Command Control Interface

01-68-03/01

Command Reference

This document describes and provides instructions for using the Command Control Interface (CCI) software to configure and perform operations on the Hitachi RAID storage systems.

HITACHI
Inspire the Next

Contents

Contents

Contents

Contents

Contents

Contents

Contents

Preface

This document describes and provides instructions for using Command Control Interface (CCI) software to configure and perform operations on Hitachi RAID storage systems.

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

Intended audience

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

Readers of this document should be familiar with the following:

- Data processing and RAID storage systems and their basic functions.

- The Hitachi RAID storage system and the *Hardware Guide* for the storage system.

- The management software for the storage system (for example, *Hitachi Device Manager - Storage Navigator*).

- The host systems attached to the Hitachi RAID storage system.

Product version

This document revision applies to CCI software version 01-68-03/01 or later.

Release notes

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

Changes in this revision

- Improved the explanations about the LDEV numbers that are displayed by using the `raidvchkscan` command and `raidcom get journal` commands.

- Added AIX support for NVMe-oF.

- Added notes on the maximum number of acceptable asynchronous commands.
- Corrected descriptions of port types.

Document conventions

This document uses the following typographic conventions:

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

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

Icon	Label	Description
📄	Note	Calls attention to additional information.
💡	Tip	Provides helpful information, guidelines, or suggestions for performing tasks more effectively.
❗	Important	Highlights information that is essential to the completion of a task.
⚠	Caution	Warns the user of adverse conditions and/or consequences (for example, disruptive operations, data loss, or a system crash).
⚠	CAUTION	Warns the user of a hazardous situation that, if not avoided, could result in major or minor injury.
⚠	WARNING	Warns the user of a hazardous situation which, if not avoided, could result in death or serious injury.

Conventions for storage capacity values

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

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

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

Logical capacity unit	Value
1 block	512 bytes

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

Getting help

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

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

Comments

Please send comments to doc.comments@hitachivantara.com. Include the document title and number, including the revision level (for example, -07), and refer to specific sections and paragraphs whenever possible. All comments become the property of Hitachi Vantara LLC.

Thank you!

Chapter 1: Differences between CCI and Device Manager - Storage Navigator

This information describes the differences between CCI and Storage Navigator/Device Manager - Storage Navigator.

Supported characters

Supported characters for CCI

There are some differences between the supported characters for CCI and the supported characters for Device Manager - Storage Navigator. Because of this, you should always use only characters that are supported by both CCI and Device Manager - Storage Navigator.

The following figure shows the characters that can be used in CCI commands and their ASCII codes. Note that "\" can be used only on Windows, whereas "/" can be used only on UNIX. For CCI, do not use characters that are not supported by the operating system in which the command is executed.

> ⚠️ **Caution:** When you execute a command to change the configuration, make sure you specify the parameters correctly. If you specify a parameter that is not contained in the syntax of the command, or you do not specify necessary parameters, the result might be different from the one you expected.

Low 4 bits	High 3 bits							
	0	1	2	3	4	5	6	7
0	NUL	DLE	SP	0	@	P	`	p
1	SOH	DC1	!	1	A	Q	a	q
2	STX	DC2	"	2	B	R	b	r
3	ETX	DC3	#	3	C	S	c	s
4	EOT	DC4	$	4	D	T	d	t
5	ENQ	NAC	%	5	E	U	e	u
6	ACK	SYN	&	6	F	V	f	v
7	BEL	ETB	'	7	G	W	g	w
8	BS	CAN	(	8	H	X	h	x
9	HT	EM	)	9	I	Y	i	y
A	LF/NL	SUB	*	:	J	Z	j	z
B	VT	ESC	+	;	K	[	k	{
C	FF	FS	,	<	L	\	l	|
D	CR	GS	-	=	M	]	m	}
E	SO	RS	.	>	N	^	n	~
F	SI	US	/	?	O	_	o	DEL

: usable

: unusable

⚠ **Caution:** Observe the following restrictions for using characters:

- A hyphen can be used in a name, but it cannot be used at the beginning of the name. You can neither use some character codes according to the OS on which you execute the command.

- A space (SP) can be used in a name, but if you specify a space for a user name, reference commands will fail, because a space is used as a delimiter and the commands cannot distinguish whether the space is for user name or delimiter.

- A space can only be used between other characters. If you use a space at the beginning or end of a name, the space is omitted automatically. Also, you cannot use a name consisting only of one or more space characters.

  Examples of using space characters (SP) in a name:

  - "group name" (different from "groupname")

  - "g r o u p n a m e" (different from "groupname" and "group name")

  Examples in which you cannot use a space:

  - " name" (same as "name")

  - "name " (same as "name")

  - " " (invalid)

- The characters that can be used in Storage Advisor Embedded are the same as the ones that can be used in CCI.

Supported characters for Device Manager - Storage Navigator

The following figure shows the characters that can be used in Device Manager - Storage Navigator commands and their ASCII codes.

📄 **Note:** Characters that can be used vary depending on names. For details, see the System Administrator Guide of your system.

Low 4 bits	High 3 bits							
	0	1	2	3	4	5	6	7
0	NUL	DLE	SP	0	@	P	`	p
1	SOH	DC1	!	1	A	Q	a	q
2	STX	DC2	"	2	B	R	b	r
3	ETX	DC3	#	3	C	S	c	s
4	EOT	DC4	$	4	D	T	d	t
5	ENQ	NAC	%	5	E	U	e	u
6	ACK	SYN	&	6	F	V	f	v
7	BEL	ETB	'	7	G	W	g	w
8	BS	CAN	(	8	H	X	h	x
9	HT	EM	)	9	I	Y	i	y
A	LF/NL	SUB	*	:	J	Z	j	z
B	VT	ESC	+	;	K	[	k	{
C	FF	FS	,	<	L	\	l	|
D	CR	GS	-	=	M	]	m	}
E	SO	RS	.	>	N	^	n	~
F	SI	US	/	?	O	_	o	DEL

☐ : usable

☐ : unusable

When using CCI and Device Manager - Storage Navigator together, and setting the name from CCI, see the notes in the following table for characters that can be used in each name.

Name	Parameter to be set	Characters that can be used in CCI	Note when using Device Manager - Storage Navigator together with CCI
Host group name (For FC)	-host_grp_name	See figure in Supported	Do not use the following characters:*

Name	Parameter to be set	Characters that can be used in CCI	Note when using Device Manager - Storage Navigator together with CCI
		characters for CCI (on page 14)	\ / : ,
LDEV nickname	-ldev_name	See figure in Supported characters for CCI (on page 14)	Do not use the following characters:* \ / : ,
WWN nickname	-wwn_nickname	See figure in Supported characters for CCI (on page 14)	Do not use the following characters:* \ / : ,
Pool name	-pool_name	See figure in Supported characters for CCI (on page 14)	Do not use the following characters:* \ / : ,
Resource group name	-resource_name	See figure in Supported characters for CCI (on page 14)	Do not use the following characters:* \ / : ,
SPM name	-spm_name	See figure in Supported characters for CCI (on page 14)	Do not use the following characters:* \ / : ,
iSCSI name	-iscsi_nickname	See figure in Supported characters for CCI (on page 14)	Do not use the following characters:* \ / : ,
CHAP user name	-target_chap_user	See figure in Supported characters for CCI (on page 14)	Do not use the following characters:* \
	-initiator_chap_user	See figure in Supported characters for CCI (on page 14)	Do not use the following characters:* \

Chapter 1: Differences between CCI and Device Manager - Storage Navigator

Name	Parameter to be set	Characters that can be used in CCI	Note when using Device Manager - Storage Navigator together with CCI
* If you use these characters, you might not be able to use some Device Manager - Storage Navigator functions such as configuration reports.			

Maximum number of characters

There are some differences between the maximum number of characters for names in CCI commands and the maximum number of characters for names in Device Manager - Storage Navigator. Because of this, you should always try to use the number of characters that can be used in both CCI and Device Manager - Storage Navigator.

If you enter more than the maximum number of characters for a name in a CCI command, the name is truncated and only the allowable number of characters is used. For example, if the maximum number of characters is 32 and you enter 35 characters, only the first 32 characters are used.

The following table specifies the maximum number of characters for names in CCI and Device Manager - Storage Navigator. To ensure that names are usable in both CCI and Device Manager - Storage Navigator, observe the following requirements:

- WWN nickname: Do not use more than 64 characters.

- User ID: Do not use more than 63 characters.

- Password: Do not use more than 63 characters.

Name	CCI command	Device Manager - Storage Navigator
Host group name	64 characters	64 characters
Device group name	32 characters	-
Device name	32 characters	-
LDEV nickname	32 characters	32 characters
WWN nickname	64 characters	64 characters
Copy group name	32 characters	32 characters
Pool name	32 characters	32 characters
Resource group name	32 characters	32 characters
User name (User ID)	63 characters	256 characters

Chapter 1: Differences between CCI and Device Manager - Storage Navigator

Name	CCI command	Device Manager - Storage Navigator
User name (Password)	63 characters	256 characters

Operational differences

The following table lists the operational differences between CCI and Device Manager - Storage Navigator.

Operation	CCI	Device Manager - Storage Navigator
Adding or changing path for an external volume (UVM)	The path definition is required for each external volume in the path group.	You can specify the path groups on the screen and execute the add/change path at one time.
Operations when the software product is not installed	You can change and delete existing resources for this software product, but you cannot add new resources.	You cannot change or delete existing resources or add new resources.
Displaying WWN	The WWN is displayed only when LUN security is enabled.	The WWN is displayed when LUN security is enabled or disabled.
Setting Pool ID	Optional	Required
Creating LDEV	You can specify the LDEV size in GB, LBA, or cylinders. When you specify GB, CCI can perform with or without size correction. If the capacity of LDEVs that are created by each GUI and CLI is the same, a copy pair might not be created. To create a pair with the LDEV that is created by GUI, create an LDEV by specifying LBA.	You can specify the LDEV size in GB, LBA, or cylinders. When you specify GB, size correction is performed.
Expanding the capacity of Dynamic Provisioning or Dynamic Provisioning for Mainframe virtual volume	You need to specify the capacity to be added to the volume.	You need to specify the total capacity after the volume is expanded.

Operation	CCI	Device Manager - Storage Navigator
Moving the CLPR assigned to the LUSE configuration volumes or the CLPR of parity groups containing LUSE configuration volumes	CLPRs can be moved, but it is not recommended.	CLPRs cannot be moved.
Moving the CLPR assigned to the journal volumes	CLPRs cannot be moved.	CLPRs can be moved if you specify all LDEVs in the journal.
Deleting SPM name	The SPM name of WWN is deleted, and the registration of the SPM name in the specified port is released.	The SPM name of WWN is deleted, but the SPM registration is maintained.
Deleting SPM group	WWN is deleted from the group, and the registration of SPM in the group of the specified port is released.	WWN is deleted from the group, but the SPM registration is maintained.
Moving the parity groups which configure the distributed parity group between the CLPRs	Parity groups cannot be moved.	All parity groups which are concatenated are moved.

Operations that cannot be used with Device Manager - Storage Navigator

You cannnot use CCI together with Device Manager - Storage Navigator for some operations.

The following table shows operations that cannot be used with Device Manager - Storage Navigator.

Function	Operation
ShadowImage	Device Manager - Storage Navigator cannnot split a pair created as a consistency group by CCI.
Universal Replicator	- Device Manager - Storage Navigator cannnot split, resynchronize, or delete a pair created by CCI. - CCI cannnot split, resynchronize, or delete a pair created byDevice Manager - Storage Navigator.

Chapter 2: Data management commands

This information provides specifications for the CCI data replication and protection management commands.

> 📄 **Note:** For the following commands marked with an asterisk (*), if the HORCC_CHECK_CHARACTER_CODE environment variable is set when executing the command, the command checks the codes of the characters used for the command options and parameters. If the HORCC_CHECK_CHARACTER_CODE environment variable is set, when you specify a character code of 0x1F or less, or 0x7F or more as a command option or parameter, then the command stops the processing and responds EX_INVARG.

paircreate*

> ⚠️ **Warning:** Use the `paircreate` command with caution. The command starts the initial copy operation, which overwrites all data on the secondary (target) volume of the pair. If the primary and secondary volumes are not identified correctly, or if the wrong options are specified (for example, vl instead of vr), data is copied to the wrong volume overwriting and the data in the target of transferring is overwritten.

The `paircreate` command is used to create a new volume pair from two unpaired volumes. The `paircreate` command can create either a paired logical volume or a group of paired volumes. The `paircreate` command allows you to specify the direction (local or remote) of the pair generation (see the following figure for pair creation). If local (`-vl` option) is specified, the server issuing the `paircreate` command has the primary volume. If remote (`-vr` option) is specified, the remote server has the primary volume. The `-split` option of the `paircreate` command (ShadowImage and Copy-on-Write Snapshot only) allows you to simultaneously create and split pairs using a replication command only. When `-split` is used, the pair status changes from COPY to PSUS (instead of PAIR) when the initial copy operation is complete.
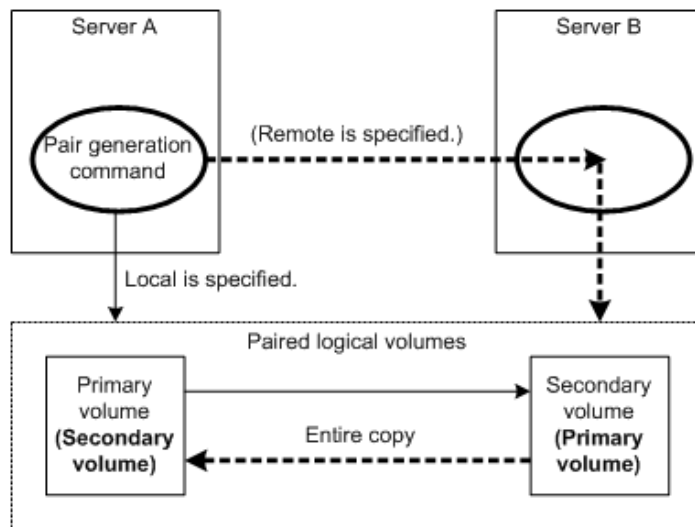
Before issuing the **paircreate** command, make sure that the secondary volume is not mounted on any system. If the secondary volume is found to be mounted after **paircreate**, delete the pair (**pairsplit -S**), unmount the secondary volume, and then reissue the **paircreate** command.

When you use TrueCopy, TrueCopy for Mainframe, TrueCopy Async, Hitachi TrueCopy Asynchronous for Mainframe, Universal Replicator, Universal Replicator for Mainframe, or global-active device, if a failure that requires maintenance operations, such as blocking the controller board or the cache memory, occurs on the primary storage system, the **paircreate** command cannot be executed.

When you use global-active device, if a failure that requires maintenance operations occurs on the secondary storage system, the **paircreate** command cannot be executed as well.

The **paircreate** command terminates before the initial copy operation is complete (except when the nocopy option is specified). Use the **pairevtwait** or **pairdisplay** command to verify that the initial copy operation completed successfully (status changes from COPY to PAIR, or from COPY to PSUS if the -split option was specified).

📄 **Note:** (For VSP G1x00 and VSP F1500, VSP E series, VSP G130, G/F350, G/F370, G/F700, G/F900) When a consistency group is created with TC, GAD, and UR by using the CCI paircreate or pairresync command, the same CTG ID cannot be used between different software. For example, if a consistency group including a GAD pair whose CTG ID is 0 exists, a consistency group including a UR pair whose CTG ID is 0 cannot be created. Therefore, if a consistency group is created with multiple software products by using CCI, the maximum number of consistency groups supported by the storage system might not be created. The following example shows the consistency groups created with multiple software products by using CCI. Make sure to use Device Manager - Storage Navigator to create the maximum number of consistency groups supported by the storage system.

The following example lists the allowed maximum numbers and the range of assigned CTG IDs of consistency groups if they are created in 3DC configurations with GAD and UR for VSP G1x00 and VSP F1500.

- GAD: 85 (CTG ID 0 through CTG ID 84)

- UR: 85 (CTG ID 85 through CTG ID 169)

- UR delta resync: 85 (CTG ID 170 through CTG ID 254)

Syntax

```
paircreate { -h | -q | -z[x] | -I[H][M][instance#] or
   -I[TC][SI][instance#]| -g <group>| -d <pair Vol> | -d[g] <raw_device>
   [MU#][-n] | -FHORC [MU#] or -FCA[MU#] | -d[g] <seq#> <ldev#>
   [MU#]| -f[g] <fence> [<CTG ID> | -startctg <START CTG ID>] | -v | -c <size> | -
nocopy | -nomsg |
   -split | [-m <mode>] | -jp <id> | -jq <id> | -js <id> | -pid <PID> |
   -fq <mode>| -cto <o-time> [c-time] [r-time] | -pvol(svol)[ldevgrp] | -nocsus | -pr
<io preference> }
```

Options and parameters

-h

Displays help/usage and version information.

-q

Terminates the interactive mode and exits the command.

-z or -zx

Makes the **paircreate** command enter interactive mode. The -zx option monitors if the HORCM is operating in interactive mode. When this option detects a HORCM shutdown, interactive mode terminates.

OpenVMS cannot use the -zx option.

-I[H][M] [instance#] or -I[TC][SI] [instance#]

Specifies the command as [HORC]/[HOMRCF] or [TrueCopy]/[ShadowImage], and used to specify the CCI instance number.

-g <group>

> Specifies a group name defined in the configuration definition file. The command is executed for the specified group unless the `-d <pair Vol>` option is specified.

-d <pair Vol>

> Specifies paired logical volume name defined in the configuration definition file. When this option is specified, the command is executed for the specified paired logical volume.

-d[g] <raw_device> [MU#] [-n]

> Searches whether the specified raw_device is included in a group on the configuration definition file (local instance). If it is in the group, the target volume is executed as the paired logical volume (-d) or group (-dg). This option is effective without specification of `-g <group>` option. If the specified raw_device is contained in two or more groups, the command is executed for the first group.
>
> This `-n` option is specified to use a device file for NVMe-oF. Specifying this option enables the use of both device files for NVMe-oF and SCSI.
>
> If the HORCC_NVME environment variable is specified, both device files for NVMe-oF and SCSI can be used without the `-n` option specified.
>
> To specify this option, a command device must be enabled in a namespace on the storage system. For more information about whether a command device can be enabled in a namespace on the storage system, see the *Provisioning Guide* for your storage system.

-d[g] <seq#> <ldev#> [MU#]

> Searches whether the specified LDEV is included in a group on the configuration definition file (local instance). If it is in the group, the target volume is executed as the paired logical volume (-d) or group (-dg). This option is effective without specification of '-g <group>' option. If the specified LDEV is contained in two or more groups, the command is executed on the first group. The <seq#> <LDEV#> values can be specified in hexadecimal (add the '0x' prefix) or decimal notation.
>
> ▪ When specifying the serial number for VSP 5000 series, add a "5" at the beginning of the serial number. For example, for serial number 12345, enter 512345.
>
> ▪ When specifying the serial number for VSP G1x00 and VSP F1500, add a "3" at the beginning of the serial number. For example, for serial number 12345, enter 312345.

-f[g] <fence> [<CTG ID> | -startctg <START CTG ID>]

> TrueCopy/TrueCopy Async/Universal Replicator/High Availability Manager/global-active device only.
>
> Specifies the fence level for assuring the consistency of paired volume data. A fence level of 'data', 'status', 'never', or 'async' must be specified. This option must always be specified. Fence level '-f async' can be specified only for TrueCopy Async/Universal Replicator. The '-fg' option is used to create a TrueCopy Sync CTG volume, and fence level must be specified as '-fg data', '-fg status', or '-fg never'.
>
> Specifies the fence level of '-f never' for HAM.

Specifies the fence level of '-f never' or '-fg never' for GAD. The '-fg never' option is used to create a CTG volume.

'CTG ID' is assigned as follows:

- When no CTG ID is assigned to the other device in the specified volume group:

  The CTG ID specified by "CTGID" option.

  If the "CTGID" option is omitted, new CTG ID.

- When the CTG ID is already assigned to another device in the specified volume group:

  The CTG ID which has been assigned to the other device in the volume group.

  Note that the CTG ID specified by "CTGID" option is invalid.

If 'CTGID' is not specified (with '-f async' or '-fg' option) and the maximum number of consistency groups already exists, an EX_ENOCTG error is returned. Therefore, the 'CTG ID' option can forcibly assign an existing CTG ID to a volume group on the RAID storage systems only when no CTG ID has been assigned to the volume group. The CTGID option is ignored unless you specify the '-f async' or '-fg' option.

The "-startctg" option is valid only when "-fg" is specified. START CTG ID is a value indicating the search start position of CTG automatic allocation. Using "START CTG ID", an unassigned CTG ID to another device is searched from the value specified by the "START CTG ID" option to the maximum value of the consistency group in ascending order. If the unassigned CTG ID is found, the search ends and the CTG ID is assigned to CTG. If the unassigned CTG ID is not found, an EX_ENOCTG error is returned.

-vl or -vr ; -pvol [ldevgrp] or -svol [ldevgrp]

Specifies the data flow direction and must always be specified. The -vl (-pvol) option specifies 'local' and the host that issues the command possesses the primary volume. The -vr (-svol) option specifies 'remote' and the remote host possesses the primary volume while the local host possesses the secondary volume. [ldevgrp] configures the specified LDEV group as the second volume.

-c <size>

TrueCopy/TrueCopy for Mainframe/TrueCopy Async/Hitachi TrueCopy Asynchronous for Mainframe/ShadowImage/global-active device only.

Specifies the track size of extents (1 to 15) to be used for the copy operation. If you specify a large number, the time for copy operation will be shortened. When you want to copy in a short time by stopping the writing of data to the P-VOL, specify the maximum value 15. If this option is not specified, the default value (3) is used.

The relationship between the track size and the copy pace is shown as follows. If the copy pace is fast, the host I/O performance might be degraded. To reduce the effect on host I/O performance, lower the copy pace.

- TrueCopy/ Async/global-active device

  - When the track size is 1, the copy pace is slow.

  - When the track size is 2 or 3, the copy pace is medium.

  - When the track size is 4 or more, the copy pace is fast.

- ShadowImage

  - When the track size is 1 or 2, the copy pace is slow.

  - When the track size is 3, the copy pace is medium.

  - When the track size is 4 or more, the copy pace is fast.

- TrueCopy for Mainframe/Hitachi TrueCopy Asynchronous for Mainframe

  - When the track size is 1 to 3, the number of the tracks copied simultaneously is 3 and the copy pace is slow.

  - When the track size is 4 or more, the number of the tracks copied simultaneously is 15 and the copy pace is fast.

> **Note:** This option is disabled in Universal Replicator, Universal Replicator for Mainframe, or ShadowImage for Mainframe. To change the copy pace in Universal Replicator or Universal Replicator for Mainframe, use the `raidcom modify journal` command for JNL option changes, or Storage Navigator. ShadowImage for Mainframe works only at medium pace.

-nocopy

Creates paired volumes without the initial copy operation when the data consistency of simplex volumes is assured by the user.

> **Note:**
> This option cannot be specified for ShadowImage or ShadowImage for Mainframe.

-nomsg

Suppresses messages to be displayed when this command is executed. It is used to execute this command from a user program. This option must be specified at the beginning of a command argument. The command execution log is not affected by this option.

-split

ShadowImage/Copy-on-Write Snapshot only.

Splits the paired volume after the initial copy operation is complete. This option will return after changed the state in P-VOL_PSUS & S-VOL_COPY immediately, and S-VOL state is changed to 'SVOL_SSUS' after all data is copied.

-m <mode>

Specifies the mode.

mode = noread (ShadowImage only): Specifies the noread mode for hiding the secondary volume. The secondary volume cannot be read when this mode option is specified. The secondary volume can be read when this mode option is omitted.

> **Note:**
>
> The primary volume becomes read-disabled during a reverse resync operation (restore option of `pairresync` command).

mode = cyl (TrueCopy/TrueCopy Async/Universal Replicator/High Availability Manager/global-active device only): Set this value when the difference of bitmap is managed by a cylinder unit.

mode = trk (TrueCopy/TrueCopy Async/Universal Replicator/High Availability Manager/global-active device only): Set this value when the difference of bitmap is managed by a track unit.

> **Note:**
>
> When this mode (cyl or trk) is not specified, the default value is used. About the default value, see the user guide of the storage system or software product.

> **Note:**
>
> When the storage system does not support the bitmap management by a unit of cylinder, the bitmap difference is managed by track even if you specify "cyl". For details about the bitmap management which is supported by the storage system, see the user guide of the storage system or software product.

> **Note:**
>
> When the connection destination storage system does not support the bitmap management by a unit of cylinder, the bitmap difference is managed by track even if you specify "cyl". For details about the bitmap management which is supported by the storage system, see the user guide of the storage system or software product.

mode = grp [CTG ID] (ShadowImage, Thin Image, or Copy-on-Write Snapshot only). Creates a group for splitting all ShadowImage pairs specified in a group. Like a TrueCopy Async or Universal Replicator consistency group, ShadowImage guarantees data consistency in a group at a single point in time when splitting a pair using the 'pairsplit -g <group>' command (except '-S' or '-E' option).

'CTG ID' is assigned as follows:

▪ When no CTG ID is assigned to other devices in the specified volume group:

  The CTG ID specified by "CTG ID" option.

  If the "CTG ID" option is omitted, new CTG ID.

▪ When the CTG ID is already assigned to another device in the specified volume group:

  The CTG ID which has been assigned to the other device in the volume group.

  Note that the CTG ID specified by "CTG ID" option is invalid.

If 'CTG ID' is not specified and the number of consistency groups reaches the maximum, an EX_ENOCTG error is returned. Therefore, the 'CTG ID' option can forcibly assign an existing CTG ID to a volume group (for example, 0 to 127 on 9900V) only when no CTG ID has been assigned to the volume group.

For detail, please refer to Restrictions on specified volumes with `-m grp` option.

> 📋 **Note:**
>
> This option cannot be specified with `-split` option in the same command.

mode = cc (Volume Migration only): Specifies Volume Migration.

The `-vl` option specifies 'local', and copies data from the local instance LU (P-VOL) to the remote instance LU (S-VOL). The original volume as the local instance LU is migrated from P-VOL to S-VOL, and the physical volume mapping between P-VOL and S-VOL is switched after copied.

The `-vr` option specifies 'remote', and copies data from the remote instance LU (P-VOL) to the local instance LU (S-VOL). The original volume as the remote instance LU is migrated from P-VOL to S-VOL, and the physical volume mapping between P-VOL and S-VOL is switched after copied. During maintenance work on the storage system (Device Manager - Storage Navigator or the maintenance utility is in modify mode), this operation cannot be completed.

> 📋 **Note:**
>
> This option cannot be specified with the `-split` option in the same command.
>
> This option ignores the '-c <size>' option.

-jp <id> or -jq <id>

Universal Replicator, the HAM configuration or the GAD configuration only.

- For Universal Replicator:

  You can use -jp <id> option when specifying a journal ID for P-VOL.

  The -jp <id> option is valid when the fence level is set to 'ASYNC', and a journal ID is automatically bound to the CTG ID.

- For the HAM configuration or the GAD configuration:

  You can create a HAM/GAD pair by using either -jp <id> or -jq <id> option.

  Specify the quorum ID with -f fence (never) option and quorum ID when creating a HAM pair or a GAD pair. You do not need to check the resource group of the quorum volume.

The following conditions must be met for the HAM configuration.

- It is not a consistency group.

- The fence level is set to 'Never'.

- The RCU path is set to 'CU Free'.

The following conditions must be met for the GAD configuration.

- The fence level is set to 'Never'.

- The RCU path is set to 'CU Free'.

-js <id>

Universal Replicator only.

This option is used when specifying a journal ID for S-VOL.

Both the -jp <id> and -js <id> options are valid when the fence level is set to 'ASYNC', and each journal ID is automatically bound to the CTG ID.

