0        0
```

f. Check the pool name.

```
raidcom get pool -key opt -IH2
PID    POLS    U(%)    POOL_NAME    Seq#    Num    LDEV#
H(%)    VCAP(%)    TYPE    PM
001    POLN    0    VOL_POOL    344444    0    30583    80
-    OPEN    N
```

3. Create the UR S-VOL. Create a virtual volume (DP-VOLs) whose capacity is 1,024,000 blocks and LDEV ID is 0x6666 in HDP pool 0. After executing the command, confirm that the asynchronous command processing ended normally.

```
raidcom add ldev -pool 1 -ldev_id 0x6666 -capacity 1024000 -IH2
raidcom get command_status -IH2
HANDLE    SSB1    SSB2    ERR_CNT    Serial#    Description
00c3      -       -       0          344444     -
```

4. Check the information of the volumes (LDEV ID: 0x6666).

```
raidcom get ldev -ldev_id 0x6666 -fx -IH2
Serial# : 344444
LDEV : 6666 VIR_LDEV : 6666
```

```
(Omitted)
VOL_Capacity(BLK) : 41943040
(Omitted)
VOL_ATTR : CVS : HDP
B_POOLID : 0
(Omitted)
```

5. Create a host group. In port CL1-A, create a host group whose host group ID is 1 and host group name is REMOTE. Set the host mode for the host group as necessary.

```
raidcom add host_grp -port CL1-A-1 -host_grp_name REMOTE -IH2
raidcom modify host_grp -port CL1-A-1 -host_mode WIN -IH2
```

6. Check the information of the host group set in the port (CL1-A).

```
raidcom get host_grp -port CL1-A -IH2
PORT    GID    GROUP_NAME    Serial#    HMD    HMO_BITs
CL1-A   0      1C-G00        344444     WIN
CL1-A   1      REMOTE       344444     WIN
```

7. Add an LU path to the UR S-VOL. When adding an LU path to UR S-VOL 0x6666, specify host group CL1-A-1 and LU 1.

```
raidcom add lun -port CL1-A-1 -lun_id 1 -ldev_id 0x6666 -IH2
```

8. Check the LU path information defined in host group CL1-A-0.

```
raidcom get lun -port CL1-A-1 -fx -IH2
PORT    GID    HMD    LUN    NUM    LDEV    CM    Serial#
HMO_BITs
CL1-A    1    WIN    1    1    6666    -    344444
```

Updating the CCI configuration definition files

Update the CCI configuration definition file to create a UR delta resync pair. The following shows an example of command execution and the description example of the configuration definition file.

1. Stop the CCI instances at the GAD secondary site and UR secondary site.
   - Server connected to the GAD secondary storage system

     ```
     horcmshutdown 1
     inst 1:
     horcmshutdown inst 1 !!!
     ```

   - Server connected to the UR secondary storage system

     ```
     horcmshutdown 2
     inst 2:
     horcmshutdown inst 2 !!!
     ```

2. Edit the configuration definition files at the GAD secondary site and UR secondary site. Add the underlined parts below to the configuration definition files at the GAD secondary site and UR secondary site.

Note: Make sure to specify the actual LDEV IDs for the GAD pair volumes, not the virtual LDEV IDs.

When specifying the serial number for VSP G1000 using CCI, add a "3" at the beginning of the serial number. For example, for serial number 22222, enter 322222.

- Example of editing the configuration definition file at the GAD secondary site: horcm1.conf

```
HORCM_MON
#ip_address    service    poll(10ms)    timeout(10ms)
localhost        31001    -1        3000

HORCM_CMD
\\.\PhysicalDrive1

HORCM_LDEV
#GRP    DEV    SERIAL    LDEV#    MU#
oraHA    dev1    322222    44:44    h0
oraDELTA    dev3    322222    44:44    h2

HORCM_INST
#GPR    IP ADR        PORT#
oraHA    localhost        31000
oraDELTA    localhost        31002
```

- Example of editing the configuration definition file at the UR secondary site: horcm2.conf

```
HORCM_MON
#ip_address    service    poll(10ms)    timeout(10ms)
localhost        31002    -1        3000

HORCM_CMD
\\.\PhysicalDrive2

HORCM_LDEV
#GRP    DEV    SERIAL    LDEV#    MU#
oraDELTA    dev3    344444    66:66    h2

HORCM_INST
#GPR    IP ADR        PORT#
oraDELTA    localhost        31001
```

3. Start the CCI instances at the GAD secondary site and UR secondary site.

- Server connected to the GAD secondary storage system

```
horcmstart 1
starting HORCM inst 1
HORCM inst 1 starts successfully.
```

- Server connected to the UR secondary storage system

```
horcmstart 2
starting HORCM inst 2
HORCM inst 2 starts successfully.
```

Creating the UR delta resync pairs

Create a UR delta resync pair. The following shows the example of command execution.

Note: Specify a different mirror ID for the GAD pair and the UR pair sharing a volume.

1. Create a UR delta resync pair from the GAD secondary storage system. Specify the group name oraDELTA, journal ID 0 of the GAD secondary site

(delta UR P-VOL), journal ID 0 of the UR secondary site (delta UR S-VOL), and the UR delta resync pair creation option (-nocsus).

```
paircreate -g oraDELTA -f async -vl -jp 0 -js 0 -nocsus -IH1
```

2. Confirm that the UR delta resync pair creation is completed. In CCI, the pair status of P-VOL is displayed as PSUE, and the mirror status of the journal is displayed as PJNS.

   The following figure shows the configuration after UR delta resync pair creation is completed.

```
pairdisplay -g oraDELTA -fxce -IH1
Group    PairVol(L/R)    (Port#,TID, LU),Seq#,LDEV#.P/
S,Status,Fence, %,P-LDEV#    M    CTG    JID    AP    EM    E-
Seq#    E-LDEV#    R/W
oraDELTA    dev3(L)        (CL1-A-1, 0, 1)322222 4444.P-VOL PSUE
ASYNC ,  0 6666    -    0    0    -    -    -    -    -/-
oraDELTA    dev3(R)        (CL1-A-1, 0, 1)344444 6666.S-VOL SSUS
ASYNC ,  0 4444    -    0    0    -    -    -    -    -/-
pairdisplay -g oraDELTA -v jnl -IH1
JID    MU    CTG    JNLS    AP    U(%)    Q-Marker    Q-CNT    D-
SZ(BLK)    Seq#    Num    LDEV#
000    1    1    PJNS    4    21    43216fde    30    512345
62500    1    39321
```

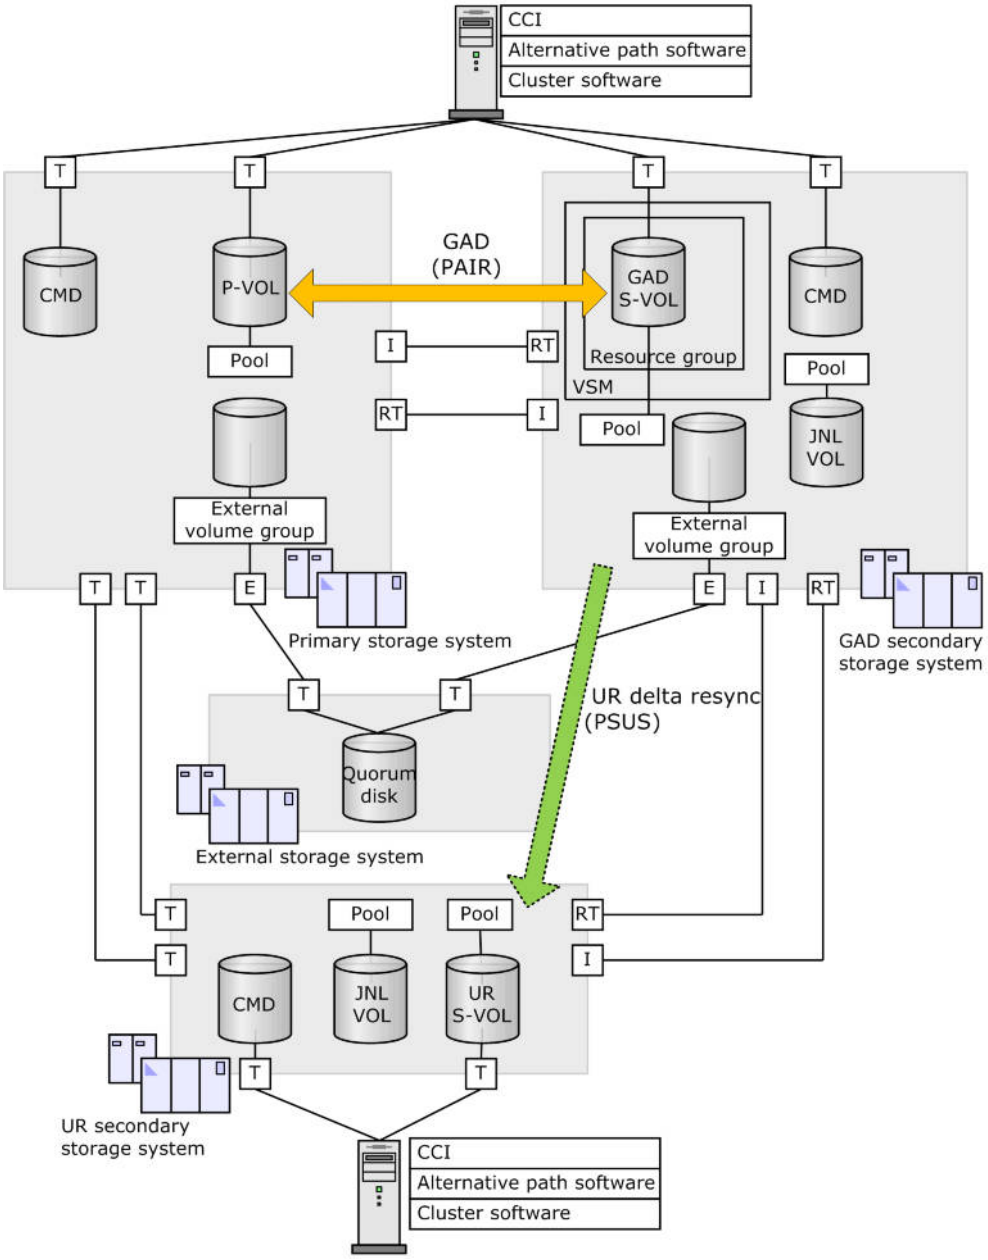

Creating the UR pairs

This section describes the procedure for creating the UR pair in a GAD+UR configuration:

1. Setting the port attributes on page 5-25
2. Adding the remote connections on page 5-25
3. Setting up the primary storage system on page 5-26
4. Updating the CCI configuration definition files on page 5-28
5. Creating the UR pairs on page 5-29

Setting the port attributes

To establish the bi-directional remote connections between the primary storage system and UR secondary storage system, first set the port attributes on each storage system. The following shows the example of command execution.

1. Set the attributes of Initiator port and RCU Target port in the primary storage system and UR secondary storage system.
   - Server connected to the primary storage system

     ```
     raidcom modify port -port CL6-A -port_attribute MCU -IH0
     raidcom modify port -port CL7-A -port_attribute RCU -IH0
     ```

   - Server connected to the UR secondary storage system

     ```
     raidcom modify port -port CL6-A -port_attribute RCU -IH2
     raidcom modify port -port CL7-A -port_attribute MCU -IH2
     ```

   Use the same procedure to change the port attributes for the alternate paths.
2. Check the setting of the port attribute.
   - Server connected to the primary storage system

     ```
     raidcom get port -IH0
     PORT     TYPE    ATTR    SPD    LPID    FAB    CONN    SSW
     SL    Serial#    WWN         PHY_PORT
     (Omitted)
     CL6-A    FIBRE    MCU    AUT    E8    N    FCAL    N    0
     311111    50060e80072b6720        -
     (Omitted)
     CL7-A    FIBRE    RCU    AUT    97    N    FCAL    N    0
     311111    50060e80072b6730        -
     (Omitted)
     ```

   - Server connected to the UR secondary storage system

     ```
     raidcom get port -IH2
     PORT     TYPE    ATTR    SPD    LPID    FAB    CONN    SSW
     SL    Serial#    WWN         PHY_PORT
     (Omitted)
     CL6-A    FIBRE    RCU    AUT    E8    N    FCAL    N    0
     344444    50060e80072b6720        -
     (Omitted)
     CL7-A    FIBRE    MCU    AUT    97    N    FCAL    N    0
     344444    50060e80072b6730        -
     (Omitted)
     ```

Adding the remote connections

After setting the port attributes, add the remote connections between the primary storage system and UR secondary storage system. The following shows the example of command execution.

Note:

- When specifying the serial number for VSP G1000 using CCI, add a "3" at the beginning of the serial number. For example, for serial number 11111, enter 311111.

- To specify the VSP G1000, enter `R800`.
- The VSP G1000 is displayed as `R8` in command output.

1. Add a bi-directional remote connection between the GAD secondary storage system and UR secondary storage system. Add a remote connection whose path group ID is 0 from port CL3-A in the primary storage system to port CL3-A in the UR secondary storage system, and add a remote connection whose path group ID is 0 from port CL3-B in the UR secondary storage system to port CL3-B in the primary storage system. After executing the command, confirm that the asynchronous command processing ended normally.

   - Server connected to the primary storage system

     ```
     raidcom add rcu -cu_free 344444 R800 0 -mcu_port CL6-A -
     rcu_port CL6-A -IH0
     raidcom get command_status -IH0
     HANDLE    SSB1    SSB2    ERR_CNT    Serial#    Description
     00c3    -    -    0    311111    -
     ```

   - Server connected to the UR secondary storage system

     ```
     raidcom add rcu -cu_free 311111 R800 0 -mcu_port CL7-A -
     rcu_port CL7-A -IH2
     raidcom get command_status -IH2
     HANDLE    SSB1    SSB2    ERR_CNT    Serial#    Description
     00c3    -    -    0    344444    -
     ```

   Add the alternate paths between the storage systems using the `raidcom add rcu_path` command.

2. Check the settings of the remote connection.

   - Server connected to the primary storage system

     ```
     raidcom get rcu -cu_free 344444 R800 0 -IH0
     Serial#    ID    PID    MCU    RCU    M/R    T    PNO
     MPORT    RPORT    STS_CD    SSIDs
     344444    R8    0    -    -    RCU    F    0    CL6-A    CL6-
     A    NML_01    -
     ```

   - Server connected to the UR secondary storage system

     ```
     raidcom get rcu -cu_free 311111 R800 0 -IH2
     Serial#    ID    PID    MCU    RCU    M/R    T    PNO
     MPORT    RPORT    STS_CD    SSIDs
     311111    R8    0    -    -    RCU    F    0    CL7-A    CL7-
     A    NML_01    -
     ```

Setting up the primary storage system

On the primary storage system, create a pool volume of HDP pool for journal, create an HDP pool for journal, create a journal volume, and create a journal. The following shows the example of command execution.

1. Create a pool volume. Specify Parity group 2-1 and create a volume whose capacity is 100 GB and whose LDEV ID is 0x9999. After executing the command, confirm that the asynchronous command processing ended normally.

   ```
   raidcom add ldev -ldev_id 0x9999 -parity_grp_id 2-1 -capacity
   100G -IH0
   ```

```
raidcom get command_status -IH0
HANDLE    SSB1    SSB2    ERR_CNT    Serial#    Description
00c3    -    -    0    311111    -
```

2. Perform a quick format of the created volume. After executing the command, confirm that the asynchronous command processing ended normally.

```
raidcom initialize ldev -operation qfmt -ldev_id 0x9999 -IH0
raidcom get command_status -IH0
HANDLE    SSB1    SSB2    ERR_CNT    Serial#    Description
00c3    -    -    0    311111    -
```

3. Create a pool. Specify the created volume 0x9999 in the pool volume and create an HDP pool whose pool ID is 1 and whose pool name is PJNL_POOL. After executing the command, confirm that the asynchronous command processing ended normally.

```
raidcom add dp_pool -pool_id 1 -pool_name PJNL_POOL -ldev_id
0x9999 -IH0
raidcom get command_status -IH0
HANDLE    SSB1    SSB2    ERR_CNT    Serial#    Description
00c3    -    -    0    311111    -
```

4. Confirm that the volume (LDEV ID: 0x9999) is set in the pool volume in the pool (Pool ID: 1).

```
raidcom get ldev -ldev_id 0x9999 -fx -IH0
Serial# : 311111
LDEV : 9999
(Omitted)
F_POOLID : 1
VOL_ATTR : CVS : POOL
(Omitted)
```

5. Check the pool capacity.

```
raidcom get dp_pool -IH0
PID    POLS    U(%)    AV_CAP(MB)    TP_CAP(MB)    W(%)
H(%)    Num    LDEV#    LCNT    TL_CAP(MB)
001    POLN    0    98196    98196    70    80    1
39321    0    0
```

6. Check the pool name.

```
raidcom get pool -key opt -IH0
PID    POLS    U(%)    POOL_NAME    Seq#    Num    LDEV#
H(%)    VCAP(%)    TYPE    PM
001    POLN    0    PJNL_POOL    311111    1    39321    80
-    OPEN    N
```

7. Create a journal volume. Create a virtual volume (DP-VOL) whose capacity is 30 GB and whose LDEV ID is 0x8888 in HDP pool ID 1. After executing the command, confirm that the asynchronous command processing ended normally.

```
raidcom add ldev -pool 1 -ldev_id 0x8888 -capacity 30G -IH0
raidcom get command_status -IH0
HANDLE    SSB1    SSB2    ERR_CNT    Serial#    Description
00c3    -    -    0    311111    -
```

8. Check the created volume.

```
raidcom get ldev -ldev_id 0x8888 -fx -IH0
Serial# : 311111
LDEV : 8888
(omitted)
```

```
VOL_ATTR : CVS : HDP
(omitted)
```

9. Create a journal. Specify the volume whose LDEV ID is 0x8888 to create a journal whose journal ID is 0. After executing the command, confirm that the asynchronous command processing ended normally.

```
raidcom add journal -journal_id 0 -ldev_id 0x8888 -IH0
raidcom get command_status -IH0
HANDLE    SSB1    SSB2    ERR_CNT    Serial#    Description
00c3    -    -    0    311111    -
```

10. Confirm that the volume (LDEV ID: 0x8888) is registered to the journal (journal ID: 0) and check the journal capacity.

```
raidcom get journal -IH0
JID    MU    CTG    JNLS    AP    U(%)    Q-Marker    Q-CNT    D-
SZ(BLK)        Seq#    Num    LDEV#
000    0    1    PJNN    4    21    43216fde    30
62914560        311111    1    34952
```

Updating the CCI configuration definition files

Update the configuration definition file to create a UR pair. The following shows the example of command execution and description example of the configuration definition file.

1. Stop the CCI instances at the primary and UR secondary sites.
   - Server connected to the primary storage system

     ```
     horcmshutdown 0
     inst 0:
     horcmshutdown inst 0 !!!
     ```

   - Server connected to the UR secondary storage system

     ```
     horcmshutdown 2
     inst 2:
     horcmshutdown inst 2 !!!
     ```

2. Edit the configuration definition files at the primary and UR secondary sites. Add the underlined parts below to the configuration definition files at the primary and UR secondary sites.

Note: Make sure to specify the actual LDEV IDs for the GAD pair volumes, not the virtual LDEV IDs.

When specifying the serial number for VSP G1000 using CCI, add a "3" at the beginning of the serial number. For example, for serial number 11111, enter 311111.

   - Example of editing the configuration definition file at the primary site: horcm0.conf

     ```
     HORCM_MON
     #ip_address    service    poll(10ms)    timeout(10ms)
     localhost        31000    -1        3000

     HORCM_CMD
     \\.\PhysicalDrive0

     HORCM_LDEV
     ```

```
#GRP    DEV    SERIAL    LDEV#    MU#
oraHA    dev1    311111    22:22    h0
oraREMOTE    dev2    311111    22:22    h1

HORCM_INST
#GPR    IP ADR        PORT#
oraHA    localhost        31001
oraREMOTE    localhost        31002
```

- Example of editing the configuration definition file at the UR secondary site: horcm2.conf

```
HORCM_MON
#ip_address    service    poll(10ms)    timeout(10ms)
localhost        31004    -1        3000

HORCM_CMD
\\.\PhysicalDrive2

HORCM_LDEV
#GRP    DEV    SERIAL    LDEV#    MU#
oraDELTA    dev3    344444    66:66    h2
oraREMOTE    dev2    344444    66:66    h1

HORCM_INST
#GPR    IP ADR        PORT#
oraDELTA    localhost        31001
oraREMOTE    localhost        31000
```

3. Start the CCI instances at the GAD secondary site and UR secondary site.

   - Server connected to the primary storage system

   ```
   horcmstart 0
   starting HORCM inst 0
   HORCM inst 0 starts successfully.
   ```

   - Server connected to the UR secondary storage system

   ```
   horcmstart 2
   starting HORCM inst 2
   HORCM inst 2 starts successfully.
   ```

Creating the UR pairs

Create a UR pair. The following shows the example of command execution.

1. Create a UR pair. Specify the group name URGPR, journal ID 0 of the GAD secondary site (UR P-VOL), and journal ID 0 of the UR secondary site (UR S-VOL).

   ```
   paircreate -g oraREMOTE -f async -vl -jp 0 -js 0 -IH0
   ```

2. Confirm that UR pair creation is completed.

   ```
   pairdisplay -g oraREMOTE -fxce -IH0
   Group    PairVol(L/R)    (Port#,TID, LU),Seq#,LDEV#.P/
   S,Status,Fence, %,P-LDEV#    M    CTG    JID    AP    EM    E-
   Seq#    E-LDEV#    R/W
   oraREMOTE    dev2(L)        (CL1-A-1, 0, 1)311111 2222.P-VOL
   PAIR ASYNC ,  0 6666    -    0    0    -    -    -    -    -/-
   oraREMOTE    dev2(R)        (CL1-A-1, 0, 1)344444 6666.S-VOL
   PAIR ASYNC ,  0 2222    -    0    0    -    -    -    -    -/-
   ```

The following figure shows the configuration after UR pair creation is completed.

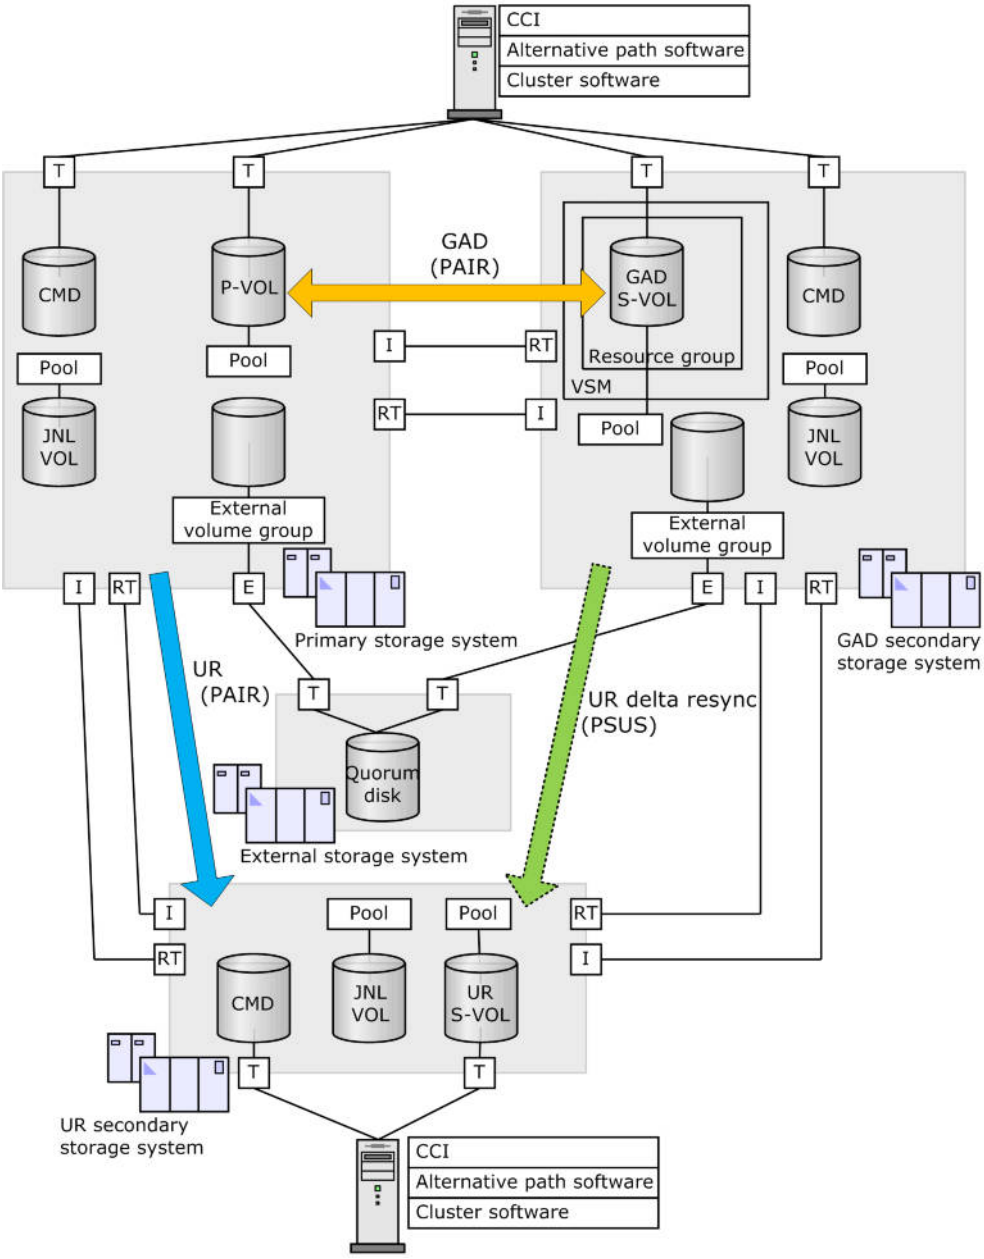

Starting GAD+UR operations

This section describes the procedure for starting operations to have the normal operation status of GAD+UR (UR delta resync pair status changes from PSUE to PSUS) after each pair of GAD+UR is created. After the following procedure, the GAD+UR configuration is ready for normal operation as shown in Figure 5-1 Sample GAD+UR configuration on page 5-2.

1. **Issuing update I/O to P-VOL/GAD S-VOL:** To have the normal GAD +UR operation status, keep updating I/O to P-VOL or GAD S-VOL and wait for two minutes.

   If update I/O is issued to P-VOL or GAD S-VOL, the status information of GAD+UR is reported among the primary, GAD secondary, and UR secondary storage systems. At the time, the UR delta resync pair status is changed from PSUE to PSUS.

2. **Checking UR delta resync pair status and mirror status:** Confirm in the GAD secondary storage system the UR delta resync pair status and the mirror status of the journal. Confirm that the UR delta resync pair status is HOLD. In CCI, the pair status of the P-VOL is displayed as PSUS, and the mirror status of the journal is displayed as PJNS.

```
pairdisplay -g oraDELTA -fxce -IH1
Group    PairVol(L/R)    (Port#,TID, LU),Seq#,LDEV#.P/
S,Status,Fence, %,P-LDEV#    M    CTG    JID    AP    EM    E-
Seq#    E-LDEV#    R/W
oraDELTA    dev3(L)         (CL1-A-1, 0, 1)322222 4444.P-VOL PSUS
ASYNC ,100 6666    -    0    0    -    -    -    -    L/M
oraDELTA    dev3(R)         (CL1-A-1, 0, 1)344444 6666.S-VOL SSUS
ASYNC ,100 4444    -    0    0    -    -    -    -    L/M
pairdisplay -g oraDELTA -v jnl -IH1
JID    MU    CTG    JNLS    AP    U(%)    Q-Marker    Q-CNT    D-
SZ(BLK)    Seq#    Num    LDEV#
000    1    1    PJNS    4    21    43216fde    30    512345
62500    1    39321
```

Note: When checking the pair status in Hitachi Command Suite, make sure to refresh the system (click **Refresh Storage System**) to update the information being displayed. The status of the UR delta resync pairs changes from HOLDING to HOLD.

3. Confirm that no failure SIMs are displayed.

6

GAD pair operations

This chapter describes and provides instructions for performing GAD pair operations using CCI.

Overview of GAD pair operations

The GAD pair operations are:

- Creating GAD pairs on page 6-2
- Suspending GAD pairs on page 6-3
  - Suspending GAD pairs by consistency group
- Resynchronizing GAD pairs on page 6-4
  - Resynchronizing GAD pairs by consistency group
- Deleting GAD pairs on page 6-5
- Managing pairs in a GAD 3DC delta resync environment on page 6-7

Caution:

- Pair operations cannot be performed on volumes that do not have an LU path. Before performing pair operations, make sure that the volumes to be assigned to pairs have at least one LU path defined.
- Pair operations cannot be performed during microcode exchange processing. Before performing pair operations, make sure that microcode exchange processing is complete.
- Pair operations cannot be performed when microcode exchange processing has been interrupted (for example, due to user cancellation or error). Make sure that microcode exchange processing has completed normally before performing pair operations.

Caution: If the following status continues while the GAD pair mirroring, the GAD pair might be suspended to prioritize the update I/O than mirroring of the GAD pair.

- The availability ratio of the processor in the MP blade/unit to which the primary volume belongs is equal to or more than 70% on the storage system at the primary site.
- There is a large amount of inward traffic of update I/O to the primary volumes on the storage system at the primary site.
- The Write Pending of the MP blade/unit to which the secondary volume belongs is equal to or more than 65% on the storage system at the secondary site.

When you create or resynchronize the GAD pair, consider the above load status of the storage system at each site.

Creating GAD pairs

Pair creation copies the data in a volume in the primary storage system to a volume in the secondary storage system. Before a pair is created, the GAD reserve attribute must be applied to the volume that will become the S-VOL.

If you want to use consistency groups to manage pairs, you must assign each GAD pair to the appropriate consistency group when you create the pair. For

details about storage system support (models, microcode) for consistency groups, see Requirements and restrictions on page 2-2.

WARNING: Pair creation is a destructive operation. When a pair is created, the data in the S-VOL is overwritten by the data in the P-VOL. Before you create a pair, you are responsible for backing up the data in the volume that will become an S-VOL.

Prerequisite

- The GAD reserve attribute must be set to the volume to be used as an S-VOL. You can verify that the reserve attribute is set by using the `raidcom get ldev` command. If the reserve attribute is set, the virtual LDEV ID (`VIR_LDEV`) is displayed as `ffff`.
- The P-VOL capacity and S-VOL capacity must be the same size (same number of blocks). To view the capacity in blocks, click **Options > Capacity Unit > block** in the **Logical Devices** window. If the capacity is displayed in GB or TB, a slight difference in P-VOL and S-VOL capacity might not be displayed.

Command example

```
paircreate -g oraHA -f never -vl -jq 0 -IH0
```

Suspending GAD pairs

Pair suspension stops write data from being copied to the S-VOL. When you suspend a pair, you can specify the volume (P-VOL or S-VOL) that will receive update data from the host while the pair is suspended. If you specify the S-VOL, the data written to the S-VOL while the pair is suspended will be copied to the P-VOL when the pair is resynchronized.

When suspending a pair that is registered to a consistency group, you can also suspend all GAD pairs in the consistency group to which that pair is registered. When you suspend GAD pairs by consistency group, the suspend pair operations are completed first, and then the GAD pair statuses change to Suspended. After you perform the pair suspension, confirm that the status of all pairs in the consistency group has changed to Suspended. When many pairs are registered to the consistency group, it might take a few minutes for all pair statuses to change to Suspended. If the pair statuses do not change to Suspended after a few minutes, the pair status transitions might have failed due to some error condition.

For details about storage system support (models, microcode) for consistency groups, see Requirements and restrictions on page 2-2.

Note: When you suspend GAD pairs by consistency group, the consistency group status is displayed as Suspending on Device Manager - Storage Navigator.

Command example

```
pairsplit -g oraHA -r -IH0
```

Enabling the ALUA mode of an existing GAD pair (VSP G1000 only)

If you need to enable the ALUA mode of an existing GAD pair, you must suspend the pair, enable the ALUA mode, restart the server while the pair is suspended, and then resynchronize the pair. When you resynchronize the pair, the ALUA mode of the S-VOL is also enabled.

Caution: This procedure requires the host server to be restarted to recognize the new ALUA mode setting.

Prerequisite

- The ALUA mode can be used on a server.

Procedure

1. Suspend the GAD pair.
   ```
   pairsplit -g oraHA -r -IH0
   ```
2. Enable the ALUA mode.
   ```
   raidcom modify ldev -ldev_id 0x2222 -alua enable
   ```
3. Restart the server.

Note: If you do not restart the server, the server might not be able to recognize the ALUA mode setting.

4. Resynchronize the GAD pair.
   ```
   pairresync -g oraHA -IH0
   ```

Resynchronizing GAD pairs

Pair resynchronization updates the S-VOL (or P-VOL) by copying the differential data accumulated since the pair was suspended. The volume that was not receiving update data while the pair was suspended is resynchronized with the volume that was receiving update data. When resynchronization completes, the host can read from and write directly to the P-VOL or the S-VOL.

When resynchronizing a pair that is registered to a consistency group, you can also resynchronize all GAD pairs in the consistency group to which that pair is registered. When resynchronizing a pair that is not registered to a consistency group, you can register that pair to a consistency group.

When resynchronizing GAD pairs by consistency group, the resynchronize pair operations are completed first, and then the GAD pair statuses change to Mirroring or Mirrored. After you perform the pair resynchronization, confirm

that the status of all pairs in the consistency group has changed to Mirroring or Mirrored. When many pairs are registered to the consistency group, it might take a few minutes for all pair statuses to change to Mirroring or Mirrored. If the pair statuses do not change to Mirroring or Mirrored after a few minutes, the pair status transitions might have failed due to some error condition.

For details about storage system support (models, microcode) for consistency groups, see Requirements and restrictions on page 2-2.

Note: When you resynchronize GAD pairs by consistency group, the consistency group status is displayed as Resynchronizing on Device Manager - Storage Navigator.

Prerequisite

- The GAD pair must be suspended (the status of both pair volumes must be PSUE). For instructions on resynchronizing a pair when the S-VOL status is PAIR, see Resynchronizing a pair when the S-VOL status is PAIR (VSP G1000 only) on page 6-5.

Command example

```
pairresync -g oraHA -IH0
```

Resynchronizing a pair when the S-VOL status is PAIR (VSP G1000 only)

To resynchronize a pair, the status of both pair volumes must be PSUE. If the P-VOL pair status is PSUE but the S-VOL pair status is PAIR (for example, due to quorum disk blockade), the resynchronize operation will fail. To resynchronize a pair when the P-VOL status is PSUE and the S-VOL status is PAIR:

1. Suspend the GAD pair by specifying the S-VOL (swap suspend). The S-VOL pair status changes to PSUE.
2. Resynchronize the GAD pair by specifying the P-VOL. The P-VOL and the S-VOL pair statuses change to PAIR.

Deleting GAD pairs

Pair deletion deletes the pair relationship between the P-VOL and the S-VOL. The data in each volume is not affected. When you delete a pair, you can specify the volume (P-VOL or S-VOL) that will receive update data from the host after the pair is deleted. The virtual LDEV ID of the unspecified volume is deleted, and the GAD reserve attribute is set for the specified volume.

The following table specifies the required conditions for the volume that will continue to receive update data from the host after pair deletion.

Volume to receive I/O after pair deletion	Required conditions
P-VOL	• Pair status: PSUS or PSUE • I/O mode: Local
S-VOL	• Pair status: SSWS • I/O mode: Local

After the GAD pair is deleted, the data on the P-VOL and S-VOL is not synchronized. To prevent viewing a duplicated volume with the same virtual LDEV ID but asynchronous data on the server, the virtual LDEV ID of the LDEV that does not continue I/O is deleted. When the virtual LDEV ID is deleted and the GAD reserve attribute is assigned to the volume, the server cannot recognize the volume.

If you want to re-create a GAD pair using a volume that was deleted from a pair, re-create the GAD pair from the storage system with the volume that was specified when you deleted the GAD pair. For example, if you deleted the GAD pair by specifying the P-VOL, re-create the GAD pair from the primary storage system. If you deleted the GAD pair by specifying the S-VOL, re-create the GAD pair from the secondary storage system.

Prerequisite

- The GAD pair must be suspended.

Command example

```
pairsplit -g oraHA -S -IH0
```

Changing preferred path settings (VSP G1000 only)

You can change the asymmetric access status setting required for specifying a preferred path regardless of the GAD pair status, even during I/O processing.

Note: Depending on the operating system of the server to which the storage system is connected, the asymmetric access status settings might not be recognized. If a path is not connected according to the asymmetric access status settings, the operating system of the server might not recognize the asymmetric access status. If this occurs, make the server recognize the device again.

Command example

```
raidcom modify lun -port CL1-A-0 -lun_id all -
asymmetric_access_state non_optimized
```

Managing pairs in a GAD 3DC delta resync environment

This section describes GAD and UR pair operations in a GAD 3DC delta resync (GAD+UR) environment.

For details about storage system support (models, microcode) for GAD+UR operations, see Requirements and restrictions on page 2-2.

Executing delta resync

When you specify an S-VOL and suspend (swap suspend) a GAD pair, the GAD S-VOL pair status changes from PAIR to SSWS. After the pair status changes to SSWS, the UR delta resync pair changes to a UR pair, and the copying from the GAD S-VOL to the UR S-VOL starts. This change in the copy source of the UR pair is the delta resync operation.

For details about storage system support (models, microcode) for GAD+UR operations, see Requirements and restrictions on page 2-2.

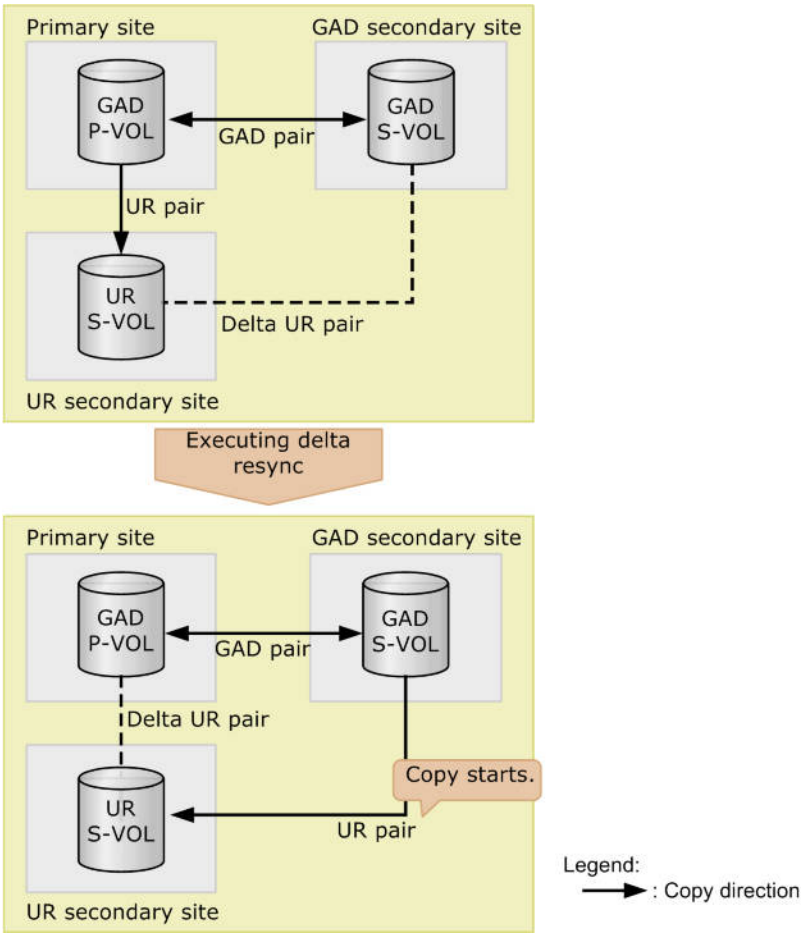

Prerequisites

- Pair status and mirror status must be as follows:

Pair type	Pair status		Mirror status	
	P-VOL	S-VOL	Master journal	Restore journal
GAD pair	PAIR	PAIR	Not applicable	Not applicable
UR pair	PAIR	PAIR	PJNN	SJNN
UR delta resync pair	PSUS*	SSUS*	PJNS	SJNS
* If you use Device Manager - Storage Navigator to suspend (swap suspend) the pair, make sure the status of the UR delta resync pair is HOLD.				

- All the differential data of the UR pairs is stored in the primary site's journal.
- Data must be able to be transferred between the UR P-VOL and S-VOL.
- The number of S-VOLs in the UR pairs is the same as that in the UR delta resync pairs.
- There is no failure in the remote path between the secondary site of GAD and the secondary site of UR.
- After creating the UR pair, you must keep updating I/O from the server to the GAD pair's P-VOL or S-VOL for about two minutes.

Command example

```
pairsplit -g oraHA -RS -IH1
```

Notes on delta resync

- When a UR pair has not been suspended and resynchronized for a long time, the data in the restore journal might exceed 70% of capacity. If this happens, old journal data is automatically deleted. In this case, the P-VOL and S-VOL are not synchronized completely by just copying the data, and delta resync will fail. In case of delta resync failure, resynchronize the UR pair.
- Journal data might be deleted in the following cases, even if the data in the restore journal does not exceed 70% of capacity:
  - When you update the P-VOL after resynchronizing the GAD pair.
  - When you update the P-VOL after resynchronizing the UR pair between the primary site and UR secondary site.
  - When retry-processing occurs because of a delay of the P-VOL update.
  - When the update of a GAD S-VOL is delayed.
- If the pair status of the UR delta resync pair does not change after the delta resync operation, the prerequisites for delta resync might not be satisfied. Review the prerequisites for the pair status of the GAD pair, UR pair, and UR delta resync pair.

In case of delta resync failure

If delta resync fails, the UR delta resync pair changes to a UR pair. The status of each pair changes as follows:

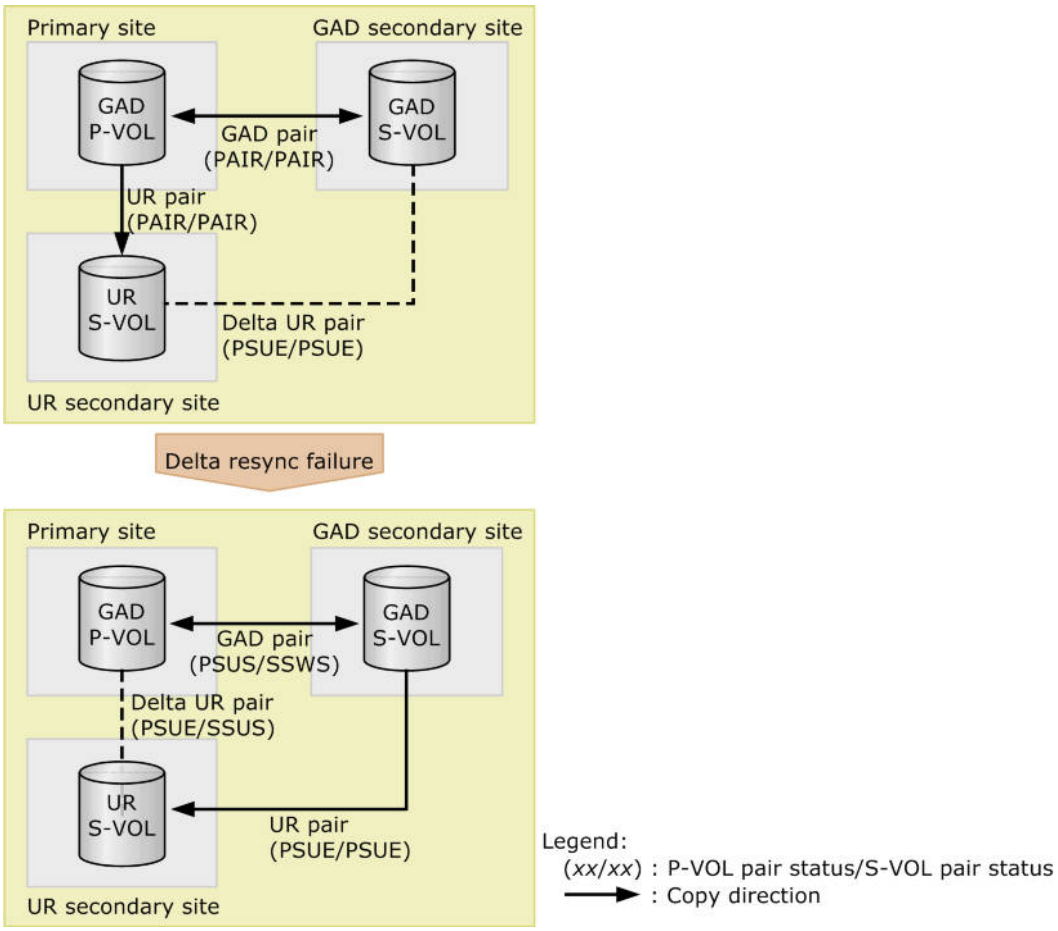

In case of delta resync failure, confirm the following two conditions:

- The system requirements for the GAD+UR configuration are met.
- The system has no failure causes.

If you resynchronize the UR pair after delta resync failure, the initial copy is performed for the GAD pair's S-VOL data to the UR pair's S-VOL.

Note: In a GAD+UR configuration, the data is not copied automatically even if you specify **Entire Copy** for **Delta Resync Failure** in the **Edit Mirror Options** window.

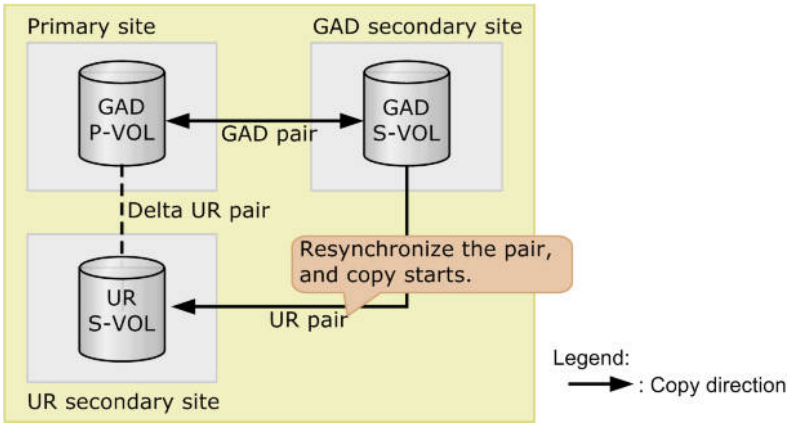

Resynchronizing GAD pairs in a GAD 3DC delta resync environment

To resynchronize a GAD pair by specifying the S-VOL (swap resync), the conditions specified below must be met.

For details about storage system support (models, microcode) for GAD+UR operations, see Requirements and restrictions on page 2-2.

Prerequisites

- Pair status must be as follows:

Pair type	Pair status	
	P-VOL	S-VOL
GAD pair	PSUS	SSWS
UR pair	PSUE or PAIR	PSUE or PAIR
UR delta resync pair	No condition	No condition

- The UR pair whose volume shared by GAD S-VOL after swap resync must be a UR delta resync pair.

Command example