-pid <PID>

Copy-on-Write Snapshot and Thin Image only.

This option is used to specify the ID of the pool in which snapshot data of the pair to be created is stored.

If you omit this option, 0 is set by default as PID.

If the specified pool is for Thin Image or Dynamic Provisioning, a Thin Image pair is created.

If the specified pool is for Copy-on-Write Snapshot, a Copy-on-Write Snapshot pair is created.

Use LDEV with the virtual volume attribute for S-VOL. With paircreate, you cannot create a snapshot pair using a volume with the HDP attribute as S-VOL. Use the raidcom add snapshot command to use a volume with the HDP attribute as S-VOL.

-fq <mode>

ShadowImage only.

Chapter 2: Data management commands

This option is used when specifying the mode whether -split is performed or not as 'QUICK'.

- **mode = normal:** The '**paircreate -split**' is performed as non-quick mode regardless of setting of $HORCC_SPLT environment variable or the system option mode 122.

- **mode = quick:** The '**paircreate -split**' is performed as Quick Split regardless of setting of $HORCC_SPLT environment variable or the system option mode 122.

If this option is not specified, then the performing of the 'Split' depends on $HORCC_SPLT environment variable or the system option mode setting through the SVP, whether the paircreate operation is Quick Split or not.

The relationship between -fq option and $HORCC_SPLT is as shown below:

-fq option	$HORCC_SPLT	Behavior
quick	Invalid	Quick Split
normal	Invalid	Normal Split
Omitted	QUICK	Quick Split
Omitted	NORMAL	Normal Split
Omitted	Omitted	- For VSP 5000 series, VSP E series, VSP G130, G/F350, G/F370, G/F700, G/F900: Quick Split - For other models: Determined by system option mode 122

📄 **Note:**

- This -fq option is also valid for TrueCopy-TrueCopy/ShadowImage cascading operation using the -FBC [MU#] option.

- The -fq option is applied to the following storage systems:
  - VSP 5000 series
  - VSP G1x00 and VSP F1500
  - VSP E series
  - VSP G130, G/F350, G/F370, G/F700, G/F900
  - VSP G200, G400, G600, G800 and VSP F400, F600, F800

  This option is ignored to maintain the compatibility on 9900V so that you can add this option to the same script.

-FHORC [MU#] or -FCA [MU#]

This option is used to create the cascading configuration with `-g <group>` and `-gs <group>` options from the local node (takeover node).

`-g <group>` is used when specifying the cascaded P-VOL, and also -gs <group> option is used when specifying the cascaded S-VOL. This operation ignores the -vl or vr option, because S-VOL is specified with `-gs <group>` option.

`-gs <group>`: This 's' option is used when specifying a group name for cascading S-VOL (defined in the configuration definition file). The command is executed for the specified group unless the -ds <pair Vol> option shown below is specified.

`-ds <pair Vol>`: The following options can be specified for cascading S-VOL.

`-d[g]s <raw_device> [MU#] ...`

`-d[g]s <seq#> <LDEV#> [MU#]`

-cto <o-time> [c-time] [r-time]

TrueCopy Async and Universal Replicator only.

If you specify -cto <o-time> [c-time] [r-time] option on TrueCopy Sync, it is ignored.

`o-time:`

This option is used when setting offloading timer for controlling inflow of write I/Os to the specified consistency group.

For TrueCopy Async, an integer from 1 through 255 (seconds) must be specified. If this option is not specified, 90 seconds is set by default.

For Universal Replicator, an integer from 1 through 255 (seconds) must be specified. If this option is not specified, the value set for the journal on the P-VOL side is applied. Use the **raidcom modify journal** command to change the value range to 256 to 600.

If o-time=0 is specified, the inflow of write I/O becomes invalid.

When the sidefile capacity exceeds the limit of the sidefile area, write I/Os from the host wait, within the specified timeout period, until the space which is large enough to store next new data becomes available. As the timeout period, you can specify a value from 1 through 255 (seconds) for TrueCopy Async, 1 through 600 (seconds) for Universal Replicator. The default timeout values are 90 seconds for TrueCopy Async, and 60 seconds for Universal Replicator. If the timeout happens occurs during this waiting state, then the pair status changes from PAIR to PSUS of sidefile (Journal) Full, and its host side Write I/Os continue and data in waiting state is managed by BITMAP mode.

Therefore the o-time timeout value must be less than the I/O timeout value of the host system

[c-time](TrueCopy Async only): This option is used when setting Copy Pending timer to the specified consistency group. Specify an integer from 1 to 15 (minutes) for c-time. If this option is not specified, then this value is set as follows:

- If a consistency group is created, 5 (minutes) is set by default.

- If a consistency group is not created, the setting is not changed.

Chapter 2: Data management commands

[r-time] (TrueCopy Async only): This option is used when specifying RCU Ready timer for the specified consistency group. r-time can be set from 1 to 10 minutes. If this option is not specified, this value is set as follows:

- If a consistency group is created, 5 (minutes) is set by default.

- If a consistency group is not created, the setting is not changed.

> **📄 Note:**
>
> For TrueCopy Async, settings changed by these options are invalid if a consistency group already exists. These parameters are also forwarded to the S-VOL side with the `paircreate` command, and are used when S-VOL is changed to P-VOL. These parameters are maintained and become valid until and when the pair-volumes are changed to SMPL.
>
> For Universal Replicator, these parameters can be set and changed when a pair is created in a journal volume in the P(S)JSN status. The parameters are set for the journal on the P-VOL side when the pair is created. To set parameters for the journals on both P-VOL and S-VOL sides, specify as follows:
>
> 1. `paircreate -g <group> -vr -f async -nocopy -jp <id> -js <id> -cto <o-time>`
> 2. `pairsplit -g <group> -S`
> 3. `paircreate -g <group> -vl -f async -jp <id> -js <id> -cto <o-time>`
>
> These parameters are maintained on each journal. Therefore, if you set the value for offloading timer, execute the `raidcom modify journal` command on both P-VOL and S-VOL sides.

-nocsus

Universal Replicator only.

This option is used to create the suspended journal volumes without copying data in order to make the delta-resync between DC2 (Sync-S-VOL) and DC3 (Universal Replicator-S-VOL).

-pr <io preference>

This option is only available for global-active device.

Set the I/O preference mode (I/O preference mode when remote path failed) when a failure occurs in the remote path between the storage systems of the primary site and the secondary site and communication becomes impossible. For details of the I/O preference mode when remote path failed, see the Global-Active Device User Guide.

The following value can be specified for <io preference>.

pvol: Primary Volume. When remote path failed, the primary volume continues host I/O preferentially.

If this option is omitted, the I/O Preference Mode When Remote Path Failed is disabled.

Returned values

The `paircreate` command sets either of the following returned values in exit (), which allows users to check the execution results using a user program.

▪ **Normal termination:**

  • **0:** When creating groups, 0 = normal termination for all pairs.

▪ **Abnormal termination:**

  • **other than 0:** Refer to the error code for error details.

Error codes

Unrecoverable errors are not resolved, even after re-executing the command. If the command fails, the detailed status is logged in the CCI command log ($HORCC_LOG).

📄 **Note:** When an option that is valid for only specific volumes is specified for other volumes, the error code might be EX_UNWOPT or EX_UNWCMD. Check whether the specified option is valid for the volume.

Category	Error Code	Error Message	Recommended Action	Value
Volume status (Unrecoverable)	EX_ENQVOL	Unmatched volume status within the group	Confirm status using the `pairdisplay` command. Make sure all volumes in the group have the same fence level and volume attributes.	236
	EX_INCSTG	Inconsistent status in group	Confirm pair status using `pairdisplay`.	229
	EX_INVVOL	Invalid volume status	Confirm volume status/attribute using `pairdisplay` (-l option) or 'raidvchkdsp -v aou'.	222
	EX_INVSTP	Invalid pair status	Confirm pair status using `pairdisplay`.	228
	EX_ENQSIZ	Unmatched volume size for pairing	Confirm volume size or number of LUSE volume using raidscan -f, and make sure volume sizes are identical.	212
Resource (Unrecoverable)	EX_ENOCTG	Not enough consistency groups in the RAID	Choose an existing CTG ID (pairvolchk displays CTG IDs). Use '-f async <CTG ID>' or '-m grp <CTG ID>' option of `paircreate` to force the pair into a pre-existing CTG ID.	217

Chapter 2: Data management commands

Category	Error Code	Error Message	Recommended Action	Value
	EX_ENXCTG	No consistency groups left for OPEN Vol use.	Confirm whether all consistency groups are already used by TrueCopy/TrueCopy Async/GAD or ShadowImage.	215
	EX_ENOPOL	Not enough Pool in RAID	Unable to retain the pool for executing a command because the threshold rate has been exceeded. Delete unnecessary/earlier generations paired volume, or re-synchronize unnecessary/earlier generations split volume.	206

Chapter 2: Data management commands

Restrictions on specified volumes with `-m grp` option

- Volume group definition

  - Volume groups specified with `-m grp` cannot be defined across the storage systems.

  - When multiple groups of CCI are contained within the same consistency group (CTG ID), pairs with the specified group are operated for the entire consistency group.

  - When ShadowImage, Thin Image, or Copy-on-Write Snapshot volumes are cascaded by TrueCopy, TrueCopy Async, Universal Replicator, or GAD volumes, the data consistency is not guaranteed by the **pairsplit** command (including **pairsplit - FMRCF**) for which I/Os are being processed.

- Registration and limitations of the number of CTG IDs

  When you create a pair, CCI maps and assigns the configuration definition file group to a CTG ID managed by the storage system. The maximum number of consistency groups that can be registered to a storage system is shown below. If you register over the maximum number of consistency groups, pair creation terminates with EX_ENOCTG error.

  - HUS VM, VSP: 256 (CTG ID 0 - CTG ID 255)

  - VSP G1x00 and VSP F1500: 256 (CTG ID 0 - CTG ID 255)

  - VSP G800, VSP F800, 9900V: 128 (CTG ID 0 - CTG ID 127)

  - VSP G400, VSP G600, VSP F400, VSP F600: 64 (CTG ID 0 - CTG ID 63)

  - VSP G200: 16 (CTG ID 0 - CTG ID 15)

  - VSP E series, VSP G130, G/F350, G/F370, G/F700, G/F900 (ShadowImage): 128 (CTG ID 0 to CTG ID 127)

  - VSP E series, VSP G130, G/F350, G/F370, G/F700, G/F900 (Thin Image): 2048 (CTG ID 0 to CTG ID 2047)

  - VSP G/F350, G/F370, G/F700, VSP E590, VSP E790: (TrueCopy/ Universal Replicator/global-active device): 128 (CTG ID 0 to CTG ID 127)

  - VSP G/F900, VSP E990, VSP E1090: (TrueCopy/ Universal Replicator/global-active device): 256 (CTG ID 0 to CTG ID 255)

  - VSP 5000 series: (except global-active device): 256 (CTG ID 0 - CTG ID 255)

  - VSP 5000 series: (global-active device): 1024 (CTG ID 0 - CTG ID 1023)

Example 1

The following figure shows an example of creating a cascading configuration with `-g <group>` and `-gs <group>` option from the local node (takeover node).
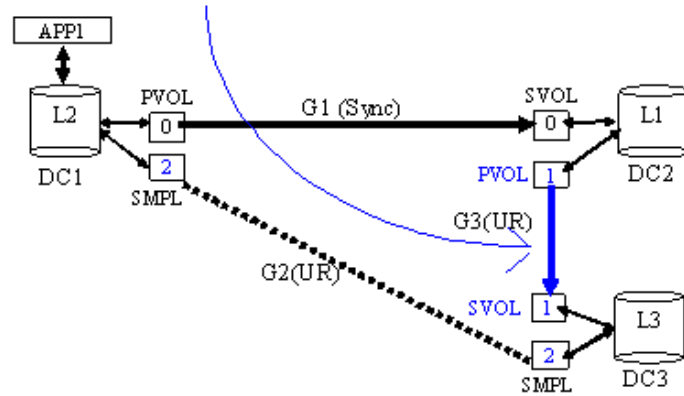
Example 2

The following figure shows an example for creating a suspended journal volume.

On DC1 side:

```
paircreate -g G1 -gs G2 -FHORC 2 -nocsus -f async <CTG ID> -jp <id>
-js <id>
```

On DC2 side:

```
paircreate -g G3 -vl  -nocsus -f async <CTG ID>  -jp <id> -js <id>
```

On DC3 side:

```
paircreate -g G3 -vr  -nocsus -f async <CTG ID>  -jp <id> -js <id>
```

Note: The journal ID for the shared Universal Replicator-SVOL must be specified the same journal ID for S-VOL currently. The CTG ID for `paircreate` can be specified the same consistency group for S-VOL currently.
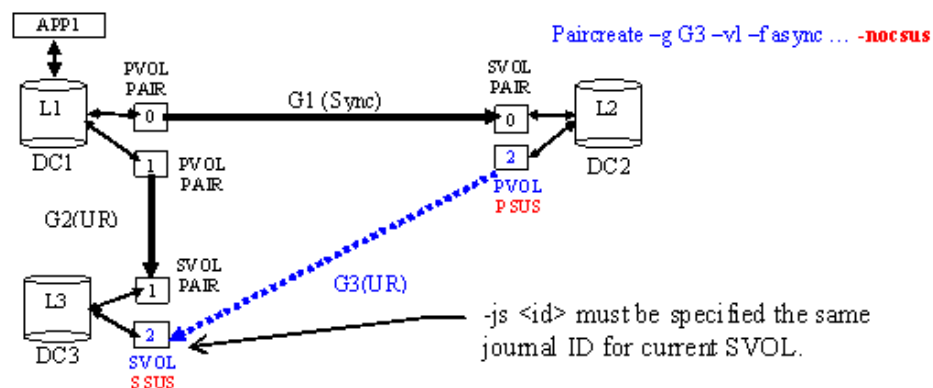
Example 3

The following figure shows a takeover example used to suspend a journal volume.

Note: The `pairresync` command must be issued after TC_Sync volume became SVOL_SSWS or PVOL_COPY/PAIR through the `horctakeover` command.
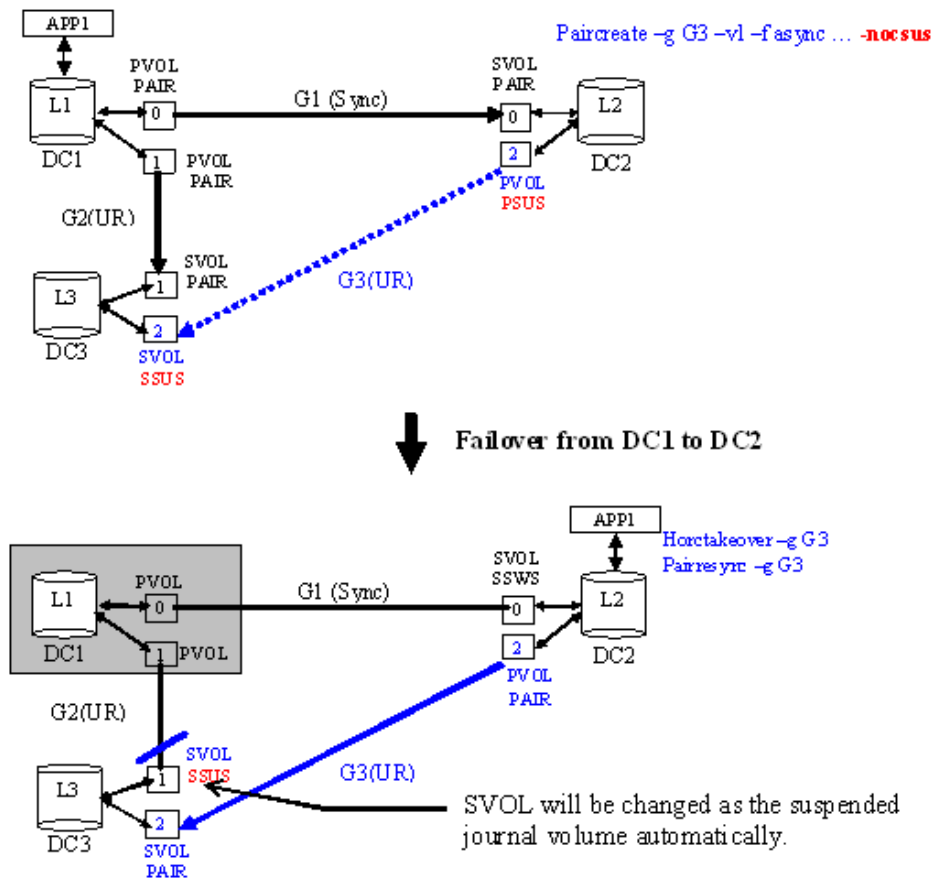
pairsplit*

The `pairsplit` command is used to split or delete volume pairs. This command stops updates to the secondary volume of a pair and can either maintain (status = PSUS) or delete (status = SMPL) the pairing status of the volumes. The `pairsplit` command can be applied to a paired logical volume or a group of paired volumes. The `pairsplit` command allows read access or read/write access to the secondary volume, depending on the selected options. When the `pairsplit` command is specified, acceptance of write requests to the primary volume depends on the fence level of the pair (data, status, never, or async). For Volume Migration, only the -S option is accepted.

The following figure illustrates the pair splitting process.
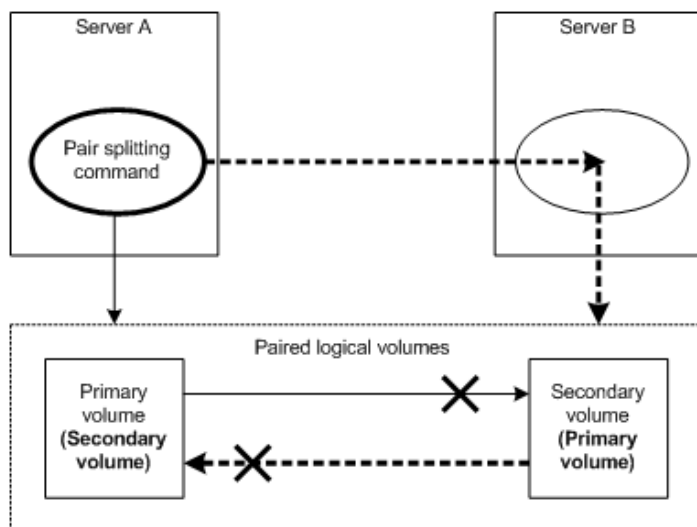
The primary volume's server is automatically detected by the **pairsplit** command, so the server does not need to be specified in the **pairsplit** command parameters. If the -S option (simplex) is used, the volume pair is deleted, the volumes are returned to the simplex state, and the primary and secondary volume status is lost. Paired volumes are split as soon as the **pairsplit** command is issued. If you want to synchronize the volumes, the **pairsplit** command must be issued after write I/Os to the paired volume have completed.

You can create and split ShadowImage pairs simultaneously using the -split option of the **paircreate** command.

You can delete pairs by using the -S option of the **pairsplit** command. When the **pairsplit -S** command is issued, the specified pair is deleted, and each volume is changed to SMPL (simplex) mode. If you want to re-establish a pair that has been deleted, you must use the **paircreate** command (not **pairresync**).

For Volume Migration, options other than -S are rejected.

Note on Quick Split:

If '$HORCC_SPLT=QUICK' environment variable is set for any of the following storage systems, the **pairsplit** and **paircreate -split** operations are performed as Quick Split regardless of the system option mode 122 setting:

- VSP 5000 series
- VSP G1x00 and VSP F1500
- VSP E series
- VSP G130, G/F350, G/F370, G/F700, G/F900
- VSP G200, G400, G600, G800 and VSP F400, F600, F800
- VSP
- HUS VM

The $HORCC_SPLT=QUICK environment variable is ignored for 9900V.

Chapter 2: Data management commands

Syntax

```
pairsplit {-h | -q | -z[x] | -I[H][M][instance#] or -I[TC][SI]
[instance#] | -g <group> | -d <pair Vol> | -d[g] <raw_device>
[MU#][-n] | -FHORC [MU#] | -FMRCF [MU#] | -d[g] <seq#> <LDEV#> [MU#]
| -r |  -rw | -S | -SF[V] | -R[S|B] | -RF[V] | -P | -l |
-t <timeout> | -nomsg | -C <size> | -E | -fq <mode>} | -iomd <mode>
```

Options and parameters

Only one **pairsplit** option (`-r`, `-rw`, `-S`, `-R`, `-P`, `-C`, or `-E`) can be specified unless stated in the description of each option. If more than one option is specified, only the last specified option will be executed, unless stated in the description of each option.

-h

Displays help/usage and version information.

-q

Terminates the interactive mode and exits this command.

-z or -zx

Makes the **pairsplit** command enter the interactive mode. The `-zx` option guards performing of the HORCM in the interactive mode. When this option is specified and a HORCM shutdown is detected, interactive mode terminates.

OpenVMS cannot use the `-zx` option

-I [H] [M] [instance#] or -I [TC] [SI] [instance#]

Specifies the command as [HORC]/[HOMRCF] or [TrueCopy]/[ShadowImage], and used when specifying the CCI instance number.

-g <group>

Specifies a group name defined in the configuration definition file. The command is executed for the specified group unless the `-d <pair Vol>` option is specified.

If the HORCC_CHECK_PAIRSPLIT_GOPTION environment variable is set, when you specify more than one parameter in the -g option, then the command suspends the processing and respond with EX_INVARG.

-d <pair Vol>

Specifies the paired logical volume name defined in the configuration definition file. When this option is specified, the command is executed for the specified paired logical volumes. This option is valid when the `-g <group>` option is specified.

-d[g] <raw_device> [MU#][-n]

Searches whether the specified raw_device is included in a group on the configuration definition file (local instance). If it is in the group, the target volume is executed as the paired logical volume (-d) or group (-dg). This option is effective without specification of `-g <group>` option. If the specified raw_device is contained in two or more groups, the command is executed for the first group.

This `-n` option is specified to use a device file for NVMe-oF. Specifying this option enables the use of both device files for NVMe-oF and SCSI.

If the HORCC_NVME environment variable is specified, both device files for NVMe-oF and SCSI can be used without the `-n` option specified.

To specify this option, a command device must be enabled in a namespace on the storage system. For more information about whether a command device can be enabled in a namespace on the storage system, see the *Provisioning Guide* for your storage system.

-d[g] <seq#> <LDEV#> [MU#]

Searches whether the specified LDEV is included in a group on the configuration definition file (local instance). If it is in the group, the target volume is executed as the paired logical volume (-d) or group (-dg). This option is effective without specification of '-g <group>' option. If the specified LDEV is contained in two or more groups, the command is executed on the first group. The <seq#> <LDEV#> values can be specified in hexadecimal (add the '0x' prefix) or decimal notation.

- When specifying the serial number for VSP 5000 series, add a "5" at the beginning of the serial number. For example, for serial number 12345, enter 512345.

- When specifying the serial number for VSP G1x00 and VSP F1500, add a "3" at the beginning of the serial number. For example, for serial number 12345, enter 312345.

-r or -rw

TrueCopy, TrueCopy for Mainframe, TrueCopy Async, Universal Replicator, and Universal Replicator for Mainframe only.

Specifies a mode of access to the S-VOL after paired volumes are split. If this option is omitted, `-r` is set. The `-r` option enables read-only access to the S-VOL. For mainframe volumes, both read and write are disabled.

The `-rw` option enables read and write access to the S-VOL.

-S

Specify this option when deleting a pair to change the P-VOL and S-VOL back to the simplex status.

> 📄 **Note:**
> - Either volume of a pair might be unpaired due to a failure. To recover from this failure, delete the pair by using the `pairsplit -S` command to cancel the P-VOL and S-VOL relation, and then create the pair again by using the `paircreate` command.
>
> - Due to a path failure, the S-VOL might not be unpaired even if you execute the `pairsplit -S` command. If the S-VOL cannot be unpaired, CCI unpairs the P-VOL.

-SF[V]

Specify the $-SF[V]$ option only when you must split a GAD pair forcibly. This option unpairs only the P-VOL of the GAD pair. It does not change the S-VOL pair status.

- $-SFV$ option: Specify this option to access the P-VOL from the host after deleting a GAD pair. This option leaves the virtual LDEV ID of the P-VOL.

- $-SF$ option: Specify this option if you do not access the P-VOL from the host after deleting a GAD pair. This option deletes the virtual LDEV ID of the P-VOL.

> ⚠️ **Caution:**
>
> Use the $-SF[V]$ option only if you cannot change the P-VOL status to Simplex by using the $-S$ option. Specify this option only when the I/O mode of both P-VOL and S-Vol is blocked. Contact technical support if you want to split GAD pair forcibly when the I/O mode is other than blocked. To confirm the I/O mode of the P-VOL and S-Vol, check the R/W row output with **pairdisplay -fe** command. B/B in the R/W row indicates the I/O mode of the volume is blocked.

> 📄 **Note:**
>
> The data of P-VOL and S-VOL becomes inconsistent and it might cause errors if the volume status is changed to Simplex by using $-SFV$ or $-RFV$ option when the server can access both P-VOL and S-VOL. Follow the steps below when you delete GAD pair forcibly.
>
> 1. Stop the access to either the P-VOL or the S-VOL from the server.
> 2. Change the volume to which you stopped access to Simplex status. For example, if you stopped the access to P-VOL, change the P-VOL status to Simplex by specifying the $-SF$ option. If you stopped the access to S-VOL, change the S-VOL status to Simplex by specifying the $-RF$ option. When the volume status is changed to Simplex by specifying the $-SF$ or $-RF$ option, the virtual LDEV ID is deleted and the GAD reserve attribute is set. The volume becomes inaccessible from the server when the reserve attribute is set.
> 3. Change the volume to which you did not stop access to Simplex status. For example, if you did not stop access to the P-VOL, change the P-VOL status to Simplex by specifying the $-SFV$ option. If you did not stop access to the S-VOL, change the S-VOL status to Simplex by specifying the $-RFV$ option.

> 📄 **Note:**
>
> In some cases (incorrect configuration definition file or other reasons), the paired volumes that are the operation target of the **pairsplit** command might both become P-VOLs. If the **pairsplit -SF[V]** command is issued under these conditions, the CCI instance that executes the command directs the storage system to delete the GAD pair of only the volume managed by that CCI instance.

-R[S|B]

TrueCopy, TrueCopy Async, Universal Replicator, and GAD only.

This option issues a command from the host at the secondary site if a failure or downtime occurs on the host accessing the P-VOL.

The `-R` option changes the S-VOL to the unpaired status (simplex volume).

- For TrueCopy, TrueCopy Async, and Universal Replicator:

  The `-R` option changes the S-VOL status, but it does not change the P-VOL pair status. However, if communication between MCU and RCU is established, the status of the P-VOL is PSUS, and the S-VOL is SMPL.

- For global-active device:

  The `-R` option deletes a GAD pair, and it changes the pair status of both the S-VOL and the P-VOL to SMPL. If you use this option to delete a GAD pair, the virtual LDEV ID of the S-VOL remains, but the virtual LDEV ID of the P-VOL is deleted.

  This option is available only when the pair status of the S-VOL is SSWS, and the I/O mode is local (L/L).

When a temporary failure such as a link failure occurs and you switch the control from the primary site host to the secondary site host, if you specify the `-RS` option, you can change the S-VOL pair status to SSWS.

- `-R`: Unpairs the S-VOL.

- `-RS`: Changes the S-VOL status to SSWS.

- `-RB`: Changes the S-VOL status from SSWS back to PSUS (PSUE) (SSUS).

Use the `-RB` option when the S-VOL status is SSWS and you want to resynchronize the pair from the P-VOL to the S-VOL. This option cannot be used for GAD.

-RF[V]

Specify the `-RF[V]` option only when you must delete a GAD pair forcibly. This option unpairs only the S-VOL of the GAD pair. It does not change the P-VOL pair status.

- `-RFV`: Specify this option to access the S-VOL from the host after deleting a GAD pair. This option leaves the virtual LDEV ID of the S-VOL.

- `-RF`: Specify this option if you do not access the S-VOL from the host after deleting a GAD pair. This option deletes the virtual LDEV ID of the S-VOL.

> ⚠️ **Caution:** Use the option only if you cannot change the S-VOL status to Simplex by using the `-R` option. Specify this option only when the I/O mode of both P-VOL and S-Vol is blocked. Contact technical support if you want to split GAD pair forcibly when the I/O mode is other than blocked. To confirm the I/O mode of the P-VOL and S-Vol, check the R/W row output with **`pairdisplay -fe`** command. B/B in the R/W row indicates the I/O mode of the volume is blocked.

> 📄 **Note:** The data of P-VOL and S-VOL becomes inconsistent and it might cause errors if the volume status is changed to Simplex by using `-SFV` or `-RFV` option when the server can access both P-VOL and S-VOL. Follow the steps below when you delete GAD pair forcibly.
>
> 1. Stop the access to either the P-VOL or the S-VOL from the server.
> 2. Change the volume to which you stopped access to Simplex status. For example, if you stopped the access to P-VOL, change the P-VOL status to Simplex by specifying the `-SF` option. If you stopped the access to S-VOL, change the S-VOL status to Simplex by specifying the `-RF` option. When the volume status is changed to Simplex by specifying the `-SF` or `-RF` option, the virtual LDEV ID is deleted and the GAD reserve attribute is set. The volume becomes inaccessible from the server when the reserve attribute is set.
> 3. Change the volume to which you did not stop access to Simplex status. For example, if you did not stop access to the P-VOL, change the P-VOL status to Simplex by specifying the `-SFV`option. If you did not stop access to the S-VOL, change the S-VOL status to Simplex by specifying the `-RFV` option.

> 📄 **Note:** In some cases (incorrect configuration definition file or other reasons), the paired volumes that are the operation target of the **pairsplit** command might both become S-VOLs. If the **pairsplit -RF[V]** command is issued under these conditions, the CCI instance that executes the command directs the storage system to delete the GAD pair of only the volume managed by that CCI instance.

-P

TrueCopy, TrueCopy Async, and Universal Replicator only.

For TrueCopy Sync, this option is used to bring the primary volume forcibly into write disabled mode like PSUE with 'fence=data'. It is issued by the secondary host to disable P-VOL data changes by the host possessing the primary volume.

For TrueCopy Async and Universal Replicator, this option is used to suspend and purge data remaining in the sidefile or journal without updating S-VOL like a link failure (PSUE). This option stops journal operations forcibly when the journal utilization traffic becomes high. This is the same as the failure that S-VOL data is not updated, but it allows write access if you specify `-rw -P`. In that situation, if you are using the S-VOL as a file system (that is, UFS, NTFS, HANFS), then an FSCK(CHKDSK) is necessary before the volume is mounted even after the P-VOL is unmounted.

-l

When the remote host cannot be used due to host down, this option enables a **pairsplit** operation by a local host only. Except the `-R` option, the target volume of a local host must be P-VOL. (ShadowImage or Copy-on-Write Snapshot volumes are able to split only S-VOL.)

-t \<timeout\>

This option is for Universal Replicator. Specifies the maximum time to wait (in seconds) for splitting or deleting a volume pair. If timeout occurs, the pairsplit

command fails with EX_EWSTOT. To avoid the timeout, specify the time required to split or delete a volume pair for this value. If this option is omitted, the value is set to the default value of 7,200 seconds (2 hours).

-nomsg

Suppresses messages to be displayed when this command is executed. It is used to execute a command from a user program. This option must be specified at the beginning of a command argument. The command execution log is not affected by this option.

-C <size>

ShadowImage only.

Copies differential data retained in the primary volume into the secondary volume, then enables reading and writing from/to the secondary volume after completion of the copying (default). If not specified, the value used for `paircreate` or `pairresync` command is used. In ShadowImage, specify 1 or 2 for slow copy pace, specify 3 for medium copy pace, and specify 4 for fast copy pace.

-E

ShadowImage only.

Specify this option when suspending a paired volume forcibly. Not normally used.

-FHORC [MU#] or -FCA [MU#]

Forcibly specifies a cascading remote copy volume for specified volume pair in a local copy environment. If the `-l` option is specified, a cascading remote copy volume is split on a local host (campus distance). If the `-l` option is not specified, a cascading remote copy volume is split on a remote host (metro distance). The target HORC volume must be a P-VOL, or '`-R[S][B]`' option can be specified for S-VOL.

-FMRCF [MU#] or -FBC [MU#]

Forcibly specifies a cascading local copy volume for specified volume pair in a remote copy environment. If the `-l` option is specified, a cascading local copy volume is split on a local host (near site). If `-l` option is not specified, a cascading local copy volume is split on a remote host (far site). The target local copy volume must be a P-VOL, and the -E option cannot be specified.

-fq <mode>

ShadowImage only.

This option is used to specify the mode whether `pairsplit` is performed as 'QUICK'.

mode = normal: `pairsplit` is performed as Non quick mode regardless of setting of $HORCC_SPLT environment variable and/or the system option mode 122 via SVP.

mode = quick: `pairsplit` is performed as Quick Split regardless of setting of $HORCC_SPLT environment variable and/or the system option mode 122 via SVP.

If this option is not specified, then performing 'Quick Split' depends on $HORCC_SPLT environment variable and/or the system option mode setting through the SVP.

The relationship between '`-fq`' option and $HORCC_SPLT is as shown below.

Chapter 2: Data management commands

-fq option	$HORCC_SPLT	Behavior
quick	Invalid	Quick Split
normal	Invalid	Normal Split
Omitted	QUICK	Quick Split
Omitted	NORMAL	Normal Split
Omitted	Omitted	▪ For VSP 5000 series, VSP E series, VSP G130, G/F350, G/F370, G/F700, G/F900: Quick Split ▪ For other models: Determined by system option mode 122

📄 **Note:**

- The `-fq` option is also valid for TrueCopy-TrueCopy/ShadowImage cascading operation using '`-FBC [MU#]`' option.

- The `-fq` option is applied to the following storage systems:

  - VSP 5000 series

  - VSP G1x00 and VSP F1500

  - VSP E series

  - VSP G130, G/F350, G/F370, G/F700, G/F900

  - VSP G200, G400, G600, G800 and VSP F400, F600, F800

  - VSP

  - HUS VM

  This option is ignored for 9900V to maintain compatibility on 9900V so that you can add this option to the same script.

-iomd <mode>
Supported storage systems:

- VSP 5000 series

- VSP G1x00 and VSP F1500

- VSP E series

- VSP G130, G/F350, G/F370, G/F700, G/F900

- VSP G200, G400, G600, G800 and VSP F400, F600, F800

This option is supported by global-active device only.

Specify this option to change the I/O mode of a pair volume.

You can specify the following values for <mode>.

- local: Change the I/O mode to local.

- block: Change the I/O mode to block.

The pair volume whose I/O mode is to be changed is the pair volume described in the configuration definition file of the instance used to issue the command, among the pair volumes of the spedified pair.

In the following example, change the I/O mode of the pair volume described in the configuration definition file of "instance: 100" to Local, among the pairs in "device group: mygroup".

Example:

```
pairsplit -g mygroup -iomd local -IH100
```

Returned values

The `pairsplit` command sets the following returned values during exit allowing you to check the execution results.

- **Normal termination:**

  - **0:** When splitting groups, 0 = normal termination for all pairs.

- **Abnormal termination:**

  - **other than 0:** Refer to the error codes for error details.

Error codes

The following table lists and describes the error codes for the `pairsplit` command. Unrecoverable errors are not resolved, even after re-executing the command. If the command failed, the detailed status is logged in the CCI command log ($HORCC_LOG).

📄 **Note:** When an option that is valid for only specific volumes is specified for other volumes, the error code might be EX_UNWOPT or EX_UNWCMD. Check whether the specified option is valid for the volume.

Category	Error Code	Error Message	Recommended Action	Value
Volume status (Unrecoverable)	EX_ENQVOL	Unmatched volume status within the group	Confirm status using the `pairdisplay` command. Make sure all volumes in the group have the same fence level and volume attributes.	236

Chapter 2: Data management commands

Category	Error Code	Error Message	Recommended Action	Value
	EX_INCSTG	Inconsistent status in group	Confirm pair status using **pairdisplay**.	229
	EX_INVVOL	Invalid volume status	Confirm volume status/attribute using pairdisplay -l or 'raidvchkdsp -v aou'. "Aou" (allocation on use) refers to dynamic provisioning.	222
	EX_EVOLCE	Pair volume combination error	Confirm pair status using **pairdisplay**, and change combination of volumes.	235
	EX_INVSTP	Invalid pair status	Confirm pair status using **pairdisplay**.	228
	EX_EWSUSE	Pair suspended at WAIT state	Issue **pairresync** manually to the identified failed paired volume to try to recover it. If the trouble persists, call the customer support.	234
Timer (Recoverable)	EX_EWSTOT	Timeout waiting for specified status	Even if timeout occurs, volume pairs might be being split or deleted. Confirm pair status using pairdisplay.	233

Examples

See the following figure for example of `-FHORC` option for **pairsplit** command.
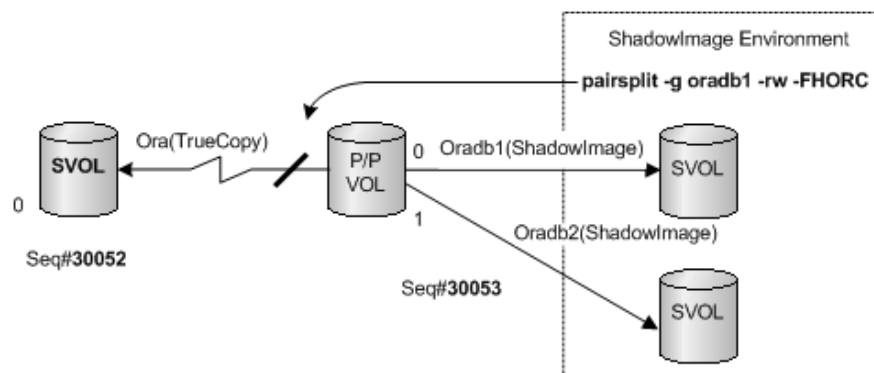
See the following figure for example of `-FMRCF` option for **pairsplit** command.
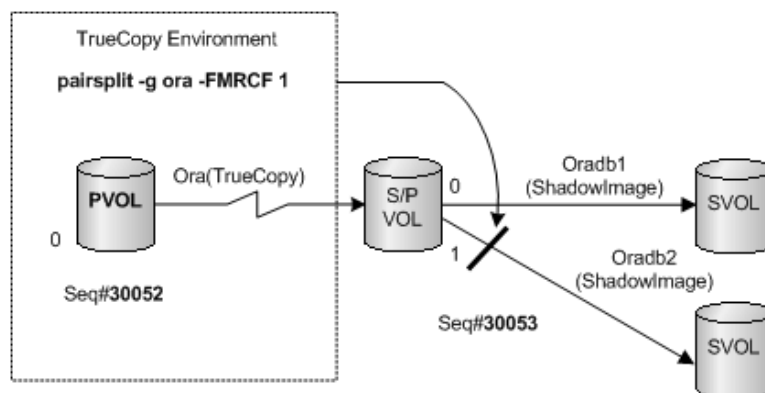
pairresync*

The **pairresync** command re-establishes a split pair volume and then restarts the update copy operations to the secondary volume. The **pairresync** command can resynchronize either a paired logical volume or a group of paired volumes. The normal direction of resynchronization is from the primary volume to the secondary volume. If the −restore option is specified (ShadowImage/Copy-on-Write Snapshot only), the pair is resynchronized in the reverse direction (that is, secondary volume to primary volume). The primary volume remains accessible during **pairresync**, except when the −restore option is specified. The secondary volume becomes write-disabled when the **pairresync** command is issued.

The **pairresync** command cannot be used for Volume Migration.

The **pairresync** command terminates before resynchronization of the secondary (or primary) volume is complete. Use the pair event waiting (pairevtwait) or **pairdisplay** command to verify that the resync operation completed successfully (status changes from COPY to PAIR). The execution log file also shows completion of the resync operation. The status transition of the paired volume is judged by the status of the primary volume. The fence level is not changed (only for TrueCopy, TrueCopy Async, Universal Replicator, or global-active device).

If no data was written to the secondary volume while the pair was split, the differential data on the primary volume is copied. If data was written to the secondary volume, the differential data on the primary volume and secondary volume is copied. This process is reversed when the ShadowImage −restore option is specified.

Before issuing the **pairresync** command (normal or reverse direction), make sure that the secondary volume is not mounted on any UNIX system. Before issuing a reverse **pairresync** command, make sure that the primary volume is not mounted on any UNIX system.

When you use TrueCopy, TrueCopy for Mainframe, TrueCopy Async, Hitachi TrueCopy Asynchronous for Mainframe, Universal Replicator, Universal Replicator for Mainframe, or global-active device, if a failure that requires maintenance operations, such as blocking the controller board or the cache memory, occurs on the primary storage system, the **paircreate** command cannot be executed.

When you use global-active device, if a failure that requires maintenance operations occurs on the secondary storage system, the **paircreate** command cannot be executed as well.

Note on Quick Resync/Restore:

If the '$HORCC_RSYN=QUICK' or the '$HORCC_REST=QUICK' environment variable is set for any of the following systems, the **pairresync** operation is performed as Quick Resync regardless of the system option mode 87/80 setting via SVP:

- VSP 5000 series

- VSP G1x00 and VSP F1500

- VSP E series

- VSP G130, G/F350, G/F370, G/F700, G/F900

- VSP G200, G400, G600, G800 and VSP F400, F600, F800

- VSP

- HUS VM

The $HORCC_RSYN and $HORCC_REST environment variables are ignored for 9900V.

TrueCopy, TrueCopy Async, Universal Replicator, and GAD only: The swaps(p) option is used to swap volumes from the S-VOL (P-VOL) to P-VOL (S-VOL) when the S-VOL (P-VOL) is suspended, and resynchronize the new S-VOL based on the new P-VOL. As a result of this operation, the volume attributes of the local host become the attributes of the new P-VOL (S-VOL).

The `swaps(p)` option:

Ignores the `-l` option.

If `-c size` option is omitted, use 3 (default value of the number of copy tracks (-c size)).

Executes the command when the pair status is PAIR including PSUS or PSUE (not applicable to COPY and SMPL).

If the target volume (remote volume for `-swapp`) is already the P-VOL, the pair operation is skipped.

The following figure shows the pair resynchronization operation.
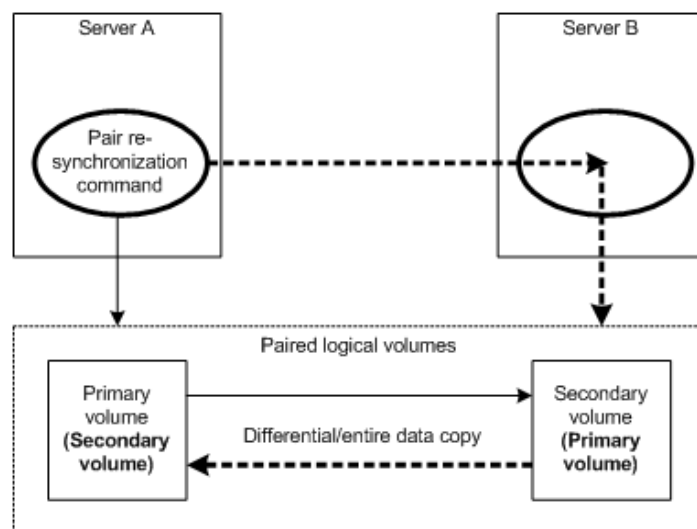
The following figure shows the Normal Resync and ShadowImage Restore Resync.

Chapter 2: Data management commands
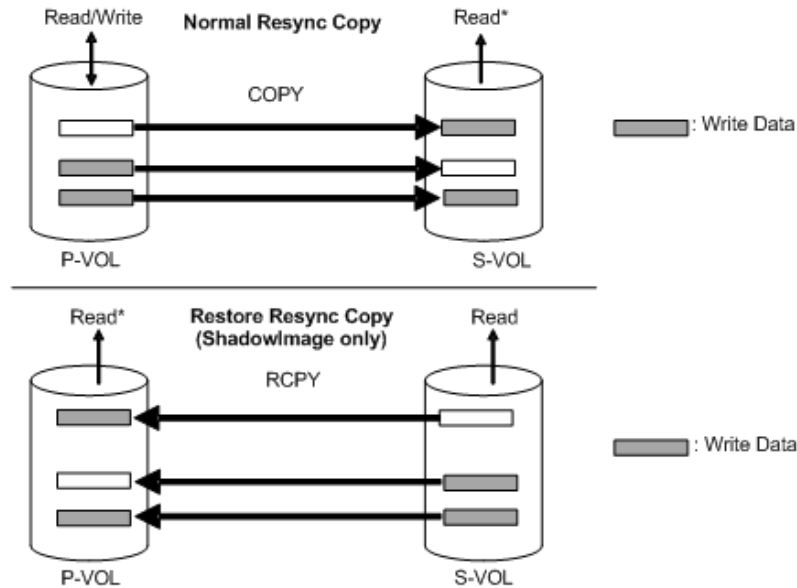
*An option for creating a pair. When you specify `-m noread`, you cannot read date in the volume.

Syntax

```
pairresync { -h | -q | -z[x] | -I[H][M][instance#] or -I[TC][SI][instance#]| -g
<group> | -d <pair Vol>| -d[g] <raw_device> [MU#][-n] | -FHORC [MU#] | -FMRCF [MU#]| -
d[g] <seq#> <LDEV#> [MU#] |-c <size> | -nomsg | -l | -restore| -swaps | -swapp | -fq
<mode>| -cto <o-time> [c-time] [r-time]| -f[g] <fence> [CTG ID] | -pr <io
preference> }
```

Options and parameters

The primary volume's server is automatically detected by the **pairresync** command, so the server does not need to be specified in the **pairresync** command options.

-h

Displays help/usage and version information.

-q

Terminates the interactive mode and exits this command.

-z or -zx

Makes the **pairresync** command enter the interactive mode. The `-zx` option guards performing of the HORCM in the interactive mode. When this option detects a HORCM shut down, interactive mode terminates.

OpenVMS cannot use the `-zx` option.

-I [H] [M] [instance#] or -I [TC] [SI] [instance#]

Specifies the command as [HORC]/[HOMRCF] or [TrueCopy]/[ShadowImage], and used when specifying the CCI instance number.

-g <group>

> Used to specify a group name defined in the configuration definition file. The command is executed for the specified group unless the `-d <pair Vol>` option is specified.

-d <pair Vol>

> Specifies a paired logical volume name defined in the configuration definition file. When this option is specified, the command is executed for the specified paired logical volumes. This option is valid when the `-g <group>` option is specified.

-d[g] <raw_device> [MU#][-n]

> Searches whether the specified raw_device is included in a group on the configuration definition file (local instance). If it is in the group, the target volume is executed as the paired logical volume (-d) or group (-dg). This option is effective without specification of `-g <group>` option. If the specified raw_device is contained in two or more groups, the command is executed for the first group.
>
> This `-n` option is specified to use a device file for NVMe-oF. Specifying this option enables the use of both device files for NVMe-oF and SCSI.
>
> If the HORCC_NVME environment variable is specified, both device files for NVMe-oF and SCSI can be used without the `-n` option specified.
>
> To specify this option, a command device must be enabled in a namespace on the storage system. For more information about whether a command device can be enabled in a namespace on the storage system, see the *Provisioning Guide* for your storage system.

-d[g] <seq#> <LDEV#> [MU#]

> Searches whether the specified LDEV is included in a group on the configuration definition file (local instance). If it is in the group, the target volume is executed as the paired logical volume (-d) or group (-dg). This option is effective without specification of '-g <group>' option. If the specified LDEV is contained in two or more groups, the command is executed on the first group. The <seq#> <LDEV#> values can be specified in hexadecimal (add the '0x' prefix) or decimal notation.
>
> - When specifying the serial number for VSP 5000 series, add a "5" at the beginning of the serial number. For example, for serial number 12345, enter 512345.
>
> - When specifying the serial number for VSP G1x00 and VSP F1500, add a "3" at the beginning of the serial number. For example, for serial number 12345, enter 312345.

-FHORC [MU#] or -FCA [MU#]

> Forcibly specifies a cascaded remote copy volume for specified pair logical volumes on local copy environment. If the `-l` option is specified, this option resyncs a cascaded remote copy volume on a local host (near site). If no `-l` option is specified, this option resyncs a cascaded remote copy volume on a remote host (far site). The target remote copy volume must be a P-VOL. The `-swapp` option cannot be specified.

-FMRCF [MU#] or -FBC [MU#]

Forcibly specifies a cascaded local copy volume for specified volume pair on remote copy environment. If the `-l` option is specified, a cascaded local copy volume is split on a local host (near site). If `-l` option is not specified, a cascaded local copy volume is split on a remote host (far site). The target local copy volume must be a P-VOL.

-swaps -FHORC [MU#]

Swaps the cascaded TrueCopy or Universal Replicator volume from the primary node for failback.

In failback operation from a 3DC cascaded site failure, if you want to failback to DC1 from DC3 directly, it is necessary to operate all cascaded volume from DC1.

In order to make this operation possible, CCI supports the '**pairresync** -swaps -FHORC' option that swaps Universal Replicator volumes on the cascaded TrueCopy Sync/Universal Replicator volume.

-c <size>

TrueCopy, TrueCopy for Mainframe, TrueCopy Async, TrueCopy Async for Mainframe, ShadowImage, and global-active device only.

Specifies the copy pace for the resync operation (range = 1 to 15 tracks). If not specified, the value used for paircreate or **pairsplit** command is used. However, when you use the `-swaps` or `-swapp` option, the default value is 3.

The relationship between the track size and the copy pace is shown as follows. If the copy pace is fast, the host I/O performance might be degraded. To reduce the effect on host I/O performance, lower the copy pace.

- TrueCopy, TrueCopy Async, and global-active device

  - When the track size is 1, the copy pace is slow.

  - When the track size is 2 or 3, the copy pace is medium.

  - When the track size is 4 or more, the copy pace is fast.

- ShadowImage

  - When the track size is 1 or 2, the copy pace is slow.

  - When the track size is 3, the copy pace is medium.

  - When the track size is 4 or more, the copy pace is fast.

- TrueCopy for Mainframe, TrueCopy Async for Mainframe

  - When the track size is 1 to 3, the number of the tracks copied simultaneously is 3 and the copy pace is slow.

  - When the track size is 4 or more, the number of the tracks copied simultaneously is 15 and the copy pace is fast.

> **Note:** This option is disabled in Universal Replicator, Universal Replicator for Mainframe, or ShadowImage for Mainframe. To change the copy pace in Universal Replicator or Universal Replicator for Mainframe, use the `raidcom modify journal` command for JNL option changes, or Device Manager - Storage Navigator. ShadowImage for Mainframe works only at medium pace.

-nomsg

Suppresses messages to be displayed when this command is executed. It is used to execute this command from a user program. This option must be specified at the beginning of a command argument. The command execution log is not affected by this option.

-l

When the remote host cannot be used for host down, this option enables the `pairresync` operation on the local host only. The target volume of the local host must be P-VOL. (For ShadowImage and Copy-on-Write Snapshot, only S-VOLs can be resynchronized.)

If you specify the `-swapp` option, this option becomes unsupported and cannot be specified at the same time.

-restore

ShadowImage and Copy-on-Write Snapshot only.

Performs reverse resync (from secondary volume to primary volume).

-swaps

TrueCopy, TrueCopy Async, Universal Replicator, and global-active device only.

Executed from the S-VOL side when there is no host on the P-VOL side to support. Typically executed in PSUS that facilitates 'fast failback' without requiring a full copy. Both `-swaps` and `-swapp` are copied to the original P-VOL based on the delta data from the original S-VOL, and then the S-VOL and the P-VOL are swapped.

-swapp

TrueCopy, TrueCopy Async, Universal Replicator, and global-active device only.

Executes the the same operation as `-swaps` from the original P-VOL, however, `-swapp` does require the cooperation of hosts at both sides.

-fq <mode>

ShadowImage and Copy-on-Write Snapshot only.

Specifies the mode whether **pairresync** or the `-restore` option is performed as 'QUICK'.

mode = normal: **pairresync** is performed as Non quick mode regardless of setting of $HORCC_RSYN environment variable and/or the system option mode 87 via SVP.

mode = quick: **pairresync** is performed as Quick Resync regardless of setting of $HORCC_RSYN environment variable and/or the system option mode 87 via SVP.

If this option is not specified, then performing **pairresync** is determined by the $HORCC_RSYN environment variable and/or the system option mode setting through the SVP regardless of whether the **pairresync** operation is Quick Split or not.

The following table shows the relationship between the `-fq` option and the $HORCC_RSYN environment variable.

-fq option	$HORCC_RSYN	Behavior
quick	Invalid	Quick Resync
normal	Invalid	Normal Copy
Omitted	QUICK	Quick Resync
Omitted	NORMAL	Normal Copy
Omitted	Omitted	▪ For VSP 5000 series, VSP E series, VSP G130, G/F350, G/F370, G/F700, G/F900: Normal Copy ▪ For other models: Determined by system option mode 87

In the case of Restore (-restore is specified):

mode = normal

The **`pairresync -restore`** command is performed as Non quick mode regardless of the setting of the $HORCC_REST environment variable and/or the system option mode 80 via SVP.

mode = quick

The **`pairresync -restore`** command is performed as Quick Restore regardless of the setting of the $HORCC_REST environment variable and/or the system option mode 80 via SVP.

If this option is not specified, then the performing of the **`pairresync`** is determined by the $HORCC_REST environment variable and/or the system option mode setting through the SVP regardless of whether the **`pairresync`** operation is Quick Restore or not.

The relationship between '-fq' option and $HORCC_REST are shown as below.

-fq option	$HORCC_REST	Behavior
quick	Invalid	Quick Restore
normal	Invalid	Normal Restore
Omitted	QUICK	Quick Restore
Omitted	NORMAL	Normal Restore
Omitted	Omitted	▪ For VSP 5000 series, VSP E series, VSP G130, G/F350, G/F370, G/F700, G/F900: Quick Restore ▪ For other models: Determined by system option mode 80

> **Note:**
> - This `-fq` option is also valid for TrueCopy-TrueCopy or ShadowImage cascaded operation using '-F BC [MU#]' option.
> - The `-fq` option is applied to the following storage systems:
>   - VSP 5000 series
>   - VSP G1x00 and VSP F1500
>   - VSP E series
>   - VSP G130, G/F350, G/F370, G/F700, G/F900
>   - VSP G200, G400, G600, G800 and VSP F400, F600, F800
>   - VSP
>   - HUS VM
>
>   This option is ignored for 9900V to maintain the compatibility on 9900V so that you can add this option to the same script.
> - If this option is combined with '-restore' option in VSP during the maintenance work in the storage system (SVP is in modify mode), this operation cannot be completed.

-cto <o-time> [c-time] [r-time]
    TrueCopy Async and Universal Replicator only.

If you specify '-cto <o-time> [c-time] [r-time]' option on TrueCopy sync, then it is ignored.

- `o-time`: Sets the offloading timer for controlling inflow of write I/O to the specified consistency group. For TrueCopy Async, specify an integer from 1 to 255 (seconds). If this option is not specified, 90 seconds is set by default. For Universal Replicator, specify an integer from 1 to 255 (seconds). If this option is not specified, the value set for the journal on the P-VOL side is applied. Use the **raidcom modify journal** command to change the value range to 256 to 600. If o-time=0 is specified, the inflow control of write I/Os is disabled. When the sidefile capacity exceeds the limit of the sidefile area, write I/Os from the host wait, within the specified timeout period, until the space which is large enough to store next new data becomes available. As the timeout period, you can specify a value from 1 through 255 (seconds) for TrueCopy Async, 1 through 600 (seconds) for Universal Replicator. The default timeout values are 90 seconds for TrueCopy Async, and 60 seconds for Universal Replicator. If the timeout occurs during this waiting state then pair status changes from PAIR to PSUS state of sidefile (Journal) Full, and its host side Write I/O is continued and data is managed by BITMAP mode. Therefore the o-time timeout value should be less than the I/O timeout value of the host system.