```
pairresync -g oraHA -swaps -IH1
```

Deleting GAD pairs in a GAD 3DC delta resync environment

For details about storage system support (models, microcode) for GAD+UR operations, see Requirements and restrictions on page 2-2.

1. Delete the UR pair.

Note: If you accidentally delete the UR delta resync pair in this step before deleting the UR pair, the UR pair might be suspended by failures.

2. Delete the UR delta resync pair.
3. Suspend the GAD pair.

4. Delete the GAD pair.

If you need to delete a GAD pair forcibly, first delete the UR pair and the UR delta resync pair, and then delete the GAD pair forcibly.

7

Using Hitachi Command Suite to set up global-active device

This chapter describes how to use Hitachi Command Suite to set up global-active device.

Setting up a global-active device environment

After the HCS management server has discovered resources, you can use the **Set up Global-Active Device** window for setup tasks that must be performed before global-active device pairs can be allocated.

About global-active device setup

The global-active device setup window provides a workflow for configuring required storage system resources before global-active device paired volumes can be allocated.

Access to the global-active device setup window is provided from:

- The **Actions** menu on the global task bar
- The **Replication** tab, from General Tasks

The setup window provides access to various HCS or Replication Manager dialog boxes. As tasks are completed, or if tasks have already been completed by other procedures, the setup window identifies completed tasks and provides a link to configuration details. If you close the setup window and reopen it later, the completed tasks remain identified.

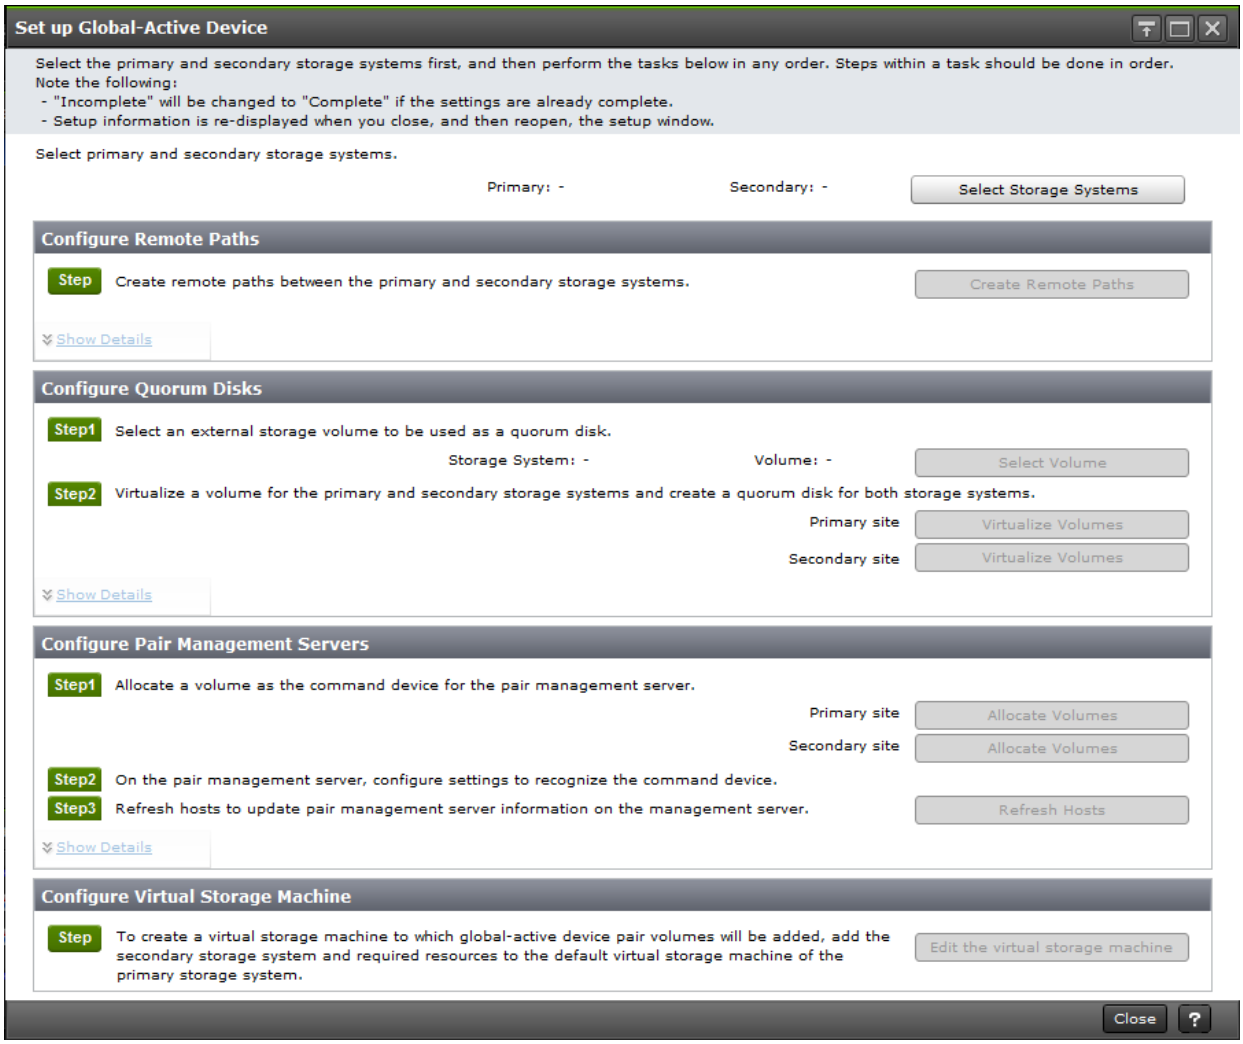

The following figure depicts the setup window configuration items in the global-active device environment:

- Select Storage Systems (see ① below). This task must be done first. The primary and secondary storage systems can be VSP G1000 or VSP G200, G400, G600, G800 storage systems.

The following setup tasks can be performed in any order, but must be completed before you can allocate global-active device pairs.

- Configure Remote Paths (see ②)
- Configure Quorum Disks (see ③)
- Configure Pair Management Servers (see ④)
- Configure Virtual Storage Machines (see ⑤)

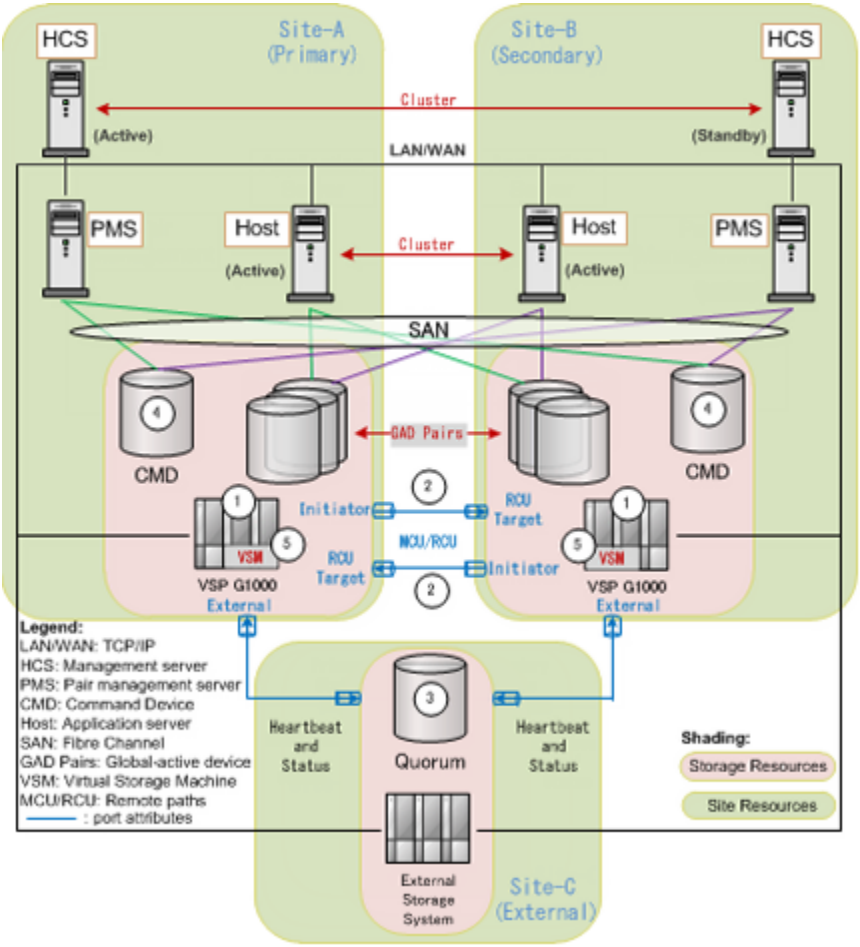

Related topics

- Prerequisites for global-active device setup using HCS on page 7-4
- Select primary and secondary storage systems on page 7-5
- Configuring remote paths on page 7-6
- Configuring the quorum disk on page 7-7
- Configuring pair management servers on page 7-9
- Configuring the virtual storage machine on page 7-11

Prerequisites for global-active device setup using HCS

Before performing global-active device setup tasks using HCS, ensure the following prerequisites are met:

- The Virtual Storage Platform G1000 microcode version for both the primary and secondary storage systems must be 80-02-01 or later to create global-active device pairs where the P-VOLs already have a virtual LDEV ID assigned and are managed by user-defined virtual storage machines. This requirement is not necessary for global-active device pairs created by adding secondary storage resources to the primary storage default virtual storage machine.

- Global-active device has been licensed (which includes Replication Manager).
- Replication Manager is installed and linked to HCS.
- A pair management server is installed with Device Manager Agent and Command Control Interface (CCI) at the primary and secondary storage sites.
- The external ports of the primary and secondary storage systems are connected to the external (quorum) storage system.
- You have registered (discovered) the primary, secondary, and external storage systems using HCS. If you are configuring quorum disks by using virtualized volumes as external volumes, registration of the quorum storage system is not necessary.

Related topics

- About global-active device setup on page 7-2

Select primary and secondary storage systems

Select primary and secondary storage systems listed by the global-active device setup. Only valid storage systems with the correct firmware revision are listed.

Procedure

1. Open the global-active device setup window:
   - From the **Actions** menu, select **Set up Global-Active Device**.
   - On the **Replication** tab, **General Tasks**, select **Set up Global-Active Device**.
2. On the **Set up Global-Active Device** window, click **Select Storage Systems**.
3. From the **Select Storage Systems** window, select the primary storage system from the list of available storage systems.
4. Select the secondary storage system from the list of available storage systems.
5. If given the option, select a default or user-defined virtual storage machine on the primary storage system in which global-active device pairs will be created.

Note: The virtual storage machine list for the primary storage system is displayed only if the microcode version of the VSP G1000 is 80-02-01 or later and the firmware version of the VSP G200, G400, G600, G800 is 83-01-2x or later.

6. Confirm the selected storage systems and virtual storage machine information are correct and click **OK**.
   The selected information is displayed on the setup window as confirmation.

Result

The storage systems and virtual storage machine are selected, and become default values for the remaining setup tasks.

Related topics

- About global-active device setup on page 7-2

Configuring remote paths

Configure remote paths between the primary and secondary storage systems.

Prerequisites

The primary and secondary storage systems have been selected.

Procedure

1. Open the global-active device setup window:
   - From the **Actions** menu, select **Set up Global-Active Device**.
   - On the **Replication** tab, **General Tasks**, select **Set up Global-Active Device**.
2. From **Configure Remote Paths**, click **Create Remote Paths**.
   The **Create Remote Path** wizard opens.
3. On the **Select Remote Storage System** page:
   a. Verify that **Copy Type** is set to **global-active device** .
   b. Verify the displayed **Primary Storage System** and **Secondary Storage System** information, and click **Next** to define paths.

Tip: The MCU/RCU check box is disabled by design for the global-active device dialog.

4. On the **Define Remote Path** page:
   a. Ensure that the **Select reverse direction path** check box is selected.
   b. (Optional) Enter a **Label** for identifying the remote paths for global-active device.
   c. (Optional) Enter a non-default **Path group ID**.
   d. Using the **Local Port** and **Remote Port** lists, select the Initiator and RCU target ports from the **Existing Ports** list for both the primary and secondary storage systems.

Tip: The port list on the left shows primary to secondary storage ports, and the port list on the right shows secondary to primary storage ports. The ports in the lists are unallocated host ports, which are target ports (port attribute of target) that are not in use by any hosts.

e. Click **Next** to review the remote path settings.

5. On the **Confirm** page:

   a. Under **Initiator/RCU Targets**, review selected ports and port attribute settings (before and after).

   b. In the **Remote Path** area, check the copy direction for the local and remote port pairs. If the port pairs are correct, select the confirmation message check box, and click **Confirm** to set the port attributes (Initiator or RCU Target).

Tip: Port pairs on the left (Local -> Remote) are primary to secondary storage remote paths. Port pairs on the right (Local <- Remote) are the secondary to primary storage remote paths.

Note: For more information about remote paths, see the *Hitachi Command Suite Replication Manager User Guide*.

Result

The remote paths between the primary and secondary storage systems are configured, and can be confirmed by clicking **Show Details**.

Related topics

- About global-active device setup on page 7-2

Configuring the quorum disk

The quorum disk detects the status of the primary and secondary storage systems. Configure the quorum disk by selecting a volume and virtualizing it as an external volume of the primary and secondary storage systems.

You can also configure the quorum disk by using a virtualized volume as an external volume of the primary or secondary storage system.

Quorum disk configuration and supported storage systems

The following table shows whether the indicated storage system supports the quorum disk settings in the initial setup dialog box for the global-active device environment.

Task	VSP G1000		VSP G200, G400, G600, G800
Quorum disk configuration	80-01-4x or later Earlier than 80-02-2x	80-02-2x or later	83-01-2x
Create a new external volume	Yes	Yes	Yes*

Task	VSP G1000		VSP G200, G400, G600, G800
Use an existing external volume	No	Yes	Yes
* The external path using iSCSI is not supported.			

Prerequisites

- The primary and secondary storage systems have been selected.
- You can create a quorum disk by using an external volume that connects using iSCSI. However, the initial settings of the HCS global-active device environment support only the display of quorum disk information and the completion (configuration) checks. Quorum disk creation is not supported. For this reason, to specify an iSCSI-connected external volume for the quorum disk, in Device Manager - Storage Navigator, create an external volume that connects using iSCSI. Then, in Configure Quorum Disks in the initial setup dialog box for the HCS global-active device environment, select the volume that was created in Device Manager - Storage Navigator.

Procedure

1. Open the global-active device setup window:
   - From the **Actions** menu, select **Set up Global-Active Device**.
   - On the **Replication** tab, **General Tasks**, select **Set up Global-Active Device**.
2. When the VSP G1000 microcode version for both the primary and secondary storage systems is 80-02-2*x* or later and the VSP G200, G400, G600, G800 firmware version for both the primary and secondary storage systems is 83-01-2*x* or later, from **Configure Quorum Disks**, do one of the following:
   - Click **Create a new external volume** to virtualize a volume as an external volume, and then continue with step 3.
   - Click **Configure an existing external volume** to use a virtualized volume as an external volume, and then click **Configure Quorum Disks** for the **Primary Site** and **Secondary Site**. On the **Quorum Disk** tab, click **Add Quorum Disks**, and then follow the instructions in the Help.
3. In the **Select Volume** window, select a **Storage system** to provide the quorum disk, and then do one of the following:
   - From the **Available Volumes** list, select a volume, click **OK**, and then go to step 12.
   - If no suitable volume is listed, click **Create Volumes**.
4. In the **Create Volumes** window, create a volume to be used as the quorum disk.

5. From the **Volume Type** list select a volume type.
6. (Optional) Click **Advanced Options** to configure volume options or specify a quorum disk label.
7. Click **Show Plan** and confirm that the information in the plan summary is correct. If changes are required, click **Back**.
8. (Optional) Update the task name and provide a description.
9. (Optional) Expand **Schedule** to specify the task schedule.
   You can schedule the task to run immediately or later. The default setting is **Now**.
10. Click **Submit.**
    If the task is scheduled to run immediately, the process begins.
11. (Optional) Check the progress and result of the task on the **Tasks & Alerts** tab. Click the task name to view details of the task.
12. From **Configure Quorum Disks**, for the primary storage, click **Virtualize Volumes**.
13. In the **Virtualize Volumes** dialog box:
    a. Verify the external and internal storage systems are correct.
    b. Click **Edit External Paths** to set or change external paths.
    c. (Optional) Select **External Path Priority Settings** to increase or decrease path priority.
    d. Select **Internal Volume Configuration** and verify that **Change Internal Volume(s) to Quorum Disk(s)** is selected by default. The default ID should be used, but can be changed to another value if necessary.
14. Click **Show Plan** and submit the task.
15. Repeat steps 12 through 14 for the secondary storage system.
    For the secondary storage system, the secondary storage system name is displayed.

Result

The quorum disk is created and virtualized for both the primary and secondary storage systems and can be confirmed by clicking **Show Details**.

Related topics

- About global-active device setup on page 7-2

Configuring pair management servers

Configure the command devices for the pair management servers from the primary and secondary storage systems. Command devices are used to communicate pair management directives.

Prerequisites

The primary and secondary storage systems have been selected.

Procedure

1. Open the global-active device setup window:
   - From the **Actions** menu, select **Set up Global-Active Device**.
   - On the **Replication** tab, **General Tasks**, select **Set up Global-Active Device**.
2. From **Configure Pair Management Server**, click **Allocate Volumes** for the primary storage system.
3. In the **Allocate Volumes** dialog box, verify that the **Storage System** is correct.
4. For **Host**, select the pair management server for the storage system.
5. Click **Advanced Options**, verify that **Command Device** is checked, and that **User Authentication** is **Enabled**.
6. Click **Show Plan** and confirm that the information in the plan summary is correct. If changes are required, click **Back**.
7. (Optional) Update the task name and provide a description.
8. (Optional) Expand **Schedule** to specify the task schedule.
   You can schedule the task to run immediately or later. The default setting is **Now**.
9. Click **Submit.**
   If the task is scheduled to run immediately, the process begins.
10. (Optional) Check the progress and result of the task on the **Tasks & Alerts** tab. Click the task name to view details of the task.
11. Repeat steps 2 through 10 for the secondary storage system. For **Host**, be sure to specify the pair management server for the secondary storage system.
12. Ensure that the pair management servers recognize the command devices, and then from **Configure Pair Management Server**, click **Refresh Hosts**.

Result

The command devices for the primary and secondary pair management servers are allocated, and pair management server information is current, and can be confirmed by clicking **Show Details**.

Related topics

- About global-active device setup on page 7-2

Configuring the virtual storage machine

To support global-active device pairs, add secondary storage system resources to the default or user-defined virtual storage machine on the primary storage system. Configuring the virtual storage machine provides a single view of the primary and secondary storage system resources.

Prerequisites

The primary and secondary storage systems have been selected in the **Set up Global-Active Device** window.

Procedure

1. Open the global-active device setup window:
   - From the **Actions** menu, select **Set up Global-Active Device**.
   - On the **Replication** tab, **General Tasks**, select **Set up Global-Active Device**.
2. From **Configure Virtual Storage Machine**, click **Edit Virtual Storage Machine**.
3. In the **Edit Virtual Storage Machine** dialog box, on the **Storage Systems** tab, verify that the displayed **Name** and **Virtual Model and Serial No.** show the primary storage system values and that the primary storage system is listed.
4. Click **Add Storage Systems**, select the secondary storage system, and then click **OK**.
5. Add secondary storage system resources, such as **Parity Groups**, **LDEV IDs**, **Storage Ports** and **Host Group/iSCSI Target Numbers** to the virtual storage machine by clicking the respective tabs.

Note: The primary storage system resources already belong to the virtual storage machine.

6. When you are finished adding resources from the secondary storage system, enter an optional task description, select whether you want to view the progress of the task, and then click **Submit**.

Result

The storage system resources are configured for the virtual storage machine.

Related topics

- About global-active device setup on page 7-2

Monitoring and managing global-active device pairs

This module describes the required tasks to allocate, monitor, and unallocate global-active device pairs for hosts.

Allocating global-active device pairs

Allocate new global-active device pairs to hosts (application servers). When you allocate global-active device pairs, an initial copy is performed for the primary and secondary volumes.

Prerequisites

You must be authenticated to the storage system and logged on to the pair management server. For details, see the *Hitachi Command Suite Administrator Guide*.

Procedure

1. On the **Resources** tab, **General Tasks**, select **Allocate Volumes**.

Tip: If you do not see **Allocate Volumes** listed, click **more...** to see all menu items.

2. In the **Allocate Volumes** dialog box, select the hosts for the global-active device pair.

Note: File-servers are not supported for global-active device pair allocations.

3. For **Allocation Type** select **global-active device**.
4. Specify the number and capacity of the volumes.
5. On the **Primary** tab, configure the following options:
   a. Select the primary **Storage system**, **Virtual Storage Machine**, and **Volume type**.
      Only valid storage systems and virtual storage machine information is displayed. For VSP G200, G400, G600, G800 and for VSP G1000 microcode 80-02-4x or earlier, you can specify only HDP volumes as the volume type. If you specify HDP or Basic as the volume type, select the **Volume location**. If you specify HDT as the volume type, the **Volume Selection** options are available.
   b. Click **Advanced Options** for further configuration.
   c. In **LUN Path Options**:
      Edit LUN paths from primary storage ports to hosts.
   d. Click **Host Group and LUN settings** for further configuration.
   e. In **Pair management Server Settings**:
      Select an available host as the pair management server. For Instance ID, click **Existing** to use a current configuration definition file and related instance ID number on the pair management server, or click **New** to create a new configuration definition file, and enter a new instance ID number which will be validated as an available instance ID. You can also specify a **UDP Port** number for the instance.
6. On the **Secondary** tab, configure the following options:

a. Select the secondary **Storage system**, **Virtual Storage Machine**, and **Volume type**.

   Virtual storage machine information defaults to the primary tab value.

b. In **Advanced options**:

   When you select Basic volumes on the Primary tab, the volumes that correspond to the settings specified on the Primary tab are displayed on the Secondary tab. Unlike the Primary tab, the Secondary tab allows you to select secondary storage system HDP pools (Pools for S-VOL) that correspond to primary storage system HDP pools configured on the primary tab (Pools for P-VOL). For example, automatic selection of HDP volumes will display primary HDP pools, and eligible secondary HDP pools for volume allocation. You can examine primary tab HDP pool and volume details with links. You can accept the HDP pools selected for secondary volumes, or make a manual selection by clicking **Select Pool** to list and select eligible HDP pools, and selecting eligible volumes from the HDP pools.

c. In **LUN Path Options**:

   Edit LUN paths from secondary storage ports to hosts. Under **LUN Path Setting Information on Primary Site**, you will see the primary storage system, and No. of LUN Paths per volume. Clicking the link will display primary storage system mapping information.

d. Click **Host Group and LUN settings** for further configuration.

e. In **Pair management Server Settings**:

   Select an available host as the pair management server. For Instance ID, click **Existing** to use a current configuration definition file and related instance ID number on the pair management server, or click **New** to create a new configuration definition file, and enter a new instance ID number which will be validated as an available instance ID. You can also specify a **UDP Port** number for the instance.

7. Click **Virtual ID Settings** for further configuration.

   These settings are displayed only in a configuration where options other than the default virtual storage machine are used on the primary side.

8. In **Pair Settings**:

   a. Select the **Quorum Disk**, and then select whether pairs should belong to a new or existing **Copy Group** for management purposes. Click **New** to assign a new copy group to the GAD pair. Select the **CTG ID** check box (VSP G1000 only) to assign a consistency group (CTG) ID to the new copy group. Click **Existing** to assign an existing copy group to the GAD pair. The existing CTG ID will be used if it exists for the selected copy group.

Note: (VSP G1000 only) To add the GAD 3DC delta resync configuration to allocated GAD pair volumes, CTG IDs must be assigned.

   b. Manually enter a name (Prefix and Start Sequence No.), or select **Automatic** for **Pair Name** and a name will be created.

9. Click **Show Plan** and confirm that the information in the plan summary is correct. If changes are required, click **Back**.
10. (Optional) Update the task name and provide a description.
11. (Optional) Expand **Schedule** to specify the task schedule.
    You can schedule the task to run immediately or later. The default setting is **Now**.
12. Click **Submit.**
    If the task is scheduled to run immediately, the process begins.
13. (Optional) Check the progress and result of the task on the **Tasks & Alerts** tab. Click the task name to view details of the task.
14. Confirm the global-active device pair is allocated.

Note: If you change the ALUA attribute of a volume for which a LUN path is set, you must rediscover the volume on the host.

Result

The global-active device pair is allocated to the target host, and initial copy is performed. For global-active device pairs that use DP pool volumes, reclaiming zero pages is also automatically performed.

Related topics

- Allocating global-active device pairs based on like volumes on page 7-14
- Allocating global-active device pairs using allocated volumes on page 7-17

Allocating global-active device pairs based on like volumes

Select an existing allocated global-active device volume, and allocate a global-active device pair using the characteristics of the selected primary or secondary volume.

Prerequisites

You must be authenticated to the storage system and logged on to the pair management server. For details, see the *Hitachi Command Suite Administrator Guide*.

Procedure

1. On the **Resources** tab, select a global-active device volume from the **Open-Allocated** volume list, or for a specific host using global-active device paired volumes, select an already allocated volume to use in a new global-active device pair, and click **Allocate Like Volumes**.

Tip: Regardless of whether the primary volume (P-VOL) or the secondary volume (S-VOL) is selected as the existing volume, the settings for the existing P-VOL are used for the newly allocated primary volume and the

settings for the existing S-VOL are used for the newly allocated secondary volume.

However, if characteristics such as drive type and drive speed of the existing volumes do not match between the primary and secondary storage systems, you should use a secondary volume whose characteristics match those of the primary volume.

2. In the **Allocate Like Volumes** dialog box, verify the information for the selected volume and make any necessary changes.
3. On the **Primary** tab, configure the following options:
   a. Select the primary **Storage system**, **Virtual Storage Machine**, and **Volume type**.

      Only valid storage and virtual storage machine information is displayed. If you specify HDP or Basic as the volume type, select the **Volume location**. If the volume type is HDT, the **Volume Selection** options are displayed.
   b. Click **Advanced Options** for further configuration.
   c. In **LUN Path Options**:

      Edit LUN paths from primary storage ports to hosts.
   d. Click **Host Group and LUN settings** for further configuration.
   e. In **Pair Management Server Settings**:

      Select an available host as the pair management server. For Instance ID, click **Existing** to use a current configuration definition file and related instance ID number on the pair management server, or click **New** to create a new configuration definition file, and enter a new instance ID number which will be validated as an available instance ID. You can also specify a **UDP Port** number for the instance.
4. On the **Secondary** tab, configure the following options:
   a. Select the secondary **Storage system**, **Virtual Storage Machine**, and **Volume type**:

      Virtual storage machine information defaults to the primary tab value.
   b. In **Advanced Options**:

      When you select Basic volumes on the Primary tab, the volumes that correspond to the settings specified on the Primary tab are displayed on the Secondary tab. Unlike the Primary tab, select HDP pools for the secondary storage system (Pools for S-VOL) that correspond to the HDP pools configured on the primary tab (Pools for P-VOL). For example, automatic selection of HDP volumes will display primary HDP pools, and eligible secondary HDP pools for volume allocation. You can examine primary tab HDP pool and volume details with links. You can accept the HDP pools selected for secondary volumes, or make a manual selection by clicking **Select Pool** to list and select eligible HDP pools, and selecting eligible volumes from the HDP pools.
   c. In **LUN Path options**:

      Edit LUN paths from secondary storage ports to hosts. **LUN Path Setting Information on Primary Site** lists the primary storage

system, and No. of LUN Paths per volume. Click the link to display primary storage system mapping information.

d. Click **Host Group and LUN settings** for further configuration.

e. Click **Pair Management Server Settings** for further configuration.

5. Click **Virtual ID Settings** for further configuration.

   These settings display only when options other than the default virtual storage machine are used on the primary storage.

6. In **Pair Settings**:

   a. Select the **Quorum Disk**, and then select whether pairs should belong to a new or existing **Copy Group** for management purposes. Click **New** to assign a new copy group to the GAD pair. Select the **CTG ID** check box (VSP G1000 only) to assign a consistency group (CTG) ID to the new copy group. Click **Existing** to assign an existing copy group to the GAD pair. The existing CTG ID will be used if it exists for the selected copy group.

   Note: (VSP G1000 only) To add the GAD 3DC delta resync configuration to allocated GAD pair volumes, CTG IDs must be assigned.

   b. Manually enter a name (Prefix and Start Sequence No.), or select **Automatic** for **Pair Name** and a name will be created.

7. Click **Show Plan** and confirm that the information in the plan summary is correct. If changes are required, click **Back**.
8. (Optional) Update the task name and provide a description.
9. (Optional) Expand **Schedule** to specify the task schedule.

   You can schedule the task to run immediately or later. The default setting is **Now**.

10. Click **Submit.**

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

11. (Optional) Check the progress and result of the task on the **Tasks & Alerts** tab. Click the task name to view details of the task.
12. Confirm the global-active device paired volumes are allocated.

Result

The global-active device pair is allocated to the target host, and initial copy is performed. For global-active device pairs that use DP pool volumes, reclaiming zero pages is also automatically performed.

Related topics

- Allocating global-active device pairs on page 7-12
- Allocating global-active device pairs using allocated volumes on page 7-17

Allocating global-active device pairs using allocated volumes

Existing open-allocated volumes in the same storage system belonging to the same host group (with the same LUN paths) can be converted to P-VOLs, and the P-VOL data migrated to S-VOLs to create a global-active device pair.

Prerequisites

You must be authenticated to the storage system and logged on to the pair management server. For details, see the *Hitachi Command Suite Administrator Guide*.

Procedure

1. On the **Resources** tab, or by searching, locate allocated volumes to convert into global-active device pairs.
2. Select the volume, and click **Change to Global-Active Device Volumes**.
3. On the **Primary** tab, configure the following options:
   a. Verify the displayed hosts, storage systems virtual storage machines for the selected volumes.
   b. **Selected Volumes** are allocated volumes eligible to become a P-VOL in a global-active device pair. Details such as WWN/iSCSI name, host group/iSCSI target, and volume characteristics are displayed.
   c. In **Pair Management Server Settings**:
      Select an available host as the pair management server. For Instance ID, select **Existing** to use a current configuration definition file and related instance ID number on the pair management server, or select **New** to create a new configuration definition file, and enter a new instance ID number which will be validated as an available instance ID. You can also specify a **UDP Port** number for the instance.
4. On the **Secondary** tab, configure the following options:
   a. Verify the displayed hosts for the selected volumes, and click **Select other Hosts** to specify more hosts that can access global-active device pairs.
   b. Select the secondary storage system where the S-VOLs will be created.
   c. Verify the displayed virtual storage machine.
   d. Click **Advanced Options** for further configuration.
   e. Click **LUN Path Options** for further configuration.
      The previously allocated primary path is displayed.
   f. Click **Host Group and LUN settings** for further configuration.
   g. In **Pair Management Server Settings**:
      Select an available host as the pair management server. For Instance ID, select **Existing** to use a current configuration definition file and related instance ID number on the pair management server, or select

New to create a new configuration definition file, and enter a new instance ID number which will be validated as an available instance ID. You can also specify a **UDP Port** number for the instance.

5. In **Pair Settings**:
   a. Select the **Quorum Disk**, and then select whether pairs should belong to a new or existing **Copy Group** for management purposes. Click **New** to assign a new copy group to the GAD pair. Select the **CTG ID** check box (VSP G1000 only) to assign a consistency group (CTG) ID to the new copy group. Click **Existing** to assign an existing copy group to the GAD pair. The existing CTG ID will be used if it exists for the selected copy group.

Note: (VSP G1000 only) To add the GAD 3DC delta resync configuration to allocated GAD pair volumes, CTG IDs must be assigned.

   b. Manually enter a name (Prefix and Start Sequence No.), or select **Automatic** for **Pair Name** and a name will be created.
6. Click **Show Plan** and confirm that the information in the plan summary is correct. If changes are required, click **Back**.
7. (Optional) Update the task name and provide a description.
8. (Optional) Expand **Schedule** to specify the task schedule.
   You can schedule the task to run immediately or later. The default setting is **Now**.
9. Click **Submit.**
   If the task is scheduled to run immediately, the process begins.
10. (Optional) Check the progress and result of the task on the **Tasks & Alerts** tab. Click the task name to view details of the task.
11. Confirm the global-active device paired volumes are created by checking volume status information.

Note: If you change the ALUA attribute of a volume for which a LUN path is set, you must rediscover the volume on the host.

Result

The global-active device pair is allocated to the target host, and initial copy is performed. For global-active device pairs that use DP pool volumes, reclaiming zero pages is also automatically performed.

Related topics

- Allocating global-active device pairs on page 7-12
- Allocating global-active device pairs based on like volumes on page 7-14

Checking global-active device pair status

You can review information (including status) of global-active device pairs or reserved volumes from the volume list and volume details windows in HCS.

Note: You can also use Replication Manager to check global-active device pair status. Open **Tools** > **Replication Manager** and locate copy groups or pairs to view pair status. For more information about pair status, conditions under which certain copy pair statuses can be changed, and about how to change the status of a copy pair by using the Change Pair Status wizard, see the online help for Replication Manager.

Procedure

1. Review global-active device status using one or more of the following options:
   - On the **Resources** tab, select **Storage Systems**, select a Virtual Storage Platform G1000, select **Volumes**, and then list **Open-Allocated** volumes to view global-active device information, as follows:
     - Volume Attribute
     - Copy Info (P-VOL)
     - Copy Info (S-VOL)
     - GAD status

     Tip: You might need to configure column settings to see global-active device information.

     Click volume links to view detailed information for global-active device pairs, Quorum disks, and command devices.
   - On the **Resources** tab, select **Hosts**, select an operating system, then select a host using global-active device pairs, and review volume status.
   - Search for hosts or volumes to locate global-active device volumes, and review volume status.

Result

The status of global-active device pairs is displayed.

Related topics

- Monitoring global-active device pair performance on page 7-19

Monitoring global-active device pair performance

Tuning Manager will collect and display performance information for global-active device pairs. If an administrator suspects a performance problem, or receives an alert indicating a performance problem, the administrator can investigate and resolve the issue.

Because global-active device pairs are kept synchronized, I/O issues on either the primary or secondary storage system can impact host performance. Status information and I/O metrics related to P-VOL and S-VOL performance can be examined to identify the problem, and resolve host I/O issues.

Prerequisites

- Tuning Manager must be installed, licensed and configured.
- When selecting one member of a global-active device pair, a single Device Manager managing the pair is required for analyzing the performance of the P-VOL and S-VOL simultaneously.

Procedure

1. On the **Analytics** tab, open the **Hosts** tree, select a host OS, and then select the target host as follows:
   a. Select the button next to the target host name, and then click **Identify Performance Problems**.
   b. Click **Host Name** to select target host volumes, and then click **Identify Performance Problems**.

Tip: When you select volumes from a host volume list or from within the **Identify Performance Problems** wizard, if you select a global-active device volume in a single Device Manager environment, the P-VOL and S-VOL will both be targets for analysis. To identify global-active device volumes, use **Column Settings** to include **Copy Info (P-VOL)**, **Copy Info (S-VOL)**, and **Virtual LDEV ID** in the host volume list or the Identify Performance Problems wizard.

   The **Identify Performance Problems** wizard is displayed (starting on the Introduction page). Review displayed messages.
2. Click **Next**.
3. On the **Target Selection** page:
   a. If prompted, select the **Storage system type/volume type (location)**, for example VSP G1000 and HDP volumes.
   b. Click **List of added volumes** to view analysis targets, including global-active device volumes.

Note: When the global-active device pair (P-VOL and S-VOL) of a selected global-active device volume are not found in HCS, the volume cannot be set to a target for analysis. These volumes are listed separately.

   c. In **Select volumes to analyze**, **Select all** selects all listed volumes for analysis. Click **Select manually** to individually select volumes or global-active device pairs (P-VOL and S-VOL) for analysis, and then click **Add** to update the **Selected Volumes** list.
   d. For **Select time period to analyze within 30 days before the current time**, configure the time period to analyze.
   e. Click **Next**.

4. On the **Overview** page:
   a. Check the analysis results and cautions. If there is a resource that exceeds a threshold value, there might be a performance issue related to the storage system.
   b. Click **Next**.
5. On the **Volume** page:
   a. Review the **Response Time** and **IOPS** graphs displaying both P-VOL and S-VOL performance metrics.

Tip: If there is a large disparity between the performance metrics for the P-VOL and S-VOL, complete the wizard steps to investigate the issue further and determine the root cause. For example, there could be significant load variances between the connected hosts for the primary/secondary storage systems, remote path latencies, or multi-path software configuration issues.

6. Continue stepping through the **Identify Performance Problems** wizard to review further information on Cache, MP Blades, and Parity Groups.

Result

The performance metrics of global-active device pairs have been displayed for analysis.

Related topics

- Checking global-active device pair status on page 7-19

Unallocating global-active device pairs

When a host with global-active device pairs is no longer needed, or the volume redundancy provided by global-active device pairs is no longer required, you can unallocate the global-active device pairs.

Depending on conditions, when you unallocate global-active device volumes, the copy pair relationship will be ended, and related copy groups and configuration definition files (horcmN.conf) could be deleted as a result.

Prerequisites

- Identify the name of the target host, and the global-active device volumes to unallocate.
- If necessary, back up data on the target volumes.
- Unmount all allocated volumes that you plan to unallocate. An IT administrator might have to perform this task.

Procedure

1. Select a host OS, then select a host using global-active device volumes and click **Unallocate Volumes**.

Tip: If you select specific volumes for the host, **Unallocate Volumes** will go directly to the **Show Plan** window.

2. From the **Unallocate Volumes** dialog box, verify that the host is correct, and then select the P-VOL or S-VOL of the global-active device pair, and select **Unallocate global-active device pair simultaneously**.
3. Click **Show Plan** and confirm that the information in the plan summary is correct. If changes are required, click **Back**.
4. (Optional) Update the task name and provide a description.
5. (Optional) Expand **Schedule** to specify the task schedule.

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

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

Result

Unallocated global-active device volumes are added back to the storage system **Open-Unallocated** volume list.

Related topics

- Unallocating individual global-active device volumes on page 7-22
- Discontinuing a global-active device environment on page 7-23

Unallocating individual global-active device volumes

Individual global-active device primary or secondary volumes can be unallocated.

Prerequisites

- Identify the name of the target host, and the global-active device primary or secondary volume to unallocate.
- If necessary, backup data for the target volume.

Procedure

1. On the **Resources** tab, select a host OS, then select a host using global-active device volumes, and click **Unallocate Volumes**.

Tip: If you select specific volumes for the host, **Unallocate Volumes** will go directly to the **Show Plan** window.

2. In the **Unallocate Volumes** dialog box, verify that the host is correct, and then select the P-VOL or S-VOL of the global-active device pair to unallocate.

Note: When unallocating the secondary volume, the primary volume is retained as a normal volume. When unallocating the primary volume, the secondary volume is retained as a normal volume.

3. Click **Show Plan** and confirm that the information in the plan summary is correct. If changes are required, click **Back**.
4. (Optional) Update the task name and provide a description.
5. (Optional) Expand **Schedule** to specify the task schedule.

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

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

Result

Unallocated global-active device primary or secondary volumes are added back to the storage system **Open-Unallocated** volume list.

Related topics

- Unallocating global-active device pairs on page 7-21
- Discontinuing a global-active device environment on page 7-23

Discontinuing a global-active device environment

To discontinue the use of a global-active device environment, you must delete remote path settings, delete and unvirtualize the quorum disk, delete the command devices for pair management servers, and delete the virtual storage machine.

Prerequisites

- Data on affected global-active device pairs has been saved.
- All global-active device pairs are unallocated, or the global-active device S-VOLs unallocated leaving normal volumes in use by hosts.

Procedure

1. Delete global-active device remote path settings.

   a. On the **Resources** tab, right-click the target storage system, and select **Remote Connections**.

b. To delete a remote path, select a remote connection to be deleted in the **Connections (To)** tab, and then select **Remove Remote Paths**.

c. To delete all remote paths allocated to the secondary storage system on the primary storage system, select the remote connections to be deleted in the **Connections (To)** tab, and then select **Remove Remote Connections**.

2. Delete and unvirtualize the global-active device quorum disk.

   a. On the **Resources** tab, in the **Storage Systems** tree, select and then right-click the target storage system. From the menu, select **Remote Connections**.

   b. In the **Remote Connections** window, on the **Quorum disk** tab, select the quorum disk, confirm it is correct, and delete it.

Note: To delete a quorum disk, all global-active device pairs that use the quorum disk must be unallocated.

   c. From **General Tasks**, select **Unvirtualize Volumes**. In the **Unvirtualize Volumes** dialog box, select the external storage system (quorum) and the internal storage system (primary or secondary storage system) and unvirtualize the external volume (quorum disk). Be sure to perform this for both the primary and secondary storage.

3. Delete the command devices for global-active device pair management servers.

   a. In the **Unallocate Volumes** dialog box, unallocate the pair management server and the command device.

   b. On the **Resources** tab, in the **Storage Systems** tree, select the primary or secondary storage system, and select **Volumes**. In the application area, click the **System GUI** link.

   c. From the volume list (LDEVs) that appear, select the command device that you want to delete.

Tip: Look for `Command Device` in the **Attribute** column.

   d. From **More Actions**, select **Edit Command Devices** and in the dialog box that appears, delete the command device.

4. Delete the global-active device virtual storage machine and remove the secondary storage system resources from the global-active device virtual storage machine.

Note: Before removing secondary storage system resources, you must unallocate the volumes on the global-active device virtual storage machine.

   a. On the **Administration** tab, select **Virtual Storage Machine**.

   b. Select the virtual storage machine related to global-active device pairs and click **Edit Virtual Storage Machine**.

   c. In the **Edit Virtual Storage Machine** window, select the secondary storage system, and then select **Remove Storage Systems**.

Result

The global-active device environment no longer exists.

Related topics

- Unallocating global-active device pairs on page 7-21
- Unallocating individual global-active device volumes on page 7-22

8

Disaster recovery of global-active device

This chapter describes the global-active device (GAD) failure locations, the SIMs issued for GAD failures, and the recovery procedures for GAD failures.

- Pair condition and recovery: secondary storage system failure
- Pair condition and recovery: path failure, primary to external storage system
- Pair condition and recovery: path failure, secondary to external storage system
- Pair condition and recovery: quorum disk failure
- Pair condition and recovery: external storage system failure
- Pair condition and recovery: other failures
- Recovering the storage systems: primary site failure with external storage system
- Reversing the P-VOL and S-VOL
- Resolving failures in multiple locations

Failure locations

The following figure and table describe the locations where GAD failures can occur, the SIMs that are issued, and whether the P-VOL and S-VOL are accessible. All GAD-related SIMs are described in SIMs related to GAD on page 8-6.

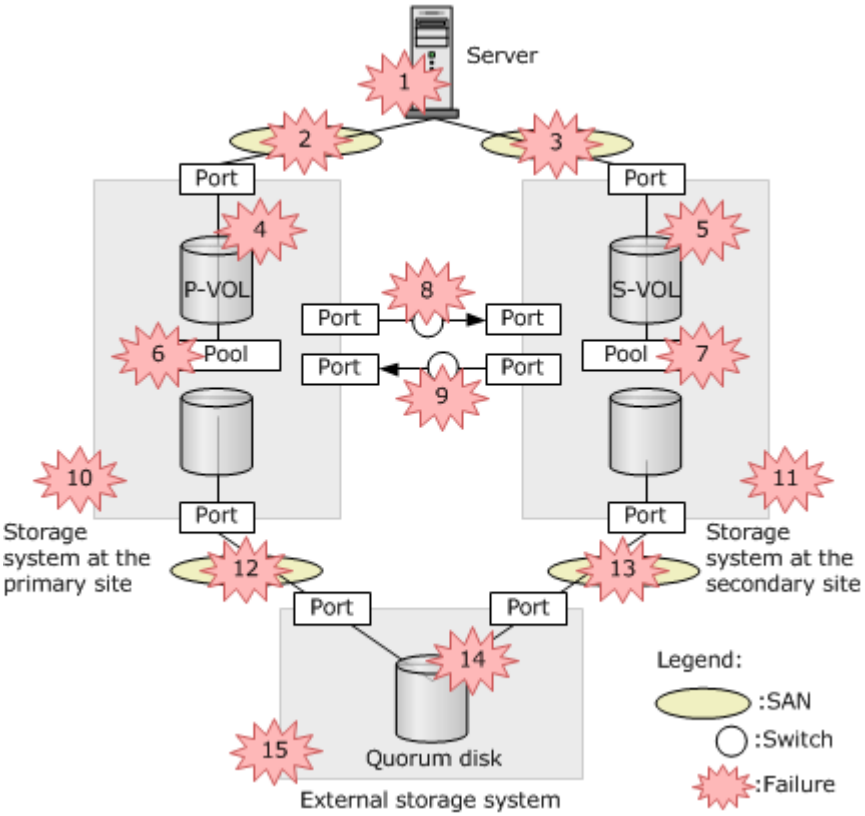

Figure 8-1 Failure locations

#	Failure location		SIM reference codes		GAD volume accessible?[1]	
			Primary storage system	Secondary storage system	P-VOL	S-VOL
1	Server		None (normal)	None (normal)	Yes	Yes
2	Path between the server and the storage system	Path between the server and the primary storage system	None (normal)	None (normal)	No	Yes[2]
3		Path between the server and the secondary storage system	None (normal)	None (normal)	Yes[3]	No
4	GAD pair volume	P-VOL	3A0*xxx* DD1*xyy*	DD1*xyy*	No	Yes[2]

#	Failure location		SIM reference codes		GAD volume accessible?[1]	
			Primary storage system	Secondary storage system	P-VOL	S-VOL
			DFA*xxx* DFB*xxx* EF9*xxx*			
5		S-VOL	DD1*xyy*	3A0*xxx* DD1*xyy* DFA*xxx* DFB*xxx* EF9*xxx*	Yes[3]	No
6	Pool for GAD pair[4]	Pool for P-VOL	622*xxx* DD1*xyy*	DD1*xyy*	No	Yes[2]
7		Pool for S-VOL	DD1*xyy*	622*xxx* DD1*xyy*	Yes[3]	No
8	Path between storage systems	Remote path from the primary to secondary storage system	2180*xx* DD0*xyy*	DD3*xyy*	Yes[3]	No
9		Remote path from the secondary to primary storage system	DD3*xyy*	2180*xx* DD0*xyy*	No	Yes[2]
10	Storage system	Primary storage system	Depends on the failure type[5]	2180*xx* DD0*xyy* DD3*xyy*	No	Yes[2]
11		Secondary storage system	2180*xx* DD0*xyy* DD3*xyy*	Depends on the failure type[5]	Yes[3]	No
12	Quorum disk	Path between the primary storage system and quorum disk	21D0*xx* 21D2*xx* DD2*xyy* DEF0*zz* EF5*xyy* EFD000 FF5*xyy*	DD2*xyy*	Yes[3]	No
13		Path between the secondary storage system and quorum disk	DD2*xyy*	21D0*xx* 21D2*xx* DD2*xyy* DEF0*zz* EF5*xyy* EFD000	Yes[3]	No

#	Failure location		SIM reference codes		GAD volume accessible?[1]	
			Primary storage system	Secondary storage system	P-VOL	S-VOL
				FF5*xyy*		
14		Quorum disk	21D0*xx* 21D2*xx* DD2*xyy* DEF0*zz* EF5*xyy* EFD000 FF5*xyy*	21D0*xx* 21D2*xx* DD2*xyy* DEF0*zz* EF5*xyy* EFD000 FF5*xyy*	Yes[3]	No
15		External storage system	21D0*xx* 21D2*xx* DD2*xyy* DEF0*zz* EF5*xyy* EFD000 FF5*xyy*	21D0*xx* 21D2*xx* DD2*xyy* DEF0*zz* EF5*xyy* EFD000 FF5*xyy*	Yes[3]	No

Notes:

1. Pairs are not suspended and do not become inaccessible for:
   - Failure in hardware used for redundancy in the storage system, such as HDD, cache, front-end director (CHA), back-end director (DKA), and MPB
   - Failure in redundant physical paths
2. The volume is not accessible if a failure occurs while the S-VOL pair status is COPY, SSUS, or PSUE.
3. The volume is not accessible if a failure occurs while the P-VOL pair status is PSUS or PSUE and the I/O mode is BLOCK.
4. A failure occurs due to a full pool for a GAD pair.
5. The SIM might not be viewable, depending on the failure (for example, all cache failure, all MP failure, storage system failure).

Related topics

- SIMs related to GAD on page 8-6
- Pair condition and recovery: server failures on page 8-7
- Pair condition and recovery: path failure between the server and storage system on page 8-8
- Pair condition and recovery: P-VOL failure (LDEV blockade) on page 8-11
- Pair condition and recovery: S-VOL failure (LDEV blockade) on page 8-18
- Pair condition and recovery: full pool for the P-VOL on page 8-25
- Pair condition and recovery: full pool for the S-VOL on page 8-28

- Pair condition and recovery: path failure, primary to secondary storage system on page 8-31
- Pair condition and recovery: path failure, secondary to primary storage system on page 8-34
- Pair condition and recovery: primary storage system failure on page 8-37
- Pair condition and recovery: secondary storage system failure on page 8-40
- Pair condition and recovery: path failure, primary to external storage system on page 8-43
- Pair condition and recovery: path failure, secondary to external storage system on page 8-46
- Pair condition and recovery: quorum disk failure on page 8-50
- Pair condition and recovery: external storage system failure on page 8-55
- Pair condition and recovery: other failures on page 8-56

SIMs related to GAD

The following table shows SIMs related to global-active device operations. All SIMs in the following table are reported to the service processor (SVP) of the storage system.

SIM reference code	Description
2180*xx*	Logical path(s) on the remote copy connections was logically blocked (due to an error condition)
21D0*xx*	External storage system connection path blocking
21D2*xx*	Threshold over by external storage system connection path response time-out
3A0*xyy*	LDEV blockade (effect of microcode error)
622*xxx*	The DP POOL FULL
DD0*xyy*	GAD for this volume was suspended (due to an unrecoverable failure on the remote copy connections)
DD1*xyy*	GAD for this volume was suspended (due to a failure on the volume)
DD2*xyy*	GAD for this volume was suspended (due to an internal error condition detected)
DD3*xyy*	Status of the P-VOL was not consistent with the S-VOL
DEE0*zz*	Quorum disk restore
DEF0*xx*	Quorum disk blocked
DFA*xxx*	LDEV blockade (drive path: boundary 0/effect of drive port blockade)

SIM reference code	Description
DFB*xxx*	LDEV blockade (drive path: boundary 1/effect of drive port blockade)
EF5*xyy*	Abnormal end of write processing in external storage system
EF9*xxx*	LDEV blockade (effect of drive blockade)
EFD000	External storage system connection device blockade
FF5*xyy*	Abnormal end of read processing in external storage system

Related topics

- Failure locations on page 8-3
- Resolving failures in multiple locations on page 8-63

Pair condition before failure

The pair status and I/O mode of a GAD pair, the accessibility of the server, and the storage location of the latest data depend on the status before a failure occurs.

The following table shows pair status and I/O mode, the volumes accessible from the server, and the location of the latest data before a failure occurs. You can compare this information with the changes that take place after a failure occurs, as described in the following topics.

Pair status and I/O mode		Volume accessible from the server		Volume with latest data
P-VOL	S-VOL	P-VOL	S-VOL	
PAIR (Mirror (RL))	PAIR (Mirror (RL))	OK	OK	Both P-VOL and S-VOL
PAIR (Mirror (RL))*	PAIR (Block)	OK	NG	Both P-VOL and S-VOL
COPY (Mirror (RL))	COPY (Block)	OK	NG	P-VOL
PSUS/PSUE (Local)	SSUS/PSUE (Block)	OK	NG	P-VOL
PSUS/PSUE (Block)	SSWS (Local)	NG	OK	S-VOL
* This row applies only to the VSP G1000.				

Pair condition and recovery: server failures

The following table shows transitions for pair status and I/O mode, the volumes that are accessible from the server, and the location of the latest data when a server failure occurs.

Before failure		After failure				
Pair status and I/O mode		Pair status and I/O mode		Volume accessible from the server[1]		Volume with latest data
P-VOL	S-VOL	P-VOL	S-VOL	P-VOL	S-VOL	
PAIR (Mirror (RL))	PAIR (Mirror (RL))	PAIR (Mirror (RL))	PAIR (Mirror (RL))	OK	OK	Both P-VOL and S-VOL
PAIR (Mirror (RL))[2]	PAIR (Block)	PAIR (Mirror (RL))	PAIR (Block)	OK	NG	Both P-VOL and S-VOL
COPY (Mirror (RL))	COPY (Block)	COPY (Mirror (RL))	COPY (Block)	OK	NG	P-VOL
PSUS/PSUE (Local)	SSUS/PSUE (Block)	PSUS/PSUE (Local)	SSUS/PSUE (Block)	OK	NG	P-VOL
PSUS/PSUE (Block)	SSWS (Local)	PSUS/PSUE (Block)	SSWS (Local)	NG	OK	S-VOL

Notes:

1. If failures occur in all servers that access the P-VOL or S-VOL, then you cannot access either volume.
2. This row applies only to the VSP G1000.

SIM

- Primary storage system: None
- Secondary storage system: None

Recovery procedure

1. Recover the server.
2. Recover the path from the server to the pair volumes.

Related topics

- Pair condition before failure on page 8-7

Pair condition and recovery: path failure between the server and storage system

If a server cannot access a pair volume whose status is PAIR, though no SIM has been issued, a failure might have occurred between the server and the storage system. The following topics provide procedures for recovering of the physical path between the server and the storage systems.

The following table shows transitions for pair status and I/O mode, the volumes that are accessible from the server, and the location of the latest

data when you can no longer use a physical path between the server and a storage system.

Before failure		After failure				
Pair status and I/O mode		Pair status and I/O mode		Volume accessible from the server[1]		Volume with latest data
P-VOL	S-VOL	P-VOL	S-VOL	P-VOL	S-VOL	
PAIR (Mirror (RL))	PAIR (Mirror (RL))	PAIR (Mirror (RL))	PAIR (Mirror (RL))	OK	OK	Both P-VOL and S-VOL
PAIR (Mirror (RL))[2]	PAIR (Block)	PAIR (Mirror (RL))	PAIR (Block)	OK	NG	Both P-VOL and S-VOL
COPY (Mirror (RL))	COPY (Block)	COPY (Mirror (RL))	COPY (Block)	OK	NG	P-VOL
PSUS/PSUE (Local)	SSUS/PSUE (Block)	PSUS/PSUE (Local)	SSUS/PSUE (Block)	OK	NG	P-VOL
PSUS/PSUE (Block)	SSWS (Local)	PSUS/PSUE (Block)	SSWS (Local)	NG	OK	S-VOL

1. If failures occur in all servers that access the P-VOL or S-VOL, then you cannot access either volume.
2. This row applies only to the VSP G1000.

SIM

- Primary storage system: None
- Secondary storage system: None

Recovery procedure

1. Recover the path between the server and the storage system.
2. Recover the path from the server to the pair volume.

Related topics

- Pair condition before failure on page 8-7

Recovering from a path failure: server to primary storage system

The following figure shows the failure area and recovery when the path between the server and the primary storage system fails.

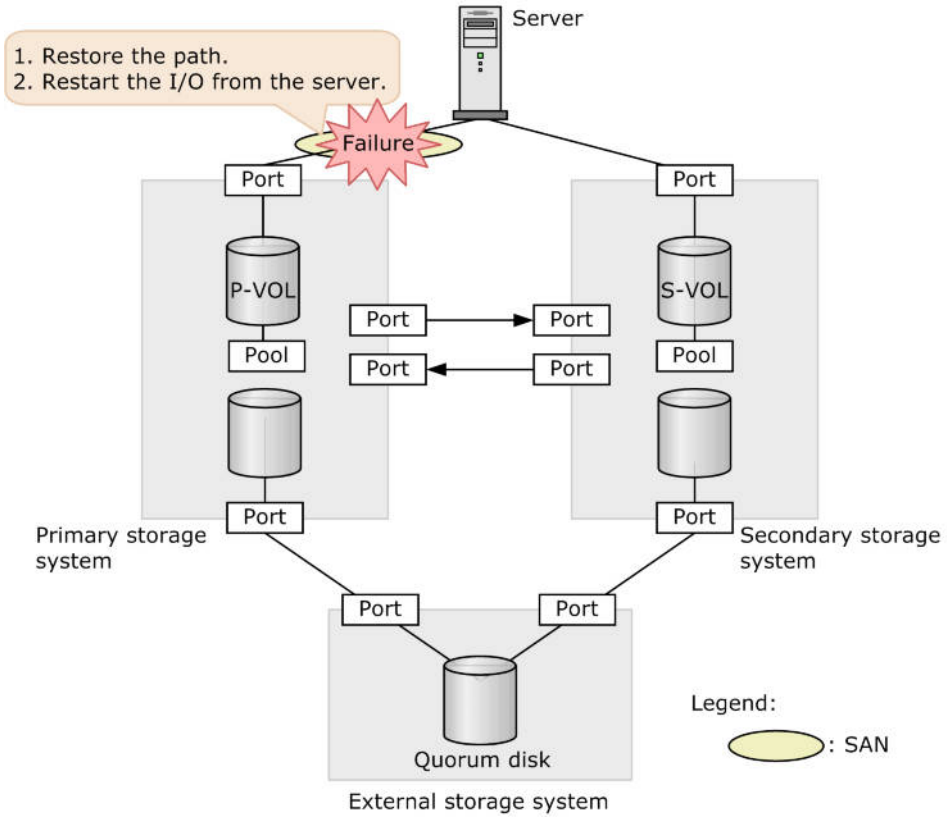

Steps for recovery

1. Recover the path.
   a. Using the alternate path software and other tools, identify the path that cannot be accessed from the server.
   b. Using the SAN management software, identify the failure location; for example, a host bus adapter, FC cable, switch, or other location.
   c. Remove the cause of failure and recover the path.
2. Using the alternate path software, resume I/O from the server to the recovered path (I/O might resume automatically).

Related topics

- Failure locations on page 8-3
- SIMs related to GAD on page 8-6

Recovering from a path failure: server to secondary storage system

The following figure shows the failure area and recovery when the path between the server and secondary storage system fails.

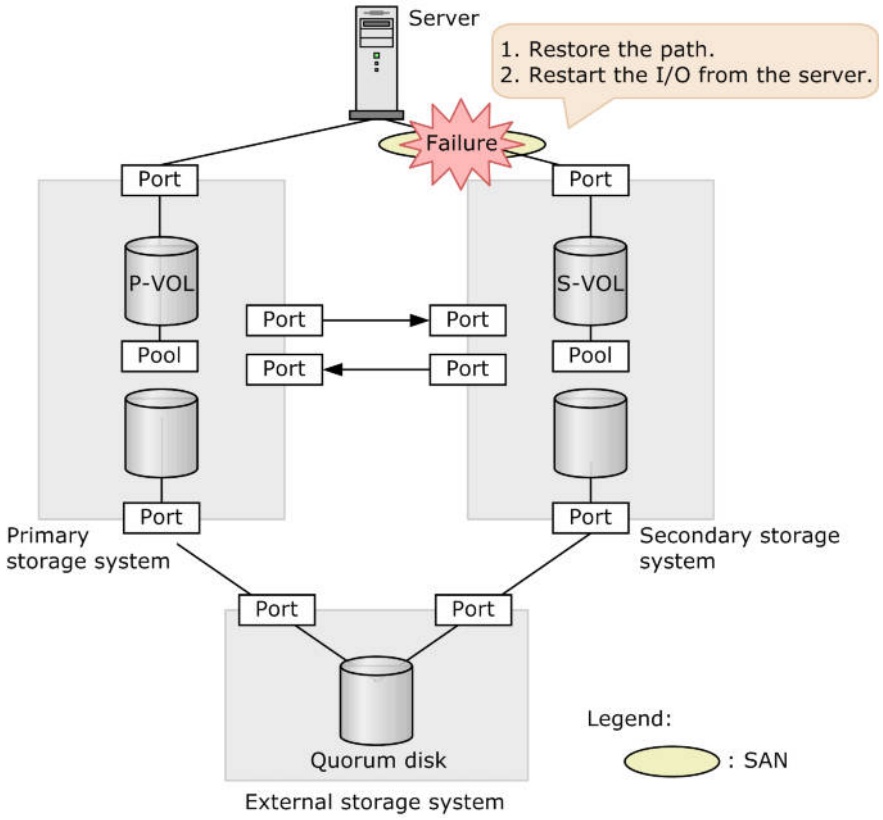

Steps for recovery

1. Recover the path.
   a. Using the alternate path software or other tools, identify the path that cannot be accessed from the server.
   b. Using SAN management software, identify the failure location; for example, a host bus adapter, FC cable, switch, or other location.
   c. Remove the cause of failure and recover the path.
2. Using the alternate path software, resume I/O from the server to the recovered path (I/O might resume automatically).

Related topics

- Failure locations on page 8-3
- SIMs related to GAD on page 8-6

Pair condition and recovery: P-VOL failure (LDEV blockade)

The following table shows transitions for pair status and I/O mode, the volumes that are accessible from the server, and the location of the latest data when you can no longer use the P-VOL due to LDEV blockade.

Before failure		After failure				
Pair status and I/O mode		Pair status and I/O mode		Volume accessible from the server[1]		Volume with latest data
P-VOL	S-VOL	P-VOL	S-VOL	P-VOL	S-VOL	
PAIR (Mirror (RL))	PAIR (Mirror (RL))	PSUE (Block)	SSWS (Local)	NG	OK	S-VOL
PAIR (Mirror (RL))[2]	PAIR (Block)	PSUE (Block)	SSWS (Local)	NG	OK	S-VOL
COPY (Mirror (RL))	COPY (Block)	PSUE (Local)	PSUE (Block)	NG	NG	None[1]
PSUS/PSUE (Local)	SSUS/PSUE (Block)	PSUS/PSUE (Local)	SSUS/PSUE (Block)	NG	NG	None[3]
PSUS/PSUE (Block)	SSWS (Local)	PSUS/PSUE (Block)	SSWS (Local)	NG	OK	S-VOL

Notes:

1. Recover the data from ShadowImage, Thin Image, or other backup data.
2. This row applies only to the VSP G1000.
3. Recover the data using the S-VOL data that is not the latest, ShadowImage, Thin Image, or other backup data.

SIM

- Primary storage system: 3A0*xyy*, DD1*xyy*, DFA*xxx*, DFB*xxx*, EF9*xxx*
- Secondary storage system: DD1*xyy*

Recovery procedure

1. Recover the P-VOL.
2. Re-create the pair.

Related topics

- Pair condition before failure on page 8-7
- Recovering the P-VOL (DP-VOL) (pair status: PAIR) on page 8-12
- Recovering the P-VOL (other than DP-VOL) (pair status: PAIR) on page 8-15

Recovering the P-VOL (DP-VOL) (pair status: PAIR)

The following figure shows the failure area and recovery when a pair is suspended due to a P-VOL failure and the P-VOL is a DP-VOL.

Note: In this example no consistency group is specified.

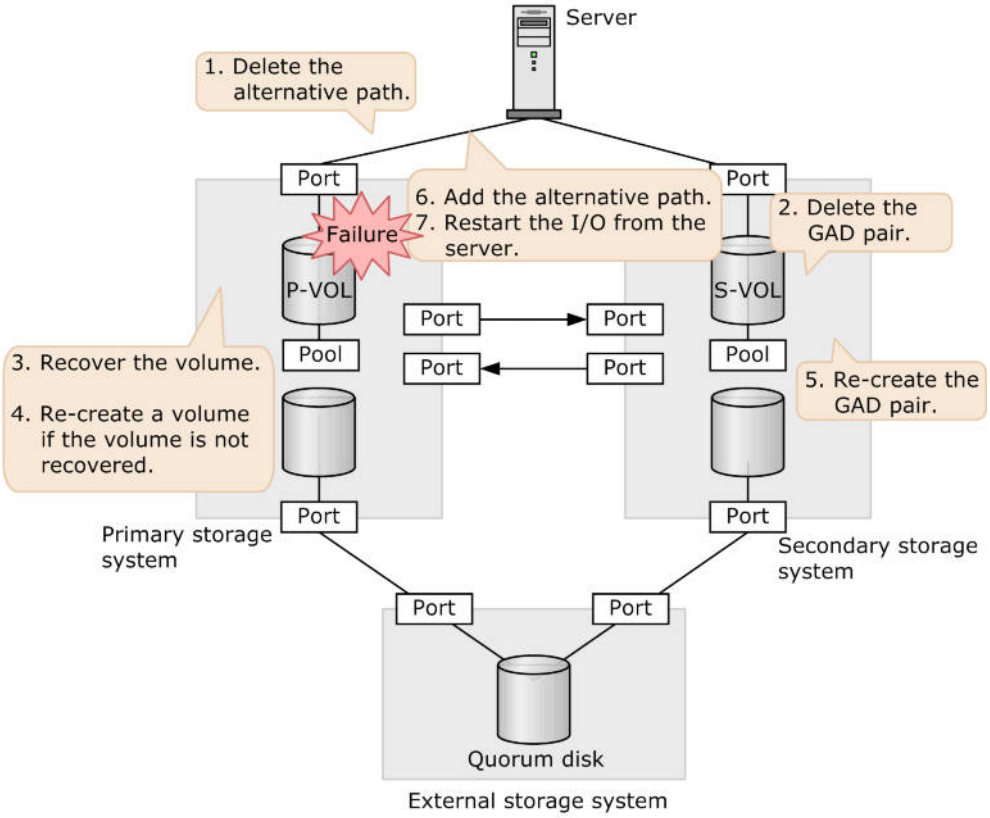

Steps for recovery

1. Delete the alternate path (logical path) to the volume that cannot be accessed from the server.

   a. Using the alternate path software, identify the volume that cannot be accessed.

   b. Confirm whether the volume (P-VOL) is blocked, and the pool ID (`B_POOLID`) of the pool to which the P-VOL is associated.

   ```
   raidcom get ldev -ldev_id 0x2222 -IH0
   (snip)
   B_POOLID : 0
   (snip)
   STS : BLK
   (snip)
   ```

   c. Display the status of the volumes configuring the pool (pool volume) to identify the blocked volume.

   ```
   raidcom get ldev -ldev_list pool -pool_id 0 -IH0
   (snip)
   LDEV : 16384
   (snip)
   STS : BLK
   (snip)
   ```

   For a blocked volume, `BLK` is indicated in the `STS` column.

d. Using the alternate path software, delete the alternate path to the volume that cannot be accessed from the server.

Go to the next step even if the alternate path cannot be deleted.

2. Delete the pair.

   a. From the secondary storage system, delete the pair specifying the actual LDEV ID of the S-VOL.

   ```
   pairsplit -g oraHA -R -d dev1 -IH1
   ```

Note: To delete the pair specifying the S-VOL, use the -R option of the pairsplit command. Specify the actual LDEV ID (device name) of the S-VOL in the -d option.

   b. Confirm that the pair is deleted.

   ```
   pairdisplay -g oraHA -fxce -IH1
   Group   PairVol(L/R) (Port#,TID, LU), Seq#,   LDEV#.P/
   S,Status,Fence,   %,    P-LDEV# M   CTG JID AP EM       E-
   Seq# E-LDEV# R/W
   oraHA   dev1(L)      (CL1-C-1, 0,   0)322222  4444.SMPL  ----
   ------,  ----- -----   -   -   -   -  -        -      -
   -/-
   oraHA   dev1(R)      (CL1-A-0, 0,   0)311111  2222.SMPL  ----
   ------,  ----- -----   -   -   -   -  -        -      -
   -/-
   ```

3. Remove the failure.

   The following example shows recovery from a pool-volume failure.

   a. Recover a pool volume that configures the P-VOL (DP-VOL).

   b. Display the status of the pool volumes to confirm that the pool volume has been recovered.

   ```
   raidcom get ldev -ldev_list pool -pool_id 0 -IH0
    (snip)
   LDEV : 16384
   (snip)
   STS : NML
   (snip)
   ```

   For a normal volume, `NML` is indicated in the `STS` column.

4. If the volume cannot be recovered, follow the procedure below to re-create the P-VOL:

   a. At the primary storage system, delete the LU path to the P-VOL.

   b. Delete the P-VOL.

   c. Create a new volume.

   d. Set an LU path to the new volume.

5. (VSP G1000 only) Confirm the quorum disk status. If the quorum disk is blocked, recover from the blockade.

6. Re-create the pair.

   a. If you created a volume in step 4, set the GAD reserve attribute to the created volume.

   ```
   raidcom map resource -ldev_id 0x2222 -virtual_ldev_id reserve
   -IH0
   ```

b. From the secondary storage system, create the pair specifying the S-VOL's actual LDEV ID.