- `[c-time]`: (TrueCopy Async only) Sets the Copy Pending timer for the specified consistency group. c-time can be specified from 1 to 15 minutes in increments of 1 minute. If this option is not specified, then this value is set as below.

  - If a consistency group is created, then 5 minutes is set as the default. If not, it is not changed.

- `[r-time]`: (TrueCopy Async only) Sets the RCU Ready timer for the specified consistency group. r-time can be specified from 1 to 10 minutes in increments of 1 minute. If this option is not specified, then the value is not changed.

Chapter 2: Data management commands

> 📄 **Note:**
>
> For TrueCopy Async, settings changed by these options are invalid if a consistency group already exists. These parameters are also forwarded to S-VOL journal with **`pairresync`** command, and are used when S-VOL is changed to P-VOL. These parameters are maintained and become valid until and when the pair-volumes are changed to SMPL.
>
> For Universal Replicator, these parameters can be set and changed when a pair is resynchronized in a journal volume in the P(S)JSN status. The parameters are forwarded to the S-VOL side and set when the pair is resynchronized. To change settings of journals on the S-VOL side, specify as follows:
>
> 1. `pairsplit -g <group>`
> 2. `pairresync -g <group> -cto <o-time>`
>
> To change settings of journals on the P-VOL side, swap them from the P-VOL to the S-VOL before specifying the above. These parameters are maintained on each journal. Therefore, if you set the value for offloading timer, execute the **`raidcom modify journal`** command on both P-VOL and S-VOL sides.

-f[g] \<fence\> [CTG ID]

TrueCopy, TrueCopyAsync, Universal Replicator, and global-active device only.

`-f[g] <fence> [CTG ID]` (TrueCopy only): This option is used to change from existing TrueCopy Sync volumes to TrueCopy Sync CTG without deleting paired-volume. It is possible to change the option (fence level) and the volume attribute as shown below. This option is valid without '-swaps' or '-swapp' option.

Changing the option (`fence level`) and the volume attribute:

Attribute	Options		
	-f fence	**-fg fence**	**-fg fence CTG ID**
Sync	Updates fence	Cmd rejected Abnormal term.	Changes to Sync CTG Updates fence*
Sync CTG	Changes to Sync Updates fence	Updates fence	Cmd rejected Abnormal term.
Async	Cmd rejected Abnormal term.	Cmd rejected Abnormal term.	Cmd rejected Abnormal term.
Explanation of terms: fence: data, status, or never ('async' cannot be specified) Cmd rejected/Abnormal term.: Command execution is rejected, and the operation terminates abnormally.			

Chapter 2: Data management commands

Attribute	Options		
	-f fence	**-fg fence**	**-fg fence CTG ID**
*If the identical CTG ID is specified with the different group name of CCI, then the command is rejected. So different CTG ID must be specified for the different group name.			

Changing the volume attribute for global-active device:

Attribute	Options		
	-f never	**-fg never**	**-fg never CTG ID**
GAD	No Updates	Cmd rejected Abnormal term.	Changes to GAD CTG
GAD CTG	Cmd rejected Abnormal term.	No Updates	Cmd rejected Abnormal term.
Explanation of terms: Cmd rejected/Abnormal term.: Command execution is rejected, and the operation terminates abnormally.			

-pr <io preference>

This option is only available for global-active device.

Set the I/O preference mode (I/O preference mode when remote path failed) when a failure occurs in the remote path between the storage systems of the primary site and the secondary site and communication becomes impossible. For details of the I/O preference mode when remote path failed, see the Global-Active Device User Guide.

The following values can be specified for <io preference>.

pvol: Primary Volume. When remote path failed, the primary volume continues host I/O preferentially.

disable: I/O Preference Mode When Remote Path Failed is disabled. Specify this value to disable I/O Preference Mode When Remote Path Failed.

If this option is omitted, the I/O Preference Mode When Remote Path Failed retains the current setting.

Pairs with different settings of I/O preference mode when remote path failed cannot be mixed in the same consistency group. Therefore, if you want to change the I/O preference mode when remote path failed of a pair that already belongs to a consistency group, make the change for the entire consistency group. If you do not specify the -d option, you can change the I/O preference mode when remote path failed for the entire consistency group.

Returned values

The **pairresync** command sets the following returned values during exit allowing you to check the execution results.

▪ **Normal termination:**

• **0:** When resynchronizing groups, 0 = normal termination for all pairs.

▪ **Abnormal termination:**

• **other than 0:** Refer to the error code for error details.

Error codes

Unrecoverable errors are fixed and not resolved, even after re-executing the command. If the command failed, the detailed status is logged in the CCI command log ($HORCC_LOG), even if the user script has no error handling.

The **pairresync** command updates both the P-VOL and the S-VOL. When you execute this command, make sure that neither the P-VOL nor the S-VOL is mounted. See the following table for specific error codes for **pairresync**.

📄 **Note:** When an option that is valid for only specific volumes is specified for other volumes, the error code might be EX_UNWOPT or EX_UNWCMD. Check whether the specified option is valid for the volume.

Category	Error Code	Error Message	Recommended Action	Value
Volume status (Unrecoverable)	EX_ENQVOL	Unmatched volume status within the group	Confirm status using the **pairdisplay** command. Make sure all volumes in the group have the same fence level and volume attributes.	236
	EX_INCSTG	Inconsistent status in group	Confirm pair status using **pairdisplay**.	229
	EX_INVVOL	Invalid volume status	Confirm pair status using **pairdisplay -l**.	222
	EX_INVSTP	Invalid pair status	Confirm pair status using **pairdisplay**.	228

Examples

See the following figure as an example of -FHORC option for pairresync
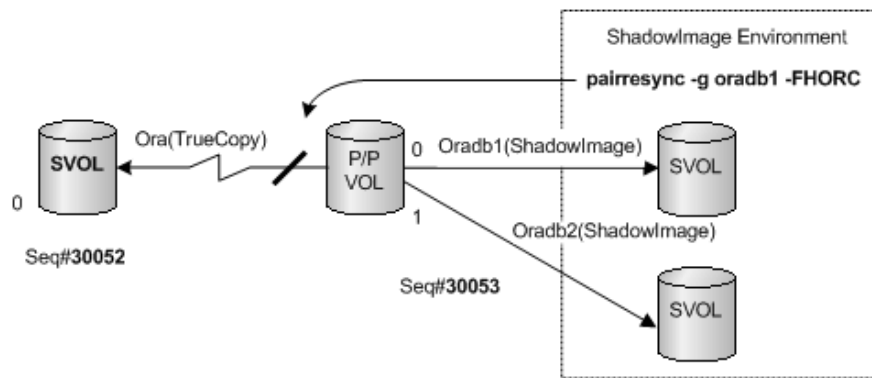
See the following figure as an example of −FMRCF option for **pairresync**
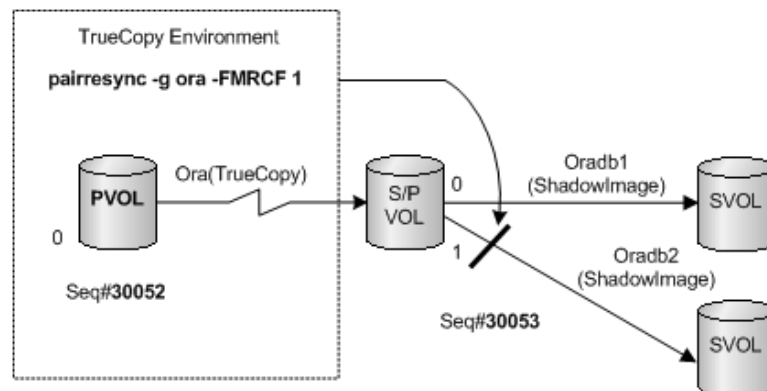
See the following figure for swap operation.
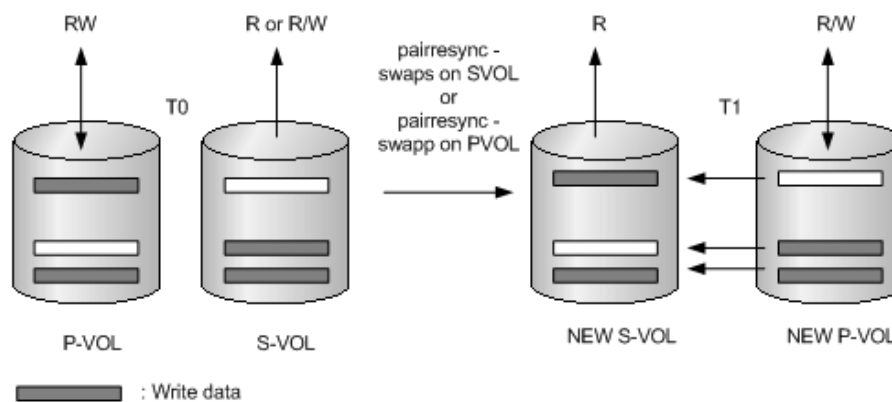
See the following figure as an example of -swaps option with -FHORC [MU#]

Chapter 2: Data management commands

PVOL
L1 0
DC1 2
SMPL

G1 (Sync)

SSW SL
0
L2
SVOL 1
DC2

G2 (UR)

PVOL 1
2
SMPL
L3
DC3

APP1

After DC1 recovery, Failback from DC3

pairresync -g G1 -FHORC 1-
pairresync -g G1 -swapp

APP1

PVOL
L1 0
DC1 2
SMPL

G1 (Sync)

SVOL
0
L2
PVOL 1
DC2

G2 (UR)

SVOL 1
2
SMPL
L3
DC3
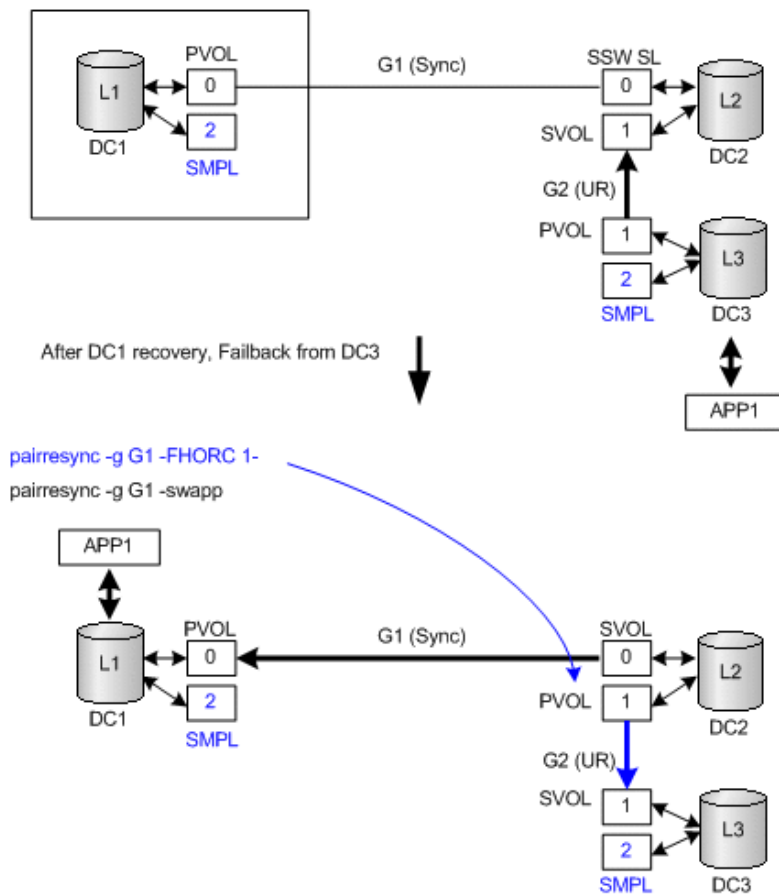
pairevtwait*

The **pairevtwait** (pair event waiting) command is used to wait for another command to complete pair operations and to confirm the result of pair operations (see the following figure for pair even waiting operation.). It waits ('sleeps') until the paired volume status becomes identical to a specified status, and terminates abnormally if an abnormal status is detected. The transition of the paired volume status is judged by the status of the volume, which is searched automatically.

The **pairevtwait** command can be executed by specifying a paired logical volume or a group of a paired volume. If the **pairevtwait** command is issued by specifying a group and the status, the command waits until the status of each volume in the group changes to the specified status. When the **pairevtwait** command with the -nowait or -nowaits option is issued for a group, the **pairevtwait** command returns the response immediately according to the status of each volume in the group. The primary and secondary volume's servers are automatically detected by the **pairevtwait** command, so the server does not need to be specified in the **pairevtwait** command parameters. For ShadowImage pairs, this command must be used to confirm a pair status transition.
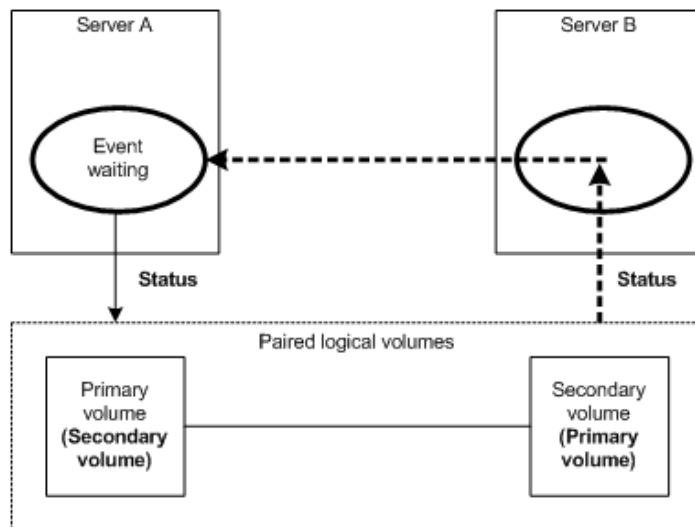
Syntax

```
pairevtwait{ -h | -q | -z[x] | -I[H][M][instance#] or
     -I[TC][SI][instance#] | -g <group> | -d <pair Vol>
     | -d[g] <raw_device> [MU#][-n] | -FHORC [MU#] | -FMRCF [MU#]
     | -d[g] <seq#> <LDEV#> [MU#] | -s [s] <status> ...
     | -t <timeout>[interval] | -nowait[s] | -l | -nomsg }
```

Options and parameters

-h

> Displays help/usage and version information.

-q

> Terminates the interactive mode and exits this command.

-z or -zx

> Makes the **pairevtwait** command enter the interactive mode. The -zx option guards performing of the HORCM in the interactive mode. When this option detects a HORCM shut down, interactive mode terminates.

> OpenVMS cannot use the -zx option.

-I[H][M] [instance#] or -I[TC][SI] [instance#]

> Specifies the command as [HORC]/[HOMRCF] or [TrueCopy]/[ShadowImage], and used when specifying the CCI instance number.

-g <group>

> Specifies a group name defined in the configuration definition file. The command is executed for the specified group unless the -d <pair Vol> option is specified.

-d <pair Vol>

> Specifies a paired logical volume name defined in the configuration definition file. When this option is specified, the command is executed for the specified paired logical volumes. This option is valid when the -g <group> option is specified.

Chapter 2: Data management commands

-d[g] <raw_device> [MU#][-n]

> Searches whether the specified raw_device is included in a group on the configuration definition file (local instance). If it is in the group, the target volume is executed as the paired logical volume (-d) or group (-dg). This option is effective without specification of `-g <group>` option. If the specified raw_device is contained in two or more groups, the command is executed for the first group.

> This `-n` option is specified to use a device file for NVMe-oF. Specifying this option enables the use of both device files for NVMe-oF and SCSI.

> If the HORCC_NVME environment variable is specified, both device files for NVMe-oF and SCSI can be used without the `-n` option specified.

> To specify this option, a command device must be enabled in a namespace on the storage system. For more information about whether a command device can be enabled in a namespace on the storage system, see the *Provisioning Guide* for your storage system.

-FHORC [MU#] or -FCA [MU#]

> Forcibly specifies a cascaded remote copy volume for specified pair logical volumes on local copy environment. If the `-l` option is specified, this option tests status of a cascaded remote copy volume on a local host (near site). If no `-l` option is specified, this option tests status of a cascaded remote copy volume on a remote host (far site).

-FMRCF [MU#] or -FBC [MU#]

> Forcibly specifies a cascaded local copy volume for specified pair logical volumes on remote copy environment. If the `-l` option is specified, this option tests status of a cascaded local copy volume on a local host (near site). If no `-l` option is specified, this option tests status of a cascaded local copy volume on a remote host (far site).

-d[g] <seq#> <LDEV#> [MU#]

> Searches whether the specified LDEV is included in a group on the configuration definition file (local instance). If it is in the group, the target volume is executed as the paired logical volume (-d) or group (-dg). This option is effective without specification of '-g <group>' option. If the specified LDEV is contained in two or more groups, the command is executed on the first group. The <seq#> <LDEV#> values can be specified in hexadecimal (add the '0x' prefix) or decimal notation.

> **Note:**
> - When specifying the serial number for VSP 5000 series, add a "5" at the beginning of the serial number. For example, for serial number 12345, enter 512345.
> - When specifying the serial number for VSP G1x00 and VSP F1500, add a "3" at the beginning of the serial number. For example, for serial number 12345, enter 312345.

-s <status>

> Specifies the waiting status, which is 'smpl', 'copy(including 'RCPY')', 'pair', 'psus', or 'psue'. If two or more statuses are specified following -s, waiting is done according to the logical OR of the specified statuses. This option is valid when the `-nowait` option is not specified.

-ss <status>

Specifies the waiting status, which is 'smpl', 'copy'('RCPY' is included), 'pair', 'ssus','psue' on S-VOL. If two or more statuses are specified following -ss, waiting is done according to the logical OR of the specified statuses. This option is valid when the `-nowaits` option is not specified.

-t <timeout> [interval]

Specifies the interval of monitoring a status specified using the `-s` and `-ss` option and the time-out period in increments of 1 second. Unless [interval] is specified, the default value (3 seconds) is used. This option is valid when the `-nowait` option is not specified. If <timeout> is specified more than 1999999, then 'WARNING' message appears.

If you execute the command in the Out-of-Band method, specify 3 seconds (default value) or more to <timeout>.

This option is valid only when the -nowait or `-nowaits` option is not specified.

-nowait

When this option is specified, the pair status at that time is reported without waiting. The pair status is set as a returned value for this command. When this option is specified, the `-t` and `-s` options are not needed.

-nowaits

When this option is specified, the pairing status on S-VOL at that time is reported without waiting. The pairing status is set as a returned value for this command. When this option is specified, the -t and -ss options are not needed.

-l

When this command cannot utilize a remote host for host down, this option executes this command by a local host only. When a paired volume used by this command meets any of the following conditions, if this option is not specified, the command fails with EX_EVOLCE:

- Both paired volumes are P-VOLs

- Both paired volumes are S-VOLs

- Either paired volume is SMPL, and the other volume is S-VOL.

-nomsg

Suppresses messages to be displayed when this command is executed. It is used to execute a command from a user program. This option must be specified at the beginning of a command argument. The command execution log is not affected by this option.

Returned values

The **pairevtwait** command sets the following returned values during exit allowing you to check the execution results. See the table below for more details.

▪ When the -nowait option is specified:

  • **Normal termination:**

    ▪ **1:** The status is SMPL.

    ▪ **2:** The status is COPY or RCPY.

    ▪ **3:** The status is PAIR.

    ▪ **4:** The status is PSUS.

    ▪ **5:** The status is PSUE.

  • When monitoring groups, 1/2/3/4/5 = normal termination for all pairs.

  • **Abnormal termination:**

    ▪ other than 0 to 127: refer to the error codes for error details.

  • The returned value of SSUS is set for PFUS.

  • The returned value of PAIR is set for PFUL.

▪ When the -nowaits option is specified:

  • **Normal termination:**

    ▪ **1:** The status is SMPL.

    ▪ **2:** The status is COPY or RCPY.

    ▪ **3:** The status is PAIR.

    ▪ **4:** The status is SSUS (Note that SVOL_PSUS appears as SSUS).

    ▪ **5:** The status is PSUE.

  • **Abnormal termination:**

    ▪ other than 0 to 127: refer to the error codes for error details.

  • The returned value of SSUS is set for SSWS/PFUS.

  • The returned value of PAIR is set for PFUL.

▪ When the -nowait and/or nowaits option is not specified:

  • **Normal termination:**

    ▪ **0:** When monitoring groups, 0 = normal termination for all pairs.

  • **Abnormal termination:**

    ▪ **other than 0 to 127:** refer to the error codes for error details.

Error codes

Unrecoverable errors are fixed and not resolved, even after re-executing the command. If the command failed, the detailed status is logged in the CCI command log ($HORCC_LOG), even if the user script has no error handling.

See the table below for specific error codes for `pairevtwait`.

Category	Error Code	Error Message	Recommended Action	Value
Volume status (Unrecoverable)	EX_ENQVOL	Unmatched volume status within the group	Confirm status using the `pairdisplay` command. Make sure all volumes in the group have the same fence level and volume attributes.	236
	EX_INCSTG	Inconsistent status in group	Confirm pair status using `pairdisplay`.	229
	EX_INVVOL	Invalid volume status	Confirm pair status using `pairdisplay -l`.	222
	EX_EVOLCE	Pair volume combination error	Confirm pair status using `pairdisplay`, and change combination of volumes.	235
	EX_EWSUSE	Pair suspended at WAIT state	Issue `pairresync` manually to the identified failed paired volume to try to recover it. If the trouble persists, call the customer support	234
Timer (Recoverable)	EX_EWSTOT	Timeout waiting for specified status	Increase timeout value using the `-t` option.	233
	EX_EWSLTO	Timeout waiting for specified status on the local host	Confirm that CCI (HORCM) on the remote host is running.	232

Examples

See the following figure as an example of `-FHORC` option for `pairevtwait`.
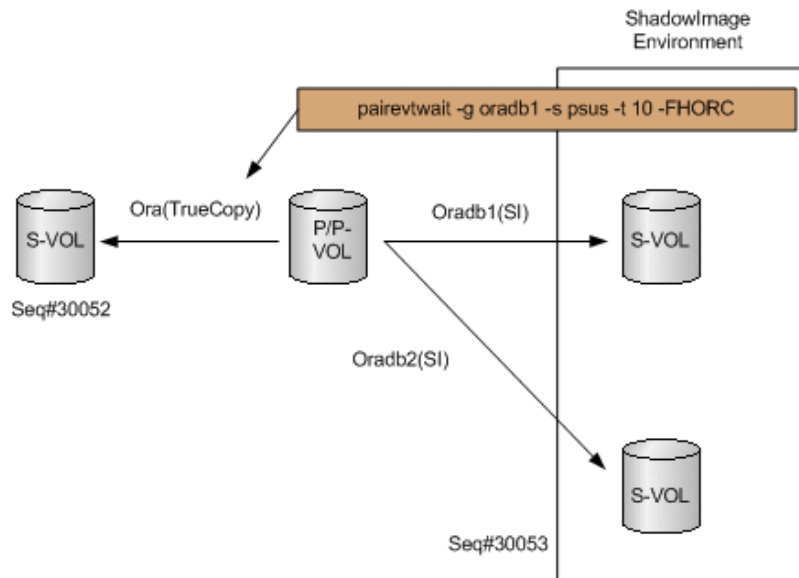
See the following figure as an example of −FMRCF option for **pairevtwait**.
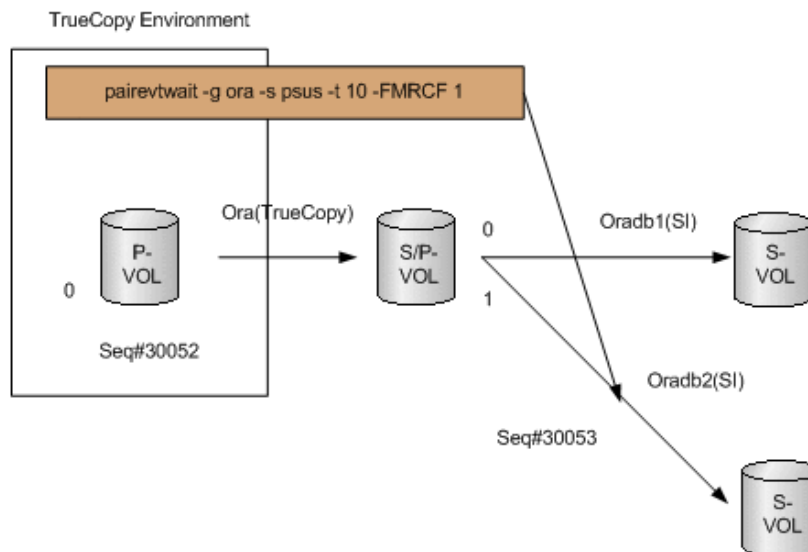
Using the -ss \<status\> ... and -nowaits options.

In PVOL_PSUS & SVOL_COPY state of ShadowImage quick mode, **pairevtwait** returns immediately even if the S-VOL is still in SVOL_COPY state because P-VOL is already in PVOL_PSUS state. If you want to wait the SVOL_SSUS state, then use the −ss \<status\> and −nowaits options to wait for the pair status on S-VOL side. This is needed for operating **pairresync −restore** or **pairsplit −S**.

The figure below shows an example of waiting on ShadowImage. The figure depicts five examples of waiting until 'PVOL_PSUS' & 'SVOL_COPY' state is changed to SVOL_SSUS.

Chapter 2: Data management commands

Pairevtwait -g G1 -ss ssus -t 600
Wait on S-VOL in communication with local and remote
Pairevtwait -g G1 -ss ssus -FHOMRCF -t 600
Wait on S-VOL in communication with remote only
Pairevtwait -g G1 -ss ssus -l -t 600
Wait on P-VOL by finding from P-VOL to S-VOL

Pairevtwait -g G1 -ss ssus -t 600
Wait on S-VOL in communication with local and remote
Pairevtwait -g G1 -ss ssus -l -t 600
Wait on S-VOL directly

The **horctakeover** command suspends G2(CA-Jnl) automatically if **horctakeover** returns 'Swap-takeover' as an exit code. In DC1 host failure, if APP1 wants to wait until DC3 becomes the suspend state, then they can verify the 'SSUS' state by using the **pairevtwait** command as shown below.

The following figure shows an example for waiting 'SSUS' on 3DC using TC/UR
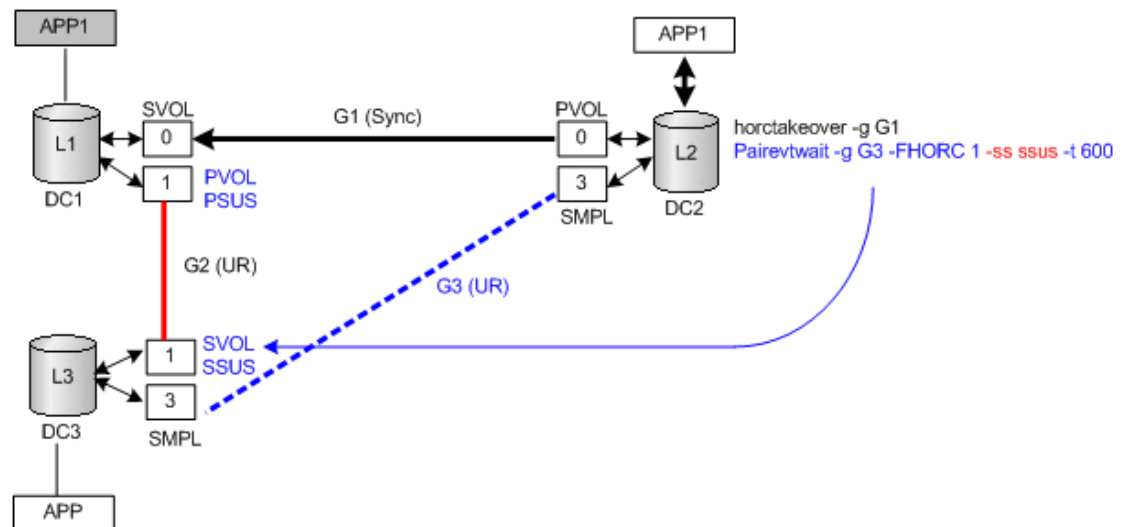
pairmon*

The **pairmon** (pair monitor) command, which is connected to the HORCM daemon, obtains the pair status transition of each volume pair and reports the status change. If the pair status changes (due to an error or a user-specified command), the **pairmon** command displays a message.