```
paircreate -g oraHA -f never -vl -jq 0 -d dev1 -IH1
```

Note: To create the pair specifying the S-VOL, specify the actual LDEV ID (device name) of the S-VOL in the -d option of the paircreate command.

The volume of the primary storage system changes to an S-VOL, and the volume of the secondary storage system changes to a P-VOL.

c. Confirm that the P-VOL and S-VOL pair statuses change to PAIR (Mirror (RL)).

```
pairdisplay -g oraHA -fxce -IH0
Group   PairVol(L/R) (Port#,TID, LU),Seq#,LDEV#.P/S,Status,
Fence,   %,P-LDEV# M CTG JID AP EM       E-Seq# E-LDEV# R/W
oraHA   dev1(L)     (CL1-A-0, 0,   0)311111  2222.P-VOL PAIR
NEVER ,  100  4444 -   -   0  -  -            -       - L/M
oraHA   dev1(R)     (CL1-C-1, 0,   0)322222  4444.S-VOL PAIR
NEVER ,  100  2222 -   -   0  -  -            -       - L/M
pairdisplay -g oraHA -fxce -IH1
Group   PairVol(L/R) (Port#,TID, LU),Seq#,LDEV#.P/S,Status,
Fence,   %,P-LDEV# M CTG JID AP EM       E-Seq# E-LDEV# R/W
oraHA   dev1(L)     (CL1-C-1, 0,   0)322222  4444.S-VOL PAIR
NEVER ,  100  2222 -   -   0  -  -            -       - L/M
oraHA   dev1(R)     (CL1-A-0, 0,   0)311111  2222.P-VOL PAIR
NEVER ,  100  4444 -   -   0  -  -            -       - L/M
```

7. Using the alternate path software, add an alternate path from the server to the S-VOL (P-VOL before the failure).
8. Using the alternate path software, resume I/O from the server to the S-VOL (P-VOL before the failure).

Note: I/O from the server might resume automatically.

9. Reverse the P-VOL and the S-VOL if necessary.

Related topics

- I/O modes on page 1-14
- Failure locations on page 8-3
- SIMs related to GAD on page 8-6
- Reversing the P-VOL and S-VOL on page 8-62
- Recovering the P-VOL (other than DP-VOL) (pair status: PAIR) on page 8-15
- Pair condition and recovery: quorum disk failure on page 8-50

Recovering the P-VOL (other than DP-VOL) (pair status: PAIR)

The following figure shows the failure area and recovery when a pair is suspended due to a P-VOL failure and the P-VOL is not a DP-VOL.

For details about storage system support (models, microcode) for volumes other than DP-VOLs, see Requirements and restrictions on page 2-2.

Note: In this example no consistency group is specified.

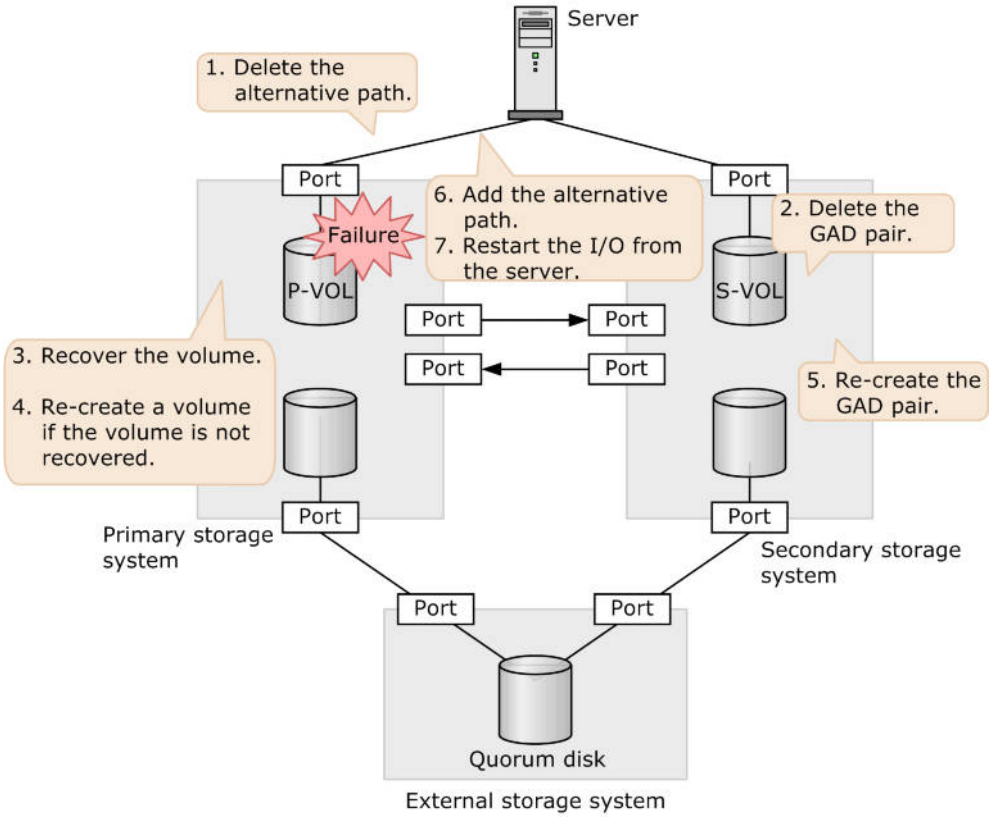

Steps for recovery

1. Delete the alternate path (logical path) to the volume that cannot be accessed from the server.
   a. Using the alternate path software, identify the volume that cannot be accessed.
   b. Confirm whether the volume (P-VOL) is blocked.

      ```
      raidcom get ldev -ldev_id 0x2222 -IH0
      (snip)
      STS : BLK
      (snip)
      ```

      For a blocked volume, BLK is indicated in the STS column.
   c. Using the alternate path software, delete the alternate path to the volume that cannot be accessed from the server.

   Go to the next step even if the alternate path cannot be deleted.
2. Delete the pair.
   a. From the secondary storage system, delete the pair specifying the actual LDEV ID of the S-VOL.

      ```
      pairsplit -g oraHA -R -d dev1 -IH1
      ```

Note: To delete the pair specifying the S-VOL, use the `-R` option of the `pairsplit` command. Specify the actual LDEV ID (device name) of the S-VOL in the `-d` option.

b. Confirm that the pair is deleted.

```
pairdisplay -g oraHA -fxce -IH1Group PairVol(L/R) (Port#,TID,
LU), Seq#, LDEV#.P/S,Status,Fence, %, P-LDEV# M CTG JID AP
        EM ESeq# E-LDEV# R/W
oraHA dev1(L) (CL1-C-1, 0, 0)322222 4444.SMPL ---- ------,
----- ----- - - - - - - -/-
oraHA dev1(R) (CL1-A-0, 0, 0)311111 2222.SMPL ---- ------,
----- ----- - - - - - - -/-
```

3. Remove the failure. The following example shows recovery from a volume failure.

   a. Recover the P-VOL.

   b. Display the status of the P-VOL to confirm that the pool volume has been recovered.

```
raidcom get ldev -ldev_id 0x2222 -IH0
(snip)
STS : NML
(snip)
```

   For a normal volume, NML is indicated in the STS column.

4. If the volume cannot be recovered, follow the procedure below to re-create the P-VOL:

   a. At the primary storage system, delete the LU path to the P-VOL.

   b. Delete the P-VOL.

   c. Create a new volume.

   d. Set an LU path to the new volume.

5. Confirm the quorum disk status. If the quorum disk is blocked, recover from the blockade.

6. Re-create the pair.

   a. If you created a volume in step 4, set the GAD reserve attribute to the created volume.

```
raidcom map resource -ldev_id 0x2222 -virtual_ldev_id reserve
-IH0
```

   b. From the secondary storage system, create the pair specifying the S-VOL's actual LDEV ID.

```
paircreate -g oraHA -f never -vl -jq 0 -d dev1 -IH1
```

Note: To create the pair specifying the S-VOL, specify the actual LDEV ID (device name) of the S-VOL in the `-d` option of the `paircreate` command. The volume in the primary storage system changes to an S-VOL, and the volume in the secondary storage system changes to a P-VOL.

   c. Confirm that the P-VOL and S-VOL pair statuses have changed to PAIR (Mirror (RL)).

```
pairdisplay -g oraHA -fxce -IH0
Group PairVol(L/R) (Port#,TID, LU),Seq#,LDEV#.P/S,Status,
Fence, %,P-LDEV# M CTG JID AP EM E-Seq# E-LDEV# R/W
oraHA dev1(L) (CL1-A-0, 0, 0)311111 2222.P-VOL PAIR NEVER ,
100 4444 - - 0 - - - - L/M
oraHA dev1(R) (CL1-C-1, 0, 0)322222 4444.S-VOL PAIR NEVER ,
100 2222 - - 0 - - - - L/M
pairdisplay -g oraHA -fxce -IH1
Group PairVol(L/R) (Port#,TID, LU),Seq#,LDEV#.P/S,Status,
Fence, %,P-LDEV# M CTG JID AP EM E-Seq# E-LDEV# R/W
oraHA dev1(L) (CL1-C-1, 0, 0)322222 4444.S-VOL PAIR NEVER ,
100 2222 - - 0 - - - - L/M
oraHA dev1(R) (CL1-A-0, 0, 0)311111 2222.P-VOL PAIR NEVER ,
100 4444 - - 0 - - - - L/M
```

7. Using the alternate path software, add an alternate path from the server to the S-VOL (P-VOL before the failure).
8. Using the alternate path software, resume I/O from the server to the S-VOL (P-VOL before the failure).

Note: I/O from the server might resume automatically.

9. Reverse the P-VOL and the S-VOL if necessary.

Related topics

- I/O modes on page 1-14
- Failure locations on page 8-3
- SIMs related to GAD on page 8-6
- Reversing the P-VOL and S-VOL on page 8-62
- Recovering the P-VOL (DP-VOL) (pair status: PAIR) on page 8-12
- Pair condition and recovery: quorum disk failure on page 8-50

Pair condition and recovery: S-VOL failure (LDEV blockade)

The following table shows the transitions for pair status and I/O mode, the volumes that are accessible from the server, and the location of the latest data when you can no longer use the S-VOL due to LDEV blockade.

Before failure		After failure				
Pair status and I/O mode		Pair status and I/O mode		Volume accessible from the server		Volume with latest data
P-VOL	S-VOL	P-VOL	S-VOL	P-VOL	S-VOL	
PAIR (Mirror (RL))	PAIR (Mirror (RL))	PSUE (Local)	PSUE (Block)	OK	NG	P-VOL
PAIR (Mirror (RL))[1]	PAIR (Block)	PSUE (Local)	PSUE (Block)	OK	NG	P-VOL

Before failure		After failure				
Pair status and I/O mode		Pair status and I/O mode		Volume accessible from the server		Volume with latest data
P-VOL	S-VOL	P-VOL	S-VOL	P-VOL	S-VOL	
COPY (Mirror (RL))	COPY (Block)	PSUE (Local)	PSUE (Block)	OK	NG	P-VOL
PSUS/PSUE (Local)	SSUS/PSUE (Block)	PSUS/PSUE (Local)	SSUS/PSUE (Block)	OK	NG	P-VOL
PSUS/PSUE (Block)	SSWS (Local)	PSUS/PSUE (Block)	SSWS (Local)	NG	NG	None[2]

1. This row applies only to VSP G1000.
2. Recover data using the P-VOL data that is not the latest, ShadowImage, Thin Image, or other backup data.

SIM

- Primary storage system: DD1*xyy*
- Secondary storage system: 3A0*xyy*, DD1*xyy*, DFA*xxx*, DFB*xxx*, EF9*xxx*

Recovery procedure

1. Recover the S-VOL.
2. Re-create the pair.

Related topics

- Pair condition before failure on page 8-7
- Recovering the S-VOL (DP-VOL) (pair status: PAIR) on page 8-19
- Recovering the S-VOL (other than DP-VOL) (pair status: PAIR) on page 8-22

Recovering the S-VOL (DP-VOL) (pair status: PAIR)

The following figure shows the failure area and recovery when a pair is suspended due to an S-VOL failure and the S-VOL is a DP-VOL.

Note: In this example no consistency group is specified.

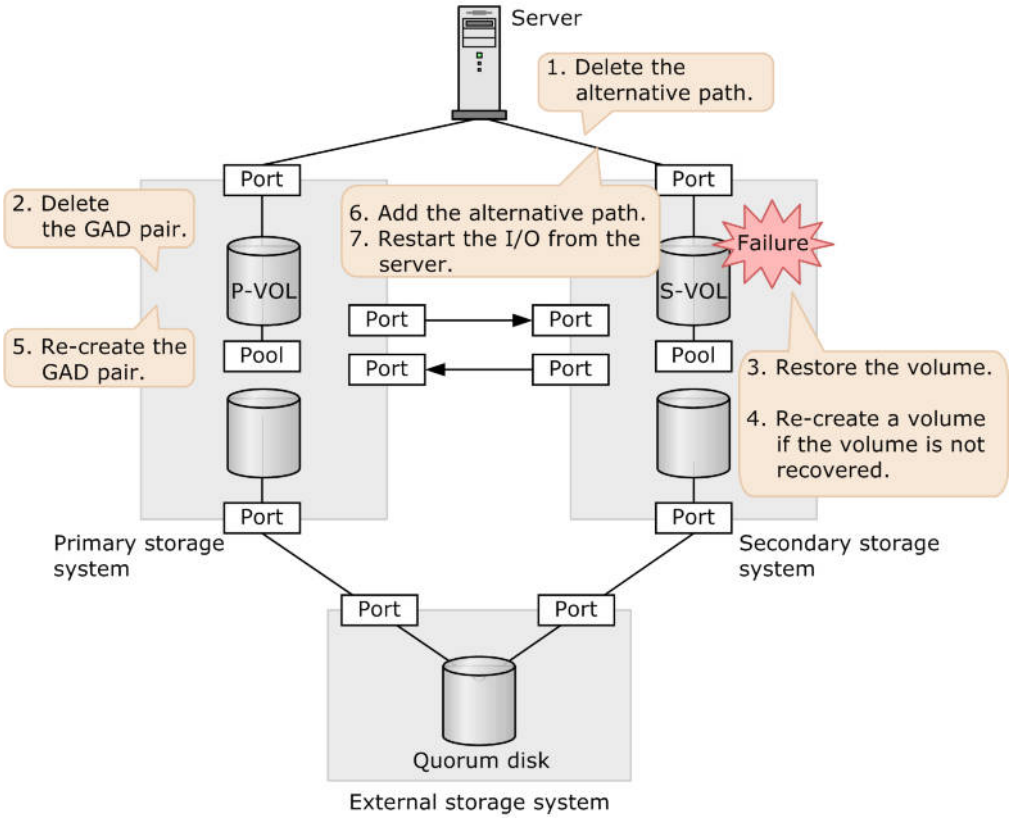

Steps for recovery

1. Delete the alternate path (logical path) to the volume that cannot be accessed from the server.

   a. Using the alternate path software, identify the volume that cannot be accessed.

   b. Confirm whether the volume (S-VOL) is blocked, and the pool ID (`B_POOLID`) of the pool to which the S-VOL is associated.

   ```
   raidcom get ldev -ldev_id 0x4444 -IH1
   (snip)
   B_POOLID : 0
   (snip)
   STS : BLK
   (snip)
   ```

   c. Display the status of the volumes configuring the pool (pool volume) to identify the blocked volume.

   ```
   raidcom get ldev -ldev_list pool -pool_id 0 -IH1
   (snip)
   LDEV : 16384
   (snip)
   STS : BLK
   (snip)
   ```

   For the blocked volume, `BLK` is indicated in the `STS` column.

   d. Using the alternate path software, delete the alternate path to the volume.

Go to the next step even if the alternate path cannot be deleted.

2. Delete the pair.

   a. From the primary storage system, delete the pair specifying the P-VOL's actual LDEV ID.

   ```
   pairsplit -g oraHA -S -d dev1 -IH0
   ```

   Note: To delete the pair specifying the P-VOL, use the -S option of the pairsplit command. Specify the actual LDEV ID (device name) of the P-VOL in the -d option.

   b. Confirm that the pair is deleted.

   ```
   pairdisplay -g oraHA -fxce -IH0
   Group   PairVol(L/R) (Port#,TID, LU), Seq#,   LDEV#.P/
   S,Status,Fence,   %,    P-LDEV# M   CTG JID AP EM       E-
   Seq# E-LDEV# R/W
   oraHA   dev1(L)      (CL1-A-0, 0,   0)311111  2222.SMPL  ----
   ------,  ----- -----   -   -   -   -  -        -      -
   -/-
   oraHA   dev1(R)      (CL1-C-1, 0,   0)322222  4444.SMPL  ----
   ------,  ----- -----   -   -   -   -  -        -      -
   -/-
   ```

3. Remove the failure.

   The following example shows recovery from a pool-volume failure.

   a. Recover a pool volume that configures the S-VOL (DP-VOL).

   b. Display the status of the pool volumes to confirm that the pool volume has been recovered.

   ```
   raidcom get ldev -ldev_list pool -pool_id 0 -IH1
   (snip)
   LDEV : 16384
   (snip)
   STS : NML
   (snip)
   ```

   For a normal volume, `NML` is indicated in the `STS` column.

4. If the volume cannot be recovered, follow the procedure below to create the S-VOL again:

   a. At the secondary storage system, delete the LU path to the S-VOL.

   b. Delete the S-VOL.

   c. Create a new volume.

   d. Set an LU path to the new volume.

5. (VSP G1000 only) Confirm the quorum disk status. If the quorum disk is blocked, recover from the blockade.

6. Re-create the pair.

   a. If you created a volume in step 4, set the GAD reserve attribute to the created volume.

   ```
   raidcom map resource -ldev_id 0x4444 -virtual_ldev_id reserve
   -IH1
   ```

   b. From the primary storage system, create the pair specifying the P-VOL's actual LDEV ID.

```
paircreate -g oraHA -f never -vl -jq 0 -d dev1 -IH0
```

Note: To create the pair specifying the P-VOL, specify the actual LDEV ID (device name) of the P-VOL in the -d option of the paircreate command.

c. Confirm that the P-VOL and S-VOL pair statuses change to PAIR (Mirror (RL)).

```
pairdisplay -g oraHA -fxce -IH0
Group   PairVol(L/R) (Port#,TID, LU),Seq#,LDEV#.P/S,Status,
Fence,   %,P-LDEV# M CTG JID AP EM       E-Seq# E-LDEV# R/W
oraHA   dev1(L)     (CL1-A-0, 0,   0)311111  2222.P-VOL PAIR
NEVER ,  100  4444 -   -   0  -  -            -       - L/M
oraHA   dev1(R)     (CL1-C-1, 0,   0)322222  4444.S-VOL PAIR
NEVER ,  100  2222 -   -   0  -  -            -       - L/M
pairdisplay -g oraHA -fxce -IH1
Group   PairVol(L/R) (Port#,TID, LU),Seq#,LDEV#.P/S,Status,
Fence,   %,P-LDEV# M CTG JID AP EM       E-Seq# E-LDEV# R/W
oraHA   dev1(L)     (CL1-C-1, 0,   0)322222  4444.S-VOL PAIR
NEVER ,  100  2222 -   -   0  -  -            -       - L/M
oraHA   dev1(R)     (CL1-A-0, 0,   0)311111  2222.P-VOL PAIR
NEVER ,  100  4444 -   -   0  -  -            -       - L/M
```

7. Using the alternate path software, add an alternate path from the server to the S-VOL.
8. Using the alternate path software, resume I/O from the server to the S-VOL.

Note: I/O from the server might resume automatically.

Related topics

- I/O modes on page 1-14
- Failure locations on page 8-3
- SIMs related to GAD on page 8-6
- Recovering the S-VOL (other than DP-VOL) (pair status: PAIR) on page 8-22
- Pair condition and recovery: quorum disk failure on page 8-50

Recovering the S-VOL (other than DP-VOL) (pair status: PAIR)

The following figure shows the failure area and recovery when a pair is suspended due to an S-VOL failure and the S-VOL is not a DP-VOL.

For details about storage system support (models, microcode) for volumes other than DP-VOLs, see Requirements and restrictions on page 2-2.

Note: In this example no consistency group is specified.

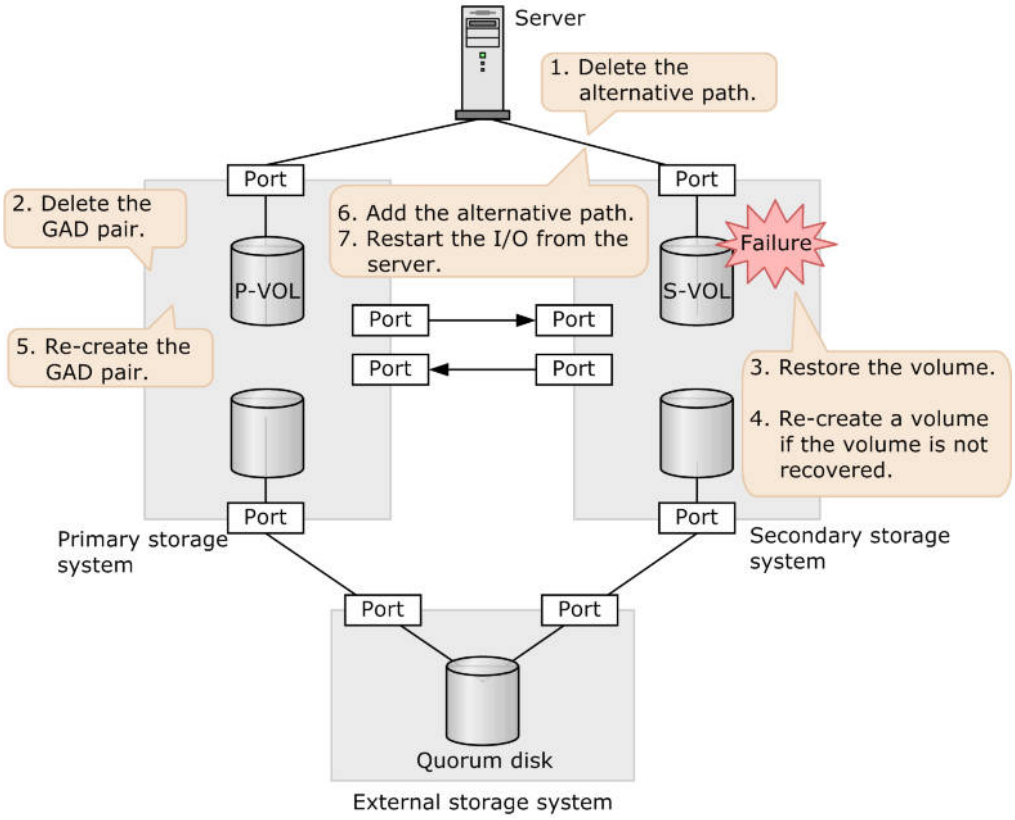

Steps for recovery

1. Delete the alternate path (logical path) to the volume that cannot be accessed from the server.
   a. Using the alternate path software, identify the volume that cannot be accessed.
   b. Confirm whether the volume (S-VOL) is blocked.

      ```
      raidcom get ldev -ldev_id 0x4444  - IH1
      (snip)
      STS : BLK
      (snip)
      ```

      For a blocked volume, `BLK` is indicated in the `STS` column.
   c. Using the alternate path software, delete the alternate path to the volume that cannot be accessed from the server.

   Go to the next step even if the alternate path cannot be deleted.
2. Delete the pair.
   a. From the primary storage system, delete the pair specifying the actual LDEV ID of the P-VOL.

      ```
      pairsplit -g oraHA -R -d dev1 -IH1
      ```

Note: To delete the pair specifying the P-VOL, use the `-S` option of the `pairsplit` command. Specify the actual LDEV ID (device name) of the P-VOL in the `-d` option.

b. Confirm that the pair is deleted.

```
pairdisplay -g oraHA -fxce  – IH0
Group PairVol(L/R) (Port#,TID, LU), Seq#, LDEV#.P/
S,Status,Fence, %, P-LDEV# M CTG JID AP EM ESeq# E-LDEV# R/W
oraHA dev1(L) (CL1-A-0, 0, 0)311111 2222.SMPL ---- ------,
----- ----- - - - - - - - -/-
oraHA dev1(R) (CL1-C-1, 0, 0)322222 4444.SMPL ---- ------,
----- ----- - - - - - - - -/-
```

3. Remove the failure. The following example shows recovery from a volume failure.

   a. Recover an S-VOL.

   b. Display the status of the P-VOL to confirm that the pool volume has been recovered.

   ```
   raidcom get ldev -ldev_id 0x4444  – IH1
   (snip)
   STS : NML
   (snip)
   ```

   For a normal volume, `NML` is indicated in the `STS` column.

4. If the volume cannot be recovered, follow the procedure below to re-create the S-VOL:

   a. At the primary storage system, delete the LU path to the S-VOL.

   b. Delete the S-VOL.

   c. Create a new volume.

   d. Set an LU path to the new volume.

5. Confirm the quorum disk status. If the quorum disk is blocked, recover from the blockade.

6. Re-create the pair.

   a. If you created a volume in step 4, set the GAD reserve attribute to the created volume.

   ```
   raidcom map resource -ldev_id 0x4444 -virtual_ldev_id reserve
   –IH1
   ```

   b. From the primary storage system, create the pair specifying the P-VOL's actual LDEV ID.

   ```
   paircreate -g oraHA -f never -vl -jq 0 -d dev1 -IH1
   ```

Note: To create the pair specifying the P-VOL, specify the actual LDEV ID (device name) of the P-VOL in the `-d` option of the `paircreate` command.

---

   c. Confirm that the P-VOL and S-VOL pair statuses have changed to PAIR (Mirror (RL)).

   ```
   pairdisplay -g oraHA -fxce -IH0
   Group PairVol(L/R) (Port#,TID, LU),Seq#,LDEV#.P/S,Status,
   Fence, %,P-LDEV# M CTG JID AP EM E-Seq# E-LDEV# R/W
   oraHA dev1(L) (CL1-A-0, 0, 0)311111 2222.P-VOL PAIR NEVER ,
   100 4444 - - 0 - - - - L/M
   oraHA dev1(R) (CL1-C-1, 0, 0)322222 4444.S-VOL PAIR NEVER ,
   100 2222 - - 0 - - - - L/M
   pairdisplay -g oraHA -fxce -IH1
   ```

```
Group PairVol(L/R) (Port#,TID, LU),Seq#,LDEV#.P/S,Status,
Fence, %,P-LDEV# M CTG JID AP EM E-Seq# E-LDEV# R/W
oraHA dev1(L) (CL1-C-1, 0, 0)322222 4444.S-VOL PAIR NEVER ,
100 2222 - - 0 - - - - L/M
oraHA dev1(R) (CL1-A-0, 0, 0)311111 2222.P-VOL PAIR NEVER ,
100 4444 - - 0 - - - - L/M
```

7. Using the alternate path software, add an alternate path from the server to the S-VOL.
8. Using the alternate path software, resume I/O from the server to the S-VOL.

Note: I/O from the server might resume automatically.

Related topics

- I/O modes on page 1-14
- Failure locations on page 8-3
- SIMs related to GAD on page 8-6
- Recovering the S-VOL (DP-VOL) (pair status: PAIR) on page 8-19
- Pair condition and recovery: quorum disk failure on page 8-50

Pair condition and recovery: full pool for the P-VOL

When the P-VOL cannot be used due to a full pool, the GAD pair is suspended.

The following table shows the transitions for pair status and I/O mode, the volumes that are accessible from the server, and the location of the latest data when you can no longer use the P-VOL due to full pool.

Before failure		After failure				
Pair status and I/O mode		Pair status and I/O mode		Volume accessible from the server		Volume with latest data
P-VOL	S-VOL	P-VOL	S-VOL	P-VOL	S-VOL	
PAIR (Mirror (RL))	PAIR (Mirror (RL))	PSUE (Block)	SSWS (Local)	NG	OK	S-VOL
PAIR (Mirror (RL))*	PAIR (Block)	PSUE (Block)	SSWS (Local)	NG	OK	S-VOL
COPY (Mirror (RL))	COPY (Block)	PSUE (Local)	PSUE (Block)	NG	NG	P-VOL
PSUS/PSUE (Local)	SSUS/PSUE (Block)	PSUS/PSUE (Local)	SSUS/PSUE (Block)	NG	NG	P-VOL

Before failure		After failure				
Pair status and I/O mode		Pair status and I/O mode		Volume accessible from the server		Volume with latest data
P-VOL	S-VOL	P-VOL	S-VOL	P-VOL	S-VOL	
PSUS/PSUE (Block)	SSWS (Local)	PSUS/PSUE (Block)	SSWS (Local)	NG	OK	S-VOL
* This row applies only to VSP G1000.						

SIM

- Primary storage system: 662*xxx*, DD1*xyy*
- Secondary storage system: DD1*xyy*

Recovery procedure

1. Increase the available pool capacity to the P-VOL.
2. Resynchronize the pair.

Related topics

- Pair condition before failure on page 8-7

Recovering a full pool for the P-VOL (pair status: PAIR)

The following figure shows the failure area and recovery when a pair is suspended due to a full pool of the P-VOL.

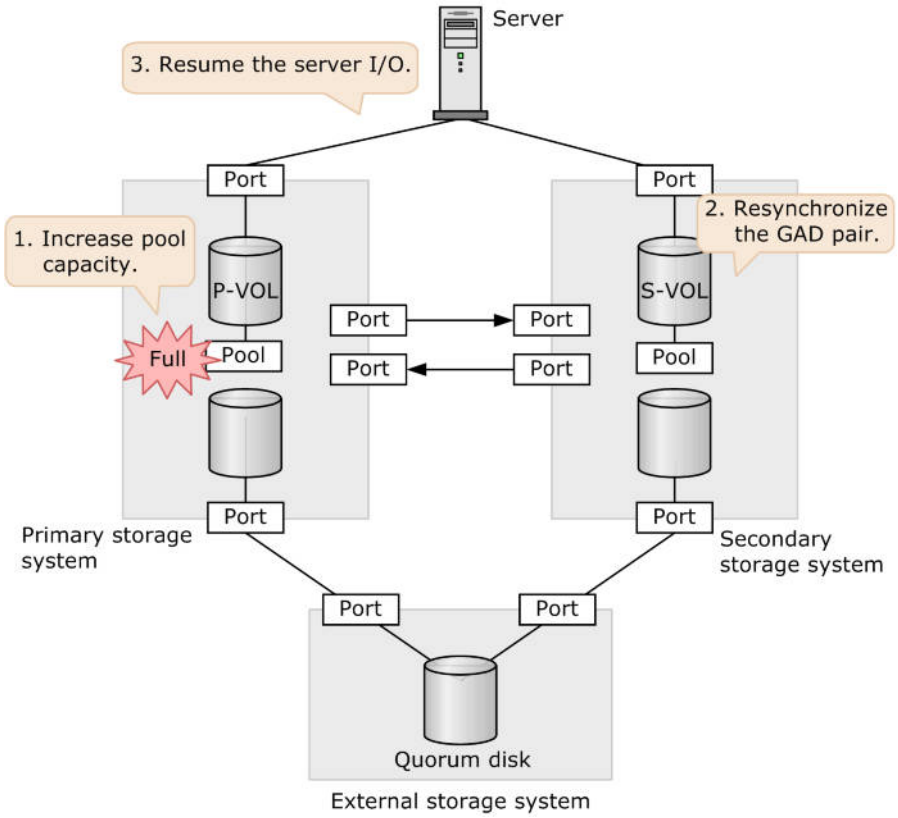

Steps for recovery

1. Increase the available capacity to the pool on which the full pool was detected.

   For details on how to increase an available pool capacity, see the *Provisioning Guide* for the storage system.

2. (VSP G1000 only) Confirm the quorum disk status. If the quorum disk is blocked, recover from the blockade.

3. Resynchronize a GAD pair.

   a. Confirm that the I/O mode of the S-VOL is Local.

   ```
   pairdisplay -g oraHA -fxce -IH1
   Group   PairVol(L/R) (Port#,TID, LU),Seq#,LDEV#.P/S,Status,
   Fence,   %,P-LDEV# M CTG JID AP EM       E-Seq# E-LDEV# R/W
   oraHA   dev1(L)     (CL1-C-1, 0,   0)322222  4444.S-VOL SSWS
   NEVER ,  100  2222 -   -   0  -  -           -       - L/L
   oraHA   dev1(R)     (CL1-A-0, 0,   0)311111  2222.P-VOL PSUE
   NEVER ,  100  4444 -   -   0  -  -           -       - B/B
   ```

   b. At the secondary storage system, resynchronize the pair.

   ```
   pairresync -g oraHA -swaps -IH1
   ```

   The volume of the primary storage system changes to an S-VOL, and the volume of the secondary storage system changes to a P-VOL.

   c. Confirm that the P-VOL and S-VOL pair statuses change to PAIR (Mirror (RL)).

   ```
   pairdisplay -g oraHA -fxce -IH0
   Group   PairVol(L/R) (Port#,TID, LU),Seq#,LDEV#.P/S,Status,
   ```

```
Fence,   %,P-LDEV# M CTG JID AP EM       E-Seq# E-LDEV# R/W
oraHA   dev1(L)     (CL1-A-0, 0,   0)311111  2222.S-VOL PAIR
NEVER ,  100  4444 -   -   0  -  -            -       - L/M
oraHA   dev1(R)     (CL1-C-1, 0,   0)322222  4444.P-VOL PAIR
NEVER ,  100  2222 -   -   0  -  -            -       - L/M
pairdisplay -g oraHA -fxce -IH1
Group   PairVol(L/R) (Port#,TID, LU),Seq#,LDEV#.P/S,Status,
Fence,   %,P-LDEV# M CTG JID AP EM       E-Seq# E-LDEV# R/W
oraHA   dev1(L)     (CL1-C-1, 0,   0)322222  4444.P-VOL PAIR
NEVER ,  100  2222 -   -   0  -  -            -       - L/M
oraHA   dev1(R)     (CL1-A-0, 0,   0)311111  2222.S-VOL PAIR
NEVER ,  100  4444 -   -   0  -  -            -       - L/M
```

4. Using the alternate path software, resume I/Os to the S-VOL that was a P-VOL before the failure (I/O might resume automatically).
5. Reverse the P-VOL and the S-VOL if necessary.

Related topics

- I/O modes on page 1-14
- Failure locations on page 8-3
- SIMs related to GAD on page 8-6
- Reversing the P-VOL and S-VOL on page 8-62
- Pair condition and recovery: quorum disk failure on page 8-50

Pair condition and recovery: full pool for the S-VOL

When the S-VOL cannot be used due to a full pool, the GAD pair is suspended.

The following table shows the transitions for pair status and I/O mode, the volumes that are accessible from the server, and the location of the latest data when you can no longer use the S-VOL due to full pool.

Before failure		After failure				
Pair status and I/O mode		Pair status and I/O mode		Volume accessible from the server		Volume with latest data
P-VOL	S-VOL	P-VOL	S-VOL	P-VOL	S-VOL	
PAIR (Mirror (RL))	PAIR (Mirror (RL))	PSUE (Local)	PSUE (Block)	OK	NG	P-VOL
PAIR (Mirror (RL))*	PAIR (Block)	PSUE (Local)	PSUE (Block)	OK	NG	P-VOL
COPY (Mirror (RL))	COPY (Block)	PSUE (Local)	PSUE (Block)	OK	NG	P-VOL
PSUS/PSUE (Local)	SSUS/PSUE (Block)	PSUS/PSUE (Local)	SSUS/PSUE (Block)	OK	NG	P-VOL

Before failure		After failure				
Pair status and I/O mode		Pair status and I/O mode		Volume accessible from the server		Volume with latest data
P-VOL	S-VOL	P-VOL	S-VOL	P-VOL	S-VOL	
PSUS/PSUE (Block)	SSWS (Local)	PSUS/PSUE (Block)	SSWS (Local)	NG	NG	S-VOL
* This row applies only to VSP G1000.						

SIM

- Primary storage system: DD1*xyy*
- Secondary storage system: 662*xxx*, DD1*xyy*

Recovery procedure

1. Increase an available pool capacity to the S-VOL.
2. Resynchronize the pair.

Related topics

- Pair condition before failure on page 8-7

Recovering a full pool for the S-VOL (pair status: PAIR)

The following figure shows the failure area and recovery when a pair is suspended due to a full pool of the S-VOL.

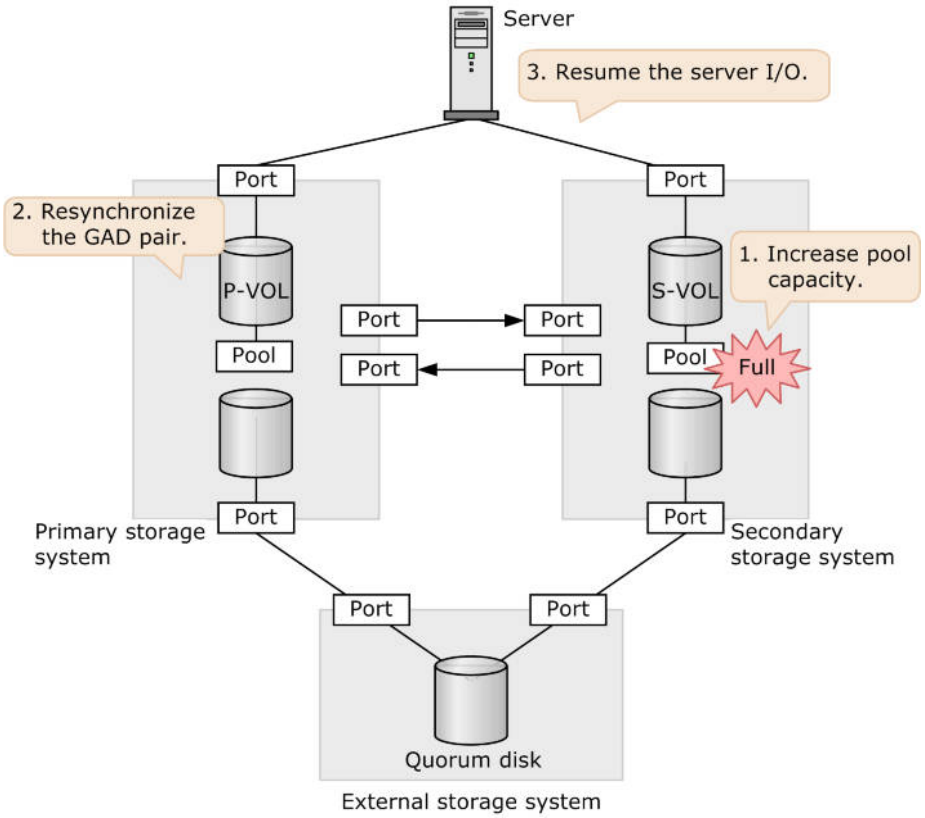

Steps for recovery

1. Increase an available capacity to the pool on which the full pool was detected.

   For details on how to increase an available pool capacity, see the *Provisioning Guide* for the storage system.

2. (VSP G1000 only) Confirm the quorum disk status. If the quorum disk is blocked, recover from the blockade.

3. Resynchronize a GAD pair.

   a. Confirm that the I/O mode of the P-VOL is Local.

   ```
   pairdisplay -g oraHA -fxce -IH0
   Group   PairVol(L/R) (Port#,TID, LU),Seq#,LDEV#.P/S,Status,
   Fence,   %,P-LDEV# M CTG JID AP EM       E-Seq# E-LDEV# R/W
   oraHA   dev1(L)     (CL1-A-0, 0,   0)311111  2222.P-VOL PSUE
   NEVER ,  100  4444 -   -   0  -  -            -       - L/L
   oraHA   dev1(R)     (CL1-C-1, 0,   0)322222  4444.S-VOL PSUE
   NEVER ,  100  2222 -   -   0  -  -            -       - B/B
   ```

   b. At the primary storage system, resynchronize the pair.

   ```
   pairresync -g oraHA -IH0
   ```

   c. Confirm that the P-VOL and S-VOL pair statuses change to PAIR (Mirror (RL)).

   ```
   pairdisplay -g oraHA -fxce -IH0
   Group   PairVol(L/R) (Port#,TID, LU),Seq#,LDEV#.P/S,Status,
   Fence,   %,P-LDEV# M CTG JID AP EM       E-Seq# E-LDEV# R/W
   oraHA   dev1(L)     (CL1-A-0, 0,   0)311111  2222.P-VOL PAIR
   ```

```
NEVER ,  100  4444 -   -   0  -  -            -       - L/M
oraHA   dev1(R)     (CL1-C-1, 0,   0)322222  4444.S-VOL PAIR
NEVER ,  100  2222 -   -   0  -  -            -       - L/M
pairdisplay -g oraHA -fxce -IH1
Group   PairVol(L/R) (Port#,TID, LU),Seq#,LDEV#.P/S,Status,
Fence,   %,P-LDEV# M CTG JID AP EM       E-Seq# E-LDEV# R/W
oraHA   dev1(L)     (CL1-C-1, 0,   0)322222  4444.S-VOL PAIR
NEVER ,  100  2222 -   -   0  -  -            -       - L/M
oraHA   dev1(R)     (CL1-A-0, 0,   0)311111  2222.P-VOL PAIR
NEVER ,  100  4444 -   -   0  -  -            -       - L/M
```

4. Using the alternate path software, resume I/O to the S-VOL (I/O might resume automatically).

Related topics

- I/O modes on page 1-14
- Failure locations on page 8-3
- SIMs related to GAD on page 8-6
- Pair condition and recovery: quorum disk failure on page 8-50

Pair condition and recovery: path failure, primary to secondary storage system

If the statuses of storage systems in both the primary and secondary sites are normal, a failure might have occurred in a physical path or switch between the storage systems.

The following table shows transitions for pair status and I/O mode, the volumes that are accessible from the server, and the location of the latest data when you can no longer use any physical path from the primary storage system to the secondary storage system.

Before failure		After failure				
Pair status and I/O mode		Pair status and I/O mode		Volume accessible from the server		Volume with latest data
P-VOL	S-VOL	P-VOL	S-VOL	P-VOL	S-VOL	
PAIR (Mirror (RL))	PAIR (Mirror (RL))	PSUE (Local)	PSUE (Block)	OK	NG	P-VOL
PAIR (Mirror (RL))*	PAIR (Block)	PSUE (Local)	PSUE (Block)	OK	NG	P-VOL
COPY (Mirror (RL))	COPY (Block)	PSUE (Local)	PSUE (Block)	OK	NG	P-VOL
PSUS/PSUE (Local)	SSUS/PSUE (Block)	PSUS/PSUE (Local)	SSUS/PSUE (Block)	OK	NG	P-VOL

Before failure		After failure				
Pair status and I/O mode		Pair status and I/O mode		Volume accessible from the server		Volume with latest data
P-VOL	S-VOL	P-VOL	S-VOL	P-VOL	S-VOL	
PSUS/PSUE (Block)	SSWS (Local)	PSUS/PSUE (Block)	SSWS (Local)	NG	OK	S-VOL
* This row applies only to VSP G1000.						

SIM

- Primary storage system: DD0*xyy*, 2180*xx*
- Secondary storage system: DD3*xyy*

Recovery procedure

1. Recover the paths from the primary storage system to the secondary storage system.
2. Resynchronize the pair.

Related topics

- Pair condition before failure on page 8-7

Recovering paths, primary to secondary storage system (pair status: PAIR)

The following figure shows the failure area and recovery when a pair is suspended due to path failure from the primary storage system to the secondary storage system.

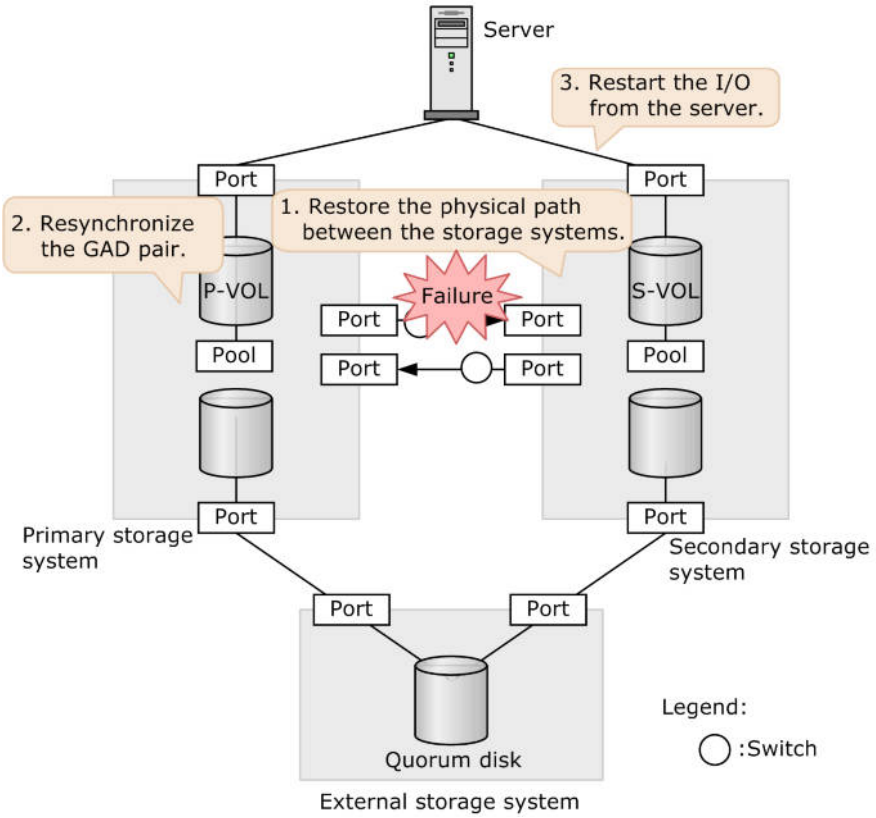

Steps for recovery

1. Reconnect the physical path or reconfigure the SAN to recover the path failure.

   When the path is recovered, the remote path is automatically recovered. If you recover the physical path but the remote path does not recover, contact Hitachi Data Systems customer support.
2. (VSP G1000 only) Confirm the quorum disk status. If the quorum disk is blocked, recover from the blockade.
3. Resynchronize the pair.

   a. Confirm that the P-VOL I/O mode is Local.

   ```
   pairdisplay -g oraHA -fxce -IH0
   Group   PairVol(L/R) (Port#,TID, LU),Seq#,LDEV#.P/S,Status,
   Fence,   %,P-LDEV# M CTG JID AP EM       E-Seq# E-LDEV# R/W
   oraHA   dev1(L)     (CL1-A-0, 0,   0)311111  2222.P-VOL PSUE
   NEVER ,  100  4444 -   -   0  -  -            -       - L/L
   oraHA   dev1(R)     (CL1-C-1, 0,   0)322222  4444.S-VOL PSUE
   NEVER ,  100  2222 -   -   0  -  -            -       - B/B
   ```

   b. At the primary storage system, resynchronize the pair.

   ```
   pairresync -g oraHA -IH0
   ```

   c. Confirm that the P-VOL and S-VOL pair statuses change to PAIR (Mirror (RL)).

   ```
   pairdisplay -g oraHA -fxce -IH0
   Group   PairVol(L/R) (Port#,TID, LU),Seq#,LDEV#.P/S,Status,
   Fence,   %,P-LDEV# M CTG JID AP EM       E-Seq# E-LDEV# R/W
   ```

```
oraHA   dev1(L)     (CL1-A-0, 0,   0)311111  2222.P-VOL PAIR
NEVER ,  100  4444 -   -   0  -  -            -       - L/M
oraHA   dev1(R)     (CL1-C-1, 0,   0)322222  4444.S-VOL PAIR
NEVER ,  100  2222 -   -   0  -  -            -       - L/M
pairdisplay -g oraHA -fxce -IH1
Group   PairVol(L/R) (Port#,TID, LU),Seq#,LDEV#.P/S,Status,
Fence,   %,P-LDEV# M CTG JID AP EM       E-Seq# E-LDEV# R/W
oraHA   dev1(L)     (CL1-C-1, 0,   0)322222  4444.S-VOL PAIR
NEVER ,  100  2222 -   -   0  -  -            -       - L/M
oraHA   dev1(R)     (CL1-A-0, 0,   0)311111  2222.P-VOL PAIR
NEVER ,  100  4444 -   -   0  -  -            -       - L/M
```

4. Using the alternate path software, resume I/O to the volume that could not be accessed from the server (I/O might resume automatically).

Related topics

- I/O modes on page 1-14
- Failure locations on page 8-3
- SIMs related to GAD on page 8-6
- Pair condition and recovery: quorum disk failure on page 8-50

Pair condition and recovery: path failure, secondary to primary storage system

If the statuses of the storage systems in both the primary and secondary sites are normal, a failure might have occurred in a physical path or switch between the storage systems.

The following table shows the transitions for pair status and I/O mode, the volumes that are accessible from the server, and the location of the latest data when you can no longer use any physical path from the secondary storage system to the primary storage system.

Before failure		After failure				
Pair status and I/O mode		Pair status and I/O mode		Volume accessible from the server		Volume with latest data
P-VOL	S-VOL	P-VOL	S-VOL	P-VOL	S-VOL	
PAIR (Mirror (RL))	PAIR (Mirror (RL))	PSUE (Block)	SSWS (Local)	NG	OK	S-VOL
PAIR (Mirror (RL))*	PAIR (Block)	PSUE (Local)	PSUE (Block)	OK	NG	P-VOL
COPY (Mirror (RL))	COPY (Block)	COPY (Mirror (RL))	COPY (Block)	OK	NG	P-VOL
PSUS/PSUE (Local)	SSUS/PSUE (Block)	PSUS/PSUE (Local)	SSUS/PSUE (Block)	OK	NG	P-VOL

Before failure		After failure				
Pair status and I/O mode		Pair status and I/O mode		Volume accessible from the server		Volume with latest data
P-VOL	S-VOL	P-VOL	S-VOL	P-VOL	S-VOL	
PSUS/PSUE (Block)	SSWS (Local)	PSUS/PSUE (Block)	SSWS (Local)	NG	OK	S-VOL
* This row applies only to VSP G1000.						

SIM

- Primary storage system: DD3*xyy*
- Secondary storage system: DD0*xyy*, 2180*xx*

Recovery procedure

1. Recover the paths from the secondary storage system to the primary storage system.
2. Resynchronize the pair.

Related topics

- Pair condition before failure on page 8-7

Recovering paths, secondary to primary storage system (pair status: PAIR)

The following figure shows the failure area and recovery when a pair is suspended due to path failure from the secondary storage system to the primary storage system.

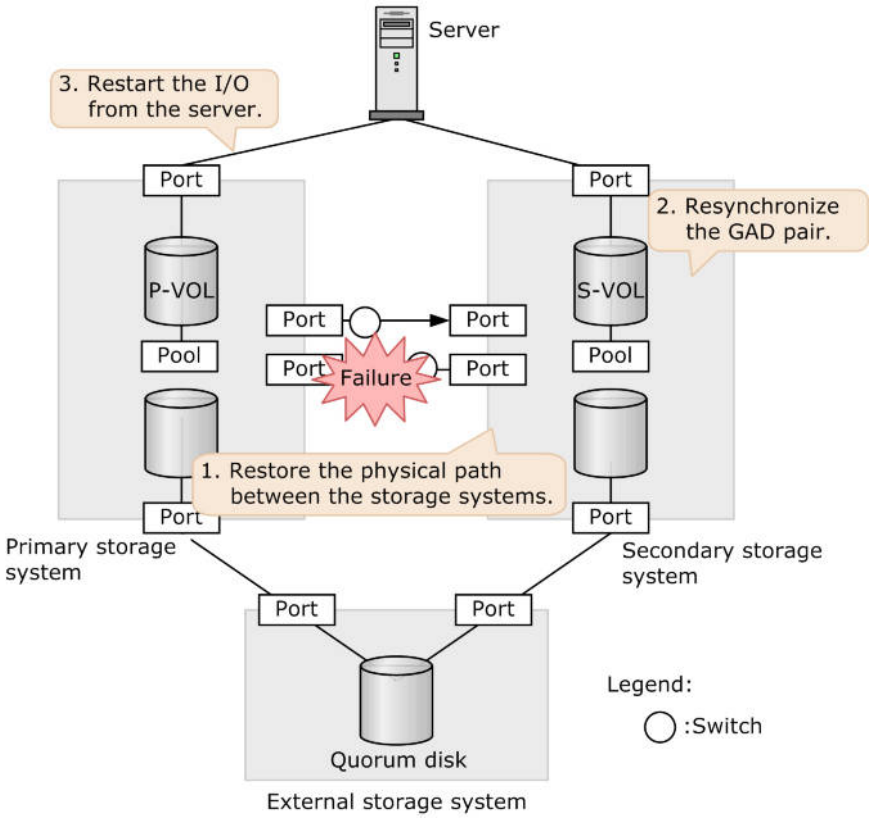

Steps for recovery

1. Reconnect the physical path or reconfigure the SAN to recover the path from the secondary storage system to the primary storage system.

   After the path is recovered, the remote path is automatically recovered. If you recover the physical path but the remote path does not recover, contact Hitachi Data Systems customer support.
2. (VSP G1000 only) Confirm the quorum disk status. If the quorum disk is blocked, recover from the blockade.
3. Resynchronize the pair.

   a. Confirm that the S-VOL I/O mode is Local.

   ```
   pairdisplay -g oraHA -fxce -IH1
   Group   PairVol(L/R) (Port#,TID, LU),Seq#,LDEV#.P/S,Status,
   Fence,   %,P-LDEV# M CTG JID AP EM       E-Seq# E-LDEV# R/W
   oraHA   dev1(L)     (CL1-C-1, 0,   0)322222  4444.S-VOL SSWS
   NEVER ,  100  2222 -   -   0  -  -            -       - L/L
   oraHA   dev1(R)     (CL1-A-0, 0,   0)311111  2222.P-VOL PSUE
   NEVER ,  100  4444 -   -   0  -  -            -       - B/B
   ```

   b. At the secondary storage system, resynchronize the pair.

   ```
   pairresync -g oraHA -swaps -IH1
   ```

   The volume on the primary storage system changes to an S-VOL, and the volume on the secondary storage system changes to a P-VOL.

   c. Confirm that the P-VOL and S-VOL pair statuses change to PAIR (Mirror (RL)).

```
pairdisplay -g oraHA -fxce -IH0
Group   PairVol(L/R) (Port#,TID, LU),Seq#,LDEV#.P/S,Status,
Fence,   %,P-LDEV# M CTG JID AP EM       E-Seq# E-LDEV# R/W
oraHA   dev1(L)     (CL1-A-0, 0,   0)311111  2222.S-VOL PAIR
NEVER ,  100  4444 -   -   0  -  -            -       - L/M
oraHA   dev1(R)     (CL1-C-1, 0,   0)322222  4444.P-VOL PAIR
NEVER ,  100  2222 -   -   0  -  -            -       - L/M
pairdisplay -g oraHA -fxce -IH1
Group   PairVol(L/R) (Port#,TID, LU),Seq#,LDEV#.P/S,Status,
Fence,   %,P-LDEV# M CTG JID AP EM       E-Seq# E-LDEV# R/W
oraHA   dev1(L)     (CL1-C-1, 0,   0)322222  4444.P-VOL PAIR
NEVER ,  100  2222 -   -   0  -  -            -       - L/M
oraHA   dev1(R)     (CL1-A-0, 0,   0)311111  2222.S-VOL PAIR
NEVER ,  100  4444 -   -   0  -  -            -       - L/M
```

4. Using the alternate path software, resume I/O to the S-VOL (P-VOL before the failure).

   I/O from the server might resume automatically.

5. Reverse the P-VOL and the S-VOL if necessary.

Related topics

- I/O modes on page 1-14
- Failure locations on page 8-3
- SIMs related to GAD on page 8-6
- Reversing the P-VOL and S-VOL on page 8-62
- Pair condition and recovery: quorum disk failure on page 8-50

Pair condition and recovery: primary storage system failure

The following table shows transitions for pair status and I/O mode, the volumes accessible from the server, and location of the latest data when you can no longer use the primary storage system due to failure.

Before failure		After failure				
Pair status and I/O mode		Pair status and I/O mode		Volume accessible from the server		Volume with latest data
P-VOL	S-VOL	P-VOL[1]	S-VOL	P-VOL	S-VOL	
PAIR (Mirror (RL))	PAIR (Mirror (RL))	PSUE (Block)	SSWS (Local)[2]	NG	OK	S-VOL
PAIR (Mirror (RL))[3]	PAIR (Block)	PSUE (Block)	PSUE (Block)	NG	NG	Both P-VOL and S-VOL
COPY (Mirror (RL))	COPY (Block)	PSUE (Local)	COPY (Block)	NG	NG	P-VOL
PSUS/PSUE (Local)	SSUS/PSUE (Block)	PSUS/PSUE (Local)	SSUS/PSUE (Block)	NG	NG	P-VOL

Before failure		After failure				
Pair status and I/O mode		Pair status and I/O mode		Volume accessible from the server		Volume with latest data
P-VOL	S-VOL	P-VOL[1]	S-VOL	P-VOL	S-VOL	
PSUS/PSUE (Block)	SSWS (Local)	PSUS/PSUE (Block)	SSWS (Local)	NG	OK	S-VOL

Notes:

1. If shared memory/control memory in the primary storage system becomes volatilized, the P-VOL status changes to SMPL, and the GAD reserve attribute is set for the volume, which prevents host access to the volume.
2. If the server does not issue write I/O, the pair status is PAIR (Mirror (RL)).
3. This row applies only to VSP G1000.

SIM

- Primary storage system: SIM varies depending on the failure type
- Secondary storage system: 2180*xx*, DD0*xyy*, DD3*xyy*

Recovery procedure

1. When the primary storage system is powered off, delete an alternate path (logical path) to the P-VOL, and then turn on the power.
   a. Using the alternate path software, distinguish the volumes which are not able to be accessed from the server.
   b. Using the alternate path software, delete the alternate paths to the P-VOL.
      If you cannot delete the alternate paths, detach all channel port connection paths (physical paths) which are connected to the server at the primary site.
2. Turn on the primary storage system.
3. Recover the primary storage system.
   For details, contact Hitachi Data Systems customer support.
4. Recover the physical path between the primary storage system and the secondary storage system.
5. If S-VOL pair status is PAIR, suspend the pair specifying the S-VOL.
6. (VSP G1000 only) Confirm the quorum disk status. If the quorum disk is blocked, recover from the blockade.
7. Resynchronize or re-create a pair using the procedure in the following table whose pair status and I/O mode match your pair's status and I/O mode.

Pair status		I/O mode		Procedure
P-VOL	S-VOL	P-VOL	S-VOL	
PSUE	COPY	Local	Block	1. Delete the pair forcibly from the P-VOL, specifying **Enable** in the **Volume Access** field of the **Delete Pairs** window. 2. Delete the pair forcibly from the S-VOL, specifying **Disable** in the **Volume Access** field of the **Delete Pairs** window. 3. Re-create the GAD pair by specifying the P-VOL.
SMPL	COPY	Not applicable	Block	1. Delete the pair forcibly from the S-VOL, specifying **Disable** in the **Volume Access** field of the **Delete Pairs** window. 2. Release the GAD reserve attribute set to the volume that was the P-VOL before the forced deletion. 3. Set the virtual LDEV ID to the volume whose reserve attribute was released. 4. Re-create the GAD pair by specifying the P-VOL.
PSUS/ PSUE	SSWS	Block	Local	Resynchronize the pair specifying the S-VOL.
SMPL	SSWS	Not applicable	Local	1. Delete the pair specifying the S-VOL. 2. When the virtual LDEV ID is set to the P-VOL, delete the virtual LDEV ID, and then set the reserve attribute to the P-VOL. 3. Re-create the pair specifying the S-VOL.
PSUS/ PSUE	SSUS/ PSUE	Local	Block	Resynchronize the pair specifying the P-VOL.
PSUE*	PSUE	Block	Block	1. Delete the pair forcibly from the primary storage system, specifying **Enable** in the **Volume Access** field of the **Delete Pairs** window. 2. Delete the pair forcibly from the secondary storage system, specifying **Disable** in the **Volume Access** field of the **Delete Pairs** window. 3. Re-create the pair specifying the P-VOL.
SMPL	SSUS/ PSUE	Not applicable	Block	1. Delete the pair forcibly from the secondary storage system, specifying **Disable** in the **Volume**

Pair status		I/O mode		Procedure
P-VOL	S-VOL	P-VOL	S-VOL	
				Access field of the **Delete Pairs** window. 2. Release the reserve attribute of the P-VOL, and then set the same virtual LDEV ID that was used before the pair was deleted. When you set the virtualization information with Device Manager - Storage Navigator, use the Edit Virtualization Management Settings window. For details, see the *Provisioning Guide* for the storage system. 3. Re-create the pair specifying the P-VOL.
* This row applies only to VSP G1000.				

8. If the alternate path to the P-VOL has been deleted, add the alternate path.
   a. If you have detached the channel port connection paths of the primary site, restore all channel port connection paths to their original status, and then add the alternate path.
   b. Using the alternate path software, add the alternate path deleted at step 1 to the P-VOL.

Related topics

- Pair condition before failure on page 8-7
- Pair condition and recovery: quorum disk failure on page 8-50

Pair condition and recovery: secondary storage system failure

The following table shows transitions for pair status and I/O mode, the volumes accessible from the server, and location of the latest data when you can no longer use the secondary storage system due to failure.

Before failure		After failure				
Pair status and I/O mode		Pair status and I/O mode		Volume accessible from the server		Volume with latest data
P-VOL	S-VOL	P-VOL	S-VOL[1]	P-VOL	S-VOL	
PAIR (Mirror (RL))	PAIR (Mirror (RL))	PSUE (Local)[2]	PSUE (Block)	OK	NG	P-VOL

Before failure		After failure				
Pair status and I/O mode		Pair status and I/O mode		Volume accessible from the server		Volume with latest data
P-VOL	S-VOL	P-VOL	S-VOL[1]	P-VOL	S-VOL	
PAIR (Mirror (RL))[3]	PAIR (Block)	PSUE (Local)	PSUE (Block)	OK	NG	P-VOL
COPY (Mirror (RL))	COPY (Block)	PSUE (Local)	PSUE (Block)	OK	NG	P-VOL
PSUS/PSUE (Local)	SSUS/PSUE (Block)	PSUS/PSUE (Local)	SSUS/PSUE (Block)	OK	NG	P-VOL
PSUS/PSUE (Block)	SSWS (Local)	PSUS/PSUE (Block)	SSWS (Local)	NG	NG	S-VOL

Notes:

1. If shared memory/control memory in the secondary storage system becomes volatilized, the S-VOL pair status changes to SMPL, and the reserve attribute is set for the volume, which prevents host access to the volume.
2. If the server does not issue write I/O, the pair status might be PAIR (Mirror (RL)).
3. This row applies only to VSP G1000.

SIM

- Primary storage system: 2180*xx*, DD0*xyy*, DD3*xyy*
- Secondary storage system: SIM varies depending on the failure type

Recovery procedure

1. When the secondary storage system is powered off, delete an alternate path (logical path) to the S-VOL, and then turn on the power.
   a. Using the alternate path software, distinguish the volumes which are not able to be accessed from the server.
   b. Using the alternate path software, delete the alternate paths to the S-VOL.
      If you cannot delete the alternate paths, detach all channel port connection paths (physical paths) which are connected to the server at the secondary site.
2. Turn on the secondary storage system.
3. Recover the secondary storage system.
   For details, contact Hitachi Data Systems customer support.
4. Recover the physical path between the primary storage system and the secondary storage system.
5. If P-VOL pair status is PAIR, suspend the pair specifying the P-VOL.

6. (VSP G1000 only) Confirm the quorum disk status. If the quorum disk is blocked, recover from the blockade.
7. Resynchronize or re-create the pair using the procedure in the following table whose pair status and I/O mode match your pair's status and I/O mode.

Pair status		I/O mode		Procedure
P-VOL	S-VOL	P-VOL	S-VOL	
PSUS/ PSUE	PSUS/ PSUE	Local	Block	Resynchronize the pair specifying the P-VOL.
PSUS/ PSUE	SMPL	Local	Not applicable	1. Delete the pair specifying the P-VOL. 2. When the virtual LDEV ID is set to the S-VOL, delete the virtual LDEV ID, and then set the reserve attribute to the S-VOL. 3. Re-create the pair specifying the P-VOL.
PSUS/ PSUE	SSWS	Block	Local	Resynchronize the pair specifying the S-VOL.
PSUS/ PSUE	SMPL	Block	Not applicable	1. Delete the pair forcibly from the primary storage system, specifying **Disable** in the **Volume Access** field (**Delete Pairs** window of Device Manager - Storage Navigator). 2. Release the reserve attribute from the S-VOL, and then set the virtual LDEV ID that was used before the pair was deleted. 3. Re-create the pair specifying the S-VOL.

8. If the alternate path to the S-VOL has been deleted, add the alternate path.
   a. If you have detached the channel port connection paths of the secondary site, restore all channel port connection paths to their original status, and then add the alternate path.
   b. Using the alternate path software, add the alternate path deleted at step 1 to the S-VOL.

Related topics

- Pair condition before failure on page 8-7
- Pair condition and recovery: quorum disk failure on page 8-50

Pair condition and recovery: path failure, primary to external storage system

If the status of the external storage system is normal, a failure might have occurred in a physical path from the primary or secondary storage system to the external storage system, or a switch. Recover from the failure that occurred in the physical path or switch.

The following table shows transitions for pair status and I/O mode, the volumes that are accessible from the server, and the location of the latest data when you can no longer use any physical path from the primary storage system to the quorum disk's external storage system.

Before failure		After failure				
Pair status and I/O mode		Pair status and I/O mode		Volume accessible from the server		Volume with latest data
P-VOL	S-VOL	P-VOL	S-VOL	P-VOL	S-VOL	
PAIR (Mirror (RL))[1]	PAIR (Mirror (RL))[1]	PSUE (Local)	PSUE (Block)	OK	NG	P-VOL
		PAIR (Mirror (RL))[2]	PAIR (Block)	OK	NG	Both P-VOL and S-VOL
PAIR (Mirror (RL))[2]	PAIR (Block)	PAIR (Mirror (RL))	PAIR (Block)	OK	NG	Both P-VOL and S-VOL
COPY (Mirror (RL))	COPY (Block)	PSUE (Local)	PSUE (Block)	OK	NG	P-VOL
PSUS/PSUE (Local)	SSUS/PSUE (Block)	PSUS/PSUE (Local)	SSUS/PSUE (Block)	OK	NG	P-VOL
PSUS/PSUE (Block)	SSWS (Local)	PSUS/PSUE (Block)	SSWS (Local)	NG	OK	S-VOL

Notes:

1. The pair status and I/O mode after failure depends on the requirement of the pair. For details, see I/O stoppage and data duplication with blocked quorum disk (VSP G1000) on page 1-21.
2. These rows apply only to VSP G1000.

SIM

- Primary storage system: 21D0*xy*, 21D2*xx*, DD2*xyy*, DEF0*zz*, EF5*xyy*, EFD000, FF5*xyy*
- Secondary storage system: DD2*xyy*

Recovery procedure

1. Recover the paths to the external storage system.
2. Resynchronize the pair.

Related topics

- Pair condition before failure on page 8-7

Recovering the path: primary to external storage system (pair status: PAIR)

The following figure shows the failure area and recovery when the GAD status of a pair changes to Suspended or Quorum disk blocked (VSP G1000 only) due to path failure from the primary storage system to the external storage system.

When the GAD status of the pair changes to Suspended, the P-VOL I/O mode changes to Local and the S-VOL I/O mode changes to Block. Server I/O continues on the P-VOL. When the GAD status of the pair changes to Quorum disk blocked (VSP G1000 only), the P-VOL I/O mode remains Mirror (RL), and the S-VOL I/O mode changes to Block. Server I/O continues on the P-VOL.

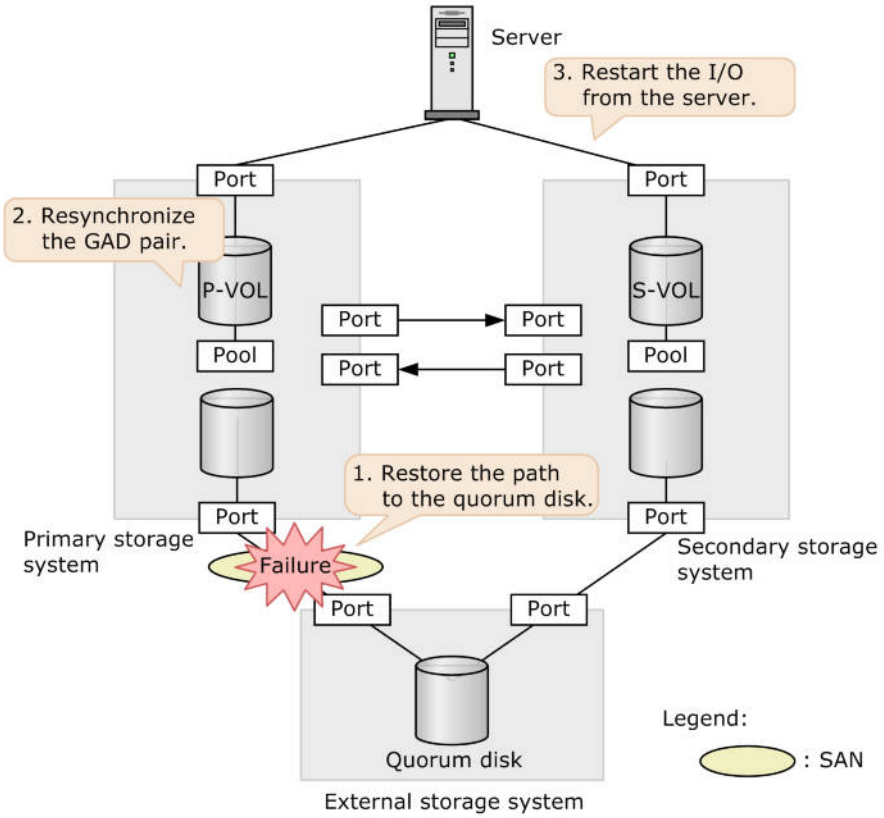

Steps for recovery

1. Recover the path to the external storage system.
   a. Reconnect the physical path or reconfigure the SAN to recover the path to the external storage system.
      After the path is recovered, the remote path is automatically recovered.
   b. Confirm that the external storage system is connected correctly.