The pair status transition events exist in the HORCM pair state transfer queue. The `-resevt` option (reset event) deletes one/all events from the HORCM pair state transfer queue. If reset event is not specified, the pair state transfer queue is maintained. If the `-s` option is not specified, **pairmon** displays all events for which it receives information from HORCM. If the `-s` option is specified, only the specified status transitions are displayed.

The CCI software supports the error monitoring and configuration confirmation commands for linkage with the system operation management of the UNIX server.

Syntax

```
pairmon { -h | -q | -z[x] | -I[H][M][instance#] or
    -I[TC][SI][instance#] | -D | -allsnd | -resevt | -nowait
    | -s <status> ... }
```

Options and parameters

-h

> Displays help/usage and version information.

-q

> Terminates the interactive mode and exits this command.

-z or -zx

> Makes the **pairmon** command enter the interactive mode. The -zx option guards performing of the HORCM in the interactive mode. When this option detects a HORCM shut down, interactive mode terminates.

> OpenVMS cannot use the -zx option.

-I[H][M] [instance#] or -I[TC][SI] [instance#]

> Specifies the command as [HORC]/[HOMRCF] or [TrueCopy]/[ShadowImage], and used when specifying the CCI instance number.

-D

> Selects the default report mode. In the default mode, if there is pair status transition information to be reported, one event is reported and the event is reset. If there is no pair status transition information to be reported, the command waits. The report mode consists of the three flags: -allsnd, -resevt, and -nowait options.

-allsnd

> Reports all events if there is pair status transition information.

-resevt

> Reports events if there is pair status transition information, and then resets all events.

-nowait

> When this option is specified, the command does not wait when there is no pair status transition information.

-s <status>...

> Specifies the pair status transition to be reported: smpl, copy (includes rcpy), pair, psus, psue. If two or more statuses are specified following -s, masking is done according to the logical OR of the specified statuses. If this option is not specified, **pairmon** displays all events which received information from HORCM.

Examples

The following shows an example of the **pairmon** command and its output.

```
# pairmon  -allsnd   -nowait
```

```
Group Pair vol Port  targ#lun#LDEV#…Oldstat code
oradb oradb1   CL1-A 1     5   145…SMPL      0x00
oradb oradb2   CL1-A 1     6   146…PAIR      0x02
```

Description of the `pairmon` command output:

Group
> The group name (dev_group) described in the configuration definition file.

Pair vol
> The paired volume name (dev_name) in the specified group that is described in the configuration definition file.

Port targ# lun#
> The port ID, TID, and LUN which is described in the configuration definition file. For further information on fibre-to-SCSI address conversion, see the *Command Control Interface Installation and Configuration Guide*.

LDEV#
> The LDEV ID for the specified device.

Oldstat
> The pair status before changing the volume status.

Newstat
> The pair status after changing the volume status.

code
> The storage system-internal code for the specified status.

The following table specifies the results of the **pairmon** command options.

-D	-nowait	-resevt	-allsnd	Actions
-D	-	-	-	When HORCM does not have an event, this option waits until an event occurs. If one or more events exist, then it reports one event and resets the event which it reported.
Invalid	-	-	-allsnd	When HORCM does not have an event, this option waits until an event occurs. If one or more events exist, then it reports all events.

Chapter 2: Data management commands

-D	-nowait	-resevt	-allsnd	Actions
Invalid	-	-resevt	-	When HORCM does not have an event, this option waits until an event occurs. If one or more events exist, then it reports one event and resets all events.
Invalid	-	-resevt	-allsnd	When HORCM does not have an event, this option waits until an event occurs. If one or more events exist, then it reports all events and resets all events.
Invalid	-nowait	-	-	When HORCM does not have an event, this option reports event nothing. If one or more events exist, then it reports one event and resets the event which it reported.
Invalid	-nowait	-	-allsnd	When HORCM does not have an event, this option reports event nothing. If one or more events exist, then it reports all events.
Invalid	-nowait	-resevt	-	When HORCM does not have an event, this option reports event nothing. If one or more events exist, then it reports one event and resets all events.
Invalid	-nowait	-resevt	-allsnd	When HORCM does not have an event, this option reports event nothing. If one or more events exist, then it reports all events and resets all events.

pairvolchk*

The **pairvolchk** command checks the attributes and status of a pair volume. It acquires and reports the attribute of a volume or group connected to the local host (issuing the command) or remote host. The volume attribute is SMPL (simplex), P-VOL (primary volume), or S-VOL (secondary volume). The -s[s] option reports the pair status in addition to the attribute.

Syntax

```
pairvolchk { -h | -q | -z[x] | -I[H][M][instance#] or
   -I[TC][SI][instance#] | -g <group> | -d <pair Vol> |
   -d[g] <raw_device> [MU#][-n] | -FHORC [MU#] | -FMRCF [MU#] |
   -d[g] <seq#> <LDEV#> [MU#] | -c | -ss |-nomsg }
```

Options and parameters

-h
　　　Displays help/usage and version information.

-q

> Terminates the interactive mode and exits the pair volume check command.

-z or -zx

> Makes the **pairvolchk** command enter the interactive mode. The `-zx` option guards performing of the HORCM in the interactive mode. When this option detects a HORCM shut down, interactive mode terminates.

> OpenVMS cannot use the `-zx` option

-I[H][M] [instance#] or -I[TC][SI] [instance#]

> Specifies the command as [HORC]/[HOMRCF] or [TrueCopy]/[ShadowImage], and used when specifying the CCI instance number.

-g <group>

> Specifies the group name defined in the configuration definition file. The command is executed for the specified group unless the `-d <pair Vol>` option is specified.

-d <pair Vol>

> Specifies the paired logical volume name defined in the configuration definition file. When this option is specified, the command is executed for the specified paired logical volumes. This option is valid when the `-g <group>` option is specified.

-d[g] <raw_device> [MU#][-n]

> Searches whether the specified raw_device is included in a group on the configuration definition file (local instance). If it is in the group, the target volume is executed as the paired logical volume (-d) or group (-dg). This option is effective without specification of `-g <group>` option. If the specified raw_device is contained in two or more groups, the command is executed for the first group.

> This `-n` option is specified to use a device file for NVMe-oF. Specifying this option enables the use of both device files for NVMe-oF and SCSI.

> If the HORCC_NVME environment variable is specified, both device files for NVMe-oF and SCSI can be used without the `-n` option specified.

> To specify this option, a command device must be enabled in a namespace on the storage system. For more information about whether a command device can be enabled in a namespace on the storage system, see the *Provisioning Guide* for your storage system.

Chapter 2: Data management commands

-d[g] <seq#> <LDEV#> [MU#]

Searches whether the specified LDEV is included in a group on the configuration definition file (local instance). If it is in the group, the target volume is executed as the paired logical volume (-d) or group (-dg). This option is effective without specification of '-g <group>' option. If the specified LDEV is contained in two or more groups, the command is executed on the first group. The <seq#> <LDEV#> values can be specified in hexadecimal (add the '0x' prefix) or decimal notation.

- When specifying the serial number for VSP 5000 series, add a "5" at the beginning of the serial number. For example, for serial number 12345, enter 512345.

- When specifying the serial number for VSP G1x00 and VSP F1500, add a "3" at the beginning of the serial number. For example, for serial number 12345, enter 312345.

-c

Checks the conform ability of the paired volumes of the local and remote hosts and reports the volume attribute of the remote host. If this option is not specified, the volume attribute of the local host is reported.

-ss

Used when acquiring the attribute of a volume and the pair status of a volume. If this option is not specified, the volume attribute of the local host is reported.

For details on the pair status to be displayed, see the table for **pairvolchk** group status.

-nomsg

Suppresses messages to be displayed when this command is executed. It is used to execute a command from a user program. This option must be specified at the beginning of a command argument. The command execution log is not affected by this option.

-FHORC [MU#] or -FCA [MU#]

Forcibly specifies a cascaded remote copy volume for specified pair logical volumes on local copy environment. If no -c option is specified, this option acquires the attributes of a cascaded remote copy volume on a local host (near site). If the -c option is specified, this option acquires the attributes of a cascaded remote copy volume on a remote host (far site).

-FMRCF [MU#] or -FBC [MU#]:

Forcibly specifies a cascaded local copy volume for specified pair logical volumes on remote copy environment. If no -c option is specified, acquires the attributes of a cascaded local copy volume on a local host (near site). If the -c option is specified, acquires the attributes of a cascaded local copy volume on a remote host (far site).

Returned values

The **pairvolchk** command sets the following returned values during exit allowing you to check the execution results. See the table below for more details.

- When the -ss option is not specified

  - **Normal termination:**

    - **1:** The volume attribute is SMPL.

    - **2:** The volume attribute is P-VOL.

    - **3:** The volume attribute is S-VOL.

  - **Abnormal termination:**

    - **other than 0 to 127:** refer to the error codes for error details.

- When the `-ss` option is specified

  - **Normal termination:**

    - **11:** The status is SMPL.

  - **Abnormal termination:**

    - specific error codes (see Error Codes) and generic error codes.

    - For TrueCopy, ShadowImage, Volume Migration, and global-active device

      - **22:** The status is P-VOL_COPY or P-VOL_RCPY.

      - **23:** The status is P-VOL_PAIR.

      - **24:** The status is P-VOL_PSUS.

      - **25:** The status is P-VOL_PSUE.

      - **26:** The status is P-VOL_PDUB (LUSE volume only).

      - **32:** The status is S-VOL_COPY or S-VOL_RCPY.

      - **33:** The status is S-VOL_PAIR.

      - **34:** The status is S-VOL_PSUS.

      - **35:** The status is S-VOL_PSUE.

      - **36:** The status is S-VOL_PDUB (LUSE volume only).

    - For TrueCopy Async and Universal Replicator

      The following values (<status-code> + 20) are returned:

      - **42:** The status is P-VOL_COPY.

      - **43:** The status is P-VOL_PAIR.

      - **44:** The status is P-VOL_PSUS.

      - **45:** The status is P-VOL_PSUE.

      - **46:** The status is P-VOL_PDUB. (LUSE volume only)

      - **47:** The status is P-VOL_PFUL.

      - **48:** The status is P-VOL_PFUS.

      - **52:** The status is S-VOL_COPY or S-VOL_RCPY.

      - **53:** The status is S-VOL_PAIR.

      - **54:** The status is S-VOL_PSUS.

      - **55:** The status is S-VOL_PSUE.

      - **56:** The status is S-VOL_PDUB. (LUSE volume only)

      - **57:** The status is S-VOL_PFUL.

      - **58:** The status is S-VOL_PFUS.

    - For Thin Image and Copy-on-Write Snapshot

      The following values are returned:

      - **22:** The status is P-VOL_COPY or P-VOL_RCPY.

Chapter 2: Data management commands

- **23:** The status is P-VOL_PAIR.
- **24:** The status is P-VOL_PSUS.
- **25:** The status is P-VOL_PSUE.
- **26:** The status is P-VOL_PDUB. (LUSE volumes only)
- **27:** The status is P-VOL_PFUL.
- **28:** The status is P-VOL_PFUS.
- **32:** The status is S-VOL_COPY or S-VOL_RCPY.
- **33:** The status is S-VOL_PAIR.
- **34:** The status is S-VOL_PSUS.
- **35:** The status is S-VOL_PSUE.
- **36:** The status is S-VOL_PDUB. (LUSE volumes only)
- **37:** The status is S-VOL_PFUL.
- **38:** The status is S-VOL_PFUS.

📄 **Note:** For pairs in the SSWS status, the command returns the same value as the S-VOL_PSUS status.

Error codes

Category	Error code	Error message	Recommended action	Value
Volume status (Unrecoverable)	EX_ENQVOL	Unmatched volume status within the group	Confirm status using the pairdisplay command. Make sure all volumes in the group have the same fence level and volume attributes.	236
	EX_EVOLCE	Pair volume combination error	Confirm pair status using the pairdisplay command, and change combination of volumes.	235

Example 1

Display example for ShadowImage/Copy-on-Write Snapshot:

```
# pairvolchk -g oradb
pairvolchk : Volstat is P-VOL.[status = PAIR ]
```

Display example for ShadowImage (specified with '-m grp' option):

```
# pairvolchk -g oradb
pairvolchk : Volstat is P-VOL.[status = PAIR, CTGID = 1 ]
```

Display example for TrueCopy:

```
# pairvolchk -g oradb
pairvolchk : Volstat is P-VOL.[status = PAIR, fence = DATA,
MINAP = 2 ]
```

Display example for TrueCopy Sync CTG:

```
# pairvolchk -g oradb
pairvolchk : Volstat is P-VOL.[status = PAIR, fence = DATA,
CTGID = 2 MINAP = 2 ]
```

Display example for TrueCopy Async:

```
# pairvolchk -g oradb
pairvolchk : Volstat is P-VOL.[status = PAIR fence = ASYNC
CTGID = 2 MINAP = 2 ]
```

Description of command output

MINAP

Displays the following two conditions (status) according to the pair status:

PVOL

This shows the minimum in Active Paths on specified group in TrueCopy/ TrueCopy Async.
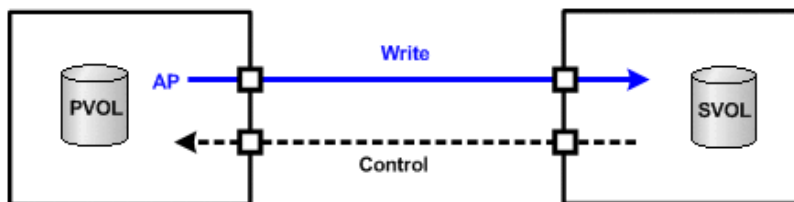
SVOL_ SSUS(SSWS)

MINAP shows the result of the suspend operation that indicates whether the remaining data on P-VOL was completely passed (synchronized) to S-VOL. If MINAP is 'one', all data is passed. Otherwise, all data is not passed from P-VOL.

📄 **Note:** If the microcode on the storage system does not support the active paths, then the 'MINAP' item is not displayed as follows:

```
pairvolchk : Volstat is P-VOL.[status = PAIR fence = ASYNC CTGID = 2]
```

Chapter 2: Data management commands

Example 2

Display example for Universal Replicator:

```
# pairvolchk -g oradb
pairvolchk : Volstat is P-VOL.[status = PAIR fence = ASYNC
CTGID = 2 MINAP = 2 ]
```

Description of command output

MINAP

Displays the following two conditions (status) according to the pair status:

PVOL or SVOL_PAIR (except SSUS(SSWS))

This shows the minimum in Active Paths on specified group in Universal Replicator.
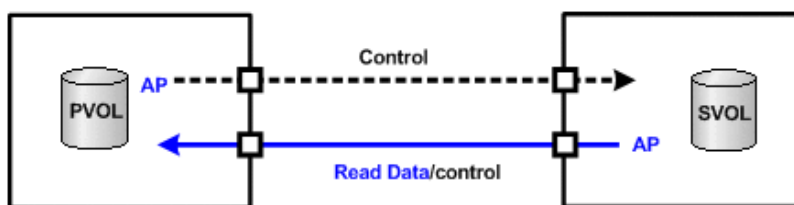
SVOL_ SSUS(SSWS)

MINAP shows the result of the suspend operation that indicates whether or not the remaining All data on P-VOL were Passed (synchronized) to S-VOL completely. If MINAP is 'one', All data were passed. If not, all data were not passed from P-VOL.

> **Note:** If the microcode on the storage system does not support the active paths, then the 'MINAP' item is not displayed as follows:
>
> ```
> pairvolchk : Volstat is P-VOL.[status = PAIR fence = ASYNC
> CTGID = 2]
> ```

Example 3

Display example for LDEV blockading:

```
# pairvolchk -g oradb
pairvolchk : Volstat is P-VOL.[status = PAIR fence = DATA
MINAP = 2  LDEV = BLOCKED]
```

Description of command output

LDEV = BLOCKED

Displays the status of LDEV blockading in order to detect a link failure of E-LUN.

Example 4

The following shows examples of the **pairvolchk** command and its output.

Chapter 2: Data management commands

TrueCopy Async:

```
# pairvolchk -g oradb
pairvolchk : Volstat is P-VOL.[status = PAIR fence = ASYNC
CTGID = 2]
```

TrueCopy:

```
# pairvolchk -g oradb
pairvolchk : Volstat is P-VOL.[status = PAIR fence = DATA ]
```

ShadowImage:

```
pairvolchk -g oradb
pairvolchk : Volstat is P-VOL.[status = PAIR ]
```

ShadowImage pair splitting with specifying the consistency group:

```
# pairvolchk -g oradb
pairvolchk : Volstat is P-VOL.[status = PAIR CTGID = 1]
```

Example 5

The following shows a pairvolchk example that acquires the status (PVOL_PSUS) of the intermediate P/Pvol through specified pair group on ShadowImage environment. The following figure shows a pairvolchk example that acquires the status (PVOL_PSUS) of the intermediate S/Pvol (MU#1) through specified pair group on TrueCopy environment.

See the following figure as an example of -FHORC option for pairvolchk.
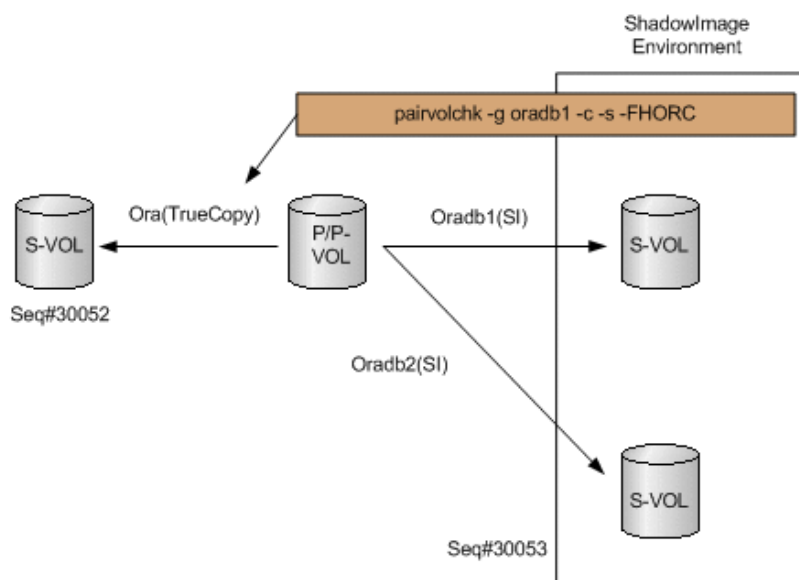
See the following figure as an example of -FMRCF option for **pairvolchk**

TrueCopy Environment
pairvolchk -g ora -c -s -FMRCF 1

PVOL — Ora(TrueCopy) → S/P VOL — Oradb1(ShadowImage) → SVOL
not-exist / exist
Oradb2(ShadowImage) → SVOL
not-exist (PVOL)
Seq#30052
Seq#30053
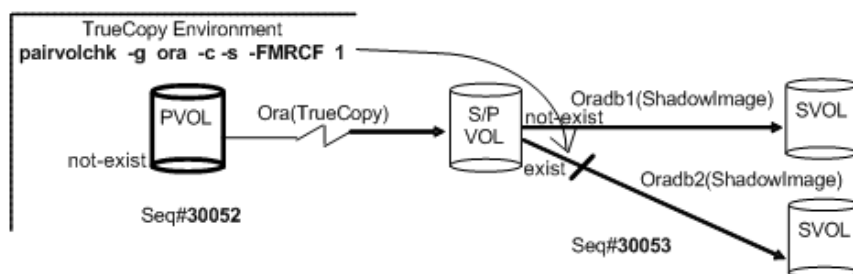
The following table shows the **pairvolchk** group status.

Option	COPY/ RCPY	PSUE	PDU B	PFUS	PSUS	PFUL	PAIR	Group status
	Status of each volume in the group							-
-s*	exist	ignore	ignore	ignore	ignore	ignore	ignore	COPY/ RCPY
	not-exist	exist	ignore	ignore	ignore	ignore	ignore	PSUE
	not-exist	not-exist	exist	ignore	ignore	ignore	ignore	PDUB
	not-exist	not-exist	not-exist	exist	ignore	ignore	ignore	PFUS
	not-exist	not-exist	not-exist	not-exist	exist	ignore	ignore	PSUS
	not-exist	not-exist	not-exist	not-exist	not-exist	exist	ignore	PFUL
	not-exist	not-exist	not-exist	not-exist	not-exist	not-exist	exist	PAIR
-ss	exist	ignore	ignore	ignore	ignore	ignore	ignore	COPY/ RCPY
	not-exist	exist	ignore	ignore	ignore	ignore	ignore	PSUE
	not-exist	not-exist	exist	ignore	ignore	ignore	ignore	PDUB
	not-exist	not-exist	not-exist	ignore	ignore	exist	ignore	PFUL
	not-exist	not-exist	not-exist	ignore	ignore	not-exist	exist	PAIR
	not-exist	not-exist	not-exist	exist	not-exist	not-exist	not-exist	PFUS

Chapter 2: Data management commands

Option	COPY/ RCPY	PSUE	PDU B	PFUS	PSUS	PFUL	PAIR	Group status
	not-exist	not-exist	not-exist	not-exist	exist	not-exist	not-exist	PSUS

*This option can be used only when **pairvolchk -s** has the variable *USE_OLD_VCHK*.

Legend:

- exist: The target status exists in the group.
- not-exist: The target status does not exist in the group.

The PFUL state appears as PAIR by all commands (except the -fc option of the **pairdisplay** command), since PFUL indicates PAIR state with sidefile at the HWM.

The PFUS state appears as PSUS by all commands (except the -fc option of the **pairdisplay** command), since PFUS indicates SUSPENDED state due to sidefile full.

The SVOL_PSUS state appears as SSUS by the **pairdisplay** command and other commands.

pairdisplay*

The **pairdisplay** command displays the pair status allowing you to verify completion of pair operations (for example, **paircreate**, **pairresync**). The **pairdisplay** command is also used to confirm the configuration of the pair volume connection path (the physical link of paired volumes and servers). The **pairdisplay** command can be used for a paired volume or a group of paired volumes.

Syntax

```
pairdisplay{ -h | -q | -z[x] | -I[H][M][instance#] or
    -I[TC][SI][instance#]
    | -g <group> | -d <pair Vol> | -d[g] <raw_device> [MU#][-n]
    | -FHORC [MU#]| -FMRCF  [MU#] | -d[g] <seq#> <LDEV#> [MU#]
    | -c | -l  | -f[xcdmnew] | -CLI | -m <mode> | -v jnl[t]
    | -v ctg  | -v pid | -v pidb [-fp]}
```

Options and parameters

-h

Displays help/usage and version information.

-q

Terminates the interactive mode and exits the pair volume check command.

-z or -zx

>
Makes the **pairdisplay** command enter interactive mode. The `-zx` option guards performing of the HORCM in the interactive mode. When this option detects a HORCM shut down, interactive mode terminates.

>
OpenVMS cannot use the `-zx` option

-I[H][M] [instance#] or -I[TC][SI] [instance#]

>
Specifies the command as [HORC]/[HOMRCF] or [TrueCopy]/[ShadowImage], and used when specifying the CCI instance number.

-g <group>

>
Specifies the group name defined in the configuration definition file. The command is executed for the specified group unless the `-d <pair Vol>` option is specified.

-d <pair Vol>

>
This option is used to specify the paired logical volume name defined in the configuration definition file. When this option is specified, the command is executed for the specified paired logical volumes. This option is valid when the `-g <group>` option is specified.

-d[g] <raw_device> [MU#] [-n]

>
Searches whether the specified raw_device is included in a group on the configuration definition file (local instance). If it is in the group, the target volume is executed as the paired logical volume (-d) or group (-dg). This option is effective without specification of `-g <group>` option. If the specified raw_device is contained in two or more groups, the command is executed for the first group.

>
This `-n` option is specified to use a device file for NVMe-oF. Specifying this option enables the use of both device files for NVMe-oF and SCSI.

>
If the HORCC_NVME environment variable is specified, both device files for NVMe-oF and SCSI can be used without the `-n` option specified.

>
To specify this option, a command device must be enabled in a namespace on the storage system. For more information about whether a command device can be enabled in a namespace on the storage system, see the *Provisioning Guide* for your storage system.

-FHORC [MU#] or -FCA [MU#]

>
Forcibly specifies a cascaded remote copy volume for specified pair logical volumes on local copy environment. If the `-l` option is specified, this option displays status of a cascaded remote copy volume on a local host (near site). If no `-l` option is specified, this option displays status of a cascaded remote copy volume on a remote host (far site). This option cannot be specified with -m <mode> option on the same command line.

-FMRCF [MU#] or -FBC [MU#]

>
Forcibly specifies a cascaded local copy volume for specified pair logical volumes on remote copy environment. If the `-l` option is specified, this option displays status of a cascaded local copy volume on a local host (near site). If no `-l` option is specified, this option displays status of a cascaded local copy volume on a remote host (far site). This option cannot be specified with -m <mode> option on the same command line.

-d[g] <seq#> <LDEV#> [MU#]

Searches whether the specified LDEV is included in a group on the configuration definition file (local instance). If it is in the group, the target volume is executed as the paired logical volume (-d) or group (-dg). This option is effective without specification of '-g <group>' option. If the specified LDEV is contained in two or more groups, the command is executed on the first group. The <seq#> <LDEV#> values can be specified in hexadecimal (add the '0x' prefix) or decimal notation.

- When specifying the serial number for VSP 5000 series, add a "5" at the beginning of the serial number. For example, for serial number 12345, enter 512345.

- When specifying the serial number for VSP G1x00 and VSP F1500, add a "3" at the beginning of the serial number. For example, for serial number 12345, enter 312345.

-c

Checks the configuration of the paired volume connection path (physical link of paired volume among the servers) and displays illegal pair configurations. If this option is not specified, the status of the specified paired volume is displayed without checking the path configuration.

-l

Displays the paired volume status of the local host (which issues this command).

-fx

Displays the LDEV ID as a hexadecimal number.

-fc

Displays copy operation progress, sidefile percentage, bitmap percentage, or Universal Replicator journal percentage. Displays PFUL/PFUS for TrueCopy Async and Universal Replicator. Used to confirm SSWS state as indication of SVOL_SSUS-takeover after.

-fd

Displays the relationship between the Device_File and the paired volumes, based on the group (as defined in the local instance configuration definition file). If Device_File column shows 'Unknown' as shown in the Display example, it means that the volume cannot be recognized by the host because the volume has not been registered when you start CCI, and pair operation are rejected (except the local option such as '-l') in protection mode.

Display example:

```
# pairdisplay -g oradb -fd
Group PairVol(L/R) Device_File M ,Seq#,LDEV#.P/S,Status, Seq#,P-LDEV# M
oradb oradev1(L) c0t3d0 0 35013 17..P-VOL COPY, 35013 18 -
oradb oradev1(R) Unknown 0 35013 ****..---- ----, ----- -- -
```

Specify ENABLE for HORCM_NVME in the configuration file. If ENABLE is not specified, a device file name for NVMe-oF is not displayed.

-fm

Displays the management units for the differential data of TC,UR,GAD, or SI pairs (SI is available only for VSP 5000 series or VSP E series). If this option is specified, the management units are displayed in M column.

-fe

Displays the serial# and LDEV# of the external LUNs mapped to the LDEV and additional information for the pair volume. This option displays the information above by adding to last column, and then ignores the format of 80 column. This option is not valid if the cascade options (-m all,-m cas) are specified.

Display example for TC/TC Async/UR/GAD:

```
# pairdisplay -g oradb -fe
Group    PairVol(L/R)    (Port#,TID, LU),    Seq#,   LDEV#.P/S, Status,   Fence,
Seq#,
P-LDEV#   M  CTG    JID   AP   EM   E-Seq#      E-LDEV# R/W     QM      DM  P  PR
Oradb    dev1(L)    (CL5-A-0,30, 0)    64568    301.P-VOL    PAIR    ASYNC,
64568
303   -      0    2   1    -    -    -  -/-    AA   D  N  P
Oradb    dev1(R)    (CL5-A-0,30, 2)    64568    303.S-VOL    PAIR    ASYNC,----
-
  301   -      0    4   1    -    -    -  -/-    AA   D  N  P
Oradb    dev2(L)    (CL5-A-0,30, 1)    64568    302.P-VOL    PAIR    ASYNC,
64568
  304   -      0    3   1    -    -    -  -/-    AA   S  N  D
Oradb    Dev2(R)    (CL5-A-0,30, 3)    64568    304.S-VOL    PAIR    ASYNC,----
-
  302   -      0    5   1    -    -    -  -/-    AA   S  N  D
```

Display example for ShadowImage/Copy-on-Write Snapshot/Volume Migration:

```
# pairdisplay -g horc0 -fe
Group ...   Seq#,LDEV#.P/S,Status, Seq#,P-LDEV# M CTG CM EM E-Seq# E-LDEV#   DM
P
```

Chapter 2: Data management commands

```
horc0 ... 63528    65.P-VOL COPY,63528     64 -   -  N  -       -     -  - -
horc0 ... 63528    64.S-VOL COPY,-----     65 -   -  N  -       -     -  - -
```

- CTG: For TC, TC Async, UR, and GAD, it displays the consistency group ID when the fence level is ASYNC. For TC Sync CTG, it displays the consistency group ID when the fence level is DATA, STATUS, or NEVER.

  > **Note:**
  >
  > If the snapshot is made by `raidcom add snapshot`, the displayed consistency group ID is not correct. To confirm the status of a snapshot that was made by `raidcom add snapshot`, use the **raidcom get snapshot** command.

- JID: The journal ID for P-VOL or S-VOL. In a HAM or GAD configuration, it shows the quorum ID and the fence level is set to 'Never'. When it is not the UR or HAM configuration, '-' is displayed.

- AP: The number of active paths for UR links on P-VOL, and it displays the number of active paths for UR links on P-VOL and S-VOL.'Unknown'is shown as '-'. Refer to 'MINAP' information that is displayed by running the **pairvolchk** command.

- CM: The Copy mode

  - N: Non SnapShot

  - S: SnapShot. In the SMPL state, this shows that pair-volume is created as SnapShot.

  - C: Cruising Copy

- EM: The external connection mode

  - H: Mapped E-lun as hidden from the host.

  - V: Mapped E-lun as visible to the host

  - ' - ': Unmapped to the E-lun

  - BH: Mapped E-lun as hidden from the host, but LDEV blockading.

  - BV: Mapped E-lun as visible to the host, but LDEV blockading

  - B: Unmapped to the E-lun, but LDEV blockading

- E-Seq#: The production (serial) number of the external LUN. 'Unknown' is shown as '-'.

  - The serial number for VSP 5000 series is displayed with a "5" at the beginning ("512345" = serial number 12345).

  - The serial number for VSP G1x00 and VSP F1500 is displayed with a "3" at the beginning ("312345" = serial number 12345).

- E-LDEV#: The LDEV# of the external LUN, 'Unknown' is shown as '-'.

- R/W: Displays the I/O mode when Read/Write is performed. For GAD configuration, the following modes are displayed:

  - L/L: Read/Write is performed on both the primary and secondary volumes.

  - L/M: Read is performed on both the primary and secondary volumes. Write is performed on the primary volume first, and then done on the secondary volume.

  - B/B: Read/Write is rejected (Illegal Request is replied). LU-undefined is returned as a response of the Inquiry command on this volume.

- QM: For GAD pairs, displays the pair operation mode when the quorum disk is blocked.

  - SP: Pair suspended.

  - AS: Pair retained (P-VOL accessible, S-VOL inaccessible).

  - AA: Pair retained (Both P-VOL and S-VOL accessible).

  - - (hyphen): Displayed for any of the following:

    - A non-GAD pair

    - A suspended GAD pair

    - A GAD pair without operation mode display support

    - A GAD pair, but the operation mode that is not supported by CCI is set

    - No volume is set for the quorum disk.

- DM: Displays the differential data management method (VSP 5000 series, VSP G1x00, VSP F1500, VSP E series, and VSP G130, G/F350, G/F370, G/F700, G/F900 only).

  - S: The differential data is managed using a shared memory.

  - D: The differential data is managed using a drive (tier difference method).

  - - (hyphen): This information is not available.

- ▪ P: Displays the status of expansion processing of DP-VOL (VSP 5000 series, VSP G1x00 and VSP F1500, VSP E series, and VSP G130, G/F350, G/F370, G/F700, G/F900 only).

  - ▪ E: Expansion is in progress.

  - ▪ N: Not processed.

  - ▪ - (hyphen): This information is not available.

    A hyphen (-) is displayed if the microcode version is not supported.

- ▪ PR: Displays the I/O preference mode when remote path failed of the GAD pair. This information is valid only for GAD pairs.

  - · P: Primary Volume.

  - · D: I/O Preference Mode When Remote Path Failed is disabled.

  - · - (hyphen): This information is not available.

    A hyphen (-) is displayed if the DKCMAIN microcode version is not supported, or the pair is not a GAD pair.

-v pidb fp

Displays the actual capacity of the pool.

> 📄 **Note:** This option displays nothing if the target volume is not a snapshot volume.

-fw

Used when displaying the WWN setting (defined by '50060E80+RAID Type,Serial#,Port#)) to the port instead of a port name. If this option is specified with the '-fe' option at the same line, then 'LUN WWN' appears as shown below. If WWN is not computed, then 'Unknown' appears (for example, SCSI).

Display example:

```
# pairdisplay -g PG01 -fw
Group PairVol(L/R) (WWN, LU-M),Seq#,LDEV#.P/S,Status, Seq#,P-LDEV# M
PG01 ora01(L) (500060e804f42001, 3-0) 62496 67. P-VOL COPY, 62496 69 -
PG01 ora01(R) (500060e804f42001, 5-0) 62496 69. S-VOL COPY, ----- 67 -
PG01 ora02(L) (500060e804f42001, 4-0) 62496 68. P-VOL COPY, 62496 64 -
PG01 ora02(R) (500060e804f42001, 6-0) 62496 64. S-VOL COPY, ----- 68 -
```

Display example:

```
# pairdisplay -IH -g PG01 -CLI -l -fwe
Group PairVol L/R WWN LU Seq# LDEV# P/S Status Fence Seq#
P-LDEV# M CTG JID AP EM E-Seq# E-LDEV# LUN-WWN  R/W QM DM P PR
PG01  oradb01 L 500060e804fa0f01 1 64015 11 P-VOL PAIR ASYNC 62496 11 - 5 - 1
- - - 60060e8005fa0f000000fa0f0000000b -/- - - - -
```

```
PG01  oradb02 L 500060e804fa0f01 2 64015 12 P-VOL PAIR ASYNC 62496 12 - 5 - 1
- - - 60060e8005fa0f000000fa0f0000000c -/- - - - -
```

- NGUID: Displays the namespace identifier (NGU ID). If NGUID is not assigned, 'Unknown' is displayed.

-fne

Used when displaying the namespace identifier (NGU ID).

Display example:

```
# pairdisplay -g nvme -I10 -fne
Group PairVol(L/R) (Port#,TID, LU),Seq#,LDEV#.P/S,Status,Fence,Seq#,P-LDEV# M
CTG
JID AP EM E-Seq# E-LDEV# NGUID R/W QM DM P PR
nvme ndev001(L) (CL7-D-0,17, 0)500054 257.SMPL ---- ------,----- ----- - - - -
- -
- 00500036000000000060e80800360101 -/- - - - -
nvme ndev001(R) (CL7-D-0,17, 1)500054 258.SMPL ---- ------,----- ----- - - - -
- -
- 00500036000000000060e80800360102 -/- - - - -
```

- NGUID: Displays the namespace identifier (NGU ID). If NGUID is not assigned, 'Unknown' is displayed.

-fwne

Used when displaying the namespace identifier (NGU ID) and LUNWWN.

Display example:

```
# pairdisplay -g ora -fwne -l
Group PairVol(L/R)(WWN , LU),Seq#,LDEV#.P/S,Status,Fence,Seq#,P-LDEV# M CTG
JID AP
EM E-Seq# E-LDEV# LUN-WWN NGUID R/W QM DM P PR
ora dev1(L) (50060e8008003622, 3)500054 20483.SMPL ---- ------,----- ----- - -
- -
- - - Unknown 00500036000000000060e80800360101 -/- - - - -
```

- NGUID: Displays the namespace identifier (NGU ID). If NGUID is not assigned, 'Unknown' is displayed.

-CLI

Used when specifying display for command line interface (CLI). This option displays to the same position that defined number of columns, and displays one header. The delimiters between columns are displayed as spaces or hyphens (-).

Display example:

```
Group PairVol L/R Port# TID LU-M  Seq# LDEV# P/S Status Seq#
P-LDEV# M
homrcf1 deva0  L   CL1-D  3  5 0 30053   271 P-VOL PAIR 30053
```

```
    263 -
homrcf1 deva1  L   CL1-D  3  5 1 30053   271   SMPL    -     -
      - -
homrcf1 deva2  L   CL1-D  3  5 2 30053   271   SMPL    -     -
      - -
```

-m <mode>

> Used when displaying a paired status of each mirror descriptors for the specified pair logical volume, and used for paired status indication of the cascaded volume. The <mode> option can be designated 'cas' or 'all':

> ▪ The 'cas' option is used when displaying a paired status of specified group that is registered to the mirror descriptor (MU#) on the cascaded configuration file.

> ▪ The 'all' option is used when displaying a paired status of all mirror descriptors (MU#).

> This option (-m <mode>) is not affected with command execution environment (TrueCopy/TrueCopy Async/Universal Replicator/GAD and ShadowImage/Copy-on-Write Snapshot /Volume Migration), and displays the paired status. This option cannot be specified with the -FHORC/-FMRCF option on the same command line.

-v jnl[t]

> Displays the JNL status for the local and remote interconnected to the group. Also finds the journal ID for each local and remote interconnected to the group via the specified group or <raw_device>, and displays information of each journal ID corresponding the local and remote. The first line shows the journal information for the local host, second line shows the journal information for the remote host. The item for displaying is the same as **raidvchkscan -v jnl[t]**. The serial number (Seq#) display varies depending on storage system.

> ▪ The serial number for VSP 5000 series is displayed with a "5" at the beginning ("512345" = serial number 12345).

> ▪ The serial number for VSP G1x00 and VSP F1500 is displayed with a "3" at the beginning ("312345" = serial number 12345).

Example:

```
# pairdisplay -g VG01 -v jnl
```

```
JID MU CTG JNLS AP U(%) Q-Marker Q-CNT D-SZ(BLK)  Seq# Nnm LDEV#
```

```
001  0   2 PJNN  4   21 43216fde   30    512345 62500   2
   265
```

```
002  0   2 SJNN  4   95 3459fd43 52000    512345 62538   3
   270
```

Chapter 2: Data management commands

Example:

```
# pairdisplay -g VG01 -v jnlt
```

```
JID MU CTG   JNLS  AP   U(%)    Q-Marker    Q-CNT   D-SZ(BLK)    Seq# DOW PBW APW
```

```
001  1   2   PJNN   4    21     43216fde      30      512345   62500  20 300  40
```

```
002  1   2   SJNN   4    95     3459fd43    52000     512345   62538  20 300  40
```

Example:

```
# pairdisplay -g VG01 -v jnl -FCA 1
```

```
JID MU CTG   JNLS  AP  U(%)   Q-Marker   Q-CNT   D-SZ(BLK)    Seq# Nnm LDEV#
```

```
003  1   2   PJNN   4   21    43216fde     30      512345   62500   2    265
```

> 📄 **Note:**
>
> - This option can be specified with following options on the same command line:
>
> ```
> {-g<group> | -d <pair Vol> | -d[g] <raw_device> [MU#] | -FCA [MU#] |
> -d[g] <seq#> <LDEV#> [MU#] | -l  | -f[x] }
> ```

- The '-FHORC [MU#]' or '-FCA [MU#]' option is used when displaying the journal information of cascaded Universal Replicator volume, and then it shows the journal information for the remote host only.

- This option displays nothing if the target volume is NOT a Universal Replicator volume.

-v ctg

This option finds the consistency group for each local and remote interconnected to the group via the specified group or <raw_device>, and displays any information of each consistency group corresponding the local and remote. The first line shows the consistency group information for the local host, second line shows the consistency group information for the remote host.

Example:

```
# pairdisplay -g VG01-v ctg
```

```
CTG P/S Status AP U(%) Q-Marker  QM-Cnt SF(%) Seq# IFC OT/s
CT/m RT/m
```

```
001  P-VOL PAIR   2   0  00000080       3   50     63528 ON   90   5   5
```

```
001  S-VOL PAIR   -   0  0000007d       -   50     63528 -    -    -   -
```

CTG: Displays the consistency group ID.

P/S: The attribute of a volume in first LDEV of the specified group.

Status: The status of the paired volume in first LDEV of the specified group.

AP: Displays the number of Active Path in Universal Replicator links on P-VOL, also displays the number of active path in Universal Replicator links on P-VOL and S-VOL, 'Unknown' is shown as '-'.

U(%): The usage sidefile/journal data, it is valid at PAIR state. For VSP G350, G370, G700, G900, VSP F350, F370, F700, F900, this indicates the actual usage rate when the -fp option is specified.

For TrueCopy Async: The sidefile percentage for consistency group in relationship to a 100% full sidefile in cache.

For Universal Replicator: The usage rate of the current journal data as 100% of the journal data space.

Q-Marker: In P-VOL, the latest sequence # of the MCU P-VOL when the write command was received. In S-VOL, the latest sequence # of the DFW on RCU. This item is valid at PAIR state.

QM-Cnt: The number of remaining Q-Marker within consistency group of the Unit. TrueCopy Async sends a token called 'dummy recordset' at regular interval time, therefore QM-Cnt always shows '2' or '3' even if host has NO writing. This item is valid at PAIR state.

SF(%):The usage of cache setting as the sidefile regardless of Universal Replicator and TrueCopy Async.

Seq#: The serial number of the RAID storage system.

- The serial number for VSP 5000 series is displayed with a "5" at the beginning ("512345" = serial number 12345).

- The serial number for VSP G1x00 and VSP F1500 is displayed with a "3" at the beginning ("312345" = serial number 12345).

IFC: Shows 'ON'/ 'OFF' for controlling Inflow.

OT/s: Displays the value of the 'offloading timer' (in seconds) for TrueCopy Async or the value of the Data Overflow Watch timer (in seconds) for Universal Replicator. Either value is displayed only for the consistency group containing the P-VOL when control of write I/O inflow (IFC) is enabled (ON). In other cases, a hyphen (-) is displayed.

CT/m: The 'Copy Pending timer' (specified in minutes) setting to consistency group for only TrueCopy Async.

RT/m: The 'RCU Ready timer' (specified in minutes) setting to consistency group for only TrueCopy Async.

> **Note:**
>
> This option displays nothing if the target volume is NOT a TrueCopy Async or Universal Replicator volume. The '-FHORC [MU#]' or '-FCA [MU#]' option is used when displaying the information of cascaded TrueCopy Async/Universal Replicator volume, and then it shows the consistency group information for the remote host only.

-v pid

Finds the pool ID for each local and remote interconnected to the group via the specified group or <raw_device>, and displays any information of each pool ID corresponding the local and remote. The first line shows the pool information for the local host, second line shows the pool information for the remote host. The displayed items are the same as for raidvchkscan -v pid.

> **Note:**
>
> - This option displays nothing if the target volume is not a snapshot volume.

Chapter 2: Data management commands

Command Control Interface Command Reference   94

Example:

```
# pairdisplay -g VG01 -v pid
```

```
PID POLS U(%) SSCNT Available(MB) Capacity(MB)  Seq#  Num LDEV# H(%)
```

```
127 POLN   0      6          3000          3000 63528    2
  864   80
```

```
127 POLN   0      6          3000          3000 63528    2
  864   80
```

```
# pairdisplay -g VG01 -v pid -l
```

```
PID POLS U(%) SSCNT Available(MB) Capacity(MB)  Seq#  Num
LDEV# H(%)
```

```
127 POLN    0    6          3000          3000 63528    2
  864   80
```

H(%): Displays the threshold rate being set to the Copy-on-Write Snapshot pool as the high water mark. 'Unknown' is shown as '-'.

-v pid -FMRCF

The '-FBC [MU#]' option is used for displaying the pool information of cascaded snapshot volumes, so that you can monitor the pool status on remote host connected to cascaded TrueCopy P-VOL to TrueCopy S-VOL/P-VOL. It shows the pool information for the remote host only.

> 📄 **Note:**
>
> - This option displays nothing if the target volume is not a snapshot volume.

Example:

```
# pairdisplay -g VG01 -v pid -FMRCF
```

```
PID POLS U(%) SSCNT Available(MB) Capacity(MB)  Seq#  Num
LDEV# H(%)
```

```
127 POLN    0    6          3000          3000 63528    2
  864   80
```

Chapter 2: Data management commands

-v pidb

> Displays basic information of the pool.

> 📄 **Note:**
>
> ▪ This option displays nothing if the target volume is not a snapshot volume.

-v pidb -fp

> Displays the actual capacity of the pool.

> If the target volume is not a snapshot volume, nothing is displayed.

Returned values

None

Error codes

None

Examples

`pairdisplay` command example for TrueCopy, TrueCopy Async, Universal Replicator, and global-active device

```
# pairdisplay -g oradb -fcx
```

```
Group  Pair Vol(L/R)  (P,T#,L#),   Seq#,  LDEV#..P/S,  Status,
Fence, Copy%, P-LDEV#  M
oradb  oradb1(L)      (CL1-B, 1,0) 1234      64..P-VOL   PAIR
Never,   75       C8  -
oradb  oradb1(R)      (CL1-A, 1,0) 5678      C8..S-VOL   PAIR
Never,  ----       64  -
```

`pairdisplay` command example for ShadowImage and Copy-on-Write Snapshot

```
# pairdisplay  -g oradb
```

```
Group  Pair Vol(L/R) (Port#,TID,LU-M),  Seq#,  LDEV#..P/S, Status,
Fence,  Seq#, P-LDEV# M
oradb      oradb1(L)      (CL1-A, 1,0)  30053      18..P-VOL  PAIR
Never, 30053     19  -
oradb      oradb1(R)      (CL1-D, 1,0)  30053      19..S-VOL  PAIR
Never,  ----      18  -
```

Examples of `pairdisplay -m`
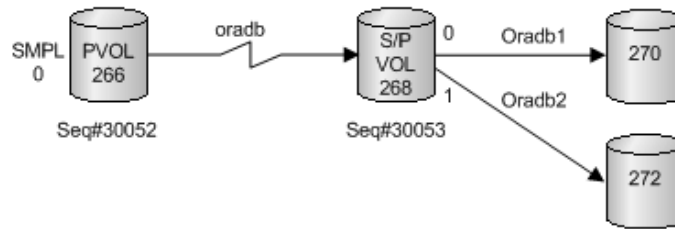
Display example for -m cas

```
# pairdisplay -g oradb -m cas
```

```
Group   PairVol(L/R) (Port#,TID,LU-M), Seq#,  LDEV#..P/S,  Status,
Seq#, P-LDEV# M
oradb   oradev1(L)   (CL1-D , 3,  0-0) 30052  266....SMPL  ----,
----- ----    -
oradb   oradev1(L)   (CL1-D , 3,  0)   30052  266....P-VOL COPY,
30053 268     -
oradb1  oradev11(R)  (CL1-D , 3,  2-0) 30053  268....P-VOL COPY,
30053 270     -
oradb2  oradev21(R)  (CL1-D , 3,  2-1) 30053  268....P-VOL PSUS,
30053 272     W
oradb   oradev1(R)   (CL1-D , 3,  2)   30053  268....S-VOL COPY,
----- 266     -
```

Display examples for -m all

```
# pairdisplay -g oradb -m all
```

```
Group   PairVol(L/R) (Port#,TID,LU-M), Seq#,  LDEV#..P/S,  Status,
Seq#, P-LDEV# M
oradb   oradev1(L)   (CL1-D , 3,  0-0) 30052  266....SMPL  ----,
---- ----     -
-----   -----(L)     (CL1-D , 3,  0-1) 30052  266....SMPL  ----,
---- ----     -
-----   -----(L)     (CL1-D , 3,  0-2) 30052  266....SMPL  ----,
---- ----     -
oradb   oradev1(L)   (CL1-D , 3,  0)   30052  266....P-VOL PAIR,
30053 268     -
oradb1  oradev11(R)  (CL1-D , 3,  2-0) 30053  268....P-VOL COPY,
30053 270     -
oradb2  oradev21(R)  (CL1-D , 3,  2-1) 30053  268....P-VOL PSUS,
30053 272     W
-----   -----(R)     (CL1-D , 3,  2-1) 30053  268....SMPL  ----,
----- ----    -
oradb   oradev1(R)   (CL1-D , 3,  2)   30053  268....S-VOL COPY,
```

Chapter 2: Data management commands

```
----- 266    -
```

```
# pairdisplay -d /dev/rdsk/c0t3d0 -l -m all
```

```
Group  PairVol(L/R) (Port#,TID,LU-M), Seq#,  LDEV#..P/S,  Status,
Seq#, P-LDEV# M
oradb  oradev1(L)    (CL1-D , 3,  0-0) 30052  266....SMPL  ----,
----   ----   -
-----  -----(L)      (CL1-D , 3,  0-1) 30052  266....SMPL  ----,
----   ----   -
-----  -----(L)      (CL1-D , 3,  0-2) 30052  266....SMPL  ----,
----   ----   -
oradb  oradev1(L)    (CL1-D , 3,  0)   30052  266....P-VOL PAIR,
30053  268    -
```

Description of the `pairdisplay` command output:

Group

Group name (dev_group) as described in the configuration definition file

Pair Vol(L/R)

Paired volume name (dev_name) as described in the configuration definition file. (L) = local host; (R) = remote host

(P,T#,L#) (TrueCopy, TrueCopy Async, Universal Replicator, global-active device)

Port, TID, and LUN as described in the configuration definition file.

(Port#,TID,LU-M) (ShadowImage, Copy-on-Write Snapshot)

Port number, TID, LUN, and MU number as described in the configuration definition file.

Seq#

Serial number of the RAID storage system.

- The serial number for VSP 5000 series is displayed with a "5" at the beginning ("512345" = serial number 12345).

- The serial number for VSP G1x00 and VSP F1500 is displayed with a "3" at the beginning ("312345" = serial number 12345).

LDEV#

Logical device number

P/S

volume attribute

Status

Status of the paired volume

Chapter 2: Data management commands

Fence (TrueCopy, TrueCopy Async, Universal Replicator, global-active device)
Fence level

%

Copy operation completion, or percent pair synchronization

Vol.	Copy	Pair	Other	Copy	Pair	Other	Copy	Pair	Pvol_psus Svol_copy	Other
	TrueCopy Async status			TrueCopy Sync/ GAD status			ShadowImage, Copy-on-Write Snapshot, or Volume Migration status			
P-VOL	CR	SF	BMP	CR	BMP	BMP	CR	CR	BMP	CR
S-VOL	-	SF	BMP	-	BMP	BMP	CR	CR	CR	CR

Volume	Copy	Pair	PSUS/SSUS (PJNS/SJNS)	Other
	Universal Replicator status			
P-VOL	CR	JF	BMP	BMP
S-VOL	-	JF	BMP	BMP
CR: Shows the copy operation rate (identical rate of a pair).				
BMP : Shows the identical percentage of BITMAP both P-VOL and S-VOL.				
SF: Shows sidefile percentage of each consistency group as sidefile 100% on cache of both P-VOL and S-VOL. Following is an arithmetic expression using the high water mark (HWM) as 100% of a sidefile space:				
HWM(%) = HWM(%)/Sidefile space (30 to 70) × 100				
JF: Shows the usage rate of the current journal data as 100% of the journal data space.				

P-LDEV#
LDEV number of the partner volume of the pair

M

- When the -fm option is not specified:

  For P-VOL and 'PSUS' state:

  - M='W' shows that S-VOL is suspending with R/W enabled through the pairsplit.

  - M='-' shows that S-VOL is suspending with read only through the pairsplit.

  For S-VOL and 'SSUS' state:

  - M='W' shows that S-VOL has been altered since entering SSUS state.

  - M='-' shows that S-VOL has NOT been altered since entering SSUS state.

  For 'COPY/RCPY/PAIR/PSUE' state:

  - M='N' shows that its volume are read-disabled through the `paircreate` '-m noread'.

  - M=other than 'N' shows that a hyphen (-) is displayed.

- When the -fm option is specified:

  - M='T' shows that the pairs are TC/UR/GAD pairs, and the management unit for their differential data is "track".

  - M='C' shows that the pairs are TC/UR/GAD pairs, and the management unit for their differential data is "cylinder".

  - M='U' shows that the system is unable to display differential data, or the pairs are SI/HTI/Copy-on-Write Snapshot pairs.

  - M='-' shows that there are no pairs.

Example of pairdisplay -v pidb

```
# pairdisplay -g VG01 -v pidb
PID POLS U(%) LCNT SSCNT Available(MB) Capacity(MB) Snap_Used(MB)
TL_CAP(MB) BM TR_CAP(MB) RCNT Seq# Num LDEV# W(%) H(%) STIP
VCAP(%)
TYPE PM PT POOL_NAME
001 POLN 0 11001 11001 46998 46998 0
2432398 NB 0 0 300050 1 0 70 80 YES - OPEN
N HDP dp_ti_pool
001 POLN 0 11001 11001 46998 46998 0
2432398 NB 0 0 300050 1 0 70 80 YES - OPEN
N HDP dp_ti_pool
```

Example of pairdisplay -v pidb -fp

```
# pairdisplay -g VG01 -v pidb -fp
PID POLS U(%) LCNT SSCNT Available(MB) Capacity(MB) Snap_Used(MB)
TL_CAP(MB) BM TR_CAP(MB) RCNT Seq# Num LDEV# W(%) H(%) STIP VCAP(%)
TYPE PM PT POOL_NAME
001 POLN 0 11001 11001 46998 46998 -
2432398 NB 0 0 300050 1 0 70 80 YES -
OPEN N HDP dp_ti_pool
```

Chapter 2: Data management commands

```
001 POLN 0 11001 11001 46998 46998 -
2432398 NB 0 0 300050 1 0 70 80 YES -
OPEN N HDP dp_ti_pool
```

Output of the pairdisplay command:

PID

Pool ID

POLS

Status of the pool

- POLN: "Pool Normal" (The pool is in the normal status.)

- POLF: "Pool Full" (The pool is in the overflow status exceeding the threshold.)

- POLS: "Pool Suspend" (The pool is in the overflow status exceeding the threshold and is blocked.)

- POLE: "Pool failure" (The pool is blocked due to a failure.) In this status, the pool information cannot be displayed.

U(%)

Usage rate of the pool

Displays the actual usage rate of the pool when the -fp option is specified.