```
raidcom get path -path_grp 1 -IH0
PHG GROUP STS CM IF MP# PORT   WWN                  PR LUN
PHS  Serial# PRODUCT_ID LB PM
  1 1-1   NML E  D    0 CL5-A  50060e8007823520     1   0
NML    33333 VSP G1000  N  M
```

Note: The `PRODUCT_ID` for the VSP G200, G400, G600, G800 is `VSP Gx00`.

c. Confirm the LDEV ID of the quorum disk by obtaining the information of the external volume from the primary storage system.

```
raidcom get external_grp -external_grp_id 1-1 -IH0
T GROUP  P_NO  LDEV#   STS         LOC_LBA
SIZE_LBA            Serial#
E 1-1       0   9999   NML  0x000000000000
0x000003c00000             333333
```

d. Confirm that the primary storage system recognizes the external volume as a quorum disk by specifying the LDEV ID of the quorum disk.

```
raidcom get ldev -ldev_id 0x9999 -fx -IH0
(snip)
QRDID : 0
QRP_Serial# : 322222
QRP_ID : R8
(snip)
```

Note: The VSP G1000 is displayed as `R8` in command output. The VSP G200, G400, G600, G800 is displayed as `M8` in command output.

2. If the GAD status of the pair is Quorum disk blocked (VSP G1000 only): The pair changes to Mirrored status automatically.

   If the GAD status of the pair is Suspended: Resynchronize the pair as follows.

   a. Confirm that the I/O mode of the P-VOL is Local.

```
pairdisplay -g oraHA -fxce -IH0
Group   PairVol(L/R) (Port#,TID, LU),Seq#,LDEV#.P/S,Status,
Fence,   %,P-LDEV# M CTG JID AP EM       E-Seq# E-LDEV# R/W
oraHA   dev1(L)     (CL1-A-0, 0,   0)311111  2222.P-VOL PSUE
NEVER ,  100  4444 -   -   0  -  -            -       - L/L
oraHA   dev1(R)     (CL1-C-1, 0,   0)322222  4444.S-VOL PSUE
NEVER ,  100  2222 -   -   0  -  -            -       - B/B
```

   b. At the primary storage system, resynchronize the pair.

```
pairresync -g oraHA -IH0
```

   c. Confirm that the P-VOL and S-VOL pair statuses change to PAIR (Mirror (RL)).

```
pairdisplay -g oraHA -fxce -IH0
Group   PairVol(L/R) (Port#,TID, LU),Seq#,LDEV#.P/S,Status,
Fence,   %,P-LDEV# M CTG JID AP EM       E-Seq# E-LDEV# R/W
oraHA   dev1(L)     (CL1-A-0, 0,   0)311111  2222.P-VOL PAIR
NEVER ,  100  4444 -   -   0  -  -            -       - L/M
oraHA   dev1(R)     (CL1-C-1, 0,   0)322222  4444.S-VOL PAIR
NEVER ,  100  2222 -   -   0  -  -            -       - L/M
pairdisplay -g oraHA -fxce -IH1
Group   PairVol(L/R) (Port#,TID, LU),Seq#,LDEV#.P/S,Status,
```

```
Fence,   %,P-LDEV# M CTG JID AP EM       E-Seq# E-LDEV# R/W
oraHA   dev1(L)     (CL1-C-1, 0,   0)322222  4444.S-VOL PAIR
NEVER ,  100  2222 -   -   0  -  -            -       - L/M
oraHA   dev1(R)     (CL1-A-0, 0,   0)311111  2222.P-VOL PAIR
NEVER ,  100  4444 -   -   0  -  -            -       - L/M
```

3. Using the alternate path software, resume I/O to the volume that could not be accessed from the server (I/O might resume automatically).

Related topics

- I/O modes on page 1-14
- Failure locations on page 8-3
- SIMs related to GAD on page 8-6

Pair condition and recovery: path failure, secondary to external storage system

If the status of the external storage system is normal, a failure might have occurred in a physical path from the primary or secondary storage system to the external storage system, or a switch. Recover from the failure that occurred in the physical path or switch.

The following table shows transitions for pair status and I/O mode, the volumes that are accessible from the server, and the location of the latest data when you can no longer use any physical path from the secondary storage system to the quorum disk's external storage system.

Before failure		After failure				
Pair status and I/O mode		Pair status and I/O mode		Volume accessible from the server		Volume with latest data
P-VOL	S-VOL	P-VOL	S-VOL[1]	P-VOL	S-VOL	
PAIR (Mirror (RL))[1]	PAIR (Mirror (RL))[1]	PSUE (Local)	PSUE (Block)	OK	NG	P-VOL
		PAIR (Mirror (RL))[2]	PAIR (Block)	OK	NG	Both P-VOL and S-VOL
PAIR (Mirror (RL))[2]	PAIR (Block)	PAIR (Mirror (RL))	PAIR (Block)	OK	NG	Both P-VOL and S-VOL
COPY (Mirror (RL))	COPY (Block)	PSUE (Local)	PSUE (Block)	OK	NG	P-VOL
PSUS/PSUE (Local)	SSUS/PSUE (Block)	PSUS/PSUE (Local)	SSUS/PSUE (Block)	OK	NG	P-VOL
PSUS/PSUE (Block)	SSWS (Local)	PSUS/PSUE (Block)	SSWS (Local)	NG	OK	S-VOL
Notes:						

Before failure		After failure				
Pair status and I/O mode		Pair status and I/O mode		Volume accessible from the server		Volume with latest data
P-VOL	S-VOL	P-VOL	S-VOL[1]	P-VOL	S-VOL	

1. The pair status and I/O mode after failure depends on the requirement of the pair. For details, see I/O stoppage and data duplication with blocked quorum disk (VSP G1000) on page 1-21.
2. This row applies only to VSP G1000.

SIM

- Primary storage system: DD2*xyy*
- Secondary storage system: 21D0*xy*, 21D2*xx*, DD2*xyy*, DEF0*zz*, EF5*xyy*, EFD000, FF5*xyy*

Recovery procedure

1. Recover the paths to the external storage system.
2. Resynchronize the pair.

Related topics

- Pair condition before failure on page 8-7

Recovering the path: secondary to external storage system (pair status: PAIR)

The following figure shows the failure area and recovery when the GAD status of a pair changes to Suspended or Quorum disk blocked (VSP G1000 only) due to path failure from the secondary storage system to the external storage system.

When the GAD status of the pair changes to Suspended, the P-VOL I/O mode changes to Local and the S-VOL I/O mode changes to Block. Server I/O continues on the P-VOL. When the GAD status of the pair changes to Quorum disk blocked (VSP G1000 only), the P-VOL I/O mode remains Mirror (RL), and the S-VOL I/O mode changes to Block. Server I/O continues on the P-VOL.

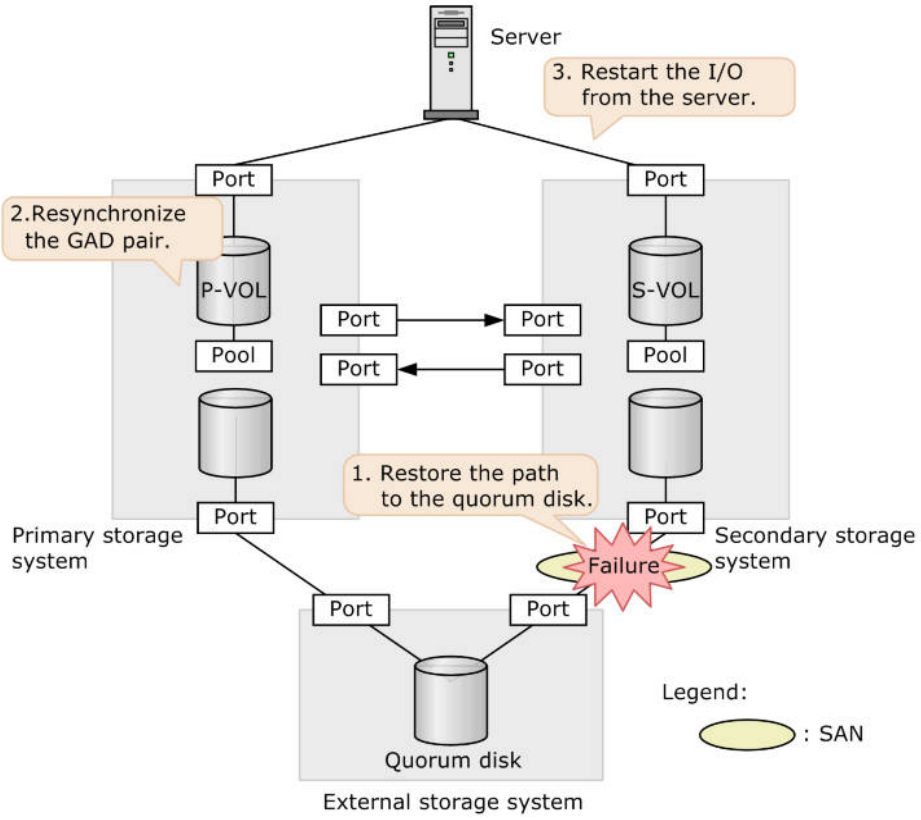

Steps for recovery

1. Recover the path to the external storage system.

   a. Reconnect the physical path or reconfigure the SAN to recover the path to the external storage system.

      After the path is recovered, the remote path is automatically recovered.

   b. Confirm that the external storage system is connected correctly.

```
raidcom get path -path_grp 1 -IH1
PHG GROUP STS CM IF MP# PORT    WWN                  PR LUN
PHS  Serial# PRODUCT_ID LB PM
  1 1-2    NML E  D     0 CL5-C  50060e8007823521     1   0
NML    33333 VSP G1000  N  M
```

Note: The `PRODUCT_ID` for the VSP G200, G400, G600, G800 is `VSP Gx00`.

   c. Confirm the LDEV ID of the quorum disk by obtaining the information of the external volume from the secondary storage system.

```
raidcom get external_grp -external_grp_id 1-2 -IH1
T GROUP  P_NO  LDEV#   STS         LOC_LBA
SIZE_LBA            Serial#
E 1-2       0   9999   NML  0x000000000000
0x000003c00000             333333
```

d. Confirm that the secondary storage system recognizes the external volume as a quorum disk by specifying the LDEV ID of the quorum disk.

```
raidcom get ldev -ldev_id 0x8888 -fx -IH1
(snip)
QRDID : 0
QRP_Serial# : 311111
QRP_ID : R8
(snip)
```

Note: The VSP G1000 is displayed as `R8` in command output. The VSP G200, G400, G600, G800 is displayed as `M8` in command output.

2. If the GAD status of the pair is Quorum disk blocked (VSP G1000 only): The pair changes to Mirrored status automatically.

   If the GAD status of the pair is Suspended: Resynchronize the pair as follows.

   a. Confirm that the P-VOL I/O mode is Local.

```
pairdisplay -g oraHA -fxce -IH0
Group   PairVol(L/R) (Port#,TID, LU),Seq#,LDEV#.P/S,Status,
Fence,   %,P-LDEV# M CTG JID AP EM       E-Seq# E-LDEV# R/W
oraHA   dev1(L)     (CL1-A-0, 0,   0)311111  2222.P-VOL PSUE
NEVER ,  100  4444 -   -   0  -  -            -       - L/L
oraHA   dev1(R)     (CL1-C-1, 0,   0)322222  4444.S-VOL PSUE
NEVER ,  100  2222 -   -   0  -  -            -       - B/B
```

   b. At the primary storage system, resynchronize the pair.

```
pairresync -g oraHA -IH0
```

   c. Confirm that the P-VOL and S-VOL pair statuses change to PAIR (Mirror (RL)).

```
pairdisplay -g oraHA -fxce -IH0
Group   PairVol(L/R) (Port#,TID, LU),Seq#,LDEV#.P/S,Status,
Fence,   %,P-LDEV# M CTG JID AP EM       E-Seq# E-LDEV# R/W
oraHA   dev1(L)     (CL1-A-0, 0,   0)311111  2222.P-VOL PAIR
NEVER ,  100  4444 -   -   0  -  -            -       - L/M
oraHA   dev1(R)     (CL1-C-1, 0,   0)322222  4444.S-VOL PAIR
NEVER ,  100  2222 -   -   0  -  -            -       - L/M
pairdisplay -g oraHA -fxce -IH1
Group   PairVol(L/R) (Port#,TID, LU),Seq#,LDEV#.P/S,Status,
Fence,   %,P-LDEV# M CTG JID AP EM       E-Seq# E-LDEV# R/W
oraHA   dev1(L)     (CL1-C-1, 0,   0)322222  4444.S-VOL PAIR
NEVER ,  100  2222 -   -   0  -  -            -       - L/M
oraHA   dev1(R)     (CL1-A-0, 0,   0)311111  2222.P-VOL PAIR
NEVER ,  100  4444 -   -   0  -  -            -       - L/M
```

3. Using the alternate path software, resume I/O to the S-VOL (I/O might resume automatically).

Related topics

- I/O modes on page 1-14
- Failure locations on page 8-3
- SIMs related to GAD on page 8-6

Pair condition and recovery: quorum disk failure

The following table shows transitions for pair status and I/O mode, the volumes that are accessible from the server, and the location of the latest data when you can no longer use the quorum disk volume.

Before failure		After failure				
Pair status and I/O mode		Pair status and I/O mode		Volume accessible from the server		Volume with latest data
P-VOL	S-VOL	P-VOL	S-VOL[1]	P-VOL	S-VOL	
PAIR (Mirror (RL))[1]	PAIR (Mirror (RL))[1]	PSUE (Local)	PSUE (Block)	OK	NG	P-VOL
		PAIR (Mirror (RL))[2]	PAIR (Block)	OK	NG	Both P-VOL and S-VOL
PAIR (Mirror (RL))[2]	PAIR (Block)	PAIR (Mirror (RL))	PAIR (Block)	OK	NG	Both P-VOL and S-VOL
COPY (Mirror (RL))	COPY (Block)	PSUE (Local)	PSUE (Block)	OK	NG	P-VOL
PSUS/PSUE (Local)	SSUS/PSUE (Block)	PSUS/PSUE (Local)	SSUS/PSUE (Block)	OK	NG	P-VOL
PSUS/PSUE (Block)	SSWS (Local)	PSUS/PSUE (Block)	SSWS (Local)	NG	OK	S-VOL

Notes:

1. The pair status and I/O mode after failure depends on the requirement of the pair. For details, see I/O stoppage and data duplication with blocked quorum disk (VSP G1000) on page 1-21.
2. This row applies only to VSP G1000.

SIM

- Primary storage system: 21D0*xy*, 21D2*xx*, DD2*xyy*, DEF0*zz*, EF5*xyy*, EFD000, FF5*xyy*
- Secondary storage system: 21D0*xy*, 21D2*xx*, DD2*xyy*, DEF0*zz*, EF5*xyy*, EFD000, FF5*xyy*

Recovery procedure

1. Recover the quorum disk.
2. Resynchronize the pair.

Related topics

- Pair condition before failure on page 8-7

Recovering the quorum disk (pair status: PAIR)

When the GAD status of the pair changes to Suspended, the P-VOL I/O mode changes to Local and the S-VOL I/O mode changes to Block. Server I/O continues on the P-VOL. When the GAD status of the pair changes to Quorum disk blocked (VSP G1000 only), the P-VOL I/O mode remains Mirror (RL), and the S-VOL I/O mode changes to Block. Server I/O continues on the P-VOL.

The following figure shows the failure area and recovery when a pair is suspended due to quorum disk failure.

Note: In this example no consistency group is specified.

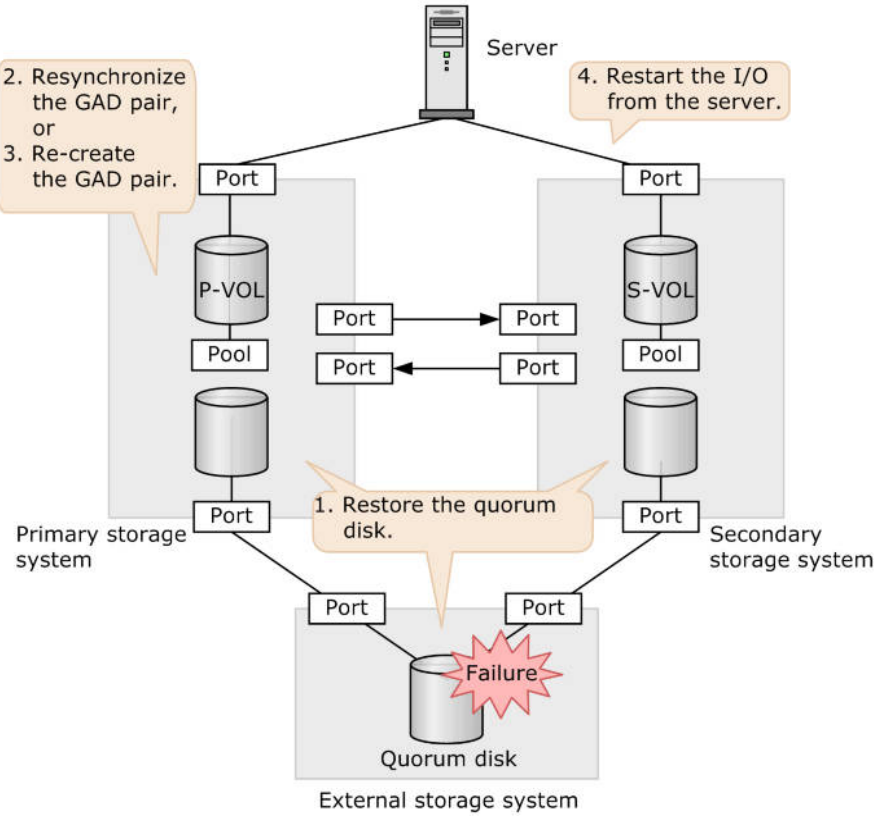

Steps for recovery

Note: The following procedure is also used for re-creating the quorum disk when it has been mistakenly reformatted.

Note: Steps 1 and 2 below describe the recovery procedure for an external storage system made by Hitachi, for example, VSP. If you are using another vendor's storage system as external storage, follow the vendor's recovery procedure for the storage system. When you complete the recovery procedure for the external storage system, start the following procedure at step 3.

1. On the external storage system, recover the quorum disk.

a. Block the quorum disk.

b. Format the quorum disk.

   If the quorum disk recovers after formatting, go to step h.

   If the quorum disk does not recover after formatting, continue to step c.

c. Confirm the following information about the quorum disk:

   - Vendor
   - Machine name
   - Volume properties
   - Device ID (if the information is valid)
   - Serial number
   - SSID
   - Product ID
   - LBA capacity (the capacity must be larger than the quorum disk before the failure occurred)
   - CVS attribute

   For details about confirming this information, see the *Hitachi Universal Volume Manager User Guide*.

   For details about confirming the CVS attribute, see Table 8-1 Confirming the CVS attribute on the external storage system on page 8-54.

d. Delete the LU path to the quorum disk.

e. Delete the volume that is used as the quorum disk.

f. Create a new volume.

   For the LDEV ID, set the same value as the LDEV ID of the quorum disk that has been used since before the failure occurred. If you cannot set the same value, go to step 3.

   Also set the same values for the following information as the values that were used before the failure occurred. If you cannot set the same value, go to step 3.

   - Vendor
   - Machine name
   - Volume properties
   - Device ID (if the information is valid)
   - Serial number
   - SSID
   - Product ID
   - LBA capacity (the capacity must be larger than the quorum disk before the failure occurred)
   - CVS attribute

   For details about confirming this information, see the *Hitachi Universal Volume Manager User Guide*. For details about confirming the CVS attribute, see Table 8-1 Confirming the CVS attribute on the external storage system on page 8-54 and Table 8-2 Conditions for the CVS

attribute for volumes created in the external storage system on page 8-55.

g. Set an LU path to the new volume.

For the LU number, set the same value as the LU number of the quorum disk that was used since before the failure occurred. If you cannot set the same value, go to step 3.

h. Reconnect the external storage system or the quorum disk to the primary and secondary storage systems.

2. If the GAD status of the pair is Quorum disk blocked (VSP G1000 only): The pair changes to the Mirrored status automatically.

If the GAD status of the pair is Suspended: Resynchronize the pair as follows.

a. Confirm that the P-VOL I/O mode is Local.

```
pairdisplay -g oraHA -fxce -IH0
Group   PairVol(L/R) (Port#,TID, LU),Seq#,LDEV#.P/S,Status,
Fence,   %,P-LDEV# M CTG JID AP EM       E-Seq# E-LDEV# R/W
oraHA   dev1(L)     (CL1-A-0, 0,   0)311111  2222.P-VOL PSUE
NEVER ,  100  4444 -   -   0  -  -            -       - L/L
oraHA   dev1(R)     (CL1-C-1, 0,   0)322222  4444.S-VOL PSUE
NEVER ,  100  2222 -   -   0  -  -            -       - B/B
```

b. On the primary storage system, resynchronize the pair.

```
pairresync -g oraHA -IH0
```

c. Confirm that the P-VOL and S-VOL pair status has changed to PAIR (Mirror (RL)). If so, go to step 4.

```
pairdisplay -g oraHA -fxce -IH0
Group   PairVol(L/R) (Port#,TID, LU),Seq#,LDEV#.P/S,Status,
Fence,   %,P-LDEV# M CTG JID AP EM       E-Seq# E-LDEV# R/W
oraHA   dev1(L)     (CL1-A-0, 0,   0)311111  2222.P-VOL PAIR
NEVER ,  100  4444 -   -   0  -  -            -       - L/M
oraHA   dev1(R)     (CL1-C-1, 0,   0)322222  4444.S-VOL PAIR
NEVER ,  100  2222 -   -   0  -  -            -       - L/M
pairdisplay -g oraHA -fxce -IH1
Group   PairVol(L/R) (Port#,TID, LU),Seq#,LDEV#.P/
S,Status,Fence,   %,P-LDEV# M CTG JID AP EM       E-Seq# E-
LDEV# R/W
oraHA   dev1(L)     (CL1-C-1, 0,   0)322222  4444.S-VOL PAIR
NEVER ,  100  2222 -   -   0  -  -            -       - L/M
oraHA   dev1(R)     (CL1-A-0, 0,   0)311111  2222.P-VOL PAIR
NEVER ,  100  4444 -   -   0  -  -            -       - L/M
```

3. Re-create the pairs.

a. On the primary storage system, delete all pairs that use the quorum disk where the failure occurred.

```
pairsplit -g oraHA -S -d dev1 -IH0
```

b. Confirm that the pairs were deleted.

```
pairdisplay -g oraHA -fxce -IH0
Group   PairVol(L/R) (Port#,TID, LU), Seq#,   LDEV#.P/
S,Status,Fence,   %,    P-LDEV# M   CTG JID AP EM       E-
Seq# E-LDEV# R/W
oraHA   dev1(L)      (CL1-A-0, 0,   0)311111  2222.SMPL  ----
------,  ----- -----   -   -   -   -  -        -      -
```

```
-/-
oraHA   dev1(R)      (CL1-C-1, 0,   0)322222  4444.SMPL  ----
------,  ----- -----   -   -   -   -  -        -      -
-/-
```

c. On the primary and secondary storage systems, delete the quorum disk.

d. On the primary and secondary storage systems, add a quorum disk.

e. On the primary storage system, create the pairs.

```
paircreate -g oraHA -f never -vl -jq 0 -d dev1 -IH0
```

f. Confirm that the P-VOL and S-VOL pair statuses change to PAIR (Mirror (RL)).

```
pairdisplay -g oraHA -fxce -IH0
Group   PairVol(L/R) (Port#,TID, LU),Seq#,LDEV#.P/S,Status,
Fence,   %,P-LDEV# M CTG JID AP EM       E-Seq# E-LDEV# R/W
oraHA   dev1(L)     (CL1-A-0, 0,   0)311111  2222.P-VOL PAIR
NEVER ,  100  4444 -   -   0  -  -            -       - L/M
oraHA   dev1(R)     (CL1-C-1, 0,   0)322222  4444.S-VOL PAIR
NEVER ,  100  2222 -   -   0  -  -            -       - L/M
pairdisplay -g oraHA -fxce -IH1
Group   PairVol(L/R) (Port#,TID, LU),Seq#,LDEV#.P/S,Status,
Fence,   %,P-LDEV# M CTG JID AP EM       E-Seq# E-LDEV# R/W
oraHA   dev1(L)     (CL1-C-1, 0,   0)322222  4444.S-VOL PAIR
NEVER ,  100  2222 -   -   0  -  -            -       - L/M
oraHA   dev1(R)     (CL1-A-0, 0,   0)311111  2222.P-VOL PAIR
NEVER ,  100  4444 -   -   0  -  -            -       - L/M
```

4. Using the alternate path software, resume I/O to the S-VOL (I/O might resume automatically).

Note: When the external storage system is installed at the primary site, if a failure occurs in both the primary storage system and the external storage system, forcibly delete the pair on the secondary storage system, and then re-create the pair. For details, see Recovering the storage systems: primary site failure with external storage system on page 8-59.

Table 8-1 Confirming the CVS attribute on the external storage system

Interface	To confirm the CVS attribute:
HCS	Open the **Logical device** window, and then confirm whether the CVS attribute is displayed in the **Emulation type** column for the LDEV that is being used as the quorum disk.
CCI	Execute the `raidcom get ldev` command from CCI to the LDEV which is used as quorum disk by external storage system, and then confirm whether the CVS attribute is output for `VOL_TYPE`. Details of the `raidcom get ldev` command; refer to *Command Control Interface Command Reference*.
Web Console*	Confirm whether the CVS attribute is displayed in the CVS column on the LUN Management window.
* Ask the maintenance personnel to operate the Web Console.	

Table 8-2 Conditions for the CVS attribute for volumes created in the external storage system

Interface	Condition			CVS attribute
HCS CCI	Internal volume or external volume			Allowed
	HDP-VOL	VSP G1000 or later VSP G200, G400, G600, G800 or later		Allowed
		VSP or earlier HUS VM or earlier	Create LDEV of maximum size	Not allowed
			Create LDEV less than maximum size	Allowed
Web Console*	The LDEV is created during the operation of the installation of Define Config & Install or ECC/LDEV, which remains the initial value of the Number of LDEVs on the Device Emulation Type Define window.			Not allowed
	Other than above			Allowed
* Ask the maintenance personnel to operate the Web Console.				

Related topics

- I/O modes on page 1-14
- Creating the quorum disk on page 4-17
- Failure locations on page 8-3
- SIMs related to GAD on page 8-6
- Configuring the quorum disk on page 7-7
- Discontinuing a global-active device environment on page 7-23

Pair condition and recovery: external storage system failure

The following table shows transitions for pair status and I/O mode, the volumes that are accessible from the server, and the location of the latest data when you can no longer use the external storage system.

Before failure		After failure				
Pair status and I/O mode		Pair status and I/O mode		Volume accessible from the server		Volume with latest data
P-VOL	S-VOL	P-VOL	S-VOL	P-VOL	S-VOL	
PAIR (Mirror (RL))[1]	PAIR (Mirror (RL))[1]	PSUE (Local)	PSUE (Block)	OK	NG	P-VOL

Before failure		After failure				
Pair status and I/O mode		Pair status and I/O mode		Volume accessible from the server		Volume with latest data
P-VOL	S-VOL	P-VOL	S-VOL	P-VOL	S-VOL	
		PAIR (Mirror (RL))[2]	PAIR (Block)	OK	NG	Both P-VOL and S-VOL
PAIR (Mirror (RL))[2]	PAIR (Block)	PAIR (Mirror (RL))	PAIR (Block)	OK	NG	Both P-VOL and S-VOL
COPY (Mirror (RL))	COPY (Block)	PSUE (Local)	PSUE (Block)	OK	NG	P-VOL
PSUS/PSUE (Local)	SSUS/PSUE (Block)	PSUS/PSUE (Local)	SSUS/PSUE (Block)	OK	NG	P-VOL
PSUS/PSUE (Block)	SSWS (Local)	PSUS/PSUE (Block)	SSWS (Local)	NG	OK	S-VOL

Notes:

1. The pair status and I/O mode after failure depends on the requirement of the pair. For details, see I/O stoppage and data duplication with blocked quorum disk (VSP G1000) on page 1-21.
2. This row applies only to VSP G1000.

SIM

- Primary storage system: 21D0*xy*, 21D2*xx*, DD2*xyy*, DEF0*zz*, EF5*xyy*, EFD000, FF5*xyy*
- Secondary storage system: 21D0*xy*, 21D2*xx*, DD2*xyy*, DEF0*zz*, EF5*xyy*, EFD000, FF5*xyy*

Recovery procedure

1. Recover the external storage system. For details, contact the vendor.
2. Re-create or resynchronize the pair.

Related topics

- Pair condition before failure on page 8-7

Pair condition and recovery: other failures

The following table shows transitions for pair status and I/O mode, the volumes that are accessible from the server, and the location of the latest data when a failure other than explained above occurs.

Before failure		After failure				
Pair status and I/O mode		Pair status and I/O mode		Volume accessible from the server		Volume with latest data
P-VOL	S-VOL	P-VOL	S-VOL	P-VOL	S-VOL	
PAIR (Mirror (RL))[1]	PAIR (Mirror (RL))[1]	PSUE (Local)	PSUE (Block)	OK	NG	P-VOL
		PSUE (Block)	SSWS (Local)	NG	OK	S-VOL
		PAIR (Mirror (RL))[4]	PAIR (Block)	OK	NG	Both P-VOL and S-VOL
PAIR (Mirror (RL))[4]	PAIR (Block)	PSUE (Block)	SSWS (Local)	NG	OK	S-VOL
COPY (Mirror (RL))	COPY (Block)	PSUE (Local)	PSUE (Block)	OK	NG	P-VOL
				NG[2]	NG	P-VOL
PSUS/PSUE (Local)	SSUS/PSUE (Block)	PSUS/PSUE (Local)	SSUS/PSUE (Block)	OK	NG	P-VOL
				NG[2]	NG	P-VOL
PSUS/PSUE (Block)	SSWS (Local)	PSUS/PSUE (Block)	SSWS (Local)	NG	OK	S-VOL
				NG	NG[3]	S-VOL

Notes:

1. The pair status and I/O mode after failure depends on the requirement of the pair. For details, see I/O stoppage and data duplication with blocked quorum disk (VSP G1000) on page 1-21.
2. Depending on the failure factor, if you cannot access the P-VOL, you cannot access the P-VOL or the S-VOL.
3. Depending on the failure factor, if you cannot access the S-VOL, you cannot access the P-VOL or the S-VOL.
4. This row applies only to VSP G1000.

SIM

- Primary storage system: SIM varies depending on the failure type
- Secondary storage system: SIM varies depending on the failure type

Recovery procedure

1. Recover the system.
2. Resynchronize the pair.

Related topics

- Pair condition before failure on page 8-7

Recovery procedure for GAD pair suspension due to other failures

A GAD pair might be suspended due to a failure other than those described in this chapter. Use the following procedure to recover a suspended pair from other types of failure.

If you are not able to restore the GAD volumes using this procedure, contact Hitachi Data Systems customer support.

1. Recover from the failure.
   a. Verify that a failure, such as a suspended GAD pair, has occurred, for example, by checking for SIMs issued by the primary or secondary storage system.
   b. When a failure has occurred, identify the failure and perform troubleshooting according to the failure type to remove the cause of the failure.
2. Confirm the quorum disk status. If the quorum disk is blocked, recover from the blockade.
3. Resynchronize the GAD pair.
   a. Check the I/O mode of the P-VOL and the S-VOL of the suspended GAD pair.

```
pairdisplay -g oraHA -fxce -IH0
Group   PairVol(L/R) (Port#,TID, LU),Seq#,LDEV#.P/S,Status,
Fence,   %,P-LDEV# M CTG JID AP EM       E-Seq# E-LDEV# R/W
oraHA   dev1(L)     (CL1-A-0, 0,   0)311111  2222.P-VOL PSUE
NEVER ,  100  4444 -   -   0  -  -            -       - L/L
oraHA   dev1(R)     (CL1-C-1, 0,   0)322222  4444.S-VOL PSUE
NEVER ,  100  2222 -   -   0  -  -            -       - B/B
pairdisplay -g oraHA -fxce -IH1
Group   PairVol(L/R) (Port#,TID, LU),Seq#,LDEV#.P/S,Status,
Fence,   %,P-LDEV# M CTG JID AP EM       E-Seq# E-LDEV# R/W
oraHA   dev1(L)     (CL1-C-1, 0,   0)322222  4444.S-VOL PSUE
NEVER ,  100  2222 -   -   0  -  -            -       - B/B
oraHA   dev1(R)     (CL1-A-0, 0,   0)311111  2222.P-VOL PSUE
NEVER ,  100  4444 -   -   0  -  -            -       - L/L
```

   b. If the I/O mode of the P-VOL is Local, resynchronize the GAD pair at the primary storage system.

```
pairresync -g oraHA -IH0
```

   c. If the I/O mode of the S-VOL is Local, resynchronize the GAD pair at the secondary storage system.

```
pairresync -g oraHA -swaps -IH1
```

      The volume in the primary storage system changes to an S-VOL, and the volume in the secondary storage system changes to a P-VOL.
   d. Confirm that the pair status of the P-VOL and the S-VOL of the GAD pair has changed to PAIR (Mirror (RL)).

```
pairdisplay -g oraHA -fxce -IH0
Group   PairVol(L/R) (Port#,TID, LU),Seq#,LDEV#.P/S,Status,
Fence,   %,P-LDEV# M CTG JID AP EM       E-Seq# E-LDEV# R/W
oraHA   dev1(L)     (CL1-A-0, 0,   0)311111  2222.P-VOL PAIR
NEVER ,  100  4444 -   -   0  -  -            -       - L/M
```

```
oraHA   dev1(R)     (CL1-C-1, 0,   0)322222  4444.S-VOL PAIR
NEVER ,  100  2222 -   -   0  -  -            -       - L/M
pairdisplay -g oraHA -fxce -IH1
Group   PairVol(L/R) (Port#,TID, LU),Seq#,LDEV#.P/S,Status,
Fence,   %,P-LDEV# M CTG JID AP EM       E-Seq# E-LDEV# R/W
oraHA   dev1(L)     (CL1-C-1, 0,   0)322222  4444.S-VOL PAIR
NEVER ,  100  2222 -   -   0  -  -            -       - L/M
oraHA   dev1(R)     (CL1-A-0, 0,   0)311111  2222.P-VOL PAIR
NEVER ,  100  4444 -   -   0  -  -            -       - L/M
```

4. Using the alternate path software, resume I/O to the S-VOL.
5. If necessary, reverse the P-VOL and the S-VOL.

Related topics

- I/O modes on page 1-14
- Reversing the P-VOL and S-VOL on page 8-62
- Pair condition and recovery: quorum disk failure on page 8-50

Recovering the storage systems: primary site failure with external storage system

If a failure occurs at the primary site in a configuration with the external storage system for the quorum disk located at the primary site, the failure might affect the primary storage system and the external storage system simultaneously. In this case, the GAD pair is suspended, and access to the GAD volumes stops.

Failure at the primary site (external storage system at the primary site)

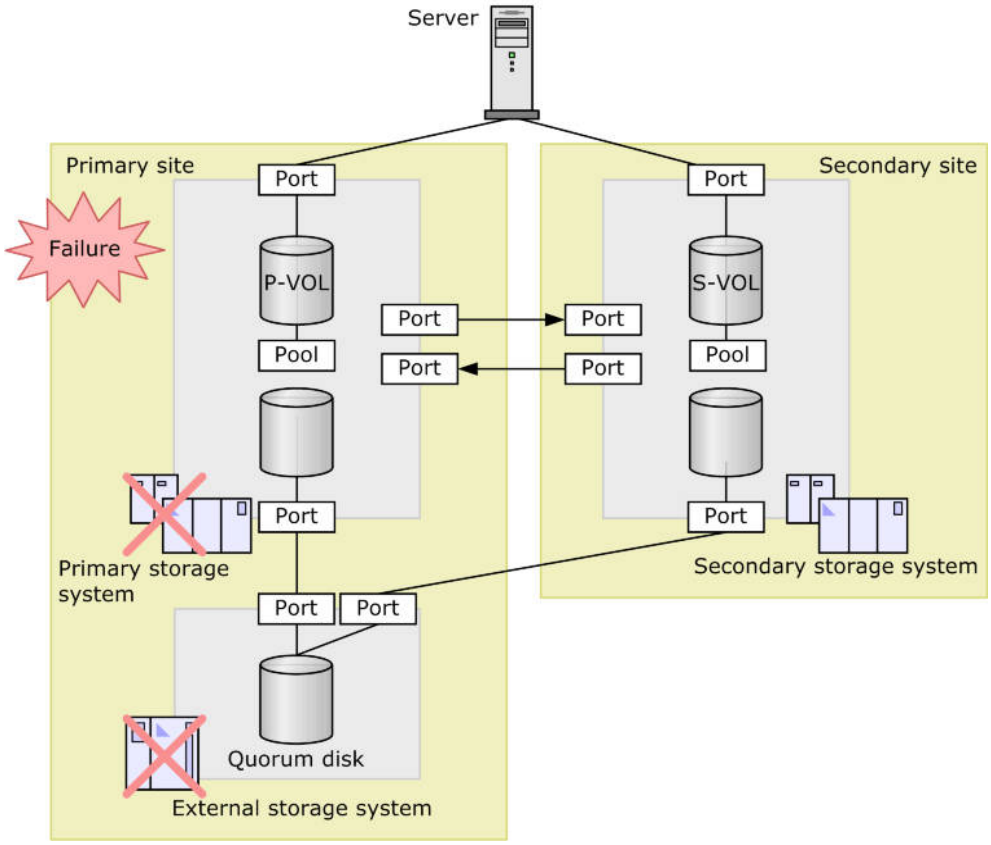

Failure locations	Reference codes of SIMs that might be issued		Can the volume access to the GAD volumes?[1]	
	Primary storage system	Secondary storage system	P-VOL	S-VOL
Both the primary storage system and the external storage system for the quorum disk	Depends on the failure type[2]	DD0*xyy* DD2*xyy* DD3*xyy* 2180*xx* 21D0*xx* 21D2*xx* EF5*xyy* EFD000 FF5*xyy* DEF0*zz*	No	No[3]

Notes:

1. Hardware such as HDD, cache, front-end director (CHA), back-end director (DKA), and MPB is redundant in the storage system configuration. Even if a failure occurs in a part of redundant hardware, the failure does not cause a GAD pair being

Failure locations	Reference codes of SIMs that might be issued		Can the volume access to the GAD volumes?[1]	
	Primary storage system	Secondary storage system	P-VOL	S-VOL
suspended, or an inaccessible GAD volume. The failure does not cause the GAD pair suspended, or the inaccessible GAD volume even if a failure occurs in a part of hardware, if the following physical paths are redundant. ○ Between a server and a storage systems of the primary and secondary sites ○ Between an external storage system and storage systems of the primary and secondary sites ○ Between storage systems of the primary and secondary sites 2. A SIM that corresponds to the failure type is issued. You might not be able to view SIMs according to the failure type. 3. You can access the S-VOL, if the pair status of the S-VOL is SSWS, even if a failure occurs.				

Steps for recovery from the failure

1. Using the alternate path software, delete the alternate path to the GAD P-VOL.
2. At the secondary storage system, delete the GAD pair forcibly.
   Select **Enable** for **Volume Access** in the **Delete Pairs** window of Device Manager - Storage Navigator.
3. Confirm that the GAD pair is deleted.
4. Using the alternate path software, resume I/Os from the server to the GAD S-VOL.
5. Restore the primary storage system from the failure.
6. At the primary storage system, delete the GAD pair forcibly.
   Select **Disable** for **Volume Access** in the **Delete Pairs** window of Device Manager - Storage Navigator.
   Depending on the failure type of the primary storage system, after the primary storage system is restored from a failure, the pair status of the P-VOL might change to SMPL, and the GAD reserve attribute might be set. In this case, you do not need to delete the GAD pair forcibly.
7. Confirm that the GAD pair is deleted.
8. Restore the external storage system from a failure.
9. From the primary and secondary storage systems, delete the quorum disk.
   Depending on the failure type of the external storage system, after the external storage system is restored from a failure, a quorum disk can be deleted. In this case, you do not need to delete the quorum disk.
10. From the primary and secondary storage systems, add a quorum disk.
11. From the secondary storage system, re-create a GAD pair.

12. Using the alternate path software, add a path to the GAD P-VOL, and then resume I/Os.
13. Reverse the P-VOL and the S-VOL if necessary.

Related topics

- I/O modes on page 1-14
- Failure locations on page 8-3
- Reversing the P-VOL and S-VOL on page 8-62
- Configuring the quorum disk on page 7-7
- Discontinuing a global-active device environment on page 7-23

Reversing the P-VOL and S-VOL

During disaster recovery operations, P-VOLs are changed to S-VOLs and S-VOLs to P-VOLs to reverse the flow of data from the secondary site to the primary site to restore the primary site. When normal operations are resumed at the primary site, the direction of copy is changed again so that the original P-VOLs become primary volumes again and the original S-VOLs become secondary volumes again with data flowing from the primary site to the secondary site.

Steps to reverse data flow

1. Using the alternate path software, stop I/O from the server to P-VOLs in the secondary storage system.

   Continue to the next step even if the alternate path cannot be deleted.
2. Confirm that the P-VOL and the S-VOL have been reversed.

```
pairdisplay -g oraHA -fxce -IH0
Group   PairVol(L/R) (Port#,TID, LU),Seq#,LDEV#.P/S,Status,
Fence,   %,P-LDEV# M CTG JID AP EM       E-Seq# E-LDEV# R/W
oraHA   dev1(L)     (CL1-A-0, 0,   0)311111  2222.S-VOL PAIR
NEVER ,  100  4444 -   -   0  -  -            -       - L/M
oraHA   dev1(R)     (CL1-C-1, 0,   0)322222  4444.P-VOL PAIR
NEVER ,  100  2222 -   -   0  -  -            -       - L/M
```

3. At the primary storage system, change the pair statuses of the S-VOLs to SSWS to suspend the pairs (swap suspension).

```
pairsplit -g oraHA -d dev1 -RS -IH0
```

4. At the secondary storage system, reverse the P-VOL and the S-VOL, and then resynchronize the pairs (swap resync).

```
pairresync -g oraHA -d dev1 -swaps -IH0
```

5. Confirm that the P-VOL and the S-VOL pair statuses change to PAIR (Mirror (RL)).