LCNT

Total number of Dynamic Provisioning virtual volumes mapped to the pool. A hyphen (-) is displayed if the information is not available for this pool.

SSCNT

Total number of snapshot data items mapped to the pool. A hyphen (-) is displayed if the information is not available for this pool.

Available(MB)

Available capacity for volume data in the pool

Displays the actual pool capacity when the -fp option is specified.

Capacity(MB)

Total capacity of the pool

Displays the actual pool capacity when the -fp option is specified.

Snap_Used(MB)

Capacity used for Thin Image data in megabytes. If the value is less than 1 MB, it is rounded up. A hyphen (-) is displayed if the information is not available for this pool.

If you entered the $-fp$ option, a hyphen (-) is displayed.

TL_CAP(MB)

Total capacity of all DP-VOLs and Thin Image pairs mapped to the pool. A hyphen (-) is displayed if the information is not available for this pool.

BM

I/O Blocking Mode of the pool:

- PF (Pool Full): If the pool is full, you cannot read from or write to the target DP-VOL. If the pool VOL is blocked, you can read from or write to the target DP-VOL.

- PB (Pool vol Blockade): If the pool VOL is blocked, you cannot read from or write to the target DP-VOL. If the pool is full, you can read from or write to the target DP-VOL.

- FB (Full or Blockade): If the pool is full or pool VOL is blocked, you cannot read from or write to the target DP-VOL.

- NB (No Blocking): If the pool is full or pool VOL is blocked, you can read from or write to the target DP-VOL.

- - (Not supported): The configuration does not support the I/O Blocking Mode.

TR_CAP(MB)

Sum of the pool capacities reserved for the volumes for which Full Allocation is enabled or reserved for Proprietary Anchor. For a configuration that does not support Full Allocation, a hyphen (-) is displayed.

RCNT

Number of volumes for which Full Allocation is enabled and mapped to the pool. For a configuration that does not support Full Allocation, a hyphen (-) is displayed.

Seq#

Serial number of the RAID storage system.

- The serial number for VSP 5000 series is displayed with a "5" at the beginning ("512345" = serial number 12345).

- The serial number for VSP G1x00 and VSP F1500 is displayed with a "3" at the beginning ("312345" = serial number 12345).

Num

Number of LDEVs belonging to the pool

LDEV#

First number of LDEVs in the pool. "65535 (ffff)" is displayed while the pool is being created.

W(%)

Threshold value of WARNING set for the pool. A hyphen (-) is displayed if the information is not available for this pool.

H(%)

Threshold value set for the pool as the high water mark

STIP

Setting for suspending Thin Image pairs when the high water mark threshold is exceeded:

- YES: Thin Image pairs are suspended.

- NO: Thin Image pairs are not suspended.

- - (hyphen): The information is not available for this pool.

VCAP(%)

The percentage of the subscription limit of V-VOL and Thin Image pair to the pool capacity:

▪ UNLIMITED: Unlimited.

▪ - (hyphen): The information is not available for this pool.

For VSP G130, G/F350, G/F370, G/F700, G/F900, a hyphen (-) indicating invalid is displayed.

TYPE

Platform type of the pool:

▪ OPEN: pool for open systems

▪ M/F: pool for mainframe systems

PM

Pool status:

▪ N: Normal status

▪ S: Shrinking or rebalancing

▪ NT: The pool for Thin Image is in the normal status.

▪ ST: The pool for Thin Image is shrinking or rebalancing.

PT

Pool type:

▪ HDP: Pool for Dynamic Provisioning

▪ HDT: Pool for Dynamic Tiering

▪ RT: Pool for active flash

▪ TI: Pool for Thin Image

▪ CW: Pool for Copy-on-Write Snapshot

▪ DM: Pool for Dynamic Provisioning with the data direct mapping attribute

POOL_NAME

Pool name

paircurchk* (for TrueCopy/global-active device)

The `paircurchk` command is used to check the currency of the TrueCopy secondary volume(s) by evaluating the data consistency based on pair status and fence level.

The following table specifies the data consistency for each possible state of a TrueCopy volume. A paired volume or group can be specified as the target of the `paircurchk` command. The `paircurchk` command assumes that the target is an S-VOL. If the `paircurchk` command is specified for a group, the data consistency of each volume in the group is checked, and all inconsistent volumes are found in the execution log file and displayed. The `paircurchk` command is also executed as part of the `horctakeover` command .

Chapter 2: Data management commands

Object Volume			Currency	
Attribute	**Status**	**Fence**	**Paircurchk**	**SVOL_takeover**
SMPL	-	-	To be confirmed	-
P-VOL	-	-	To be confirmed	-
S-VOL	COPY	Data	Inconsistent	Inconsistent
		Status		
		Never		
		Async	Inconsistent	Inconsistent
	PAIR	Data	OK	OK
		Status	OK	OK
		Never	To be analyzed	To be analyzed
	PAIR	Async	To be analyzed	OK (assumption)
	PFUL		To be analyzed	OK (assumption)
	PSUS	Data	Suspected	Suspected
		Status	Suspected	Suspected
		Never	Suspected	Suspected
	PSUS	Async	Suspected	Suspected
	PFUS		Suspected	Suspected
	PSUE	Data	OK	OK
	PDUB	Status	Suspected	Suspected
		Never	Suspected	Suspected
		Async	Suspected	OK (assumption)
	SSWS	Data	Suspected	-
		Status	Suspected	
		Never	Suspected	
		Async	Suspected	

Legend:

- To be confirmed = It is necessary to check the object volume, since it is not the secondary volume.

- Inconsistent = Data in the volume is inconsistent because it was being copied.

Object Volume			Currency	
Attribute	**Status**	**Fence**	**Paircurchk**	**SVOL_takeover**
▪ To be analyzed = It cannot be judged from the status of the secondary volume whether data is consistent or not. It is OK if the status of the primary volume is PAIR. It is Suspected if the status is PSUS or PSUE. ▪ Suspected = The primary volume data and secondary volume data are not consistent. ▪ OK (assumption) = Mirroring consistency is not assured, but as S-VOL of TrueCopy Async or Universal Replicator, the sequence of write data is ensured.				

Syntax

```
paircurchk { -h | -q | -z[x] | -I[H][M][instance#] or
    -I[TC][SI][instance#]
    | -g <group> | -d <pair Vol> | -d[g] <raw_device> [MU#][-n]
    | -d[g] <seq#> <LDEV#> [MU#] | -nomsg }
```

Options and parameters

-h

Displays help/usage and version information.

-q

Terminates the interactive mode and exits the command.

-z or -zx

Makes the **paircurchk** command enter the interactive mode. The -zx option guards performing of the HORCM in the interactive mode. When this option detects a HORCM shut down, interactive mode terminates.

OpenVMS cannot use the -zx option.

-I[H][M] [instance#] or -I[TC][SI] [instance#]

Specifies the command as [HORC]/[HOMRCF] or [TrueCopy]/[ShadowImage], and used when specifying the CCI instance number.

-g <group>

Specifies a group name defined in the configuration definition file. The command is executed for the specified group unless the -d <pair Vol> option is specified.

-d <pair Vol>

Specifies paired logical volume name defined in the configuration definition file. When this option is specified, the command is executed for the specified paired logical volume.

-d[g] <raw_device> [MU#][-n]

Searches whether the specified raw_device is included in a group on the configuration definition file (local instance). If it is in the group, the target volume is executed as the paired logical volume (-d) or group (-dg). This option is effective without specification of -g <group> option. If the specified raw_device is contained in two or more groups, the command is executed for the first group.

This -n option is specified to use a device file for NVMe-oF. Specifying this option enables the use of both device files for NVMe-oF and SCSI.

If the HORCC_NVME environment variable is specified, both device files for NVMe-oF and SCSI can be used without the -n option specified.

To specify this option, a command device must be enabled in a namespace on the storage system. For more information about whether a command device can be enabled in a namespace on the storage system, see the *Provisioning Guide* for your storage system.

-d[g] <seq#> <LDEV#> [MU#]

Searches whether the specified LDEV is included in a group on the configuration definition file (local instance). If it is in the group, the target volume is executed as the paired logical volume (-d) or group (-dg). This option is effective without specification of '-g <group>' option. If the specified LDEV is contained in two or more groups, the command is executed on the first group. The <seq#> <LDEV#> values can be specified in hexadecimal (add the '0x' prefix) or decimal notation.

- When specifying the serial number for VSP 5000 series, add a "5" at the beginning of the serial number. For example, for serial number 12345, enter 512345.

- When specifying the serial number for VSP G1x00 and VSP F1500, add a "3" at the beginning of the serial number. For example, for serial number 12345, enter 312345.

-nomsg

Suppresses messages to be displayed when this command is executed. This option must be specified at the beginning of a command argument. The command execution log is not affected by this option.

Returned values

The **paircurchk** command sets the following returned values during exit allowing you to check the execution results.

- **Normal termination:**

  - **0:** Data is consistent

- **Abnormal termination:**

  - **other than 0:** Refer to the error code for error details.

Error codes

The following table lists and describes the error codes for the **paircurchk** command. Unrecoverable errors are fixed and is not resolved, even after re-executing the command. If the command fails, the detailed status is logged in the CCI command log ($HORCC_LOG), even if the user script has no error handling.

Category	Error Code	Error Message	Recommended Action	Value
Volume status (Unrecoverable)	EX_VOLCUR	S-VOL currency error	Check volume list to see if an operation was directed to the wrong S-VOL.	225

Examples

The following shows an example of the **paircurchk** command for a group and the resulting display of inconsistent volumes in the specified group.

```
# paircurchk -g oradb
```

```
Group  Pair vol   Port   targ#  lun#  LDEV#  Volstatus  Status  Fence
To be...
oradb  oradb1     CL1-A  1      5     145    S-VOL      PAIR    NEVER   Analyzed
oradb  oradb2     CL1-A  1      6     146    S-VOL      PSUS    STATUS  Suspected
```

Output of the paircurchk command:

- Group: Displays the group name (dev_group) defined in the configuration definition file.

- Pair vol: Displays the pair volume name (dev_name) for the group defined in the configuration definition file.

- Port targ# lun#: Displays the port ID, TID, and LUN defined in the configuration definition file.

- LDEV#: Displays the LDEV number from the storage system.

- Volstat: Displays the volume attributes (P-VOL, S-VOL, SMPL).

- Status: Displays the status of a pair volume.

- Fence: Displays the fence level of a pair volume.

- To be...: Displays S-VOL's currency which was confirmed by mirror consistency check for S-VOL.

pairsyncwait*

The **pairsyncwait** command is used to confirm data consistency between the TrueCopy Async/Universal Replicator P-VOL and S-VOL by confirming that required writing was stored in the DFW area of RCU, and confirming whether the last writing just before this command reached the RCU DFW area.

This command obtains the MCU sidefile sequence number (the sequence number of the P-VOL within the consistency group ID), and the sequence number of RCU DFW in the CTG ID for <group> or <raw_device> specified by **pairsyncwait**. Both sequence numbers are the latest when the command is received. After that, this command compares the MCU and RCU sequence numbers which are the latest when the command is received at a certain interval.

If the RCU sequence number is over the value of the MCU sequence number within the term specified by **pairsyncwait**, this command displays the return code 0, indicating completion of synchronization. The `-nowait` option shows the latest sequence numbers (Q-marker) of the MCU P-VOL and CTG ID that are obtained when the command is received. The marker is a 10-digit hexadecimal number.

When a client issues the **pairsyncwait** command, this command is placed on the queue buffer for waiting in the HORCM daemon as a command request. HORCM gets the latest sequence number obtained from the MCU sidefile when the command is received and the sequence number whose block was transferred and stored in the DFW area of RCU with data consistency, and compares the latest sequence number of MCU sidefile with the sequence number of the RCU DFW area within the term. HORCM replies return codes to this command, when the write of MCU sidefile was stored in RCU DFW area.

Using this function, a client can confirm that a commit() has reached the remote site, and also the backup utility on a remote site can split the cascaded ShadowImage volumes (TrueCopy Async/Universal Replicator to TrueCopy Async/ShadowImage/Universal Replicator) without splitting TrueCopy Async/Universal Replicator.

More robust systems need to confirm the data consistency between the TrueCopy Async/Universal Replicator P-VOL and S-VOL. In DB operations (for example, Oracle), the commit() of DB transaction (see the following figure showing the synchronization for TrueCopy Async/Universal Replicator) is needed to confirm that a last writing for the commit() on a local site reached the remote site by using CCI-unique API command.

Note: Write() shows that synchronous writing or commit() of DB is used.
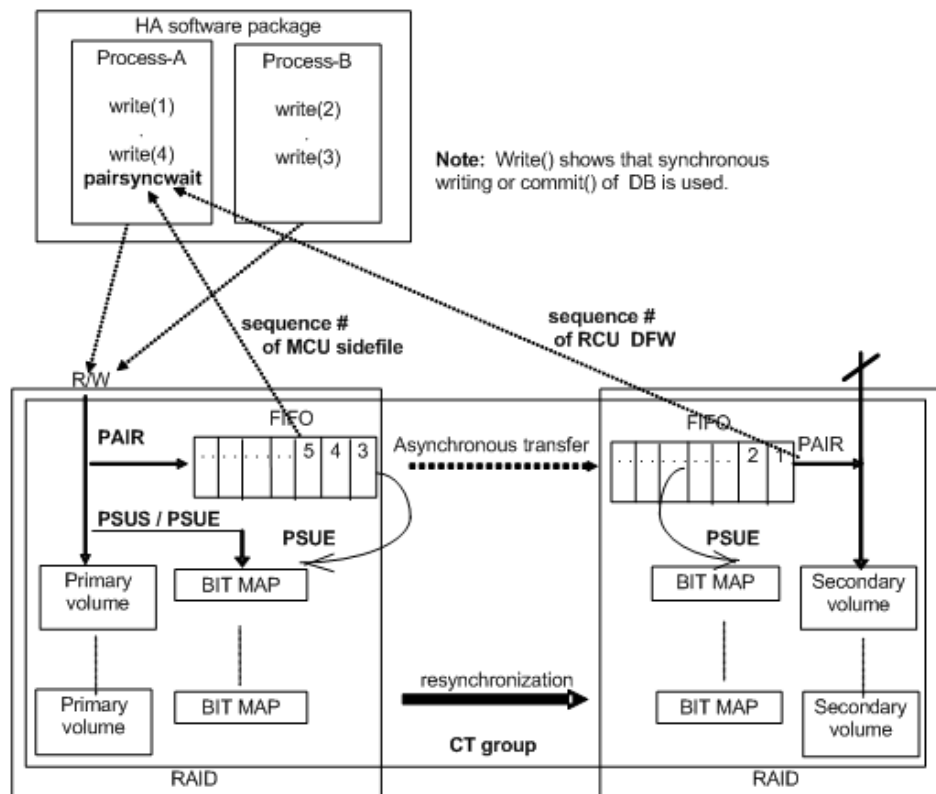
> ⚠️ **Caution:** If an extended consistency group is used in Universal Replicator, the **pairsyncwait** command cannot be used.

Syntax

```
pairsyncwait{ -h | -q | -z[x] | -I[H][M][instance#] or -I[TC][SI][instance#]| -g
<group> | -d <pair Vol>| -d[g] <raw_device> [MU#][-n] | -d[g] <seq#> <LDEV#> [MU#]| -
m <marker> | -t <timeout> | -nowait | -nomsg | -fq }
```

Options and parameters

-h

> Displays help/usage and version information.

-q

> Terminates the interactive mode and exits the command.

-z or -zx

> Makes the **pairsyncwait** command enter the interactive mode. The -zx option guards performing of the HORCM in the interactive mode. When this option detects a HORCM shut down, interactive mode terminates.

> OpenVMS cannot use the -zx option.

-I [H] [M] [instance#] or -I [TC] [SI] [instance#]

> Specifies the command as [HORC]/[HOMRCF] or [TrueCopy]/[ShadowImage], and used when specifying the CCI instance number.

-g <group>

> Specifies a group name defined in the configuration definition file. The command is executed for the specified group unless the `-d <pair Vol>` option is specified.

-d <pair Vol>

> Specifies paired logical volume name defined in the configuration definition file. When this option is specified, the command is executed for the specified paired logical volume.

-d[g] <raw_device> [MU#][-n]

> Searches whether the specified raw_device is included in a group on the configuration definition file (local instance). If it is in the group, the target volume is executed as the paired logical volume (-d) or group (-dg). This option is effective without specification of `-g <group>` option. If the specified raw_device is contained in two or more groups, the command is executed for the first group.
>
> This `-n` option is specified to use a device file for NVMe-oF. Specifying this option enables the use of both device files for NVMe-oF and SCSI.
>
> If the HORCC_NVME environment variable is specified, both device files for NVMe-oF and SCSI can be used without the `-n` option specified.
>
> To specify this option, a command device must be enabled in a namespace on the storage system. For more information about whether a command device can be enabled in a namespace on the storage system, see the *Provisioning Guide* for your storage system.

-d[g] <seq#> <LDEV#> [MU#]

> Searches whether the specified LDEV is included in a group on the configuration definition file (local instance). If it is in the group, the target volume is executed as the paired logical volume (-d) or group (-dg). This option is effective without specification of '-g <group>' option. If the specified LDEV is contained in two or more groups, the command is executed on the first group. The <seq#> <LDEV#> values can be specified in hexadecimal (add the '0x' prefix) or decimal notation.
>
> - When specifying the serial number for VSP 5000 series, add a "5" at the beginning of the serial number. For example, for serial number 12345, enter 512345.
>
> - When specifying the serial number for VSP G1x00 and VSP F1500, add a "3" at the beginning of the serial number. For example, for serial number 12345, enter 312345.

-m <marker>

> Specifies the sequence number of the MCU P-VOL in 10-digit hexadecimal, called the Q-marker. If the application gets Q-marker as the result of execution of **pairsyncwait** because of timeout or **'-nowait'**, the application can reconfirm the completion of Async transfer by using **pairsyncwait** with Q-marker. If the application does not specify Q-marker, CCI uses the latest sequence number obtained when CCI receives **pairsyncwait**. It is also possible to wait for the completion from S-VOL side with this option.
>
> Q-Marker format: iissssssss, where ii = regeneration number of pair volume, and ssssssss = latest sequence number obtained when the command is received on the side of the P-VOL.

-t <timeout>

  Specifies the timeout value to wait for the completion of RCU DFW area. The unit is 100 ms. MCU gets the latest sequence number obtained from RCU when the command is received at regular interval.

-nowait

  Gets the latest sequence numbers of the MCU P-VOL and CTG ID obtained when the command is received without waiting. When this option is specified, the latest sequence number of the MCU P-VOL obtained when the command is received is reported immediately, and `-t <timeout>`option is ignored.

-nomsg

  Suppresses messages to be displayed when this command is executed from a user program. This option must be specified at the beginning of the command arguments.

-fq

  Displays the number of remaining Q-Markers within the consistency group by adding 'QM-Cnt' to the last column. 'QM-Cnt' is shown as follows:

- When specifying '-nowait -fq'

  'QM-Cnt' is shown as the number of remaining Q-Marker at this time within consistency group.

- When specifying '-nowait -m <marker> -fq'

  'QM-Cnt' is shown as the number of remaining Q-Marker from the specified <marker> within consistency group.

- When specifying 'TIMEOUT' without '-nowait'

  'QM-Cnt' is shown as the number of remaining Q-Marker at this timeout within consistency group.

- 'QM-Cnt' is shown as '-', if the status for Q-Marker is invalid (that is, status is 'BROKEN' or 'CHANGED').

Example:

```
# pairsyncwait -g oradb -nowait -fq
```

```
UnitID  CTGID     Q-Marker  Status    Q-Num   QM-Cnt
```

```
    0      3    01003408ef  NOWAIT       2      120
```

```
# pairsyncwait -g oradb  -nowait -m 01003408e0 -fq
```

```
UnitID  CTGID     Q-Marker  Status    Q-Num   QM-Cnt
```

```
    0      3    01003408e0  NOWAIT       2      105
```

```
# pairsyncwait -g oradb  -t 50 -fq
```

```
UnitID  CTGID     Q-Marker  Status    Q-Num   QM-Cnt
```

```
    0      3    01003408ef  TIMEOUT      2        5
```

Restriction

Specified <group> volume must be P-VOL with status PAIR. Other cases reply with error (EX_INVVOL). It is possible to issue **pairsyncwait** from S-VOL side with -m <marker>.

Returned values

The `pairsyncwait` command sets the following returned values during exit allowing you to check the execution results.

- When the `-nowait` option is specified

  - **Normal termination:**

    - **0:** The status is NOWAIT.

  - **Abnormal termination:**

    - **other than 0 to 127:** Refer to the error code for error details.

- When the `-nowait` option is not specified

  - **Normal termination:**

    - **0:** The status is DONE (completion of synchronization).

    - **1:** The status is TIMEOUT (timeout).

    - **2:** The status is BROKEN (Q-marker synchronized process is rejected).

    - **3:** The status is CHANGED (Q-marker is invalid due to resynchronize).

  - **Abnormal termination:**

    - **other than 0 to 127:** Refer to the error code for error details.

Error codes

Unrecoverable errors are fixed and not resolved, even after re-executing the command. If the command failed, the detailed status is logged in the CCI command log ($HORCC_LOG), even if the user script has no error handling.

Category	Error Code	Error Message	Recommended Action	Value
Volume status (Unrecoverable)	EX_INVVOL	Invalid volume status	Confirm pair status using pairdisplay -l.	222

Examples

If the latest sequence number is not synchronized within 10 seconds when the command is received:

```
# pairsyncwait -g oradb -t 100
UnitID    CTGID    Q-Marker      Status      Q-Num
0         3    01003408ef    TIMEOUT        3
```

If the latest sequence number is synchronized within 10 seconds when the command is received (restored in S-VOLs):

```
# pairsyncwait -g oradb -t 100
```

Chapter 2: Data management commands

```
UnitID    CTGID    Q-Marker     Status     Q-Num
0             3    01003408ef    DONE         2
```

When you check the synchronization status for a specific checkpoint, specify the Q-Maker value acquired by **-nowait** with the -m option, and then check the Status.

If you obtain the latest sequence number when the command is received:

```
 # pairsyncwait -g oradb -nowait
UnitID    CTGID    Q-Marker     Status     Q-Num
0             3    02006811DE   NOWAIT        2
```

If the specified sequence number (02006811DE) is not synchronized within 10 seconds:

```
# pairsyncwait -g oradb -t 100 -m 02006811DE
UnitID    CTGID    Q-Marker     Status     Q-Num
0             3    0200682000   TIMEOUT       3
```

If the specified sequence number (02006811DE) is synchronized within 10 seconds:

```
# pairsyncwait -g oradb -t 100 -m 02006811DE
UnitID    CTGID    Q-Marker     Status     Q-Num
0             3    0200682F8D   DONE          0
```

Description of the `pairsyncwait` command output:

UnitID

Unit ID in the case of multiple storage system connection

CTGID

Consistency group ID within Unit ID

Q-Marker

The latest sequence # of MCU PVol (Marker) when the command is received.

Status

The status after the execution of command.

Q-Num

The number of process queues to wait for synchronization within the consistency group (CTG).

QM-Cnt

The number of remaining Q-Markers within consistency group of the Unit. TrueCopy Async/Universal Replicator sends a token called 'dummy recordset' at regular intervals, therefore QM-Cnt always shows '2' or '3' even if host has NO writing.

Arithmetic expression for determining the remaining data in a consistency group (CTG):

```
Remaining data in CTG = sidefile capacity × sidefile percentage ÷ 100
```

Chapter 2: Data management commands

Sidefile percentage is the rate showed to '%' column with 'PAIR' state by Pairdisplay command. Sidefile capacity is the capacity within 30% to 70% of the cache setting as the sidefile.

Arithmetic expression for determining the average data per Q-Marker in a consistency group (CTG):

```
Data per Q-Marker = Remaining data in CTG ÷ QM-Cnt
```

horctakeover*

The `horctakeover` command is a scripted command for executing several takeover operations. The `horctakeover` command checks the specified volume's or group's attributes (`paircurchk`), decides the takeover function based on the attributes, executes the chosen takeover function, and returns the result. The four takeover functions designed for HA software operation are (see `horctakeover` command functions in the *User and Reference Guide*): takeover-switch, swap-takeover, PVOL-takeover, and SVOL-takeover. A paired volume or a group can be specified as the target of the TrueCopy takeover command. If SVOL-takeover is specified for a group, the data consistency check is executed for all volumes in the group, and all inconsistent volumes are found in the execution log file and displayed (same as `paircurchk` command). You can execute the `horctakeover` command only for remote copy pairs.

The `horctakeover` command allows swapping of the primary and secondary volumes, so that if the primary or secondary volume is switched due to a server error or package transfer, duplex operations can be continued using the reversed volumes. When control is handed over to the current node, swapping the volumes again eliminates the need to copy them. The `horctakeover` command also allows the secondary volume to be separated for disaster recovery operations.

Syntax

```
horctakeover { -h | -q | -z[x] | -I[H][M][instance#] or
    -I[TC][SI][instance#] | -g <group> | -d <pair Vol>
    | -d[g] <raw_device> [MU#][-n] | -d[g] <seq#> <LDEV#> [MU#]
    | -S | -l | -t <timeout> | -nomsg }
```

Options and parameters

-h

Displays help/usage and version information.

-q

Terminates the interactive mode and exits the command.

-z or -zx

Makes the `horctakeover` command enter the interactive mode. The `-zx` option guards performing of the HORCM in the interactive mode. When this option detects a HORCM shut down, interactive mode terminates.

OpenVMS cannot use the -zx option.

-I[H][M] [instance#]or -I[TC][SI] [instance#]

Specifies the command as [HORC]/[HOMRCF] or [TrueCopy]/[ShadowImage], and used to specify the CCI instance number.

-g <group>

Specifies a group name defined in the configuration definition file. The command is executed for the specified group unless the -d <pair Vol> option is specified.

-d <pair Vol>

Specifies paired logical volume name defined in the configuration definition file. When this option is specified, the command is executed for the specified paired logical volume.

-d[g] <raw_device> [MU#][-n]

Searches whether the specified raw_device is included in a group on the configuration definition file (local instance). If it is in the group, the target volume is executed as the paired logical volume (-d) or group (-dg). This option is effective without specification of -g <group> option. If the specified raw_device is contained in two or more groups, the command is executed for the first group.

This -n option is specified to use a device file for NVMe-oF. Specifying this option enables the use of both device files for NVMe-oF and SCSI.

If the HORCC_NVME environment variable is specified, both device files for NVMe-oF and SCSI can be used without the -n option specified.

To specify this option, a command device must be enabled in a namespace on the storage system. For more information about whether a command device can be enabled in a namespace on the storage system, see the *Provisioning Guide* for your storage system.

-d[g] <seq#> <LDEV#> [MU#]

Searches whether the specified LDEV is included in a group on the configuration definition file (local instance). If it is in the group, the target volume is executed as the paired logical volume (-d) or group (-dg). This option is effective without specification of '-g <group>' option. If the specified LDEV is contained in two or more groups, the command is executed on the first group. The <seq#> <LDEV#> values can be specified in hexadecimal (add the '0x' prefix) or decimal notation.

- When specifying the serial number for VSP 5000 series, add a "5" at the beginning of the serial number. For example, for serial number 12345, enter 512345.

- When specifying the serial number for VSP G1x00 and VSP F1500, add a "3" at the beginning of the serial number. For example, for serial number 12345, enter 312345.

-S

Selects and executes SVOL-takeover. The target volume of the local host must be an S-VOL. If this option is specified, then the following '-l' option is invalid.

-l

>Enables read and write to the primary volume(s) by a local host only without a remote host, and executes PVOL-takeover when the primary volume cannot be used because it is fenced (fence = DATA or STATUS, state = PSUE or PDUB, or PSUE or PDUB volume is contained in the group). If the primary volume can be accessed, nop-takeover is executed. The target volume of the local host must be a P-VOL.