```
pairdisplay -g oraHA -fxce -IH0
Group   PairVol(L/R) (Port#,TID, LU),Seq#,LDEV#.P/S,Status,
Fence,   %,P-LDEV# M CTG JID AP EM       E-Seq# E-LDEV# R/W
oraHA   dev1(L)     (CL1-A-0, 0,   0)311111  2222.P-VOL PAIR
NEVER ,  100  4444 -   -   0  -  -            -       - L/M
```

```
oraHA   dev1(R)     (CL1-C-1, 0,   0)322222  4444.S-VOL PAIR
NEVER ,  100  2222 -   -   0  -  -            -       - L/M
pairdisplay -g oraHA -fxce -IH1
Group   PairVol(L/R) (Port#,TID, LU),Seq#,LDEV#.P/S,Status,
Fence,   %,P-LDEV# M CTG JID AP EM       E-Seq# E-LDEV# R/W
oraHA   dev1(L)     (CL1-C-1, 0,   0)322222  4444.S-VOL PAIR
NEVER ,  100  2222 -   -   0  -  -            -       - L/M
oraHA   dev1(R)     (CL1-A-0, 0,   0)311111  2222.P-VOL PAIR
NEVER ,  100  4444 -   -   0  -  -            -       - L/M
```

6. Using the alternate path software, restart I/Os from the server to S-VOLs in the secondary storage system.

Related topics

- I/O modes on page 1-14

Resolving failures in multiple locations

If failures occur in multiple locations, use the following recovery procedure:

1. Identify the failure locations from the SIMs issued by the primary and secondary storage systems and using the SAN management software, and then recover from the failures.
2. If data has been lost from both volumes, recover from the backup data using ShadowImage or Thin Image volumes, or backup software.
3. If I/O is stopped, resume I/O from the server.
4. If GAD pairs are suspended, resynchronize the pairs.

   If the pairs cannot be resynchronized, delete the pairs and then create them again.

Related topics

- SIMs related to GAD on page 8-6

Pair condition and recovery: quorum disk and primary-to-secondary path failure (VSP G1000 only)

When a failure occurs on the quorum disk, the S-VOL pair status changes from PAIR (Mirror(RL)) to PAIR (Block). Another failure then occurs on the path from the primary storage system to the secondary storage system, and the P-VOL pair status changes from PAIR (Mirror(RL)) to PSUE(Local), as shown in the following figure.

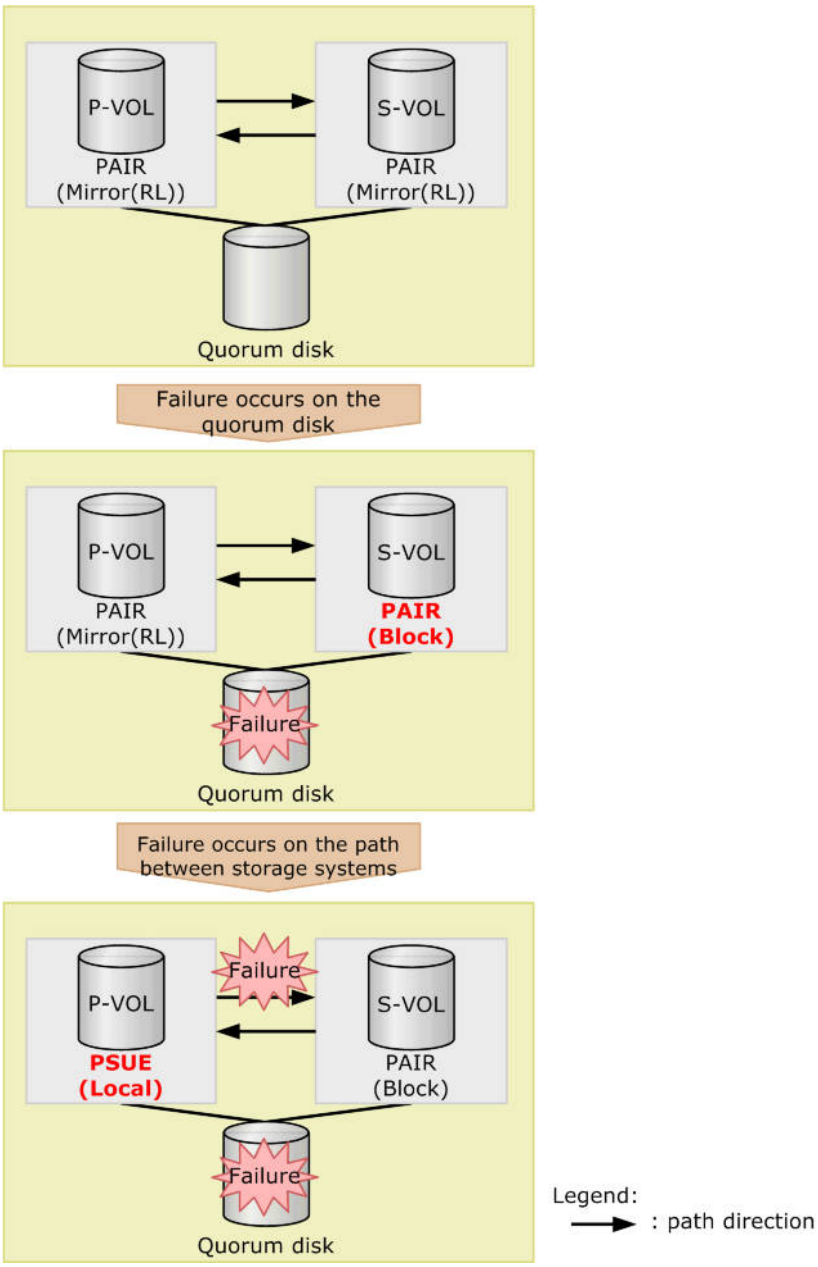

The following table shows transitions for pair status and I/O mode, the volumes that are accessible from the server, and the location of the latest data when you can no longer use any physical path from the primary storage system to the secondary storage system after the quorum disk failure.

After quorum disk failure		After primary-to-secondary path failure				
Pair status and I/O mode		Pair status and I/O mode		Volume accessible from the server		Volume with latest data
P-VOL	S-VOL	P-VOL	S-VOL	P-VOL	S-VOL	
PAIR (Mirror (RL))	PAIR (Block)	PSUE (Local)	PAIR (Block)	OK	NG	P-VOL

SIMs

- Primary storage system: 21d0*xy*, 21d2*xx*, dd2*xyy*, def0*zz*, ef5*xyy*, efd000, ff5*xyy*, dd0*xyy*, 2180*xx*
- Secondary storage system: 21d0*xy*, 21d2*xx*, dd2*xyy*, def0*zz*, ef5*xyy*, efd000, ff5*xyy*, dd3*xyy*, dd0*xyy*, dd1*xyy*

Recovery procedure

1. Recover the quorum disk failure and the path to the external storage system.
2. Recover the path from the primary storage system to the secondary storage system.
3. Resynchronize the GAD pair suspended due to the failure.
4. Confirm the pair status.

   When the pair status of the P-VOL and S-VOL is PAIR (Mirror(RL)), the recovery is completed.

   When the P-VOL pair status is PSUE (Local) and the S-VOL pair status is PAIR (Block), proceed to step 5.
5. Suspend the GAD pair by specifying the S-VOL (swap suspend).

   The pair suspension operation fails, but the S-VOL pair status changes to PSUE (Block).
6. Resynchronize the GAD pair by specifying the P-VOL.

   The pair status of the P-VOL and S-VOL changes to PAIR (Mirror(RL)).

Recovering the quorum disk and primary-to-secondary path failure

The following figure shows the failure area and recovery from the path failure from the primary storage system to the secondary storage system after the GAD status changes to Quorum disk blocked (VSP G1000 only).

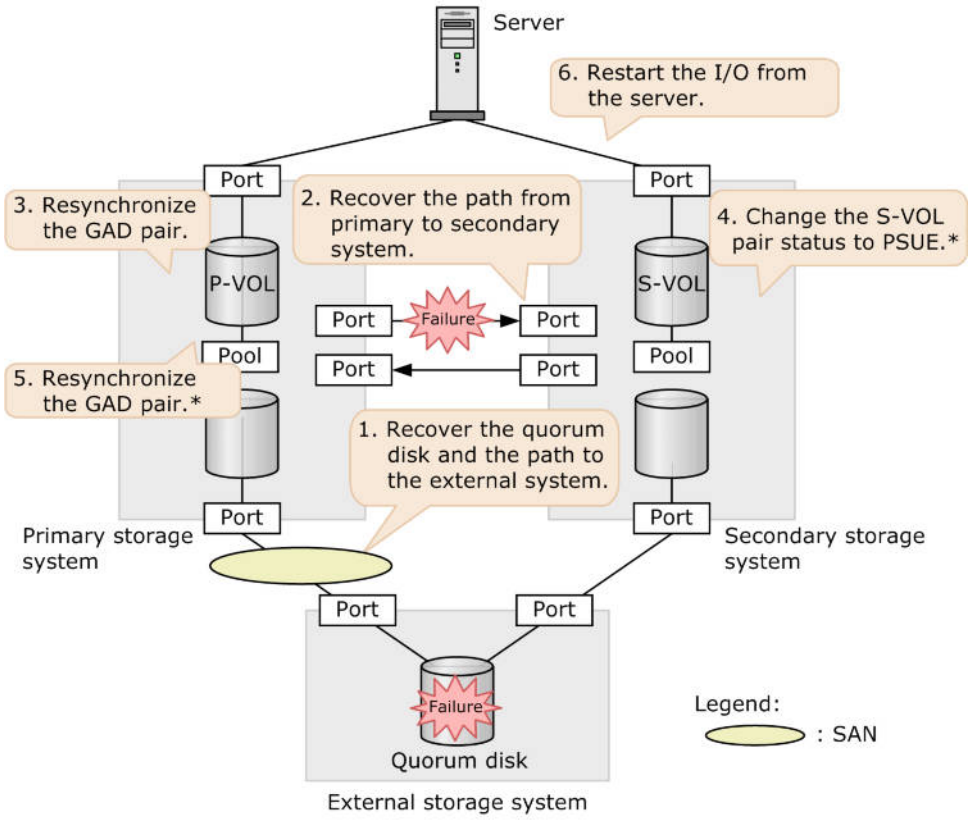

*These steps are necessary only when the GAD pair resynchronization fails in step 3.

Recovery procedure

1. Recover the quorum disk failure and the path to the external storage system.

   a. Recover the quorum disk.

   b. Reconnect the physical path or reconfigure the SAN to recover the path to the external storage system. When the path is recovered, the external path is automatically recovered.

   c. Confirm that the external storage system is connected correctly.

```
raidcom get path -path_grp 1 -IH0
PHG GROUP STS CM IF MP# PORT WWN PR LUN PHS Serial#
PRODUCT_ID LB PM
1 1-1 NML E D 0 CL5-A 50060e8007823520 1 0 NML 33333 VSP
G1000 N M
```

Note: The `PRODUCT_ID` for the VSP G200, G400, G600, G800 is `VSP Gx00`.

   d. Confirm the LDEV ID of the quorum disk by obtaining the information of the external volume from the primary storage system.

```
raidcom get external_grp -external_grp_id 1-1 -IH0
T GROUP P_NO LDEV# STS LOC_LBA SIZE_LBA Serial#
E 1-1 0 9999 NML 0x000000000000 0x000003c00000 333333
```

e. Confirm that the primary storage system recognizes the external volume as a quorum disk by specifying the LDEV ID of the quorum disk.

```
raidcom get ldev -ldev_id 0x9999 -fx -IH0
(snip)
QRDID : 0
QRP_Serial# : 322222
QRP_ID : R8
(snip)
```

Note: The VSP G1000 is displayed as `R8` in command output. The VSP G200, G400, G600, G800 is displayed as `M8` in command output.

2. Reconnect the physical path or reconfigure the SAN to recover the path failure from the primary to secondary storage system.

   When the path is recovered, the remote path is automatically recovered. If you recover the physical path but the remote path does not recover, contact Hitachi Data Systems customer support.

3. Resynchronize the GAD pair whose GAD status is Suspended.

   a. Confirm that the P-VOL I/O mode is Local.

```
pairdisplay -g oraHA -fxce -IH0
Group PairVol(L/R) (Port#,TID, LU), Seq#,LDEV#.P/
S,Status,Fence, %,P-LDEV# M CTG JID AP EM E-Seq# E-LDEV# R/W
oraHA dev1(L) (CL1-A-0, 0, 0)311111 2222.P-VOL PSUE NEVER ,
100 4444 - - 0 - - - - L/L
oraHA dev1(R) (CL1-C-1, 0, 0)322222 4444.S-VOL PSUE NEVER ,
100 2222 - - 0 - - - - B/B
```

   b. At the primary storage system, resynchronize the pair.

```
pairresync -g oraHA -IH0
```

Note: When the P-VOL pair status is PSUE (Local) and the S-VOL pair status is PAIR(Block), the pair resynchronization fails. The result of the pair resynchronization depends on whether the GAD pair is registered to the consistency group.

- When the GAD pair is registered to the consistency group, the pair resynchronization operation fails.
- When the GAD pair is not registered to the consistency group, the pair resynchronization operation succeeds, but the pair resynchronization process fails. The pair status of the P-VOL after the pair resynchronization remains PSUE.

   c. Confirm that the pair status of the P-VOL and S-VOL changes to PAIR (Mirror (RL)).

```
pairdisplay -g oraHA -fxce -IH0
Group PairVol(L/R) (Port#,TID, LU), Seq#, LDEV#.P/
S,Status,Fence, %,P-LDEV# M CTG JID AP EM E-Seq# E-LDEV# R/W
oraHA dev1(L) (CL1-A-0, 0, 0)311111 2222.P-VOL PAIR NEVER ,
100 4444 - - 0 - - - - L/M
oraHA dev1(R) (CL1-C-1, 0, 0)322222 4444.S-VOL PAIR NEVER ,
100 2222 - - 0 - - - - L/M
pairdisplay -g oraHA -fxce -IH1
Group PairVol(L/R) (Port#,TID, LU), Seq#, LDEV#.P/
S,Status,Fence, %,P-LDEV# M CTG JID AP EM E-Seq# E-LDEV# R/W
```

```
oraHA dev1(L) (CL1-C-1, 0, 0)322222 4444.S-VOL PAIR NEVER ,
100 2222 - - 0 - - - - L/M
oraHA dev1(R) (CL1-A-0, 0, 0)311111 2222.P-VOL PAIR NEVER ,
100 4444 - - 0 - - - - L/M
```

Note: When the pair whose P-VOL pair status is PSUE (Local) and S-VOL pair status is PAIR (Block) exists, go to step 4. When no pairs meet this condition, go to step 6.

4. Suspend all GAD pairs whose P-VOL pair status is PSUE (Local) and S-VOL pair status is PAIR (Block) by specifying the S-VOL (swap suspend).

   ```
   pairsplit -g oraHA -RS -d dev1 -IH0
   ```

   The pair suspension operation fails, but the S-VOL pair status changes to PSUE.

Note:

- Even if the pairs are registered to a consistency group, swap-suspend the pairs by pair.
- The following SIMs might be issued, but you do not need to address these SIMs: dd0*xyy*, dd1*xyy*, dd2*xyy*, dd3*xyy*

5. At the primary storage system, resynchronize the GAD pair.

   ```
   pairresync -g oraHA -IH0
   ```

6. Using the alternate path software, resume I/O to the S-VOL (I/O might resume automatically).

9

Disaster recovery in a GAD 3DC delta resync environment

This chapter provides information and instructions specifically for disaster recovery operations in a GAD 3DC delta resync (GAD+UR) environment.

For details about storage system support (models, microcode) for GAD+UR operations, see Requirements and restrictions on page 2-2.

Status before failure

The following figure shows the status in a GAD 3DC delta resync (GAD+UR) environment before a failure occurs. The data is copied from the GAD P-VOL to the UR S-VOL.

Disaster recovery procedures for GAD+UR failure conditions:

- Recovering from a primary site failure (GAD+UR) on page 9-2
- Recovering from a failure at the secondary site (GAD+UR) on page 9-5
- Recovering from a failure at the primary and secondary sites (GAD+UR) on page 9-7
- Recovering from a failure on the UR delta resync pair on page 9-10
- Recovering from a quorum disk failure (GAD+UR) on page 9-11

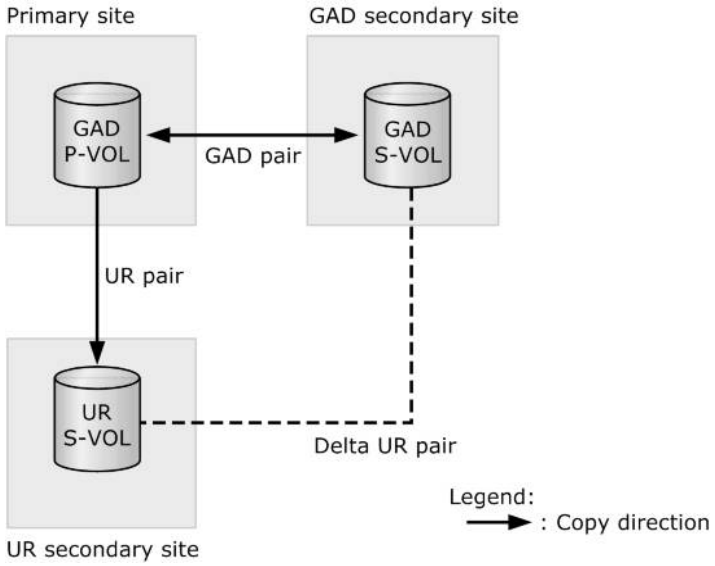

Recovering from a primary site failure (GAD+UR)

Use the following procedure to recover from a failure on the P-VOL at the primary site when sharing GAD volumes with UR in a GAD 3DC delta resync configuration.

When a failure occurs at the primary site, the status of the GAD pair changes to PSUE/SSWS. In this case, delta resync is automatically performed and the pair statuses change as follows:

Before failure			After failure		
Pair type	Pair status		Pair type	Pair status	
	P-VOL	S-VOL		P-VOL	S-VOL
GAD pair	PAIR	PAIR	GAD pair	PSUE	SSWS

Before failure			After failure		
Pair type	Pair status		Pair type	Pair status	
	P-VOL	S-VOL		P-VOL	S-VOL
UR pair	PAIR	PAIR	UR delta resync pair	PSUE	SSUS
UR delta resync pair	PSUS	SSUS	UR pair	COPY	COPY

The UR delta resync pair changes to a UR pair, and copying from the GAD S-VOL to the UR S-VOL starts. When the UR pair copying is completed, the status of the P-VOL and the status of the S-VOL of the UR pair change to PAIR.

Overview of failure recovery

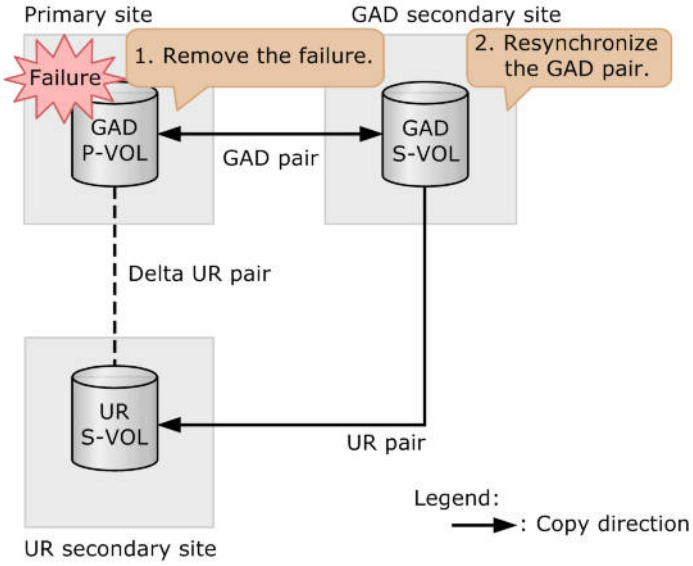

Steps for recovery

1. Remove the failure on the P-VOL.
2. Reverse the P-VOL and the S-VOL, and then resynchronize the GAD pairs (swap resync) on the storage system at the GAD secondary site.

   Command example: `pairresync -g oraHA -swaps -IH1`

   The volume on the primary storage system changes to an S-VOL, and the volume on the GAD secondary storage system changes to a P-VOL.
3. Confirm that the GAD P-VOL and S-VOL pair statuses change to PAIR.

   Command example:

```
pairdisplay -g oraHA -fxce -IH0
Group   PairVol(L/R) (Port#,TID, LU),Seq#,LDEV#.P/S,Status,
Fence,   %,P-LDEV# M CTG JID AP EM       E-Seq# E-LDEV# R/W
oraHA   dev1(L)     (CL1-A-0, 0,   0)311111  2222.P-VOL PAIR
NEVER ,  100  4444 -   -   0  -  -            -       - L/M
```

```
oraHA   dev1(R)     (CL1-C-1, 0,   0)322222  4444.S-VOL PAIR
NEVER ,  100  2222 -   -   0  -  -            -       - L/M
pairdisplay -g oraHA -fxce -IH1
Group   PairVol(L/R) (Port#,TID, LU),Seq#,LDEV#.P/S,Status,
Fence,   %,P-LDEV# M CTG JID AP EM       E-Seq# E-LDEV# R/W
oraHA   dev1(L)     (CL1-C-1, 0,   0)322222  4444.S-VOL PAIR
NEVER ,  100  2222 -   -   0  -  -            -       - L/M
oraHA   dev1(R)     (CL1-A-0, 0,   0)311111  2222.P-VOL PAIR
NEVER ,  100  4444 -   -   0  -  -            -       - L/M
```

4. Keep updating I/O from the server to the P-VOL or S-VOL of the GAD pair for about two minutes.
5. Confirm that the delta UR P-VOL pair status is PSUS.

   Command example:

```
pairdisplay -g oraREMOTE -fxce -IH0
Group     PairVol(L/R)  (Port#,TID, LU),Seq#,  LDEV#.P/S,
Status,Fence, %,  P-LDEV#  M  CTG  JID  AP  EM  E-Seq#  E-LDEV#
R/W
oraDELTA  dev2(L)       (CL1-A-1, 0, 1) 311111 2222. P-VOL
PSUS   ASYNC ,0   6666     -  0    0    -   -   -       -
-/-
oraDELTA  dev2(R)       (CL1-A-1, 0, 1) 344444 6666. S-VOL
SSUS   ASYNC ,0   2222     -  0    0    -   -   -       -
-/-
```

Note: To check the status of a pair in Device Manager - Storage Navigator, select **Refresh All** in the **File** menu, update the information displayed on Device Manager - Storage Navigator, and then view the pair status. The status of the UR delta resync pairs changes from HOLDING to HOLD.

6. Confirm that the mirror status of the journal of the UR delta resync pair is PJNS.

   Command example:

```
pairdisplay -g oraREMOTE -v jnl -IH0
JID  MU  CTG  JNLS  AP  U(%)  Q-Marker  Q-CNT  D-SZ(BLK)  Seq#
Num  LDEV#
000  1   1    PJNS  4   21    43216fde  30     512345     62500
1    39321
```

7. Confirm that no failure SIMs are displayed.
8. Reverse the GAD P-VOL and the S-VOL if necessary. For instructions, see Reversing the GAD P-VOL and S-VOL (GAD+UR) on page 9-4.

Reversing the GAD P-VOL and S-VOL (GAD+UR)

Use the following procedure to reverse the GAD P-VOL and S-VOL when sharing GAD volumes with UR in a GAD 3DC delta resync (GAD+UR) configuration.

1. Suspend the GAD pair by specifying the S-VOL (swap suspend).

   Command example: `pairsplit -g oraHA -RS -IH0`
2. Resynchronize the GAD pair by specifying the S-VOL (swap resync).

   Command example: `pairresync -g oraHA -swaps -IH0`

The volume on the primary storage system changes to a P-VOL, and the volume on the GAD secondary storage system changes to an S-VOL.

3. Keep updating I/O from the server to the P-VOL or S-VOL of the GAD pair for about two minutes.
4. Confirm that the delta UR P-VOL pair status is PSUS.

   Command example:

```
pairdisplay -g oraDELTA -fxce -IH0
Group     PairVol(L/R)  (Port#,TID, LU),Seq#,  LDEV#.P/S,
Status,Fence, %,  P-LDEV#  M  CTG  JID  AP  EM  E-Seq#  E-LDEV#
R/W
oraDELTA  dev2(L)       (CL1-A-1, 0, 1) 311111 2222. P-VOL
PSUS   ASYNC ,0   6666     -  0    0    -   -   -       -
-/-
oraDELTA  dev2(R)       (CL1-A-1, 0, 1) 344444 6666. S-VOL
SSUS   ASYNC ,0   2222     -  0    0    -   -   -       -
-/-
```

Note: To check the status of a pair in Device Manager - Storage Navigator, select **Refresh All** in the **File** menu, update the information displayed on Device Manager - Storage Navigator, and then view the pair status. The status of the UR delta resync pairs changes from HOLDING to HOLD.

5. Confirm that the mirror status of the journal of the UR delta resync pair is PJNS.

   Command example:

```
pairdisplay -g oraDELTA -v jnl -IH0
JID  MU  CTG  JNLS  AP  U(%)  Q-Marker  Q-CNT  D-SZ(BLK)  Seq#
Num  LDEV#
000  1   1    PJNS  4   21    43216fde  30     512345     62500
1    39321
```

6. Confirm that no failure SIMs are displayed.

Recovering from a failure at the secondary site (GAD+UR)

Use the following procedure to recover from a failure on the S-VOL at the secondary site when sharing GAD volumes with UR in a GAD 3DC delta resync (GAD+UR) configuration.

When a failure occurs at the secondary site, the status of the P-VOL and S-VOL of the GAD pair changes to PSUE.

Overview of failure recovery

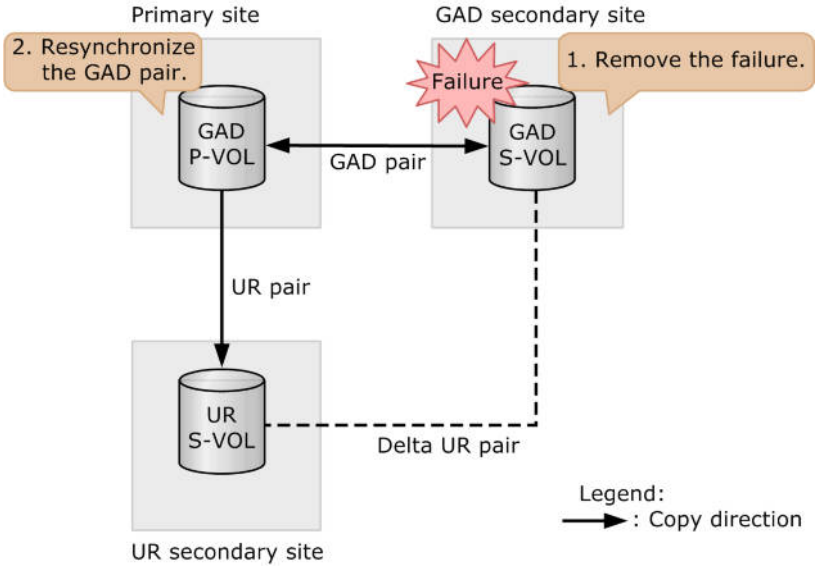

Steps for recovery

1. Remove the failure on the S-VOL.
2. Resynchronize the GAD pair at the primary storage system.
   Command example: `pairresync -g oraHA -IH0`
3. Confirm that the GAD P-VOL and S-VOL pair statuses change to PAIR.
   Command example:

```
pairdisplay -g oraHA -fxce -IH0
Group   PairVol(L/R) (Port#,TID, LU),Seq#,LDEV#.P/S,Status,
Fence,   %,P-LDEV# M CTG JID AP EM       E-Seq# E-LDEV# R/W
oraHA   dev1(L)     (CL1-A-0, 0,   0)311111  2222.P-VOL PAIR
NEVER ,  100  4444 -   -   0  -  -            -       - L/M
oraHA   dev1(R)     (CL1-C-1, 0,   0)322222  4444.S-VOL PAIR
NEVER ,  100  2222 -   -   0  -  -            -       - L/M
pairdisplay -g oraHA -fxce -IH1
Group   PairVol(L/R) (Port#,TID, LU),Seq#,LDEV#.P/S,Status,
Fence,   %,P-LDEV# M CTG JID AP EM       E-Seq# E-LDEV# R/W
oraHA   dev1(L)     (CL1-C-1, 0,   0)322222  4444.S-VOL PAIR
NEVER ,  100  2222 -   -   0  -  -            -       - L/M
oraHA   dev1(R)     (CL1-A-0, 0,   0)311111  2222.P-VOL PAIR
NEVER ,  100  4444 -   -   0  -  -            -       - L/M
```

4. Keep updating I/O from the server to the P-VOL or S-VOL of the GAD pair for about two minutes.
5. Confirm that the pair status of the delta UR P-VOL is PSUS.
   Command example:

```
pairdisplay -g oraDELTA -fxce -IH1
Group     PairVol(L/R)  (Port#,TID, LU),Seq#,  LDEV#.P/S,
Status,Fence, %,  P-LDEV#  M  CTG  JID  AP  EM  E-Seq#  E-LDEV#
R/W
oraDELTA  dev3(L)       (CL1-A-1, 0, 1) 322222 4444. P-VOL
PSUS   ASYNC ,0   6666     -  0    0    -   -   -       -
-/-
oraDELTA  dev3(R)       (CL1-A-1, 0, 1) 344444 6666. S-VOL
```

```
SSUS   ASYNC ,0   4444     -  0    0    -   -   -       -
-/-
```

Note: To check the status of a pair in Device Manager - Storage Navigator, select **Refresh All** in the **File** menu, update the information displayed on Device Manager - Storage Navigator, and then view the pair status. The status of the UR delta resync pairs changes from HOLDING to HOLD.

6. Confirm that the mirror status of the journal of the UR delta resync pair is PJNS using the CCI instance for managing the GAD secondary storage system.

   Command example:

```
pairdisplay -g oraDELTA -v jnl -IH1
JID  MU  CTG  JNLS  AP  U(%)  Q-Marker  Q-CNT  D-SZ(BLK)  Seq#
Num  LDEV#
000  1   1    PJNS  4   21    43216fde  30     512345     62500
1    39321
```

7. Confirm that no failure SIMs are displayed.

Recovering from a failure at the primary and secondary sites (GAD+UR)

Use the following procedure to recover from a failure at the primary and secondary sites when sharing GAD volumes with UR in a GAD 3DC delta resync (GAD+UR) configuration. You need to delete and re-create all pairs.

Steps for recovery

1. Suspend the UR pair by specifying the S-VOL (swap suspend).

   Command example: `pairsplit -g oraREMOTE -RS -IH2`

   You can resume operations at this point by using the S-VOL at the UR secondary site.

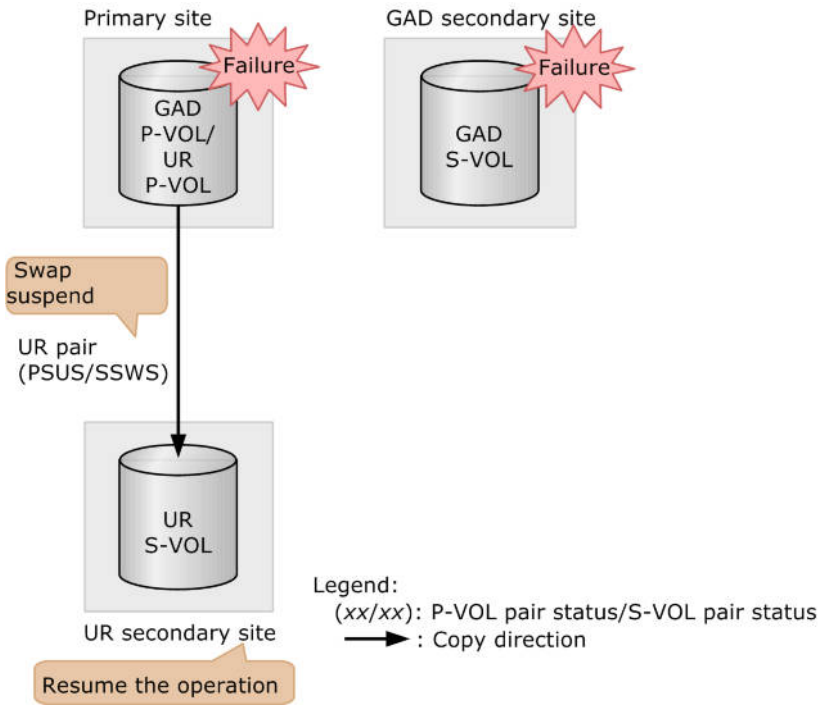

2. Remove the failure at the primary and secondary sites.
3. Delete the UR pair.
4. Delete the UR delta resync pair.
5. Delete the GAD pair.

   When the I/O mode of both the primary and secondary volumes is Block, select **Force** in **Delete Mode** in the **Delete Pairs** window of Device Manager - Storage Navigator.

Caution: You can forcibly delete a pair only when the I/O mode of both the primary volume and secondary volume is Block. If you want to forcibly delete a pair when the I/O mode is not Block, call Hitachi Data Systems customer support.

6. Create a UR pair between the UR secondary and primary sites.

   Command example: `paircreate -g oraREMOTE -f async -vl -jp 0 -js 0 -IH2`
7. Suspend the UR pair by specifying the S-VOL (swap suspend).

   Command example: `pairsplit -g oraREMOTE -RS -IH0`

   You can resume operations at this point by using the S-VOL at the primary site.

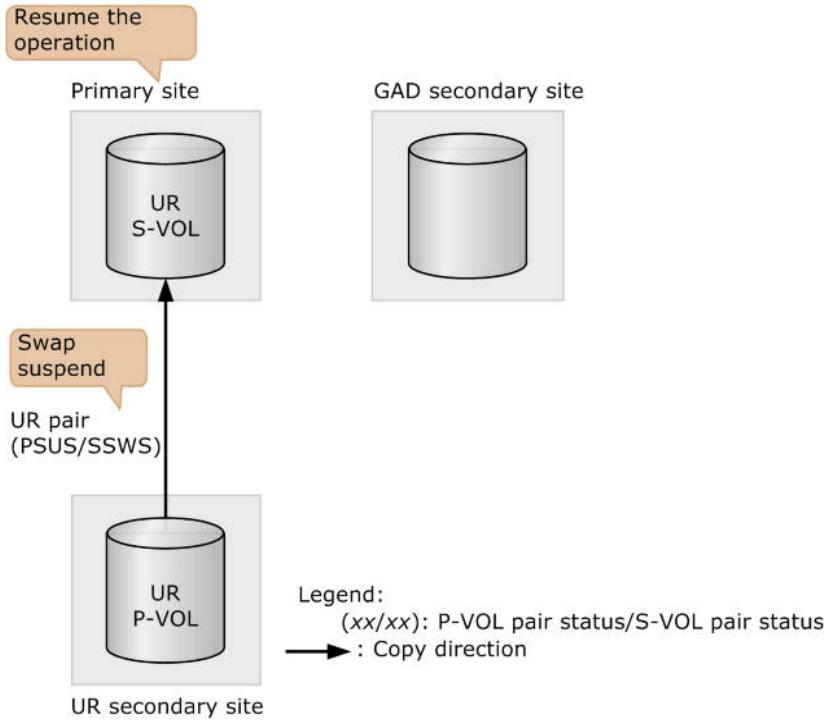

8. Resynchronize the UR pair by specifying the S-VOL (swap resync).
   Command example: `pairresync -g oraREMOTE -swaps -IH0`
   The volume on the primary storage system changes to a P-VOL, and the volume on the UR secondary storage system changes to an S-VOL.

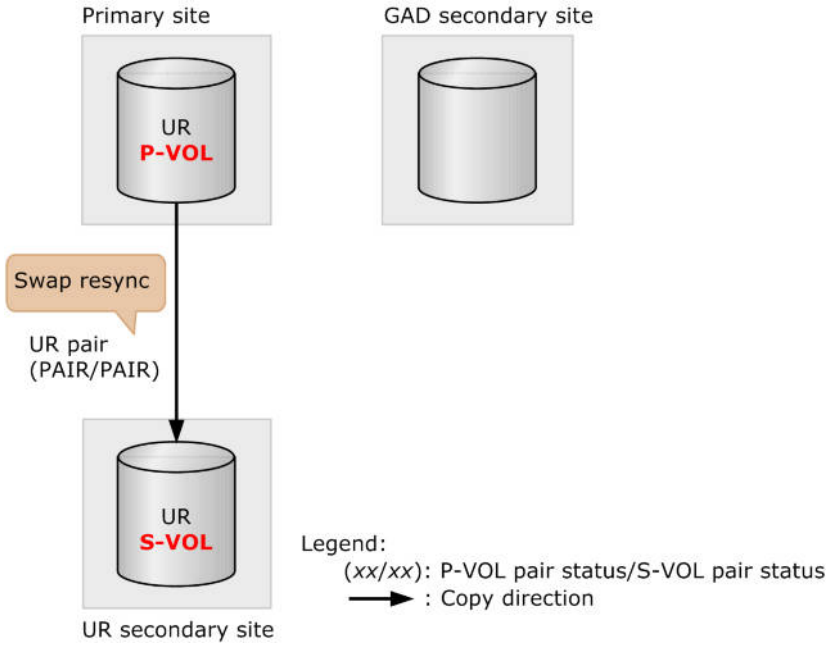

9. Delete the UR pair.
   Command example: `pairsplit -g oraREMOTE -S -IH0`
10. Re-create the GAD pair.
    Command example: `paircreate -g oraHA -fg never 2 -vl -jq 0 -IH0`

11. Re-create the UR delta resync pair.
    Command example: `paircreate -g oraDELTA -f async -vl -jp 0 -js 0 -nocsus -IH1`
12. Re-create the UR pair.
    Command example: `paircreate -g oraREMOTE -f async -vl -jp 0 -js 0 -IH0`

Recovering from a failure on the UR delta resync pair

Use the following procedure to recover from a failure on the UR delta resync pair due to a journal full condition on the storage system at the GAD secondary site.

When a failure occurs on the UR delta resync pair, a SIM is displayed.

Overview of failure recovery

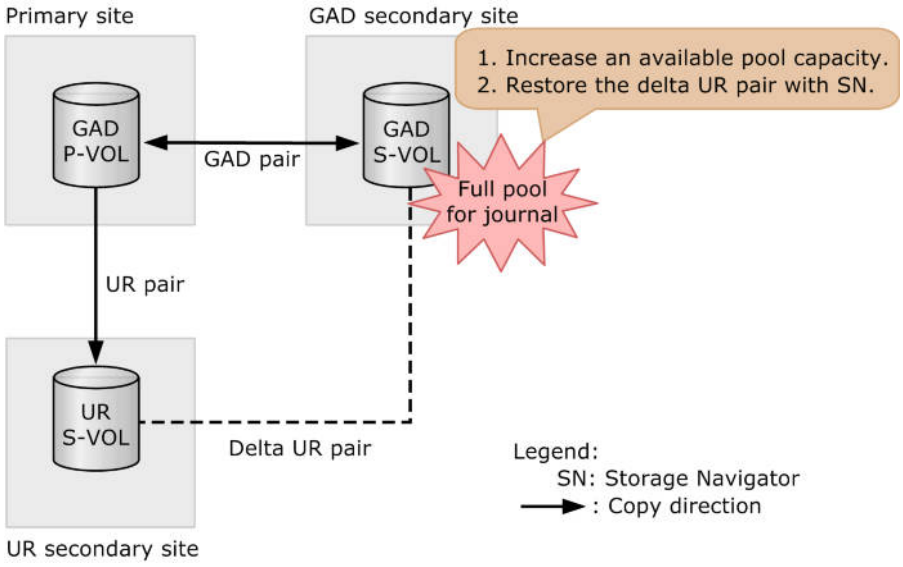

Steps for recovery

1. Increase the available capacity of the pool used by the journal at the GAD secondary storage system.
2. Restore the UR delta resync pair.
   Specify **Return to standby** in **Resync Mode** of the **Resync Mirrors** window of Device Manager - Storage Navigator.
3. Confirm that the pair status of the delta UR P-VOL is PSUS.
   Command example:

```
pairdisplay -g oraDELTA -fxce -IH1
Group     PairVol(L/R)  (Port#,TID, LU),Seq#,  LDEV#.P/S,
Status,Fence, %,  P-LDEV#  M  CTG  JID  AP  EM  E-Seq#  E-LDEV#
R/W
oraDELTA  dev3(L)       (CL1-A-1, 0, 1) 322222 4444. P-VOL
PSUS   ASYNC ,0   6666     -  0    0    -   -   -       -
```

```
-/-
oraDELTA  dev3(R)        (CL1-A-1, 0, 1) 344444 6666. S-VOL
SSUS   ASYNC ,0   4444     -  0    0    -   -   -       -
-/-
```

4. Confirm that the mirror status of the journal of the UR delta resync pair is PJNS using the CCI instance for managing the GAD secondary storage system.

   Command example:

```
pairdisplay -g oraDELTA -v jnl -IH1
JID  MU  CTG  JNLS  AP  U(%)  Q-Marker  Q-CNT  D-SZ(BLK)  Seq#
Num  LDEV#
000  1   1    PJNS  4   21    43216fde  30     512345     62500
1    39321
```

5. Confirm that no failure SIMs are displayed.

Recovering from a quorum disk failure (GAD+UR)

Use the following procedure to recover from a quorum disk failure when sharing GAD volumes with UR in a GAD 3DC delta resync (GAD+UR) configuration. The procedure depends on the following status:

- Able to access either volume of the GAD pair from the server
- Not able to access either volume of the GAD pair from the server

To check whether you can access the volumes of a GAD pair from the server, use I/O mode for the GAD pair.

Command example

```
pairdisplay -g oraHA -fxce -IH0
Group   PairVol(L/R) (Port#,TID, LU),Seq#,LDEV#.P/S,Status,
Fence,   %,P-LDEV# M CTG JID AP EM       E-Seq# E-LDEV# R/W
oraHA   dev1(L)     (CL1-A-0, 0,   0)311111  2222.S-VOL PSUS
NEVER ,  100  4444 -   -   0  -  -           -       - L/L
oraHA   dev1(R)     (CL1-C-1, 0,   0)322222  4444.P-VOL PAIR
NEVER ,  100  2222 -   -   0  -  -           -       - B/B
```

The server can access a volume whose I/O mode (`R/W`) is `L/L` (Local), and cannot access a volume whose I/O mode (`R/W`) is `B/B` (Block).

Related topics

- I/O modes on page 1-14
- Recovering from a quorum disk failure when one GAD volume is accessible on page 9-12
- Recovering from a quorum disk failure when neither GAD volume is accessible on page 9-12

Recovering from a quorum disk failure when one GAD volume is accessible

Use the following procedure to recover a failure on the quorum disk when you can access either volume of a GAD pair from the server.

Overview of failure recovery

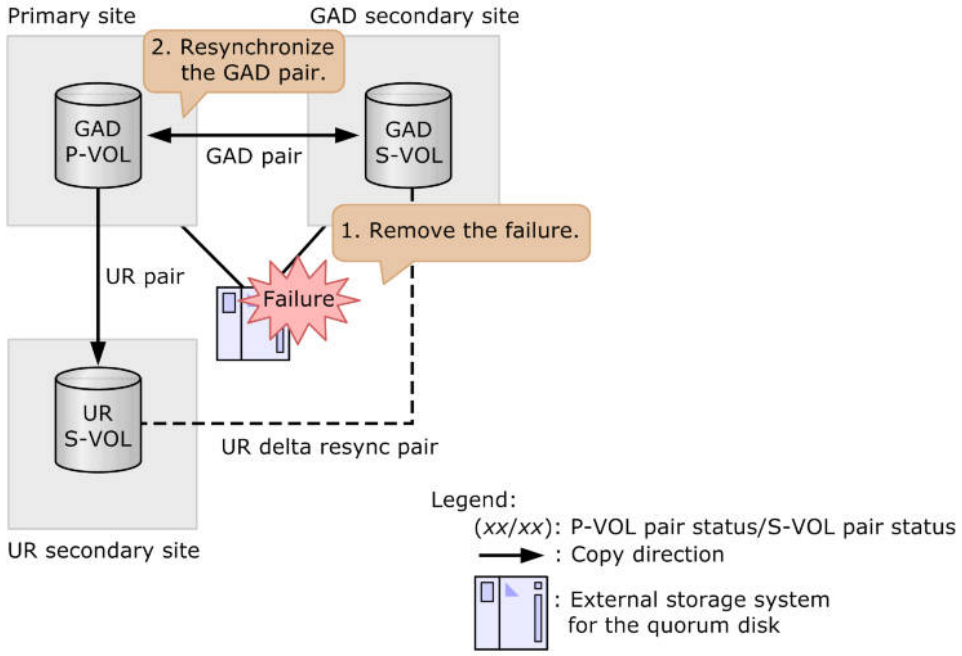

Steps for recovery

1. Remove the failure on the quorum disk.
2. If the GAD status of the pair is Suspended, resynchronize the pair.
   If the GAD status of the pair is Quorum disk blocked (VSP G1000 only), the pair changes to the Mirrored status automatically.

Related topics

- Pair condition and recovery: quorum disk failure on page 8-50

Recovering from a quorum disk failure when neither GAD volume is accessible

When a failure occurs on the quorum disk and you cannot access either volume of a GAD pair from the server, you need to delete the UR and GAD pairs, then recover from the failure by using the volume at the UR secondary site, and then re-create the pairs.

Steps for recovery

1. Delete the UR pair.

Command example: `pairsplit -g oraREMOTE -S -IH0`

2. Delete the UR delta resync pair.
   Command example: `pairsplit -g oraDELTA -S -IH1`
3. Delete the GAD pair.
   When the I/O mode of both the primary and secondary volumes is Block, select **Force** in **Delete Mode** in the **Delete Pairs** window of Device Manager - Storage Navigator.

Caution: You can forcibly delete a pair only when the I/O mode of both the primary volume and secondary volume is Block. If you want to forcibly delete a pair when the I/O mode is not Block, call Hitachi Data Systems customer support.

4. Create a UR pair between the UR secondary and primary sites.
   Command example: `paircreate -g oraREMOTE -f async -vl -jp 0 -js 0 -IH2`
   You can resume operations at this point by using the P-VOL at the UR secondary site.

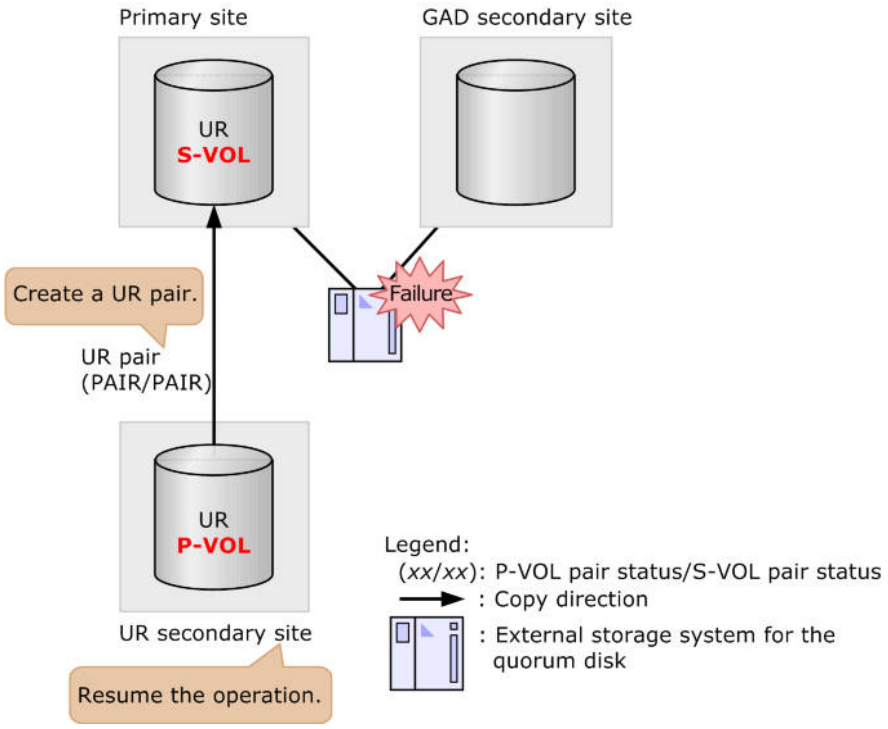

5. Remove the failure on the quorum disk.
6. Suspend the UR pair by specifying the S-VOL (swap suspend).
   Command example: `pairsplit -g oraREMOTE -RS -IH0`
   You can resume operations at this point by using the S-VOL at the primary site.

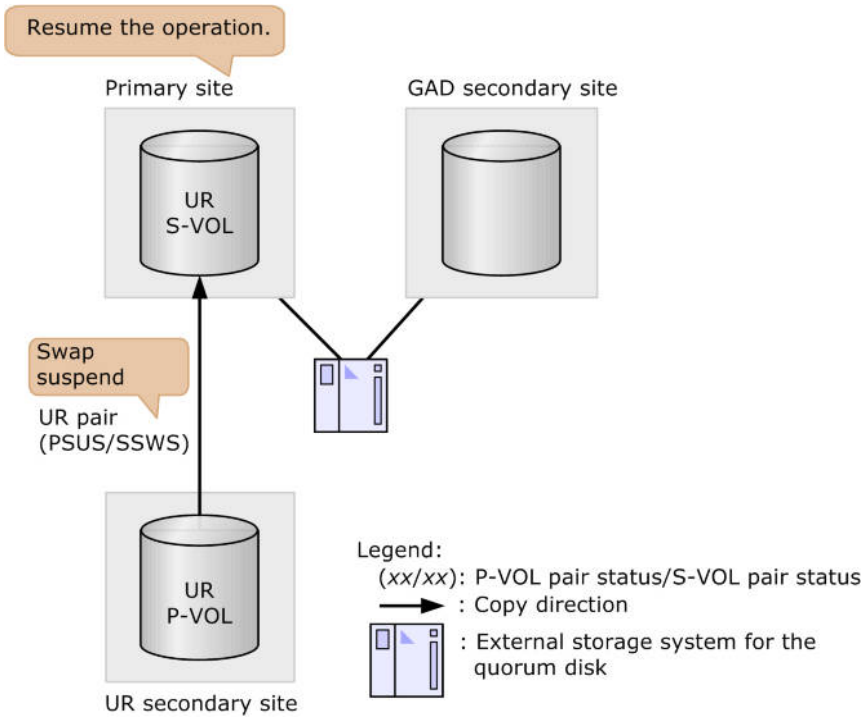

7. Resynchronize the UR pair by specifying the S-VOL (swap resync).
   Command example: `pairresync -g oraREMOTE -swaps -IH0`
   The volume on the primary storage system changes to a P-VOL, and the volume on the UR secondary storage system changes to an S-VOL.

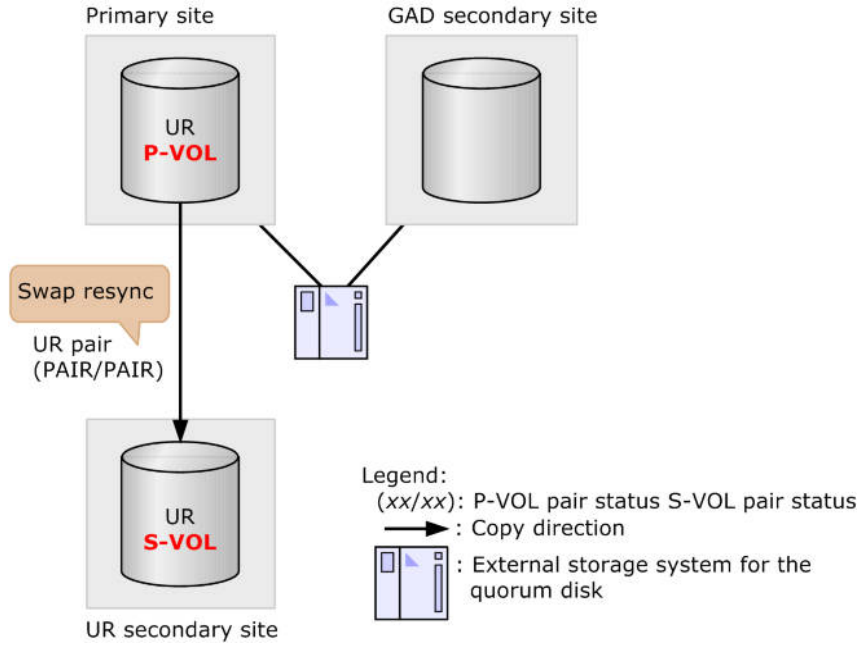

8. Delete the UR pair.
   Command example: `pairsplit -g oraREMOTE -S -IH0`
9. Re-create the GAD pair.
   Command example: `paircreate -g oraHA -fg never 2 -vl -jq 0 -IH0`

10. Re-create the UR delta resync pair.
    Command example: `paircreate -g oraDELTA -f async -vl -jp 0 -js 0 -nocsus -IH1`
11. Re-create the UR pair.
    Command example: `paircreate -g oraREMOTE -f async -vl -jp 0 -js 0 -IH0`

10

Planned outage of global-active device storage systems

This chapter describes and provides instructions for performing planned outages of global-active device (GAD) storage systems.

Planned power off/on: primary storage system

Powering off the primary storage system

1. Direct server I/O to the storage system at the secondary site.
   Using the alternate path software, stop I/O from servers to the storage system at the primary site.
2. On the storage system at the secondary site, suspend the GAD pairs to change the pair status of the S-VOLs to SSWS (swap suspension).
   ```
   pairsplit -g oraHA -RS -IH1
   ```
3. Verify that the pair status of P-VOLs of theGAD pairs has changed to PSUS(Block) and that the pair status of the S-VOLs has changed to SSWS(Local).
   ```
   pairdisplay -g oraHA -fcxe -IH1
   ```
4. Power off the storage system at the primary site.

Powering on the primary storage system

1. Power on the storage system at the primary site.
2. Confirm that the primary and secondary storage systems and the external storage system for the quorum disk do not have any blocked parts.
   If any parts are blocked, recover them.
3. Check for SIMs about path blockage, and delete the SIMs.
4. Confirm that the pair status of the GAD P-VOLs is PSUS(Block) and that the pair status of the S-VOLs is SSWS(Local).
   ```
   pairdisplay -g oraHA -fcxe -IH1
   ```
5. On the storage system at the secondary site, resynchronize the GAD pairs by reversing the primary and secondary volumes (swap resync).
   ```
   pairresync -g oraHA -swaps -IH1
   ```
6. Confirm that the pair status of the P-VOLs and S-VOLs of the GAD pairs has changed to PAIR (Mirror (RL)).
   ```
   pairdisplay -g oraHA -fcxe -IH1
   ```
7. Using the alternate path software, resume I/O from the servers to the storage system at the primary site.
8. If necessary, reverse the primary and secondary volumes.

Planned power off/on: secondary storage system

Powering off the secondary storage system

1. Direct server I/O to the storage system at the primary site.
   Using the alternate path software, stop I/O from servers to the storage system at the secondary site.

2. On the storage system at the primary site, suspend the GAD pairs by specifying the primary volume.
   ```
   pairsplit -g oraHA -r -IH0
   ```
3. Confirm that the pair status of P-VOLs of GAD pairs has changed to PSUS(Local) and the pair status of the S-VOLs has changed to SSUS(Block).
   ```
   pairdisplay -g oraHA -fcxe -IH0
   ```
4. Power off the storage system at the secondary site.

Powering on the secondary storage system

1. Power on the storage system at the secondary site.
2. Confirm that the primary and secondary storage systems and the external storage system for the quorum disk do not have any blocked parts.
   If any parts are blocked, recover them.
3. Check for SIMs about path blockage, and delete the SIMs.
4. Confirm that the pair status of the P-VOLs of the GAD pairs is PSUS(Local) and that the pair status of the S-VOLs is SSUS(Block).
   ```
   pairdisplay -g oraHA -fcxe -IH0
   ```
5. On the storage system at the primary site, resynchronize the GAD pairs by specifying the primary volume.
   ```
   pairresync -g oraHA -IH0
   ```
6. Confirm that the pair status of the P-VOLs and S-VOLs of the GAD pairs has changed to PAIR (Mirror (RL)).
   ```
   pairdisplay -g oraHA -fcxe -IH0
   ```
7. Using the alternate path software, resume I/O from the servers to the storage system at the secondary site.

Planned power off/on of the external storage system, I/O at primary site

Powering off the external storage system, I/O at primary site

1. Direct server I/O to the storage system at the primary site.
   Using the alternate path software, stop I/O from the servers to the storage system at the secondary site.
2. On the storage system at the primary site, suspend the GAD pairs by specifying the primary volume.
   ```
   pairsplit -g oraHA -r -IH0
   ```
3. Confirm that the pair status of P-VOLs of the GAD pairs has changed to PSUS(Local) and that the pair status of the S-VOLs has changed to SSUS(Block).
   ```
   pairdisplay -g oraHA -fcxe -IH0
   ```
4. On the primary and secondary storage systems, disconnect the quorum disks.

```
raidcom disconnect external_grp -ldev_id 0x9999 -IH0
```

Note: When you disconnect a quorum disk, SIM (def0*zz*) (quorum disk blockade) might be issued (*zz*: quorum disk ID). If this SIM is issued, delete the SIM after powering on the storage system and reconnecting the quorum disk.

5. On the primary and secondary storage systems, confirm that the quorum disk has been disconnected.
   ```
   raidcom get path -path_grp 1 -IH0
   ```
6. Power off the external storage system.

Powering on the external storage system, I/O at primary site

1. Power on the external storage system.
2. Confirm that the primary and secondary storage systems and the external storage system for the quorum disk do not have any blocked parts.
   If any parts are blocked, recover them.
3. Check for SIMs about path blockage, and delete the SIMs.
4. On the primary and secondary storage systems, establish the connections to the quorum disks.
   ```
   raidcom check_ext_storage external_grp -ldev_id 0x9999 -IH0
   ```
5. On the primary and secondary storage systems, confirm that the connections to the quorum disks have been established.
   ```
   raidcom get external_grp -external_grp_id 1-1 -IH0
   ```
6. Confirm that the external volumes of the primary and secondary storage systems are recognized as quorum disks.
   ```
   raidcom get ldev -ldev_id 0x9999 -IH0
   ```
7. Check for SIMs about quorum disk blockade, and delete the SIMs.
8. Confirm that the pair status of the P-VOLs of the GAD pairs is PSUS(Local) and that the pair status of the S-VOLs is SSUS(Block).
   ```
   pairdisplay -g oraHA -fcxe -IH0
   ```
9. Resynchronize the GAD pairs on the primary storage system by specifying the primary volume.
   ```
   pairresync -g oraHA -IH0
   ```
10. Confirm that the pair status of the P-VOLs and S-VOLs of the GAD pairs has changed to PAIR (Mirror (RL)).
    ```
    pairdisplay -g oraHA -fcxe -IH0
    ```
11. Using the alternate path software, resume server I/O to the storage system at the secondary site.

Planned power off/on of the external storage system, I/O at secondary site

Powering off the external storage system, I/O at secondary site

1. Direct server I/O to the storage system at the secondary site.
   Using the alternate path software, stop I/O from the servers to the storage system at the primary site.
2. On the secondary storage system, suspend the GAD pairs to change the pair status of the S-VOLs to SSWS (swap suspension).
   ```
   pairsplit -g oraHA -RS -IH1
   ```
3. Verify that the pair status of the P-VOLs of the GAD pairs has changed to PSUS(Block) and that the pair status of the S-VOLs has changed to SSWS(Local).
   ```
   pairdisplay -g oraHA -fcxe -IH1
   ```
4. On the primary and secondary storage systems, disconnect the quorum disks.
   ```
   raidcom disconnect external_grp -ldev_id 0x8888 -IH1
   ```

Note: When you disconnect a quorum disk, SIM (def0*zz*) (quorum disk blockade) might be issued (*zz*: quorum disk ID). If this SIM is issued, delete the SIM after powering on the storage system and reconnecting the quorum disk.

5. On the primary and secondary storage systems, confirm that the quorum disks have been disconnected.
   ```
   raidcom get path -path_grp 1 -IH1
   ```
6. Power off the external storage system.

Powering on the external storage system, I/O at secondary site

1. Power on the external storage system.
2. Confirm that the primary and secondary storage systems and the external storage systems for the quorum disks do not have any blocked parts.
   If any parts are blocked, recover them.
3. Check for SIMs about path blockage, and delete the SIMs.
4. On the primary and secondary storage systems, establish connections to the quorum disks.
   ```
   raidcom check_ext_storage external_grp -ldev_id 0x8888 -IH1
   ```
5. On the primary and secondary storage systems, confirm that the connections to the quorum disks have been established.
   ```
   raidcom get external_grp -external_grp_id 1-2 -IH1
   ```
6. Confirm that the external volumes of the primary and secondary storage systems are recognized as quorum disks.
   ```
   raidcom get ldev -ldev_id 0x8888 -IH1
   ```
7. Check for SIMs about quorum disk blockade, and delete the SIMs.

8. Confirm that the pair status of the P-VOLs of the GAD pairs is PSUS(Block) and that the pair status of the S-VOLs is SSWS(Local).
   ```
   pairdisplay -g oraHA -fcxe -IH1
   ```
9. Resynchronize the GAD pairs from the secondary storage system by reversing the primary and secondary volumes (swap resync).
   ```
   pairresync -g oraHA -swaps -IH1
   ```
10. Confirm that the pair status of the P-VOls and S-VOLs of the GAD pairs has changed to PAIR (Mirror (RL)).
    ```
    pairdisplay -g oraHA -fcxe -IH1
    ```
11. Using the alternate path software, resume I/O from the servers to the storage system at the primary site.
12. If necessary, reverse the primary and secondary volumes.

Planned power off/on: primary and secondary storage systems

Powering off the primary and secondary storage systems

1. Stop server I/O to the primary and secondary storage systems.
2. On the primary storage system, suspend the GAD pairs by specifying the primary volume.
   ```
   pairsplit -g oraHA -r -IH0
   ```
3. Confirm that the pair status of the P-VOLs of the GAD pairs has changed to PSUS(Local) and that the pair status of the S-VOLs has changed to SSUS(Block).
   ```
   pairdisplay -g oraHA -fcxe -IH0
   ```
4. Power off the primary and secondary storage systems.

Powering on the primary and secondary storage systems

1. Power on the primary and secondary storage systems.
2. Confirm that the primary and secondary storage systems and the external storage system for the quorum disks do not have any blocked parts.
   If any parts are blocked, recover them.
3. Check for SIMs about path blockage, and delete the SIMs.
4. Confirm that the pair status of the P-VOLs of the GAD pairs is PSUS(Local) and that the pair status of the S-VOLs is SSUS(Block).
   ```
   pairdisplay -g oraHA -fcxe -IH0
   ```
5. On the primary storage system, resynchronize the GAD pairs by specifying the primary volume.
   ```
   pairresync -g oraHA -IH0
   ```
6. Confirm that the pair status of the P-VOLs and S-VOLs of the GAD pairs has changed to PAIR (Mirror (RL)).
   ```
   pairdisplay -g oraHA -fcxe -IH0
   ```

7. Resume I/O from the servers to the primary and secondary storage systems.

Planned power off/on of the primary and external storage systems

Powering off the primary and external storage systems

1. Direct server I/O to the secondary storage system.

   Using the alternate path software, stop I/O from the servers to the primary storage system.
2. On the secondary storage system, swap and suspend the GAD pairs (swap suspension).

   ```
   pairsplit -g oraHA -RS -IH1
   ```
3. Confirm that the pair status of the P-VOLs of the GAD pairs has changed to PSUS(Block) and that the pair status of the S-VOLs has changed to SSWS(Local).

   ```
   pairdisplay -g oraHA -fcxe -IH1
   ```
4. On the primary and secondary storage systems, disconnect the quorum disks.

   ```
   raidcom disconnect external_grp -ldev_id 0x8888 -IH1
   ```

Note: When you disconnect a quorum disk, SIM (def0*zz*) (quorum disk blockade) might be issued (*zz*: quorum disk ID). If this SIM is issued, delete the SIM after powering on the storage system and reconnecting the quorum disk.

5. On the primary and secondary storage systems, verify that the quorum disks have been disconnected.

   ```
   raidcom get path -path_grp 1 -IH1
   ```
6. Power off the storage system at the primary site and the external storage system.

Powering on the primary and external storage systems

1. Power on the storage system at the primary site and the external storage system.
2. Confirm that the primary and secondary storage systems and the external storage system for the quorum disks do not have any blocked parts.

   If any parts are blocked, recover them.
3. Check for SIMs about path blockage, and delete the SIMs.
4. On the primary and secondary storage systems, establish connections to the quorum disks.

   ```
   raidcom check_ext_storage external_grp -ldev_id 0x8888 -IH1
   ```
5. On the primary and secondary storage systems, verify that the connections to the quorum disks have been established.

```
raidcom get external_grp -external_grp_id 1-2 -IH1
```

6. Confirm that external volumes of the primary and secondary storage systems are recognized as quorum disks.
   ```
   raidcom get ldev -ldev_id 0x8888 -IH1
   ```
7. Check for SIMs about quorum disk blockade, and delete the SIMs.
8. Confirm that the pair status of the P-VOLs of the GAD pairs is PSUS(Block) and that the pair status of the S-VOLs is SSWS(Local).
   ```
   pairdisplay -g oraHA -fcxe -IH1
   ```
9. Resynchronize GAD pairs from the secondary storage system by reversing the primary and secondary volumes (swap resync).
   ```
   pairresync -g oraHA -swaps -IH1
   ```
10. Confirm that the pair status of the P-VOLs and S-VOLs of the GAD pairs has changed to PAIR (Mirror (RL)).
    ```
    pairdisplay -g oraHA -fcxe -IH1
    ```
11. Using the alternate path software, resume I/O from the servers to the storage system at the primary site.
12. If necessary, reverse the primary and secondary volumes.

Planned power off/on of the secondary and external storage systems

Powering off the secondary and external storage systems

1. Direct server I/O to the storage system at the primary site.
   Using the alternate path software, stop I/O from the servers to the secondary storage system.
2. On the primary storage system, suspend the GAD pairs by specifying the primary volume.
   ```
   pairsplit -g oraHA -r -IH0
   ```
3. Confirm that the pair status of the P-VOLs of the GAD pairs has changed to PSUS(Local) and that the pair status of the S-VOLs has changed to SSUS(Block).
   ```
   pairdisplay -g oraHA -fcxe -IH0
   ```
4. On the primary and secondary storage systems, disconnect the quorum disks.
   ```
   raidcom disconnect external_grp -ldev_id 0x9999 -IH0
   ```

Note: When you disconnect a quorum disk, SIM (def0*zz*) (quorum disk blockade) might be issued (*zz*: quorum disk ID). If this SIM is issued, delete the SIM after powering on the storage system and reconnecting the quorum disk.

5. On the primary and secondary storage systems, verify that the quorum disks have been disconnected.
   ```
   raidcom get path -path_grp 1 -IH0
   ```

6. Power off the storage system at the secondary site and the external storage system.

Powering on the secondary and external storage systems

1. Power on the storage system at the secondary site and the external storage system.
2. Confirm that the primary and secondary storage systems and the external storage system for the quorum disks do not have any blocked parts.

   If any parts are blocked, recover them.
3. Check for SIMs about path blockage, and delete the SIMs.
4. On the primary and secondary storage systems, establish connections to the quorum disks.

   ```
   raidcom check_ext_storage external_grp -ldev_id 0x9999 -IH0
   ```
5. On the primary and secondary storage systems, verify that the connections to the quorum disks have been established.

   ```
   raidcom get external_grp -external_grp_id 1-1 -IH0
   ```
6. Confirm that the external volumes of the primary and secondary storage systems are recognized as quorum disks.

   ```
   raidcom get ldev -ldev_id 0x9999 -IH0
   ```
7. Check for SIMs about quorum disk blockade, and delete the SIMs.
8. Confirm that the pair status of the P-VOLs of the GAD pairs is PSUS(Local) and that the pair status of the S-VOLs is SSUS(Block).

   ```
   pairdisplay -g oraHA -fcxe -IH0
   ```
9. Resynchronize the GAD pairs from the primary storage system by specifying the primary volume.

   ```
   pairresync -g oraHA -IH0
   ```
10. Confirm that the pair status of the P-VOLs and S-VOLs of the GAD pairs has changed to PAIR (Mirror (RL)).

    ```
    pairdisplay -g oraHA -fcxe -IH0
    ```
11. Using the alternate path software, resume I/O from the servers to the storage system at the secondary site.

Planned power off/on of all GAD storage systems

Powering off the primary, secondary, and external storage systems

1. Using the alternate path software, stop server I/O to the primary and secondary storage systems.
2. On the primary storage system, suspend the GAD pairs.

   ```
   pairsplit -g oraHA -r -IH0
   ```
3. Confirm that the pair status of the P-VOLs of the GAD pairs has changed to PSUS(Local) and that the pair status of the S-VOLs has changed to SSUS(Block).

   ```
   pairdisplay -g oraHA -fcxe -IH0
   ```

4. On the primary and secondary storage systems, disconnect the quorum disks.
   ```
   raidcom disconnect external_grp -ldev_id 0x9999 -IH0
   ```

Note: When you disconnect a quorum disk, SIM (def0*zz*) (quorum disk blockade) might be issued (*zz*: quorum disk ID). If this SIM is issued, delete the SIM after powering on the storage system and reconnecting the quorum disk.

5. On the primary and secondary storage systems, verify that the quorum disks have been disconnected.
   ```
   raidcom get path -path_grp 1 -IH0
   ```
6. Power off the primary and secondary storage systems and the external storage system.

Powering on the primary, secondary, and external storage systems

1. Power on the primary and secondary storage systems and the external storage system.
2. Confirm that the primary and secondary storage systems and the external storage system for the quorum disks do not have any blocked parts.
   If any parts are blocked, recover them.
3. Check for SIMs about path blockage, and delete the SIMs.
4. On the primary and secondary storage systems, establish connections to the quorum disks.
   ```
   raidcom check_ext_storage external_grp -ldev_id 0x9999 -IH0
   ```
5. On the primary and secondary storage systems, verify that the connections to the quorum disks have been established.
   ```
   raidcom get external_grp -external_grp_id 1-1 -IH0
   ```
6. Confirm that the external volumes of the primary and secondary storage systems are recognized as quorum disks.
   ```
   raidcom get ldev -ldev_id 0x9999 -IH0
   ```
7. Check for SIMs about quorum disk blockade, and delete the SIMs.
8. Confirm that the pair status of the P-VOLs of the GAD pairs is PSUS(Local) and that the pair status of the S-VOLs is SSUS(Block).
   ```
   pairdisplay -g oraHA -fcxe -IH0
   ```
9. Resynchronize the GAD pairs from the primary storage system by specifying the primary volume.
   ```
   pairresync -g oraHA -IH0
   ```
10. Confirm that the pair status of the P-VOLs and S-VOLs of the GAD pairs has changed to PAIR (Mirror (RL)).
    ```
    pairdisplay -g oraHA -fcxe -IH0
    ```
11. Using the alternate path software, resume I/O from the servers to the storage systems at the primary and secondary sites.

Planned outage of storage systems in a GAD 3DC delta resync environment

This topic describes and provides instructions for performing planned outages of storage systems in a GAD 3DC delta resync (GAD+UR) environment.

For details about storage system support (models, microcode) for GAD+UR operations, see Requirements and restrictions on page 2-2.

Planned power off/on: primary GAD+UR storage system

Powering off the GAD+UR primary storage system

Use the following procedure to power off the primary storage system.

1. Using the alternate path software, stop I/O from servers to the primary storage system.
   Direct server I/O to the GAD secondary storage system.
2. On the GAD secondary storage system, suspend the GAD pairs by specifying the S-VOL (swap suspension).
   Command example: `pairsplit -g oraHA -RS -IH1`
3. Confirm that the pair types and the pair statuses are changed as shown in the following table.

Before the GAD pair suspension			After the GAD pair suspension		
Pair type	Pair status		Pair type	Pair status	
	P-VOL	S-VOL		P-VOL	S-VOL
GAD pair	PAIR (Mirror(RL))	PAIR (Mirror(RL))	GAD pair	PSUS (Block)	SSWS (Local)
UR pair	PAIR	PAIR	UR delta resync pair	PSUE[1]	SSUS[2]
UR delta resync pair	PSUS[2]	SSUS[2]	UR pair	PAIR	PAIR

Notes:

1. Device Manager - Storage Navigator displays HOLDING.
2. Device Manager - Storage Navigator displays HOLD.

4. Power off the primary storage system.

Related topics

- I/O modes on page 1-14

Powering on the GAD+UR primary storage system

Use the following procedure to power on the primary storage system.

1. Power on the primary storage system.
2. Confirm that the storage systems do not have any blocked parts.
   If any parts are blocked, recover them.
3. Check for SIMs about path blockage, and delete the SIMs.
4. On the GAD secondary storage system, resynchronize the GAD pair by specifying the S-VOL (swap resync).
   Command example: `pairresync -g oraHA -swaps -IH1`
   A volume on the primary storage system changes to an S-VOL, and a volume on the secondary storage system changes to a P-VOL.
5. Confirm that the pair statuses are changed as shown in the following table.

Pair type	Before the GAD pair resync		After the GAD pair resync	
	P-VOL	S-VOL	P-VOL	S-VOL
GAD pair	PSUS (Block)	SSWS (Local)	PAIR (Mirror(RL))	PAIR (Mirror(RL))
UR delta resync pair	PSUE[1]	SSUS[2]	PSUE[1]	SSUS[2]
UR pair	PAIR	PAIR	PAIR	PAIR

Notes:
1. Device Manager - Storage Navigator displays HOLDING.
2. Device Manager - Storage Navigator displays HOLD.

6. Keep updating I/O to the P-VOL or GAD S-VOL for two minutes.
   The information on the system configuration is reported among the primary, GAD secondary, and UR secondary storage systems. At the time, the UR delta resync pair status is changed from PSUE to PSUS. The mirror status of the journal is displayed as PJNS. Confirm that no failure SIMs are displayed.

Note: The GAD P-VOL and S-VOL are now reversed compared to the original configuration before the planned outage. To restore the original configuration, reverse the GAD P-VOL and S-VOL.

Related topics

- I/O modes on page 1-14
- Reversing the GAD P-VOL and S-VOL (GAD+UR) on page 9-4

Planned power off/on: GAD secondary storage system

Powering off the GAD secondary storage system

Use the following procedure to power off the GAD secondary storage system.

1. Using the alternate path software, stop I/O from servers to the GAD secondary storage system.

Direct server I/O to the primary storage system.

2. On the primary storage system, resynchronize the GAD pair by specifying the P-VOL.
   Command example: `pairsplit -g oraHA -IH0`
3. Confirm that the pair statuses are changed as shown in the following table.

Pair type	Before the GAD pair suspension		After the GAD pair suspension	
	P-VOL	S-VOL	P-VOL	S-VOL
GAD pair	PAIR (Mirror(RL))	PAIR (Mirror(RL))	PSUS (Local)	PSUS (Block)
UR pair	PAIR	PAIR	PAIR	PAIR
UR delta resync pair	PSUS[1]	SSUS[1]	PSUE[2]	SSUS[1]

Notes:
1. Device Manager - Storage Navigator displays HOLD.
2. Device Manager - Storage Navigator displays HOLDING.

4. Power off the GAD secondary storage system.

Related topics

- I/O modes on page 1-14

Powering on the GAD secondary storage system

Use the following procedure to power on the GAD secondary storage system.

1. Power on the GAD secondary storage system.
2. Confirm that the primary and secondary storage systems and the external storage system for the quorum disk do not have any blocked parts.
   If any parts are blocked, recover them.
3. Check for SIMs about path blockage, and delete the SIMs.
4. On the primary storage system, resynchronize the GAD pairs by specifying the P-VOL.
   Command example: `pairresync -g oraHA -IH0`
5. Confirm that the pair statuses are changed as shown in the following table.

Pair type	Before the GAD pair resync		After the GAD pair resync	
	P-VOL	S-VOL	P-VOL	S-VOL
GAD pair	PSUS (Local)	PSUS (Block)	PAIR (Mirror(RL))	PAIR (Mirror(RL))
UR pair	PAIR	PAIR	PAIR	PAIR

Pair type	Before the GAD pair resync		After the GAD pair resync	
	P-VOL	S-VOL	P-VOL	S-VOL
UR delta resync pair	PSUE[1]	SSUS[2]	PSUE[1]	SSUS[2]

Notes:

1. Device Manager - Storage Navigator displays HOLDING.
2. Device Manager - Storage Navigator displays HOLD.

6. Keep updating I/O to P-VOL or GAD S-VOL for two minutes.

   The information on the system configuration is reported among the primary, GAD secondary, and UR secondary storage systems. At the time, the UR delta resync pair status is changed from PSUE to PSUS. The mirror status of the journal is displayed as PJNS. Confirm that no failure SIMs are displayed.

Related topics

- I/O modes on page 1-14

Planned power off/on: UR secondary storage system

Powering off the UR secondary storage system

Use the following procedure to power off the UR secondary storage system.

1. On the primary storage system, suspend the UR pairs by specifying the P-VOL.

   Command example: `pairsplit -g oraREMOTE -IH0`

2. Confirm that the pair statuses are changed as shown in the following table.

Pair type	Before the GAD pair suspension		After the GAD pair suspension	
	P-VOL	S-VOL	P-VOL	S-VOL
GAD pair	PAIR (Mirror(RL))	PAIR (Mirror(RL))	PAIR (Mirror(RL))	PAIR (Mirror(RL))
UR pair	PAIR	PAIR	PSUS	PSUS
UR delta resync pair	PSUS[1]	SSUS[1]	PSUE[2]	SSUS[1]

Notes:

1. Device Manager - Storage Navigator displays HOLD.
2. Device Manager - Storage Navigator displays HOLDING.

3. Power off the UR secondary storage system.

Related topics

- I/O modes on page 1-14

Powering on the UR secondary storage system

Use the following procedure to power on the UR secondary storage system.

1. Power on the UR secondary storage system.
2. Confirm that the primary and secondary storage systems and the external storage system for the quorum disk do not have any blocked parts.
   If any parts are blocked, recover them.
3. Check for SIMs about path blockage, and delete the SIMs.
4. On the primary storage system, resynchronize the UR pairs by specifying the P-VOL.
   Command example: `pairresync -g oraREMOTE -IH0`
5. Confirm that the pair statuses are changed as shown in the following table.

Pair type	Before the GAD pair resync		After the GAD pair resync	
	P-VOL	S-VOL	P-VOL	S-VOL
GAD pair	PAIR (Mirror(RL))	PAIR (Mirror(RL))	PAIR (Mirror(RL))	PAIR (Mirror(RL))
UR pair	PSUS	PSUS	PAIR	PAIR
UR delta resync pair	PSUE[1]	SSUS[2]	PSUE[1]	SSUS[2]

Notes:

1. Device Manager - Storage Navigator displays HOLDING.
2. Device Manager - Storage Navigator displays HOLD.

6. Keep updating I/O to P-VOL or GAD S-VOL for two minutes.
   The information on the system configuration is reported among the primary, GAD secondary, and UR secondary storage systems. At the time, the UR delta resync pair status is changed from PSUE to PSUS. The mirror status of the journal is displayed as PJNS. Confirm that no failure SIMs are displayed.

Related topics

- I/O modes on page 1-14

Planned power off/on: external storage system, I/O at primary site (GAD+UR)

Powering off the external storage system, I/O at primary site (GAD+UR)

Use the following procedure to power off the external storage system when server I/O continues to the primary storage system.

1. Using the alternate path software, stop I/O from servers to the GAD secondary storage system.

   Direct server I/O to the primary storage system.
2. On the primary storage system, resynchronize the GAD pair by specifying the P-VOL.

   Command example: `pairsplit -g oraHA -IH0`
3. Confirm that the pair statuses are changed as shown in the following table.

Pair type	Before the GAD pair suspension		After the GAD pair suspension	
	P-VOL	S-VOL	P-VOL	S-VOL
GAD pair	PAIR (Mirror(RL))	PAIR (Mirror(RL))	PSUS (Local)	PSUS (Block)
UR pair	PAIR	PAIR	PAIR	PAIR
UR delta resync pair	PSUS[1]	SSUS[1]	PSUE[2]	SSUS[1]

Notes:

1. Device Manager - Storage Navigator displays HOLD.
2. Device Manager - Storage Navigator displays HOLDING.

4. On the primary and GAD secondary storage systems, disconnect the quorum disks.

   Command example: `raidcom disconnect external_grp -ldev_id 0x9999 -IH0`

Note: When you disconnect a quorum disk, SIM (def0*zz*) (quorum disk blockade) might be issued (*zz*: quorum disk ID). If this SIM is issued, delete the SIM after powering on the storage system and reconnecting the quorum disk.

5. Power off the external storage system.

Related topics

- I/O modes on page 1-14

Powering on the external storage system, I/O at primary site (GAD+UR)

Use the following procedure to power on the external storage system when server I/O continues to the primary storage system.

1. Power on the external storage system.
2. Confirm that the primary and secondary storage systems and the external storage system for the quorum disk do not have any blocked parts.

   If any parts are blocked, recover them.
3. Check for SIMs about path blockage, and delete the SIMs.
4. On the primary and GAD secondary storage systems, establish the connections to the quorum disks.

Command example: `raidcom check_ext_storage external_grp -ldev_id 0x9999 -IH0`

5. On the primary and GAD secondary storage systems, confirm that the connections to the quorum disks have been established.
   Command example: `raidcom get external_grp -external_grp_id 1-1 -IH0`
6. Confirm that the external volumes of the primary and GAD secondary storage systems are recognized as quorum disks.
   Command example: `raidcom get ldev -ldev_id 0x9999 -IH0`
7. On the primary and GAD secondary storage systems, check for SIMs about quorum disk blockade, and delete the SIMs.
8. Resynchronize the GAD pairs on the primary storage system by specifying the P-VOL.
   Command example: `pairresync -g oraHA -IH0`
9. Confirm that the pair statuses are changed as shown in the following table.

Pair type	Before the GAD pair resync		After the GAD pair resync	
	P-VOL	S-VOL	P-VOL	S-VOL
GAD pair	PSUS (Local)	PSUS (Block)	PAIR (Mirror(RL))	PAIR (Mirror(RL))
UR pair	PAIR	PAIR	PAIR	PAIR
UR delta resync pair	PSUE[1]	SSUS[2]	PSUE[1]	SSUS[2]

Notes:
1. Device Manager - Storage Navigator displays HOLDING.
2. Device Manager - Storage Navigator displays HOLD.

10. Keep updating I/O to P-VOL or GAD S-VOL for two minutes.
    The information on the system configuration is reported among the primary, GAD secondary, and UR secondary storage systems. At the time, the UR delta resync pair status is changed from PSUE to PSUS. The mirror status of the journal is displayed as PJNS. Confirm that no failure SIMs are displayed.

Related topics

- I/O modes on page 1-14

Planned power off/on: external storage system, I/O at secondary site (GAD+UR)

Powering off external storage system, I/O at secondary site (GAD+UR)

Use the following procedure to power off the external storage system when server I/O continues to the GAD storage system.

1. Using the alternate path software, stop I/O from servers to the primary storage system.

   Direct server I/O to the GAD secondary storage system.
2. On the GAD secondary storage system, suspend the GAD pairs by specifying the S-VOL (swap suspension).

   Command example: `pairsplit -g oraHA -RS -IH1`
3. Confirm that the pair types and the pair statuses are changed as shown in the following table.

Before the GAD pair suspension			After the GAD pair suspension		
Pair type	Pair status		Pair type	Pair status	
	P-VOL	S-VOL		P-VOL	S-VOL
GAD pair	PAIR (Mirror(RL))	PAIR (Mirror(RL))	GAD pair	PSUS (Block)	SSWS (Local)
UR pair	PAIR	PAIR	UR delta resync pair	PSUE[1]	SSUS[2]
UR delta resync pair	PSUS[2]	SSUS[2]	UR pair	PAIR	PAIR

Notes:

1. Device Manager - Storage Navigator displays HOLDING.
2. Device Manager - Storage Navigator displays HOLD.

4. On the primary and GAD secondary storage systems, disconnect the quorum disks.

   Command example: `raidcom disconnect external_grp -ldev_id 0x8888 -IH1`

Note: When you disconnect a quorum disk, SIM (def0*zz*) (quorum disk blockade) might be issued (*zz*: quorum disk ID). If this SIM is issued, delete the SIM after powering on the storage system and reconnecting the quorum disk.

5. Power off the external storage system.

Related topics

- I/O modes on page 1-14

Powering on external storage system, I/O at secondary site (GAD+UR)

Use the following procedure to power on the external storage system when server I/O continues to the GAD storage system.

1. Power on the external storage system.
2. Confirm that the primary and secondary storage systems and the external storage system for the quorum disk do not have any blocked parts.

   If any parts are blocked, recover them.

3. Check for SIMs about path blockage, and delete the SIMs.
4. On the primary and GAD secondary storage systems, establish the connections to the quorum disks.

   Command example: `raidcom check_ext_storage external_grp -ldev_id 0x8888 -IH1`
5. On the primary and GAD secondary storage systems, confirm that the connections to the quorum disks have been established.

   Command example: `raidcom get external_grp -external_grp_id 1-2 -IH1`
6. Confirm that the external volumes of the primary and GAD secondary storage systems are recognized as quorum disks.

   Command example: `raidcom get ldev -ldev_id 0x8888 -IH1`
7. On the primary and GAD secondary storage systems, check for SIMs about quorum disk blockade, and delete the SIMs.
8. Resynchronize the GAD pairs on the GAD secondary storage system by specifying the S-VOL (swap resync).

   Command example: `pairresync -g oraHA -swaps -IH1`

   A volume on the primary storage system changes to an S-VOL, and a volume on the secondary storage system changes to a P-VOL.
9. Confirm that the pair statuses are changed as shown in the following table.

Pair type	Before the GAD pair resync		After the GAD pair resync	
	P-VOL	S-VOL	P-VOL	S-VOL
GAD pair	PSUS (Block)	SSWS (Local)	PAIR (Mirror(RL))	PAIR (Mirror(RL))
UR delta resync pair	PSUE[1]	SSUS[2]	PSUE[1]	SSUS[2]
UR pair	PAIR	PAIR	PAIR	PAIR

Notes:

1. Device Manager - Storage Navigator displays HOLDING.
2. Device Manager - Storage Navigator displays HOLD.

10. Keep updating I/O to P-VOL or GAD S-VOL for two minutes.

    The information on the system configuration is reported among the primary, GAD secondary, and UR secondary storage systems. At the time, the UR delta resync pair status is changed from PSUE to PSUS. The mirror status of the journal is displayed as PJNS. Confirm that no failure SIMs are displayed.

Note: The GAD P-VOL and S-VOL are reversed compared to the original configuration (before the planned outage). To restore the original configuration, reverse the GAD P-VOL and S-VOL.

Related topics

- I/O modes on page 1-14
- Reversing the GAD P-VOL and S-VOL (GAD+UR) on page 9-4

11

Data migration and volume reuse

This chapter describes and provides instructions for performing nondisruptive data migration using global-active device (GAD) and discontinuing GAD operations after the migration so that the volumes at both sites can be reused.

Workflow for data migration

The global-active device (GAD) feature enables you to perform data migration without interrupting business operations.

The following figure shows the system configuration for data migration using global-active device. The source and target storage systems must both support global-active device.

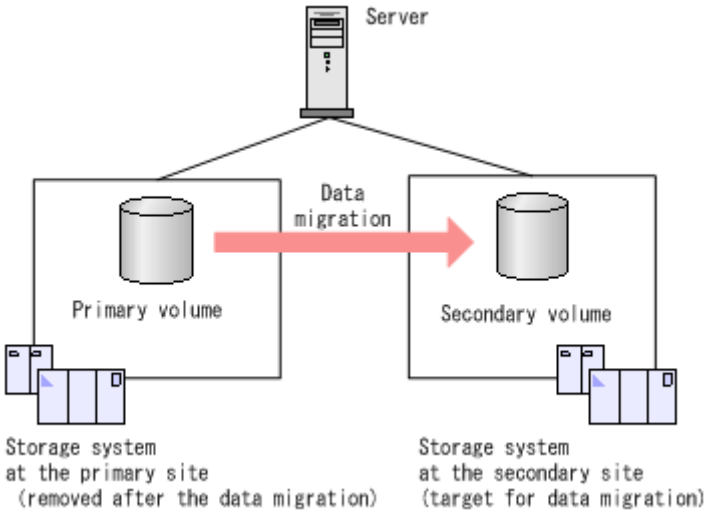

1. Create GAD pairs between the primary and secondary storage systems. The data on the volumes is duplicated, and the server issues I/O operations to volumes in both storage systems of the GAD pairs.
2. Monitor the status of the GAD pairs, and make sure that the pair status of all pairs is **PAIR** before continuing.
3. On the server, stop I/O to the primary volumes at the primary site.

   At this time, do not stop I/O to the secondary volumes at the secondary site.
4. At the secondary site, suspend the GAD pairs by specifying the S-VOLs.

   When you suspend a GAD pair by specifying the S-VOL, the pair status I/O mode of the P-VOL and S-VOL change as follows:

   - The pair status of the S-VOL changes to SSWS, and the I/O mode of the S-VOL changes to Local.
   - The pair status of the P-VOL changes to PSUS, and the I/O mode of the P-VOL changes to Block.
5. At the secondary site, delete the GAD pairs by specifying the S-VOLs.

   When you delete a GAD pair by specifying the S-VOL, the GAD reserve attribute is applied to the volume that was the P-VOL, which prevents host access to the volume. The volume that was the S-VOL keeps the virtual LDEV ID and continues to receive I/O from the server.
6. After you have deleted the GAD pairs, at the primary site delete the LU paths to the volumes that were the P-VOLs. If desired, you can now delete the volumes at the primary site, as they have been nondisruptively migrated to the secondary site.

7. On the primary and secondary storage systems, release the quorum disk settings for the external volume that was the quorum disk.
8. On the primary and secondary storage systems, disconnect the external volume that was the quorum disk.

Note: When you disconnect the quorum disk, SIM (def0*zz*) (quorum disk blockade) might be issued (*zz*: quorum disk ID). If this SIM is issued, you can delete it.

9. On the primary and secondary storage systems, delete the remote connections between the storage systems.
10. If necessary, uninstall the GAD license.
11. Remove the physical paths between the primary and secondary storage systems.
12. Stop and remove the storage system at the primary site.

Reusing volumes after data migration

This topic provides instructions for reusing volumes that were the P-VOLs and S-VOLs of GAD pairs that have been deleted.

Reusing a volume that was an S-VOL

When you delete a GAD pair by specifying the P-VOL, the GAD reserve attribute remains set for the volume that was the S-VOL. When you delete a GAD pair by specifying the S-VOL, the GAD reserve attribute is applied to the volume that was the P-VOL. When you check a volume that has the reserve attribute using the `raidcom get ldev` command, the virtual LDEV ID (`VIR_LDEV`) is displayed as `ffff`.

1. Delete the LU path to the volume that has the reserve attribute.
2. Remove the reserve attribute.
   Command example for removing the reserve attribute for LDEV ID (0x4444):
   ```
   raidcom unmap resource -ldev_id 0x4444 -virtual_ldev_id reserve
   ```
   The volume from which the reserve attribute was removed changes to a volume whose virtual LDEV ID was deleted. If you use the `raidcom get ldev` command to check a volume whose virtual LDEV ID was deleted, the virtual LDEV ID (`VIR_LDEV`) is displayed as `fffe`.
3. Reserve an LDEV ID for the resource group that will use the volume.
   Command example for reserving LDEV ID (0x4444) for resource group (#0):
   ```
   raidcom add resource -resource_name meta_resource -ldev_id 0x4444
   ```
4. Set a virtual LDEV ID for the volume.

Note: You must set a virtual LDEV ID that is unique within the storage system that uses the volume. If the same virtual LDEV ID is used in other storage systems or virtual storage machines with the same serial number

and model, identification of multiple volumes with the same virtual LDEV ID might cause problems. At worst, the server might detect an inconsistency.

Command example for setting virtual LDEV ID (0x5555) for volume (0x4444):

```
raidcom map resource -ldev_id 0x4444 -virtual_ldev_id 0x5555
```

5. Specify a new port and host group for the volume, and set an LU path.

Reusing a volume that was a P-VOL

After you delete a GAD pair by specifying the P-VOL, you can continue to use the volume that was the P-VOL of the pair. When you execute the `raidcom get ldev` command for a volume that continues to be available after pair deletion, a value other than `ffff` or `fffe` is displayed for the VIR_LDEV (virtual LDEV ID), or the VIR_LDEV is not displayed.

Use the following procedure to move the volume to another resource group (virtual storage machine) so that the server recognizes it as a different volume and it can be used.

1. Delete the LU path to the volume.
2. Delete the virtual LDEV ID.

   Command example for deleting virtual LDEV ID (0x5555) for LDEV ID (0x4444):

   ```
   raidcom unmap resource -ldev_id 0x4444 -virtual_ldev_id 0x5555
   ```

   When you delete the virtual LDEV ID, the volume changes to a volume whose virtual LDEV ID has been deleted. If you execute the `raidcom get ldev` command for a volume whose virtual LDEV ID has been deleted, `fffe` is displayed for the VIR_LDEV (virtual LDEV ID).
3. Reserve an LDEV ID for a resource group to be used for a different purpose.

   Command example for reserving LDEV ID (0x4444) for resource group (AnotherGroup) to which the volume is registered:

   ```
   raidcom add resource -resource_name AnotherGroup -ldev_id 0x4444
   ```
4. Set a virtual LDEV ID for the volume.

Note: You must set a virtual LDEV ID that is unique within the storage system that uses the volume. If the same virtual LDEV ID is used in other storage systems or virtual storage machines with the same serial number and model, identification of multiple volumes with the same virtual LDEV ID might cause problems. At worst, the server might detect an inconsistency.

Command example for setting virtual LDEV ID (0xe000) for volume (0x4444):

```
raidcom map resource -ldev_id 0x4444 -virtual_ldev_id 0xe000
```

5. Specify a new port and host group for the volume, and set an LU path.

12

Troubleshooting

This chapter provides troubleshooting information for global-active device operations and instructions for contacting Hitachi Data Systems customer support.

General troubleshooting

Problem	Recommended action
The HCS computer stops, or global-active device does not operate properly.	• Verify that there are no problems with the HCS computer, with the Ethernet connection, or with the software products, and then restart the HCS computer. Restarting the HCS computer does not affect the GAD operations that are currently running. • Confirm that all GAD requirements and restrictions (such as LU types) are met. • Confirm that the storage systems of the primary site and secondary site are powered on and that their functions are fully enabled. • Confirm that all values and parameters (such as serial number and ID of the remote storage system, path parameter, IDs of the primary volume and secondary volume) were entered correctly.
The LED on the storage system's control panel that indicates that the channel of the initiator is available is off or blinking.	Contact Hitachi Data Systems customer support.
GAD error messages are displayed on the HCS computer.	Correct the error, and then re-execute the GAD operation.
The status of a path to the remote storage system is abnormal.	Check the status of the paths in the Remote Connections window, and make the required corrections.
A timeout error occurred while creating a pair or resynchronizing a pair.	• If a timeout occurred due to a hardware error, a SIM is generated. Contact Hitachi Data Systems customer support, and after solving the problem, re-execute the GAD operation. • Large workload: If a SIM is not generated, wait for 5 to 6 minutes, and then check the status of the pair you want to create or resynchronize. If the pair status correctly changed, the failed operation completed after the timeout error message was displayed. If the pair status did not change as anticipated, the GAD operation cannot complete due to the large workload. In this case, re-execute the GAD operation when the workload of the storage system is smaller. • If a communication error between HCS and the SVP occurred, see *Hitachi Command Suite User Guide*.
A GAD volume has pinned tracks.	Recover the pinned track volume.
The monitoring switch is enabled, but the monitoring data is not updated.	• The monitoring data might not be updated because the time setting of SVP was changed. Disable the monitoring switch, and then enable it again. For details about the monitoring switch, see the *Provisioning Guide* for the storage system.

Problem	Recommended action
	• Verify that the settings for the target being monitored are correct.

Related topics

- Troubleshooting related to remote path status on page 12-3
- Procedure for recovering pinned track of a GAD volume on page 12-20

Troubleshooting related to remote path status

Remote path status and description	Status description	Recommended action
Normal Normal	This remote path is correctly set, and the path can be used for copying GAD.	The remote path status is normal. Recovery is not required.
Initialization Failed Initialization error	A physical connection between the local storage system and the remote storage system, or a connection between the local storage system and the switch does not exist. Therefore, the error occurred when the connection to the remote storage system was initialized.	Check the following, and correct them if they are not correct: • The cable between the ports of the local storage system and the remote storage system or between the ports of the local storage system and the switch of the local storage system is properly connected. • The serial number (S/N) and system ID of the remote storage system, the port number of the local storage system, and the port number of the remote storage system are correct. • The topology (Fabric, FC-AL, Point-to-point) of the ports of the local storage system and remote storage system is correctly set.
Communication Time Out Communication timeout	A timeout occurred in a communication between the local storage system and remote storage system.	Check the following, and correct them if they are not correct: • The remote storage system is powered on, and the remote storage system can be used normally. • The following network relaying devices are correctly configured, and can be properly used: ◦ Connectors

Remote path status and description	Status description	Recommended action
		◦ Cables ◦ Switches (zoning settings) ◦ Channel extenders (if channel extenders are connected) ◦ Lines and systems connected among between channel extenders (if channel extenders are connected)
Port Rejected Insufficient resources	All resources of the local storage system or remote storage system are being used for other connections. Therefore, the local storage system or remote storage system rejected the connection control function that sets remote paths.	• In the Remove Remote Paths window, remove all remote paths that are not currently used. • In the Remove Remote Connections window, remove all remote storage systems that are not currently used. • VSP G1000: Confirm that the port attribute of the local storage system is Initiator, and that the port attribute of the remote storage system is set to RCU Target is set. If these settings are incorrect, change them to the correct port attributes.
Serial Number Mismatch Mismatched serial number	The serial number of the remote storage system does not match the specified serial number.	Check the following, and correct them if they are not correct: • The serial number (S/N) and system ID of the remote storage system, the port number of the local storage system, and the port number of the remote storage system are correct. • The topology (Fabric, FC-AL, Point-to-point) of the ports of the local storage system and remote storage system is correctly set. • The following network relaying devices are correctly configured, and can be properly used: ◦ Connectors ◦ Cables ◦ Switches (zoning settings)

Remote path status and description	Status description	Recommended action
		◦ Channel extenders (if channel extenders are connected) ◦ Lines and systems connected among between channel extenders (if channel extenders are connected)
Invalid Port Invalid port	The specified port of the local storage system is in the following status: • The port is not mounted. • VSP G1000: The port attribute is not Initiator. • The remote path already exists.	Check the following, and correct them if they are not correct: • VSP G1000: The port of the local storage system is mounted, or the Initiator is set to the attribute. • No remote path with the same configuration (the same port number of the local storage system and the same port number of the remote storage system) exists. • The topology (Fabric, FC-AL, Point-to-point) of the ports of the local storage system and remote storage system is correctly set. • The following network relaying devices are correctly configured, and can be properly used: ◦ Connectors ◦ Cables ◦ Switches (zoning settings) ◦ Channel extenders (if channel extenders are connected) ◦ Lines and systems connected among between channel extenders (if channel extenders are connected) • The serial number (S/N) and system ID of the remote storage system, the port number of the local storage system, and the port number of the remote storage system are correct.
Pair-Port Number Mismatch	The specified port of the remote storage system is not physically	Check the following, and correct them if they are not correct:

Remote path status and description	Status description		Recommended action
Incorrect port number of the remote storage system	connected to the local storage system.		• The port number of the remote storage system is correct. • The cable between the ports of the local storage system and the remote storage system or between the ports of the local storage system and the switch of the local storage system is properly connected. • The topology (Fabric, FC-AL, Point-to-point) of the ports of the local storage system and remote storage system is correctly set.
Pair-Port Type Mismatch Incorrect port type of the remote storage system (VSP G1000 only)	The attribute of the specified port of the remote storage system is not set to RCU Target.		Set RCU Target to the attribute of the port of the remote storage system that was specified to RCU Target.

Communication Failed
Communication error
	The local storage system is correctly connected to the remote storage system, but a logical communication timeout occurred.		Check the following, and correct them if they are not correct:
			• The port of the remote storage system and the network relaying devices are correctly set.
			• The following network relaying devices are correctly configured, and can be properly used:
			  ◦ Connectors
			  ◦ Cables
			  ◦ Switches (zoning settings)
			  ◦ Channel extenders (if channel extenders are connected)
			  ◦ Lines and systems connected among between channel extenders (if channel extenders are connected)

Path Blockade Logical blockade	Blocked because path errors or link errors repeatedly occurred.	The port of the local storage system is out of order.	Repair the port of the local storage system. Then, recover the remote path.*
		The port of the remote storage	Repair the port of the remote storage system. Then, recover the remote path.*

Remote path status and description	Status description		Recommended action
		system is out of order.	
		A relaying device is out of order.	Repair the relaying device. Then, recover the remote path.*
		The cable is damaged.	Replace the cable. Then, recover the remote path.*
Program Error Program error	A program error was detected.		Recover the remote path.*
In Progress In progress	A remote path is being created, the remote path is being deleted, or the port attribute (VSP G1000 only) is being changed.		Wait until the processing ends.

* Recover the remote path by either of the following methods:

- To use Hitachi Command Suite (either of the following):
  - Remove the remote connection in the Remove Remote Connections window, and then register the remote connection again in the Add Remote Connection window. For instructions, see the topic "Removing remote connections" in the help.
  - Remove the remote path in the Remove Remote Paths window, and then create a remote path again in the Add Remote Paths window. For instructions, see the topic "Removing remote paths" in the help.
- To use CCI:

  Use the raidcom delete rcu_path command to remove the remote path, and then use the raidcom add rcu_path command to re-create the remote path.

If the remote path is still not recovered after these operations, contact Hitachi Data Systems customer support.

Error codes and messages

If an error occurs during a global-active device operation, GAD displays an error message that describes the error and includes an error code.

Make sure to record the error codes, so that you can report them if you need to contact Hitachi Data Systems customer support. For details about Device Manager - Storage Navigator error codes, see the *Messages* guide for the storage system. For details about HCS error codes, see *Command Suite Messages*.

Troubleshooting for CCI

If an error occurs when you are using CCI to perform a GAD operation, you might be able to determine the cause of the error by viewing the logs that are output in the CCI window or the CCI operation logs.

The following is an output example:

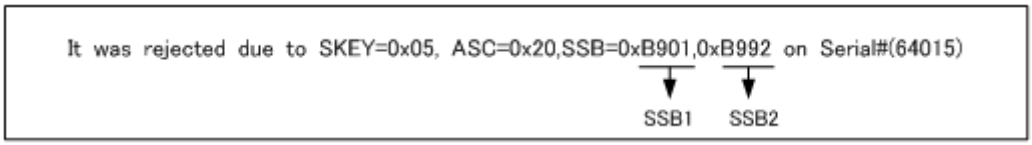

The alphanumerics after "SSB=" indicate the error code. The last four digits of the alphanumerics before the comma (,) is SSB1 (example: B901), and the last four digits of the alphanumerics after the comma is SSB2 (example: B992).

For details about the CCI error codes, see the *Command Control Interface User and Reference Guide*.

If the problem persists, send the contents of the /HORCM/log* folder to Hitachi Data Systems customer support.

Table 12-1 Error codes and details when operating CCI (when SSB1 is 2E31, B901, B90A, B90B, B912 or D004)

Error code (SSB2)	Details
9100	The command cannot be executed because user authentication has not been performed.
B952	The specified LU is not defined. A configuration of the storage system might have been changed. Restart CCI.
B9A2	You cannot create the GAD pair, because the specified volume is a command device.
B9A3	You cannot create the GAD pair, because the specified volume is a mainframe volume.
B9A4	You cannot create the GAD pair because no SCSI path is defined on the specified volume.
B9A5	You cannot create the GAD pair or perform a swap resync for the pair, because one of the following remote paths is not set: • A bidirectional remote path between the storage systems at the primary site and secondary site • A remote path from the storage system at the primary site to the storage system at the secondary site • A remote path from the storage system at the secondary site to the storage system at the primary site
B9BD	A configuration of the LDEV in the storage system might have been changed while CCI was running. Restart CCI.
B9C0	There are no free resources in the command device. Use LUN Manager to turn off and then turn on the command device.
DB89	You cannot change the status of the GAD pair, even though a request to suspend or to delete the pair has been received. This is because the volume paired with the specified volume is in an unusable status.

Error code (SSB2)	Details
DB8A	You cannot change the status of the GAD pair, even though a request to suspend or to delete the pair has been received. This is because the volume paired with the specified volume is blocked.
DB8B	You cannot change the status of the GAD pair, even though a request to suspend or to delete the pair has been received. This is because the volume paired with the specified volume is in an unusable status.
DB8D	You cannot change the status of the GAD pair, even though a request to suspend or to delete the pair has been received. This is because the number of remote paths from the storage systems at the primary site to the storage systems at the secondary site is less than the minimum number of remote paths.
FA00	You cannot create the GAD pair, because the capacity of the volume that has been specified as the primary volume is larger than the maximum capacity of a GAD pair that can be created.
FA01	You cannot create the GAD pair, because the volume that has been specified as the primary volume is being used by nondisruptive migration.
FA02	You cannot create the GAD pair, because the storage system cache at the primary site is in one of the following statuses: • One side is blocked or is transitioning to being blocked. • One side is recovering. • Recovering
FA03	You cannot create the GAD pair, because the remote paths from the storage systems at the primary site to the storage systems at the secondary site are in either of the following status: • The number of remote paths is 0 (unspecified). • The number of remote paths is less than the minimum number.
FA04	You cannot create the GAD pair, because the emulation type of the volume that has been specified as the primary volume is not OPEN-V.
FA05	You cannot create the GAD pair, because the volume that has been specified as the primary volume is a migration volume of a product of another company.
FA07	The pair status of the volume that has been specified as the primary volume is not SMPL.
FA08	The pair status of the volume that has been specified as the primary volume is not PSUS or PSUE.
FA09	There is a pinned track on the volume that has been specified as the primary volume.
FA0A	You cannot create the GAD pair, because the volume that has been specified as the primary volume is blocked.
FA0B	You cannot create the GAD pair, because the volume that has been specified as the primary volume is in one of the following statuses: • Blocked • Being formatted • Read only

Error code (SSB2)	Details
FA0C	You cannot create the GAD pair, because the volume that has been specified as the primary volume is a mainframe volume.
FA0D	You cannot create the GAD pair because the virtual emulation type of the P-VOL is not supported.
FA0E	You cannot create the GAD pair, because the volume that has been specified as the primary volume is not a virtual volume of Dynamic Provisioning, Dynamic Tiering, or active flash.
FA0F	The device type of the volume that has been specified as the primary volume is not supported.
FA10	You cannot create the GAD pair because the secondary volume is in an unusable status.
FA12	You cannot create the pair, because the GAD reserve attribute has been set for the volume that was specified as the primary volume.
FA13	The specified volume is being used by TrueCopy.
FA14	The specified volume is being used by Universal Replicator.
FA15	You cannot create the pair due to one of the following reasons: • The volume that has been specified as the primary volume of GAD is a primary volume of Thin Image which is being restored. • The volume that has been specified as the secondary volume of GAD is a primary volume of Thin Image. • The specified volume is a secondary volume of Thin Image.
FA16	The specified volume is a secondary volume of ShadowImage.
FA17	The specified volume is being used by Volume Migration.
FA18	The specified volume is a volume of a ShadowImage pair that is in the process of reverse copying.
FA1B	You cannot create the GAD pair, because the information about the virtual storage machine at the primary site disagrees with the one at the secondary site.
FA1C	The access attribute set by Data Retention Utility for the primary volume cannot be transferred to the secondary volume because Data Retention Utility is not installed in the storage system at the secondary site.
FA1D	You cannot create the GAD pair, by using the specified secondary volume because of either one of the following two reasons: • The specified secondary volume is already used for the other GAD pair. • The information about the GAD pair still remains only in the secondary volume.
FA1E	You cannot create the GAD pair, because the primary volume is a command device.
FA1F	You cannot create the GAD pair, because the secondary volume is a command device.
FA20	Internal process is taking time for the specified P-VOL. Retry the operation after a while.

Error code (SSB2)	Details
FA21	You cannot create or resynchronize the GAD pair, because the differential bitmap area cannot be secured due to either of the following reasons: • The primary storage system does not have enough available shared memory. • The DP pool of the volume specified as the P-VOL does not have enough available capacity.
FA23	Internal process is taking time for the specified S-VOL. Retry the operation after a while.
FA24	You cannot create the GAD pair, because the specified secondary storage system's microcode version does not support GAD consistency groups.
FA25	The operation on the GAD pair failed. The T10 PI attribute setting for the primary volume is not the same as for the secondary volume.
FA27	You cannot create the GAD pair, because the attribute of the volume specified as a P-VOL is not supported.
FA28	You cannot resynchronize or swap-resync the GAD pair in a GAD+UR configuration, because the following conditions are not satisfied (both must be satisfied): • The volume must be used as both a GAD P-VOL and a UR P-VOL. • The volume must be used as both a GAD S-VOL and a UR delta resync P-VOL.
FA29	You cannot create the GAD pair, because the volume that has been specified as the secondary volume is not installed or is a command device.
FA2A	You cannot create the GAD pair, because the volume that has been specified as the secondary volume is in the intervention-required condition.
FA2B	You cannot create the GAD pair, because the volume that has been specified as the secondary volume is blocked.
FA2C	The secondary volume is in an unusable status.
FA30	The pair status of the volume that has been specified as the secondary volume is not PSUS or PSUE.
FA31	The pair status of the volume that has been specified as the secondary volume is not SMPL.
FA32	There is a pinned track on the volume that has been specified as the secondary volume.
FA33	You cannot create the GAD pair, because the volume that has been specified as the secondary volume is in one of the following statuses: • Blocked • Being formatted • Read only
FA34	The operation cannot be performed because the storage system is executing internal processing or the configuration of the storage system is being changed by another operation. Wait a while, and then try again.
FA35	You cannot create the GAD pair, because the volume that has been specified as the secondary volume is blocked.

Error code (SSB2)	Details
FA36	You cannot create or resynchronize the GAD pair, because the remote storage system does not support pair creation using volumes that are not DP-VOLs.
FA37	You cannot create the GAD pair, because the volume that has been specified as the secondary volume is a migration volume of a product of another company.
FA38	You cannot create the GAD pair, because the volume that has been specified as the secondary volume is not OPEN-V.
FA3A	You cannot create the GAD pair, because the capacity of the volume that has been specified as the secondary volume is larger than the maximum capacity of a GAD pair that can be created.
FA3B	You cannot create the GAD pair, because the volume that has been specified as the secondary volume is being used by nondisruptive migration.
FA3C	The device type of the volume that has been specified as the secondary volume is not supported.
FA3D	You cannot create the GAD pair, because the attribute of the volume specified as an S-VOL is not supported.
FA3E	The controller emulation types of the storage systems at the primary site and the secondary site are inconsistent.
FA3F	The global-active device feature is not installed in the storage systems at the secondary site.
FA40	The shared memory that is required to create a GAD pair is not installed on the storage system at the secondary site.
FA41	The volume that has been specified as the secondary volume is not installed.
FA42	You cannot create the GAD pair, because the storage system cache at the secondary site is in one of the following statuses: • One side is blocked or is transitioning to being blocked • One side is recovering. • Recovering
FA43	You cannot create the GAD pair, because the remote path from the storage system at the secondary site to the storage system at the primary site is in either of the following status: • The number of remote paths is 0 (unspecified). • The number of remote paths is less than the minimum number.
FA44	You cannot create the GAD pair, because the volume that has been specified as the secondary volume is a mainframe volume.
FA46	You cannot create the GAD pair, because the volume that has been specified as the secondary volume is not a virtual volume of Dynamic Provisioning, Dynamic Tiering, or active flash.
FA48	You cannot create the GAD pair, because the P-VOL attribute differs from the S-VOL attribute.

Error code (SSB2)	Details
FA49	You cannot create the pair, because the serial numbers of the storage systems for the primary volume and the secondary volume are the same.
FA4B	You cannot create the pair, because the GAD reserve attribute is not set for the secondary volume.
FA4C	You cannot create the pair, because no virtual LDEV ID is set for the secondary volume.
FA4D	No LU path to the specified secondary volume is defined.
FA4E	You cannot create the GAD pair, because the capacities of the primary volume and the secondary volume are different.
FA4F	No LU path to the specified secondary volume is defined.
FA50	One of the following is incorrect: • A primary volume parameter (port name, host group ID, LUN ID) • A secondary volume parameter (port name, host group ID, LUN ID)
FA5B	You cannot create the GAD pair, because the remote paths from the storage system at the primary site to the storage system at the secondary site are in one of the following states: • The number of remote paths is 0 (unspecified). • The requirement for the minimum number of paths is not met.
FA60	You cannot create a pair, because the remote storage system product or its microcode version does not support the global-active device function.
FA62	You cannot create the GAD pair, because no virtual LDEV ID is set for the volume specified as the primary volume.
FA63	Failed to create the GAD pair. An external volume for which the Data Direct Mapping attribute is enabled is specified as the primary volume.
FA64	Failed to create the GAD pair. An external volume for which the Data Direct Mapping attribute is enabled is specified as the secondary volume.
FA65	Failed to create the GAD pair. The Data Direct Mapping attribute of the volume that was specified as the secondary volume is enabled, but the M-DKC firmware version does not support the functionality for mapping external volumes with a capacity exceeding 4TB.
FA66	Failed to create the GAD pair. The Data Direct Mapping attribute of the volume that was specified as the secondary volume is enabled, but the Data Direct Mapping attribute of the volume that was specified as the primary volume is disabled.
FA67	Failed to create the GAD pair. The Data Direct Mapping attribute of the volume that was specified as the secondary volume is enabled, but the R-DKC firmware version does not support the functionality for mapping external volumes with a capacity exceeding 4TB.
FB30	The command cannot be executed because the shared memory cannot be accessed.
FB41	You cannot create the GAD pair, because the operation by the consistency group is performed in the consistency group to which the specified volume is registered.

Error code (SSB2)	Details
FB42	You cannot create the GAD pair, because the operation from the remote storage system or the processing of the suspension by failure is performed in the consistency group to which the specified volume is registered.
FB43	You cannot create the GAD pair in the consistency group to which the specified volume is registered, because either of the following volumes exists on the storage system that has the secondary volume: • The primary volume whose pair status is COPY • The volume whose I/O mode is Local
FB44	You cannot create the GAD pair in the consistency group to which the specified volume is registered, because either of the following volumes exists on the storage system that has the primary volume. • The secondary volume whose pair status is COPY • The volume whose I/O mode is Block
FB45	You cannot create the GAD pair, because the suspension by failure occurs in the consistency group to which the specified volume is registered.
FB46	You cannot create the GAD pair, because you specify 256 or more as the consistency group ID.
FB47	You cannot create the GAD pair, because the number of GAD pairs registered in one consistency group or mirror has reached the maximum (8,192).
FB48	You cannot create the GAD pair, because you specify a different quorum disk ID from the ID used by the GAD pair registered in the consistency group.
FB49	You cannot create the GAD pair, because you specify a different mirror ID from the ID used by the GAD pair registered in the consistency group.
FB50	You cannot resynchronize or swap resynchronize the GAD pair, because the operation by the consistency group is performed in the consistency group to which the specified volume is registered.
FB51	You cannot resynchronize or swap resynchronize the GAD pair, because the operation from the remote storage system or the processing of the suspension by failure is performed in the consistency group to which the specified volume is registered.
FB54	You cannot resynchronize or swap resynchronize the GAD pair, because the suspension by failure occurs in the specified consistency group.
FB60	You cannot suspend the GAD pair, because the processing of the suspension by the consistency group is performed in the consistency group to which the specified volume is registered.
FB61	You cannot suspend the GAD pair, because either of the following volumes exists in the consistency group in the storage system that receives a request to suspend a pair. • The secondary volume whose pair status is COPY • The volume whose I/O mode is Block
FB62	You cannot suspend the GAD pair, because either of the following volumes exists in the consistency group in the paired storage system of the storage system that receives a request to suspend a pair.

Error code (SSB2)	Details
	• The primary volume whose pair status is COPY • The volume whose I/O mode is Local
FB63	You cannot suspend the GAD pair, because the suspension by failure occurs in the specified consistency group.
FB64	You cannot suspend the GAD pair, because the processing of the resynchronization by the consistency group is performed in the consistency group to which the specified volume is registered.
FB65	You cannot suspend the GAD pair by the consistency group, because the operation by the volume is performed in the specified consistency group.
FB6E	You cannot swap suspend the GAD pairs registered to a consistency group by volume.
FB70	You cannot delete the GAD pair, because the operation by the consistency group is performed in the consistency group to which the specified volume is registered.
FB71	You cannot delete the GAD pair, because the suspension by failure occurs in the consistency group to which the specified volume is registered.
FB80	The paircreate command or the pairresync command was rejected for one of the following reasons: • You specified the **-fg** option. When you specify the **-fg** option, confirm that the microcode version of the primary storage system supports GAD consistency groups. • The specified quorum disk ID is out of range. If none of the above applies, contact Hitachi Data Systems customer support.
FB92	You cannot create or resynchronize the GAD pair, because the differential bitmap area cannot be secured due to either of the following reasons: • The secondary storage system does not have enough available shared memory. • The DP pool of the volume specified as the S-VOL does not have enough available capacity.
FBB0	A request to delete the GAD pair was received, but the pair cannot be deleted because the volume paired with the specified volume is shared with ShadowImage.
FBB1	A request to delete the GAD pair was received, but the pair cannot be deleted because the volume paired with the specified volume is shared with Thin Image.
FBE0	The command operating on the GAD pair was rejected because the -f data or -f status option was specified for the paircreate or pairresync command.
FBE1	The command operating on the GAD pair was rejected because the -SM block or -SM unblock option was specified for the paircreate or pairresync command.
FBE8	The command operating on the GAD pair was rejected because the -P option was specified for the pairsplit command.

Error code (SSB2)	Details
FBE9	The command operating on the GAD pair was rejected because the -rw option was specified for the pairsplit command.
FBEA	The command operating on the GAD pair was rejected because the -RB option was specified for the pairsplit command.
FBEB	The command operating on the GAD pair was rejected because the -SM block or -SM unblock option was specified for the pairsplit command.
FC10	The command was rejected because the storage system of the specified volume is in one of the following states: • The storage system includes microcode that does not support GAD. • The GAD feature is not installed. • Shared memory for GAD has not been added. If none of the above applies, contact Hitachi Data Systems customer support.
FC12	The same operation or a different operation is already being executed.
FC13	An operation to suspend the pair is being processed because an error was detected.
FC14	You cannot execute the command, because you specify a different consistency group ID from the ID of the consistency group to which the specified volume is registered.
FC20	You cannot create a pair for one of the following reasons: • The primary storage system does not have enough available shared memory. • The DP pool of the volume specified as the P-VOL does not have enough available capacity. • No shared memory is installed in the primary storage system. • No Resource Partition Manager license is installed in the storage system of the specified volume. • The GAD license capacity in the storage system of the specified volume is insufficient.
FC21	You cannot create a pair for one of the following reasons: • The secondary storage system does not have enough available shared memory. • The DP pool of the volume specified as the S-VOL does not have enough available capacity. • No shared memory is installed in the secondary storage system. • No Resource Partition Manager license is installed in the storage system of the volume to be paired with the specified volume. • The GAD license capacity in the storage system of the volume to be paired with the specified volume is insufficient.
FC22	You cannot create a pair for one of the following reasons: • The pair status of the specified volume is not SMPL. • The specified volume is a single volume or is not the primary volume in the GAD pair.

Error code (SSB2)	Details
FC23	You cannot create a pair for one of the following reasons: • The pair status of the volume to be paired with the specified volume is not SMPL or COPY. • The volume to be paired with the specified volume is a single volume or is not the secondary volume in the GAD pair.
FC24	You cannot create a pair at the primary site for either one of the following two reasons: • The capacity of the specified volume is being expanded. • The pool containing the specified volume is being initialized. • The volume does not have a virtual LDEV ID.
FC25	You cannot create a pair at the secondary site for one of the following reasons: • The capacity of the volume to be paired with the specified volume is being expanded. • The pool of the volume paired with the specified volume is being initialized. • The virtual LDEV ID of the volume to be paired with the specified volume is duplicated in the virtual storage machine. • You specified the virtual LDEV ID at the primary site the same as the actual LDEV ID at the secondary site from the volume to be paired with the specified volume. However, the actual information of the virtual emulation type (including the settings for CVS and LUSE) or the virtual SSID is different from the virtual information. • The virtual LDEV ID of the volume to be paired with the specified volume is already in use. • The number of LUN paths defined in the volume to be paired with the specified volume exceeds the number of LUN paths that can be used by ALUA.
FC26	You cannot create a pair, because verification of the remote path between storage systems failed in the storage system of the specified volume.
FC27	You cannot create a pair, because verification of the remote path between storage systems failed in the storage system of the volume to be paired with the specified volume.
FC28	You cannot create a pair for one of the following reasons: • The mirror count for a single volume is depleted for the specified volume. • The specified volume is already being used by another GAD pair. • The pair management table of the specified volume is depleted.
FC29	You cannot create a pair for one of the following reasons: • The mirror count for a single volume is depleted for the volume to be paired with the specified volume. • The volume to be paired with the specified volume is already being used by another GAD pair. • The pair management table of the volume to be paired with the specified volume is depleted.

Error code (SSB2)	Details
FC2A	The operation cannot be performed, either because the storage system is executing internal processing, or because the configuration of the storage system is being changed by another operation. Wait a while, and then try again.
FC2B	The operation cannot be performed because the storage system is enabling or disabling ALUA mode. Wait a while, and then try again.
FC30	The pair resynchronization or swap resync was rejected for one of the following reasons: • The pair status of the volume specified for the pair resynchronization is not PSUS or PSUE. • The volume specified for the pair resynchronization is not the primary volume of the GAD pair. • The I/O mode of the volume specified for the pair resynchronization is Block. • The pair status of the volume paired with the volume specified for the swap resync is not PSUS or PSUE. • The I/O mode of the volume paired with the volume specified for the swap resync is not Block.
FC31	The pair resynchronization or swap resync was rejected for one of the following reasons: • The pair status of the volume specified for the swap resync is not SSWS. • The volume specified for the swap resync is not the secondary volume of the GAD pair. • The I/O mode of the volume specified for the swap resync is Block. • The pair status of the volume paired with the volume specified for the pair resynchronization is not SSUS or PSUE. • The I/O mode of the volume paired with the volume specified for the pair resynchronization is not Block.
FC32	You cannot resynchronize or swap resynchronize the GAD pair for one of the following reasons: • You execute the command with specifying the consistency group ID to the volume that is not registered to the consistency group. • You execute the command without specifying the consistency group ID to the volume that is registered to the consistency group. • You execute the swap resynchronization by the volume to the volume that is registered in the consistency group.
FC38	A request to suspend a pair was received, but the pair cannot be suspended because the specified volume meets one of the following conditions: • An instruction specifying that the primary volume be suspended is directed at the secondary volume. • An instruction specifying that the secondary volume be suspended is directed at the primary volume. • The pair status is not PAIR or COPY.

Error code (SSB2)	Details
FC39	A request to suspend a pair was received, but the pair cannot be suspended because the volume paired with the specified volume meets one of the following conditions: • The paired volume is the primary volume, but an instruction specifies that the primary volume be suspended. • The paired volume is the secondary volume, but an instruction specifies that the secondary volume be suspended. • The pair status is not PAIR or COPY.
FC3A	You cannot suspend the GAD pair for either of the following reasons: • You execute the command with specifying the consistency group ID to the volume that is not registered in the consistency group. • You execute the command without specifying the consistency group ID to the volume that is registered in the consistency group.
FC40	A request to delete a pair was received, but the pair cannot be deleted because the specified volume meets one of the following conditions: • An instruction specifying that the primary volume be deleted is directed at the secondary volume. • An instruction specifying that the secondary volume be deleted is directed at the primary volume. • The pair status is not PSUS, SSUS, SSWS, or PSUE. • The I/O mode is not Local. • The volume is shared with a Universal Replicator pair.
FC41	A request to delete a pair was received, but the pair cannot be deleted because the volume paired with the specified volume meets one of the following conditions: • The paired volume is the primary volume, but an instruction specifies that the primary volume be deleted. • The paired volume is the secondary volume, but an instruction specifies that the secondary volume be deleted. • The pair status is not PSUS, SSUS, SSWS, or PSUE. • The I/O mode is not Block. • The volume is shared with a Universal Replicator pair.
FC7E	A request to create a pair, resynchronize a pair, or perform a swap resync was received, but the request was rejected because the status of the quorum disk meets one of the following conditions: • The ID of the specified quorum disk is out of range. • The quorum disk has not been created. • The specified remote storage system is not the same as when the quorum disk was created. • The same quorum disk ID is allocated to separate external volumes in the storage systems at the primary and secondary sites. • The quorum disk is blocked. • An error occurred on the external path between the storage systems at the primary and secondary sites and the external storage system for the quorum disk.

Error code (SSB2)	Details
	• The quorum disk was used to cancel the pair.

SIM reports of GAD operations

If a storage system requires maintenance, a SIM is issued and displayed in Device Manager - Storage Navigator. A SIM is also issued when the pair status of a primary or secondary GAD volume changes. SIMs are categorized into service, moderate, serious, and acute according to their severity. The GAD operation history appears in the History window.

If SNMP is installed on the storage systems, SIMs trigger an SNMP trap that is sent to the corresponding server. For details about SNMP operations, see the *Hitachi Command Suite User Guide* or the *Hitachi SNMP Agent User Guide*.

Related topics

- SIMs related to GAD on page 8-6

Procedure for recovering pinned track of a GAD volume

To recover the pinned track and secure the entire data integrity of the pair at the same time, follow this procedure:

1. Connect to the storage system of the primary site for a GAD pair that contains the pinned track volume, and then select a correct CU.
2. Remove the GAD pair that contains the pinned track volume.
3. Perform a normal procedure for recovering data from the pinned track. Use the pinned track recovery procedure of the OS that is being used, or contact Hitachi Data Systems customer support.
4. Use the Create GAD Pairs window to create the pair. Make sure to select Entire for Initial Copy Type.

Contacting Hitachi Data Systems customer support

If you need to contact Hitachi Data Systems customer support, please provide as much information about the problem as possible, including:

- The circumstances surrounding the error or failure.
- The content of any error messages displayed on the host systems.
- The content of any error messages displayed by Device Manager - Storage Navigator.
- The Device Manager - Storage Navigator configuration information (use the Dump Tool).

- The service information messages (SIMs), including reference codes and severity levels, displayed by Device Manager - Storage Navigator.

The Hitachi Data Systems customer support staff is available 24 hours a day, seven days a week. If you need technical support, log on to Hitachi Data Systems Support Connect for contact information: https://support.hds.com/en_us/contact-us.html

A

Correspondence between GUI operations and CLI commands

This appendix describes the correspondence between the GUI operations and the CCI commands. You can perform global-active device operations using either the GUI or the CLI, whichever you prefer, but a few operations can be performed using only the GUI (for example, forcibly deleting a pair) or only the CLI (for example, disaster recovery procedures).

Correspondence between GUI operations and CCI commands

The following tables show the correspondence between CCI commands and Device Manager - Storage Navigator operations.

Table A-1 Correspondence between GUI operations and CCI commands: configuration operations

Operation	GUI	CCI	
	Supported?	Supported?	Command
Edit ports[1]	Yes	Yes	raidcom modify port
Add remote connections	Yes	Yes	raidcom add rcu
Select external path groups	Yes	Yes	raidcom add external_grp
Create external volumes	Yes	Yes	raidcom add ldev
Add quorum disks	Yes	Yes	raidcom modify ldev
Create virtual storage machines (resource groups)	No[2]	Yes	raidcom add resource
Reserve host group IDs	No	Yes	raidcom add resource
Delete virtual LDEV IDs (release reserve attribute)	Yes	Yes	raidcom unmap resource
Reserve LDEV IDs	Yes	Yes	raidcom add resource
Assign GAD reserve attribute	Yes	Yes	raidcom map resource
Create host groups	Yes	Yes	raidcom add host_grp
Create pools	Yes	Yes	raidcom add dp_pool
Create virtual volumes	Yes	Yes	raidcom add ldev
Add LU paths	Yes	Yes	raidcom add lun

Notes:

1. Applies only to VSP G1000. You do not need to set the ports on the VSP G200, G400, G600, G800.
2. You can create VSMs using the Hitachi Command Suite GUI.

Table A-2 Correspondence between GUI operations and CCI commands: pair operations

Operation	Parameter	GUI	CCI		
		Supported ?	Supported ?	Command	Option
Create GAD pairs	Fence Level	No[1]	Yes	paircreate	-f never
	Copy Pace	Yes	Yes	paircreate	-c <size>
	No initial copy	Yes	Yes	paircreate	-nocopy
	Registration to a consistency group[2, 7]	Yes	Yes	paircreate	-fg never <CTG ID>
	ALUA mode change[7]	Yes	Yes	raidcom modify ldev	-alua{enable | disable}
Suspend pairs	P-VOL specification	Yes	Yes	pairsplit	-r
	S-VOL specification	Yes[3]	Yes	pairsplit	-RS
Suspend consistency groups[2, 7]		Yes	Yes	pairsplit	None[4]
Resync pairs	P-VOL specification	Yes	Yes	pairresync	
	S-VOL specification	Yes	Yes	pairresync	-swaps
	Copy Pace	Yes	Yes	pairresync	-c <size>
	Registration to a consistency group[2, 7]	Yes	Yes	pairresync	-fg never <CTG ID>
	ALUA mode change	Yes	Yes	raidcom modify ldev	-alua{enable | disable}
Resync consistency groups[2, 5, 7]	None	Yes	Yes	pairresync	None[4]
	ALUA mode change	Yes	Yes	raidcom modify ldev	-alua{enable | disable}
Delete pairs	Normal (P-VOL specification)	Yes	Yes	pairsplit	-S
	Normal (S-VOL specification)	Yes	Yes	pairsplit	-R
	Force (**Enable** is specified for **Volume Access**)[6]	Yes	No	None	

Operation	Parameter	GUI	CCI		
		Supported ?	Supported ?	Command	Option
	Force (**Disable** is specified for **Volume Access**)[6]	Yes	No	None	

Notes:

1. The fence level is set to Never automatically.
2. Consistency group operations can be performed only on the VSP G1000.
3. If you want to suspend a pair registered to a consistency group from the storage system that contains the S-VOL, you must use the Suspend Consistency Groups window. You cannot suspend a pair registered to a consistency group from the storage system that contains the S-VOL by selecting the S-VOL on the GAD Pairs tab in the Remote Replication window and opening the Suspend Pairs window.
4. When you suspend or resynchronize pairs registered to a consistency group without specifying any option, the pairs are suspended or resynchronized by the consistency group.
5. When you resynchronize pairs not registered to a consistency group, you can register the pairs to a consistency group.
6. The operation to delete pairs forcibly can only be performed using the Delete Pairs window of Device Manager - Storage Navigator.
7. VSP G1000 only.

Table A-3 Correspondence between GUI operations and CCI commands: displaying status

Operation	Parameter	GUI	CCI		
		Supported ?	Supported ?	Command	Option
View pair properties	I/O mode	Yes	Yes	pairdisplay	-fxce or -fxde
	Status	Yes	Yes	pairdisplay	-fxc or -fxce
View consistency group properties*	Consistency group ID	Yes	Yes	pairdisplay	-fxc or -fxce
	Consistency group status	Yes	No	None	
View pair synchronization rate		Yes	Yes	pairdisplay	-fxc
View remote connection properties		Yes	Yes	raidcom get rcu	
View virtual storage machines		Yes	Yes	raidcom get resource	-key opt
View quorum disks		Yes	Yes	raidcom get ldev	

Operation	Parameter	GUI	CCI		
		Supported ?	Supported ?	Command	Option
Check the status of volumes	Existence of a virtual LDEV ID	Yes	Yes	raidcom get ldev	
	Existence of the GAD reserve attribute	Yes	Yes	raidcom get ldev	
View the ALUA mode*		Yes	Yes	raidcom get ldev	
View the asymmetric access status*		Yes	Yes	raidcom get host_grp	
* Consistency group operations can be performed only on the VSP G1000.					

Table A-4 Correspondence between GUI operations and CCI commands: changing settings

Operation	Parameter	GUI	CCI		
		Supported ?	Supported ?	Command	Option
Edit remote replica options		Yes	No	None	
Edit virtualization management settings	Virtual LDEV ID	Yes	Yes	raidcom map resource raidcom unmap resource	-virtual_ldev_id <ldev#>
	Virtual emulation type (including CVS and LUSE settings)	Yes	Yes	raidcom map resource	-emulation <emulation type>
	Virtual SSID	Yes	Yes	raidcom map resource	-ssid <ssid>
Remove quorum disks		Yes	Yes	raidcom modify ldev	
Release GAD reserve attribute		Yes	Yes	raidcom unmap resource	-virtual_ldev_id reserve
Forcibly delete pairs		Yes	No	None	
Edit remote connection options	RIO MIH time	Yes	Yes	raidcom modify rcu	-rcu_option <mpth> <rto> <rtt>
Add remote paths		Yes	Yes	raidcom add rcu_path	

Operation	Parameter	GUI	CCI		
		Supported ?	Supported ?	Command	Option
Remove remote paths		Yes	Yes	raidcom delete rcu_path	
Remove remote connections		Yes	Yes	raidcom delete rcu	

Glossary

This glossary defines the special terms used in this document. Click the letter links below to navigate.

#

3DC
three-data-center

A

administrative logical unit (ALU)
An LU used for the conglomerate LUN structure, a SCSI architecture model. In the conglomerate LUN structure, all host access is through the ALU, which functions as a gateway to sort the I/Os for the subsidiary logical units (SLUs) grouped under the ALU. The host requests I/Os by using SCSI commands to specify the ALU and the SLUs grouped under the ALU. An ALU is called a Protocol Endpoint (PE) in vSphere.
See also *subsidiary logical unit (SLU)*.

alternate path
A secondary path (port, target ID, LUN) to a logical volume, in addition to the primary path, that is used as a backup in case the primary path fails.

ALU
See *administrative logical unit*.

async
asynchronous

| # | A | B | C | D | E | F | G | H | I | J | K | L | M | N | O | P | Q | R | S | T | U | V | W | X | Y | Z |

B

B
block

BLK, blk
block

bmp
bitmap

C

C/T
See *consistency time (C/T)*.

ca
cache

CCI
Command Control Interface

CLI
command line interface

CLPR
cache logical partition

cluster
Multiple-storage servers working together to respond to multiple read and write requests.

CMD
command device

consistency group (CTG)
A group of pairs on which copy operations are performed simultaneously; the pairs' status changes at the same time. See also *extended consistency group (EXCTG)*.

consistency time (C/T)
Shows a time stamp to indicate how close the target volume is to the source volume. C/T also shows the time stamp of a journal and extended consistency group.

control memory
Another name for shared memory.

CTG
See *consistency group (CTG)*.

CU
control unit

currency of data
The synchronization of the volumes in a copy pair. When the data on the secondary volume (S-VOL) is identical to the data on the primary volume (P-VOL), the data on the S-VOL is current. When the data on the S-VOL is not identical to the data on the P-VOL, the data on the S-VOL is not current.

CVS
custom volume size

D

data consistency
When the data on the secondary volume is identical to the data on the primary volume.

data path
The physical paths used by primary storage systems to communicate with secondary storage systems in a remote replication environment.

differential data
Changed data in the primary volume not yet reflected in the copy.

DP-VOL
Dynamic Provisioning-virtual volume. A virtual volume used by Dynamic Provisioning.

dynamic provisioning
An approach to managing storage. Instead of "reserving" a fixed amount of storage, it removes capacity from the available pool when data is actually written to disk. Also called thin provisioning.

E

ERC
error reporting communications

EXG
external volume group

ext.
external

external port
A Fibre Channel port that is configured to be connected to an external storage system for Universal Volume Manager operations.

external volume
A logical volume whose data resides on drives that are physically located outside the Hitachi storage system.

G

GAD
global-active device

GUI
graphical user interface

H

HA
high availability

HCS
Hitachi Command Suite

HDLM
Hitachi Dynamic Link Manager

HDP
Hitachi Dynamic Provisioning

HDS
Hitachi Data Systems

HRpM
Hitachi Replication Manager

HTI
Hitachi Thin Image

HUR
Hitachi Universal Replicator

I

initial copy
An initial copy operation is performed when a copy pair is created. Data on the primary volume is copied to the secondary volume.

initiator port
A Fibre Channel port configured to send remote I/Os to an RCU target port on another storage system. See also *RCU target port* and *target port*.

in-system replication
The original data volume and its copy are located in the same storage system. ShadowImage in-system replication provides duplication of logical volumes; Thin Image provides "snapshots" of logical volumes that are stored and managed as virtual volumes (V-VOLs).

internal volume
A logical volume whose data resides on drives that are physically located within the storage system. See also *external volume*.

IO, I/O
input/output

IOPS
I/Os per second

J

JNL
journal

L

L
local

LDEV
logical device

local copy
See *in-system replication*.

LU
logical unit

LUN
logical unit number

LUSE
LUN Expansion

M

M-DKC
main disk controller. Refers to the primary storage system.

MIH
missing interrupt handler

MU
mirror unit

N

NVS
nonvolatile storage

P

PiT
point-in-time

pool
A set of volumes that are reserved for storing pool volumes (pool-VOLs) used by Thin Image, Dynamic Provisioning, Dynamic Provisioning for Mainframe, Dynamic Tiering, Dynamic Tiering for Mainframe, active flash, or active flash for mainframe.

pool volume (pool-VOL)
A logical volume that is reserved for storing snapshot data for Thin Image operations or write data for Dynamic Provisioning, Dynamic Provisioning for Mainframe, Dynamic Tiering, Dynamic Tiering for Mainframe, active flash, and active flash for mainframe.

port attribute
Indicates the type of Fibre Channel port: target, RCU target, or initiator.

primary site
The physical location of the storage system that contains the original data to be replicated and that is connected to one or more storage systems at the remote or secondary site via remote copy connections.

primary storage system
A storage system that contains primary volumes (P-VOLs) of copy pairs.

primary volume (P-VOL)
The volume in a copy pair that contains the original data to be replicated. The data on the primary volume is duplicated synchronously or asynchronously on the secondary volume.
See also *secondary volume (S-VOL)*.

P-VOL
See *primary volume*.

R

R
read

R/W, r/w
read/write

RCP
remote control port

RCU
remote control unit

RCU target port
A Fibre Channel port that is configured to receive remote I/Os from an initiator port on another storage system.

R-DKC
remote disk controller. Refers to the secondary storage system.

remote copy connections
The physical paths that connect a storage system at the primary site to a storage system at the secondary site. Also called data path.

remote site
See *secondary site*.

resync
resynchronize, resynchronization

RIO
remote I/O

R-JNL
restore journal

RPO
recovery point objective

RTO
recovery time objective

S

S/N, SN
serial number

secondary site
The physical location of the storage system that contains the secondary volumes of remote replication pairs at the secondary site. The secondary storage system is connected to the primary storage system via remote copy connections. See also *primary site*.

secondary storage system
A storage system that contains secondary volumes (S-VOLs) of copy pairs.

secondary volume
The volume in a copy pair that contains the copy of the data on the primary volume. See also *primary volume*.

service processor (SVP)
The computer inside a RAID storage system that hosts the Device Manager - Storage Navigator software and is used by service personnel for configuration and maintenance of the storage system.

shared volume
A volume that is being used by more than one replication function. For example, a volume that is the primary volume of a GAD pair and the primary volume of a UR pair is a shared volume.

SI
Hitachi ShadowImage

SIz
Hitachi ShadowImage for Mainframe

SIM
service information message

SLU
See *subsidiary logical unit (SLU)*.

SM
shared memory

SOM
system option mode

SSB
sense byte

SSID
(storage) subsystem identifier. SSIDs are used as an additional way to identify a control unit. Each group of 64 or 256 volumes requires one SSID, resulting in one or four SSIDs per CU image. For the VSP G1000, one SSID is associated with 256 volumes.

subsidiary logical unit (SLU)
An LU used for the conglomerate LUN structure, a SCSI architecture model. An SLU is an LU that stores actual data. You can use a DP-VOL or snapshot data (or a V-VOL allocated to snapshot data) as an SLU. All host access to SLUs is through the administrative logical unit (ALU). An SLU is called a virtual volume (VVol) in vSphere. See also *administrative logical unit (ALU)*.

S-VOL
See *secondary volume*.

SVOS
Storage Virtualization Operating System

SVP
See *service processor (SVP)*.

sync
synchronize, synchronous

system option mode (SOM)
Additional operational parameters for the RAID storage systems that enable the storage system to be tailored to unique customer operating requirements. SOMs are set on the service processor.

T

T10 PI
T10 Protection Information

target port
A Fibre Channel port that is configured to receive and process host I/Os.

TC
Hitachi TrueCopy

TCz
Hitachi TrueCopy for Mainframe

TID
target ID

TSV
tab-separated values

U

update copy
An operation that copies differential data on the primary volume of a copy pair to the secondary volume. Update copy operations are performed in response to write I/Os on the primary volume after the initial copy operation is completed.

UR
Hitachi Universal Replicator

URMF
Hitachi Universal Replicator for Mainframe

URz
Hitachi Universal Replicator for Mainframe

USP V/VM
Hitachi Universal Storage Platform V/VM

UVM
Universal Volume Manager

V

virtual volume (V-VOL)
A logical volume that has no physical storage space. Hitachi Thin Image uses V-VOLs as secondary volumes of copy pairs. In Dynamic Provisioning, Dynamic Provisioning for Mainframe, Dynamic Tiering, Dynamic Tiering for Mainframe, active flash, and active flash for mainframe, V-VOLs are referred to as DP-VOLs.

VM
virtual machine; volume migration

VOL, vol
volume

VSM
virtual storage machine

VSP G1000
Hitachi Virtual Storage Platform G1000

V-VOL
virtual volume

Index

P

Q

R

S

T

U

V

Hitachi Data Systems

Corporate Headquarters
2845 Lafayette Street
Santa Clara, California 95050-2639
U.S.A.
www.hds.com

Regional Contact Information

Americas
+1 408 970 1000
info@hds.com

Europe, Middle East, and Africa
+44 (0)1753 618000
info.emea@hds.com

Asia Pacific
+852 3189 7900
hds.marketing.apac@hds.com

Hitachi Data Systems

MK-92RD8072-08