-t <timeout>

>Can be specified for asynchronous pairs only, ignored for synchronous pairs. Specifies the maximum time to wait (in seconds) for swap-takeover and SVOL-takeover operation to synchronize the P-VOL and S-VOL. If this timeout occurs, the `horctakeover` command fails with EX_EWSTOT. To avoid timeout, set this value less than or equal to the start-up timeout value of the HA control script. This option is required for asynchronous pairs.

-nomsg

>Suppresses messages to be displayed when this command is executed. This option must be specified at beginning of a command argument. The command execution log is not affected by this option.

Returned values

The `horctakeover` command sets the following returned values during exit allowing you to check the execution results.

- **Normal termination:**

  - **0:** Nop-takeover (no operation).

  - **1:** Swap-takeover was successfully executed.

  - **2:** SVOL-takeover was successfully executed.

  - **3:** PVOL-SMPL-takeover was successfully executed.

  - **4:** PVOL-PSUE-takeover was successfully executed. (This value depends on the microcode level.)

  - **5:** SVOL-SSUS-takeover was successfully executed. (This value depends on the microcode level.)

- **Abnormal termination:**

  - **other than 0-5:** Refer to the error code for error details.

Error codes

The following table lists and describes the error codes for the `horctakeover` command. Unrecoverable errors are fixed and not resolved, even after re-executing the command. If the command fails, the detailed status is logged in the CCI command log ($HORCC_LOG), even if the user script has no error handling.

Category	Error Code	Error Message	Recommended Action	Value
Volume status (Unrecoverable)	EX_ENQVOL	Unmatched volume status within the group	Confirm status using **pairdisplay** command. Make sure all volumes in the group have the same fence level and volume attributes.	236
	EX_INCSTG	Inconsistent status in group	Confirm pair status using **pairdisplay**.	229
	EX_EVOLCE	Pair volume combination error	Confirm pair status using **pairdisplay**, and change combination of volumes.	235
	EX_VOLCUR	S-VOL currency error	Check volume list to see if an operation was directed to the wrong S-VOL.	225
	EX_VOLCUE	Local volume currency error	Confirm pair status of the local volume.	224
	EX_VOLCRE	Local and remote volume currency error	Confirm pair status of remote and local volumes using **pairdisplay** command.	223
Timer (Recoverable)	EX_EWSTOT	Timeout waiting for specified status	Increase timeout value using -t option.	233

raidscan*

The **raidscan** command displays configuration and status information for the specified port/TID(s)/LUN/MU#. The information is acquired directly from the storage system (not the configuration definition file).

Syntax

```
raidscan { -h | -q | -z[x] | -I[H][M][instance#] or
    -I[TC][SI][instance#] | -p <port#> [hgrp]|-pd[g] <raw_device>[-n]
    | -s <Seq#> | -t <targ> | -l <lun> | [ -f[xfgde] ] | -CLI
    | -find[g] [op] [MU#] | [-g <group>][-n] | -pi <strings>
    | -m <MU#> | -fw }
```

Options and parameters

-h

Displays help/usage and version information.

-q

    Terminates the interactive mode and exits the command.

-z or -zx

    Makes the **raidscan** command enter the interactive mode. The `-zx` option guards performing of the HORCM in the interactive mode. When this option detects a HORCM shut down, interactive mode terminates.

    OpenVMS cannot use the `-zx` option.

-I[H][M] [instance#]or -I[TC][SI] [instance#]

    Specifies the command as [HORC]/[HOMRCF] or [TrueCopy]/[ShadowImage], and used when specifying the CCI instance number.

-p <port#> [hgrp]

    Specifies the port ID of the port to be scanned.

- Valid ports are CL1-A to CL1-R and CL2-A to CL2-R (excluding CL1-I, CL1-O, CL2-I, CL2-O).

- Valid expanded ports are CL3-A to CL3-R (excluding CL3-I and CL3-O) and CLG-A to CLG-R (excluding CLG-I and CLG-O).

    The port is not case sensitive (for example, CL1-A = cl1-a = CL1-a = cl1-A).

    This option must be specified if '`-find`' or '`-pd <raw_device>`'option is not specified.

    [hgrp] is specified to display only the LDEVs mapped to a host group on a port.

    If only the `<port>` option is specified, the absolute LUN of CCI is displayed for LU#.

-pd[g] <raw_device> [-n]

    Specifies the raw device name. This option finds Seq# and port_name of the storage system to which the specified device can be connected, and scans the port of the storage system which corresponds with the unit ID that searches the unit ID from Seq#. This option must be specified if the '-find' option is not specified. If this option is specified, the following -s <seq#> option is not valid.

    The -pdg option is used when displaying a LUN on the host view by finding a host group.

    This `-n` option is specified to use a device file for NVMe-oF. Specifying this option enables the use of both device files for NVMe-oF and SCSI.

    If the HORCC_NVME environment variable is specified, both device files for NVMe-oF and SCSI can be used without the `-n` option specified.

    To specify this option, a command device must be enabled in a namespace on the storage system. For more information about whether a command device can be enabled in a namespace on the storage system, see the *Provisioning Guide* for your storage system.

-s <Seq#>

Used to specify the seq# (serial#) of the storage system when this option cannot specify the unit ID which is contained for '`-p <port>`' option. This option scans the port specified by '`-p <port>`' option of the storage system which corresponds with the unit ID that searches the unit ID from seq#. If this option is specified, then the unit ID that is contained in '`-p <port>`' option is invalid.

- When specifying the serial number for VSP 5000 series, add a "5" at the beginning of the serial number. For example, for serial number 12345, enter 512345.

- When specifying the serial number for VSP G1x00 and VSP F1500, add a "3" at the beginning of the serial number. For example, for serial number 12345, enter 312345.

-t <targ>

Specifies the target ID of the specified port. If this option is not specified, the command applies to all target IDs.

-l <lun>

Specifies the LUN of the specified target ID. If this option is not specified, the command applies to all LUNs. If this option is specified, the TID must also be specified.

-f or -ff

Specifies display of volume-type for a display column. If this is specified, `-f[g] [d]` option is invalid.

-fx

Displays the LDEV number in hexadecimal notation.

-fg

Specifies display of group_name for a display column. This option searches a group on the configuration definition file (local CCI instance) from the scanned LDEV, and displays a group_name when the scanned LDEV is contained in the group. If this option is specified, the `-f[f]` option is not allowed and the `-f[d]` option is invalid.

-fd

Displays the Device_File that was registered to the group of the HORCM in the output, based on the LDEV (as defined in local instance config. def. file). If this option is specified, `-f[f][g]` option is not allowed.

Specify ENABLE for HORCM_NVME in the configuration file. If ENABLE is not specified, a device file name for NVMe-oF is not displayed.

-fe

Displays the serial# (E-Seq#) and LDEV# (E-LDEV#) of the external LUNs only mapped to the LDEV. If the external LUN mapped to the LDEV on a specified port does not exist, then this option does nothing. Also if this option is specified, `-f[f][g] [d]` option is not allowed.

Chapter 2: Data management commands

Display example:

```
# raidscan -p cl1-a-0 -fe -CLI
```

```
PORT# /ALPA/C TID# LU#  Seq# Num LDEV# P/S Status Fence
E-Seq# E-LDEV#
```

```
CL1-A-0 ef  0    0 48 62468   2   256 SMPL     -    -
 30053      17
```

```
CL1-A-0 ef  0    0 49 62468   2   272 SMPL     -    -
 30053      23
```

```
CL1-A-0 ef  0    0 50 62468   1   288 SMPL     -    -
 30053      28
```

-CLI

> Specifies display for command line interface (CLI). This option displays to the same position that defined number of columns, and displays one header. The delimiters between columns are displayed as spaces or hyphens (-).

Display example:

```
Port# TargetID# Lun#  Seq# Num LDEV#   P/S Status Fence  P-
  Seq#  P-LDEV#
```

```
CL1-C       1    0 30053   1   274  SMPL     -    -
    -       -
```

```
CL1-C       2    2 30053   1   260 P-VOL   PAIR  NEVER
 30053      268
```

```
CL1-C       2    3 30053   1   261 P-VOL   PAIR  NEVER
 30053      269
```

-pi <strings>

> Changes a strings via STDIN for -find option to '<strings>'. If this option is specified, the -find option is ignored a raw device file provided via STDIN, and <strings> is used as input. A <strings> must be specified within 255 characters.

-m <MU#>

> This option is used for displaying only the specified mirror descriptor. If you want to display all mirror descriptor, specify '-m all' for displaying all MUs.

-fw

Specify this option to display the NAA identifier (LUN WWN) for the LU.

-find [op] [MU#] [-n]

Executes the specified [op] using a raw device file provided via STDIN. If the -pi <strings> option is specified, this option does not use a strings via STDIN, and -pi <strings> is used as input.

Restrictions: Special files via STDIN are specified in the following ways:

- HP-UX:
  - /dev/rdsk/*
  - /dev/rdisk/disk*

- Solaris:
  - /dev/rdsk/*s2
  - c*s2

- Linux:
  - /dev/sd...
  - /dev/rd...
  - /dev/raw/raw*
  - /dev/nvme...

- zLinux:
  - /dev/sd...
  - /dev/dasd…
  - /dev/rd...
  - /dev/raw/raw*

- AIX:
  - /dev/rhdisk*
  - /dev/hdisk*
  - hdisk*

- DIGITAL or Tru64:
  - /dev/rrz*c
  - /dev/rdisk/dsk*c
  - /dev/cport/scp*

- DYNIX:
  - /dev/rdsk/sd*
  - sd* for only unpartitioned raw device

- IRIX64:
  - /dev/rdsk/*vol
  - /dev/rdsk/node_wwn/*vol/*
  - /dev/dsk/*vol
  - /dev/dsk/node_wwn/*vol/*

- Windows:
  - hdX-Y
  - $LETALL
  - $Volume
  - $Phys
  - D:\Vol(Dms,Dmt,Dmr)X\DskY
  - \Vol(Dms,Dmt,Dmr)X\DskY
  - D:\DskX\pY
  - \DskX\pY

For further information on LDM volumes for Windows systems, see **Volume Discovery Function** in the *User and Reference Guide*.

Lines starting with '#' via STDIN are interpreted as comments.

Lines starting with 'quit' via STDIN are interpreted as exit.

This -n option is specified to use a device file for NVMe-oF. Specifying this option enables the use of both device files for NVMe-oF and SCSI.

If the HORCC_NVME environment variable is specified, both device files for NVMe-oF and SCSI can be used without the -n option specified.

To specify this option, a command device must be enabled in a namespace on the storage system. For more information about whether a command device can be enabled in a namespace on the storage system, see the *Provisioning Guide* for your storage system.

-find[g] [-n]

Displays the port, target ID, LUN (RAID storage system notation) and so on that was mapped for LDEV using a special file (raw device file) provided via STDIN. If target ID and LUN are Unknown for the target device file, you should start CCI without any description for HORCM_DEV and HORCM_INST, and should describe the shown port, target ID, and LUN for HORCM_DEV. This option also uses the -fx option to display the LDEV numbers in hexadecimal.

The -findg option is used to show a LUN on the host view by finding a host group.

This -n option is specified to use a device file for NVMe-oF. Specifying this option enables the use of both device files for NVMe-oF and SCSI.

If the HORCC_NVME environment variable is specified, both device files for NVMe-oF and SCSI can be used without the -n option specified.

Chapter 2: Data management commands

To specify this option, a command device must be enabled in a namespace on the storage system. For more information about whether a command device can be enabled in a namespace on the storage system, see the *Provisioning Guide* for your storage system.

-find inst [-n]

Registers the Device File name (raw device file provided via STDIN) to all mirror descriptors of the LDEV map table for HORCM, permits the matching volumes on horcm.conf in protection mode, and is started automatically. Therefore, you do not need to use this option normally. This option is also terminated to avoid wasteful scanning when the registration has been finished with based on HORCM. Therefore if HORCM does not need the registration any more, then nothing is done and it exits. This option can be used with '-fx' option to display LDEV numbers in hexadecimal.

Example for HP-UX:

ioscan -fun | grep rdsk | raidscan -find inst

DEVICE_FILE Group PairVol PORT TARG LUN M SERIAL LDEV

/dev/rdsk/c0t3d0 oradb oradev1 CL1-D 3 0 - 35013 17

/dev/rdsk/c0t3d0 oradb oradev1 CL1-D 3 0 0 35013 17

/dev/rdsk/c0t3d0 oradb1 oradev2 CL1-D 3 0 1 35013 17

> 📄 **Note:**
>
> When multiple device files share the same LDEV, the first device file is registered to the LDEV map table.
>
> - Group: Displays the group name (dev_group) defined in the configuration definition file.
>
> - PairVol: Displays the paired volume name (dev_name) within the group defined in the configuration definition file.
>
> - PORT: Displays the port number (port#) defined in the configuration definition file.
>
> - TARG: Displays the target ID (TargetID) defined in the configuration definition file.
>
> - LUN: Displays the LUN (LU#) defined in the configuration definition file.
>
> - M: Displays the MUN (MU#) defined in the configuration definition file. MU# for TrueCopy/GAD are shown as '-'. MU# for ShadowImage/Copy-on-Write Snapshot are shown as '0', '1', '2'.
>
> - SERIAL: Displays the production (serial#) number of the RAID storage system.
>
> - LDEV: Displays the LDEV# within the RAID storage system.

This -n option is specified to use a device file for NVMe-oF. Specifying this option enables the use of both device files for NVMe-oF and SCSI.

If the HORCC_NVME environment variable is specified, both device files for NVMe-oF and SCSI can be used without the -n option specified.

To specify this option, a command device must be enabled in a namespace on the storage system. For more information about whether a command device can be enabled in a namespace on the storage system, see the *Provisioning Guide* for your storage system.

-find verify [MU#] [-n]

Displays the relationship between group on the configuration definition file and Device_File registered to the LDEV map tables, based on the Device File name (raw device file provided via STDIN). This option can be used with '-fx' and '-fd' options. This option is affected by the command execution environment (HORCC_MRCF).

```
# ioscan -fun | grep rdsk | raidscan -find verify
```

DEVICE_FILE	Group	PairVol	PORT	TARG	LUN	M	SERIAL	LDEV
/dev/rdsk/c0t3d0	oradb	oradev1	CL1-D	3	0	0	35013	17
/dev/rdsk/c0t3d1	oradb	oradev2	CL1-D	3	1	0	35013	18
/dev/rdsk/c0t3d2	-	-	-	-	-	0	35013	19

```
# ioscan  -fun | grep rdsk | raidscan -find verify 1 -fd
```

DEVICE_FILE	Group	PairVol	Device_File	M	SERIAL	LDEV
/dev/rdsk/c0t3d0	oradb	oradev1	c0t3d0	1	35013	17
/dev/rdsk/c0t3d1	oradb	oradev2	Unknown	1	35013	18
/dev/rdsk/c0t3d2	-	-	-	1	35013	19

Chapter 2: Data management commands

> 📄 **Note:**
>
> If the device name is different between DEVICE_FILE and Device_File, then it shows shared LDEV among multiple device files.
>
> - Group: Displays the group name (dev_group) defined in the configuration definition file.
>
> - PairVol: Displays the paired volume name (dev_name) within the group defined in the configuration definition file.
>
> - PORT: Displays the port number (port#) defined in the configuration definition file.
>
> - TARG: Displays the target ID (TargetID) defined in the configuration definition file.
>
> - LUN: Displays the LUN (LU#) defined in the configuration definition file.
>
> - M: Displays the MUN (MU#) defined in the configuration definition file. MU# for TrueCopy/GAD are shown as '-'. MU# for ShadowImage/Copy-on-Write Snapshot are shown as '0', '1', '2'.
>
> - Device_File: Displays Device_File which is registered to the LDEV map tables into the CCI.
>
> - SERIAL: Displays the production (serial#) number of the RAID storage system.
>
> - LDEV: Displays the LDEV# within the RAID storage system.

This -n option is specified to use a device file for NVMe-oF. Specifying this option enables the use of both device files for NVMe-oF and SCSI.

If the HORCC_NVME environment variable is specified, both device files for NVMe-oF and SCSI can be used without the -n option specified.

To specify this option, a command device must be enabled in a namespace on the storage system. For more information about whether a command device can be enabled in a namespace on the storage system, see the *Provisioning Guide* for your storage system.

-find[g] conf [MU#] [-g <group>]

Displays the port, target ID, LUN in horcm.conf image by using a special file (raw device file) provided via STDIN. If **target ID & LUN are Unknown** for the target device file, you must start CCI without any description for HORCM_DEV and HORCM_INST, and should be described the shown port, target ID, LUN for HORCM_DEV. This option can be used with the '-fx' option.

[-g <group>] specifies the group for 'dev_group' on horcm.conf. If omitted, the group applies 'VG' as default.

Chapter 2: Data management commands

The -findg option is used when displaying a LUN on the host view by finding a host group.

```
# cat /etc/horcmperm.conf | raidscan -find conf 0 -g ORA
```

```
HORCM_DEV
```

```
#dev_group          dev_name          port#      TargetID      LU#      MU#
```

```
# /dev/rdsk/c23t0d0    SER =     61456  LDEV =  192 [ FIBRE FCTBL = 4 ]
```

```
ORA                ORA_000          CL2-J          0         0         0
```

```
# /dev/rdsk/c23t0d1    SER =     61456  LDEV =  193 [ FIBRE FCTBL = 4 ]
```

```
ORA                ORA_001          CL2-J          0         1         0
```

```
# /dev/rdsk/c23t0d2    SER =     61456  LDEV =  194 [ FIBRE FCTBL = 4 ]
```

```
ORA                ORA_002          CL2-J          0         2         0
```

```
# /dev/rdsk/c23t0d3    SER =     61456  LDEV =  195 [ FIBRE FCTBL = 4 ]
```

```
ORA                ORA_003          CL2-J          0         3         0
```

```
# ERROR [CMDDEV] /dev/rdsk/c23t0d7    SER =     61456  LDEV =  259 [ OPEN-3-
CM  ]
```

If the target device has shared an LDEV among multiple device files and an LDEV is displayed by another target device already, then its target device is suppressed as a comment as shown below:

```
# ERROR [LDEV LINK] /dev/rdsk/c24t0d3  SER =     61456  LDEV =  195 [FIBRE
FCTBL = 4]
```

Chapter 2: Data management commands

If the target device does not have a valid MU#, then its target device is suppressed as a comment as shown below:

```
# ERROR [INVALID MUN (2 < 1)] /dev/rdsk/c24t0d3  SER =  61456  LDEV =  195
[ OPEN-3 ]
```

If the target device is mixed with a different RAID TYPE, then its target device is suppressed as a comment as shown below:

```
# ERROR [MIXING RAID TYPE] /dev/rdsk/c24t0d3   SER =   61456  LDEV =  195
[ OPEN-3  ]
```

-find sync[d] [MU#] [-g <group>]

Flushes the system buffer associated to a logical drive which corresponds to a `[-g <group>]` through the KEY WORD **($Volume,$LETALL, $Physical)** provided via STDIN.

`[-g <group>]` specifies the group for 'dev_group' on horcm.conf. If this option is not specified, then flushes the system buffer associated to all groups for the local instance.

Example of flushing the system buffer associated to ORB group through $Volume (Windows):

```
echo $Volume | raidscan -find sync -g ORB    or
```

```
raidscan -pi $Volume -find sync -g ORB
```

```
[SYNC] : ORB ORB_000[-] -> \Dmt1\Dsk1 : Volume{bf48a395-0ef6-11d5-8d69-
00c00d003b1e}
```

```
[SYNC] : ORB ORB_001[-] -> \Dmt1\Dsk2 : Volume{bf48a395-0ef6-11d5-8d69-
00c00d003b1e}
```

```
[SYNC] : ORB ORB_002[-] -> \Dmt1\Dsk3 : Volume{bf48a395-0ef6-11d5-8d69-
00c00d003b1e}
```

Example of flushing the system buffer associated to all groups for the local instance (Windows):

```
echo $Volume | raidscan -find sync          or
```

```
raidscan -pi $Volume -find sync
```

```
[SYNC] : ORA ORA_000[-] -> \Vol44\Dsk0 : Volume{56e4954a-28d5-4824-a408-
3ff9a6521e5d}
```

```
[SYNC] : ORA ORA_000[-] -> \Vol45\Dsk0 : Volume{56e4954a-28d5-4824-a408-
3ff9a6521e5e}
```

```
[SYNC] : ORB ORB_000[-] -> \Dmt1\Dsk1  : Volume{bf48a395-0ef6-11d5-8d69-
00c00d003b1e}
```

```
[SYNC] : ORB ORB_001[-] -> \Dmt1\Dsk2  : Volume{bf48a395-0ef6-11d5-8d69-
00c00d003b1e}
```

```
[SYNC] : ORB ORB_002[-] -> \Dmt1\Dsk3  : Volume{bf48a395-0ef6-11d5-8d69-
00c00d003b1e}
```

> **Note:**
> 1. The option cannot be specified the device object name as follows: D:\Vol(Dms,Dmt,Dmr)X\DskY,\Vol(Dms,Dmt,Dmr)X\DskY
> 2. Sync executes the following behavior under any conditions:
>   - If the logical drive which corresponds to a [-g <group>] is not open by any applications, then sync flushes the system buffer to a drive and makes the dismount state for this drive.
>   - If the logical drive which corresponds to a [-g <group>] is already opened by any applications, then sync flushes only the system buffer to a drive. This is allowed to flush the system buffer before pairsplit with mounting the P-VOL (opening state), and indicates the behavior as [FLUSH] below: [FLUSH] : ORA ORA_000[-] -> \Vol44\Dsk0 : Volume{56e4954a-28d5-4824-a408-3ff9a6521e5d}

-fn
Specifies the namespace identifier (NGU ID).

-nguid
Displays the LDEV number and the serial number for the entered NGUID.

Returned values

None

Error codes

None

`raidscan` command examples for Fibre Channel ports

```
# raidscan -p cl1-r
```

```
PORT#/ALPA/C,TID#,LU#Num(LDEV#...)P/S,Status,Fence,LDEV#,P-Seq#P-LDEV#
CL1-R/  ce/15,15, 7 5(100,101..)P-VOL PAIR    NEVER   100,5678   200
CL1-R/  ce/15,15, 6 5(200,201..)SMPL  ----    ----  ----  ----  -----
```

```
# raidscan -p cl1-r -f
```

```
PORT#/ALPA/C,TID#,LU#Num(LDEV#...)P/S,Status,Fence,LDEV#,Vol.Type
CL1-R/  ce/15,15, 7 5(100,101..)P-VOL PAIR    NEVER   100,OPEN-3
CL1-R/  ce/15,15, 6 5(200,201..)SMPL  ----    ----   ----OPEN-3
```

Example of the `-find` option for `raidscan`

```
# ls /dev/rdsk/* | raidscan -find
```

```
DEVICE_FILE      UID S/F  PORT TARG LUN SERIAL LDEV    PRODUCT_ID
/dev/rdsk/c0t0d4  0   S CL1-M   0   4  31168  216 OPEN-3-CVS-CM
/dev/rdsk/c0t0d1  0   S CL1-M   0   1  31168  117 OPEN-3-CVS
/dev/rdsk/c1t0d1  -   - CL1-M   -   -  31170  121 OPEN-3-CVS
```

Example of the `-fw` option for `raidscan`

```
#raidscan -p cl1-e-0 -l 0 -CLI -fw -IH
```

```
PORT#   /ALPA/C TID# LU#   Seq# Num LDEV# P/S   Status Fence LUN-WWN
CL1-E-0 cd  4   16   0    493017  1   768 P-VOL   PAIR NEVER
60060e8016016400000101640000893
#raidscan -p cl1-e-0 -l 0 -fw -IM -CLI
PORT#   /ALPA/C TID# LU#   Seq# Num LDEV# P/S Status LUN-WWN
CL1-E-0 cd  4   16   0 0 493017   1   768 SMPL    - 60060e8016016400000101640000893
CL1-E-0 cd  4   16   0 1 493017   1   768 SMPL    - 60060e8016016400000101640000893
CL1-E-0 cd  4   16   0 2 493017   1   768 SMPL    - 60060e8016016400000101640000893
```

Description of the `raidscan` command output:

Port#, ALPA/C, TID#, LU#

   Port ID, arbitrated loop physical address, target ID, LUN.

MU#

   For ShadowImage, `raidscan` displays the MU# for each LUN (for example, LUN 7-0, 7-1, 7-2).

Num(LDEV#…):

   Number of LDEVs and LDEV ID for a LUSE volume

P/S

   Volume attribute

Status

   Status of the paired volume

Fence

   Fence level (TrueCopy/global-active device only)

P-Seq#

   Serial # of the storage system which contains the partner volume of the pair.

   - The serial number for VSP 5000 series is displayed with a "5" at the beginning ("512345" = serial number 12345).

   - The serial number for VSP G1x00 and VSP F1500 is displayed with a "3" at the beginning ("312345" = serial number 12345).

P-LDEV#

   LDEV number of the partner volume of the pair

Vol.Type

   Logical unit (LU) type (for example, OPEN-V, OPEN-9)

Group

   Group name (dev_group) as described in the configuration definition file

UID

   Displays the unit ID for multiple storage system configuration. If UID appears as '-', the command device for HORCM_CMD is not found.

S/F

   Displays "S" for a SCSI port or "F" for a port other than SCSI. For storage system models newer than VSP and HUS VM, confirm the port type with the TYPE field which is the execution result of the `raidcom get port` command.

PORT

   Displays the RAID storage system port number

TARG

   Displays the target ID (that was converted by the fibre conversion table)

LUN

Displays the LUN (that was converted by the fibre conversion table)

SERIAL

Displays the production (serial#) number of the RAID storage system.

- The serial number for VSP 5000 series is displayed with a "5" at the beginning ("512345" = serial number 12345).

- The serial number for VSP G1x00 and VSP F1500 is displayed with a "3" at the beginning ("312345" = serial number 12345).

LDEV

Displays the LDEV# within the RAID storage system

PRODUCT_ID

Displays product-id field in the STD inquiry page

LUN-WWN

NAA identifier (LUN WWN) of the LU.

Example of the `-fn` option for raidscan

```
#raidscan -p cl1-e-0 -l 0 -CLI -fn -IH
PORT# /ALPA/C TID# LU# Seq# Num LDEV# P/S Status Fence NGUID
CL1-E-0 cd 4 16 0 493017 1 768 P-VOL PAIR NEVER 1234…
```

Description of each column in output example:

NGUID

Displays the namespace identifier (NGU ID). If NGUID is not assigned, 'Unknown' is displayed.

Example of the LDEV number and the serial number of NGUID: 12345678901234567890012345

```
#raidscan -nguid 12345678901234567890012345
NGUID LDEV# SERIAL
12345678901234567890012345 2345 12345
```

raidar*

The `raidar` command displays configuration, status, and I/O activity information for the specified port/TID(s)/LUN at the specified time interval. The configuration information is acquired directly from the storage system (not from the configuration definition file).

The I/O activity of a TrueCopy, TrueCopyAsync, Universal Replicator, or global-active device S-VOL in the COPY or PAIR state includes TC/TC Async/UR/GAD remote I/Os (update copy operations) in addition to host-requested I/Os.

The I/O activity of a ShadowImage, Copy-on-Write Snapshot, or Volume Migration S-VOL in the COPY or PAIR state includes only host-requested I/Os (ShadowImage, Copy-on-Write Snapshot, and Volume Migration update copy operations are excluded).

The I/O activity of a P-VOL or simplex volume includes only host-requested I/Os. If the status changed into SMPL in S-VOL (COPY, PAIR) I/O activity, I/O activity is reported as the SMPL status, until the pair status is changed.

Syntax

```
raidar { -h | -q | -z[x] | -I[H][M][instance#] or -I[TC][SI][instance#] |
   -p <port#> <targ> <lun> [mun]| -pd[g] <raw_device> [mun][-n] |
   -s [interval] [count] }
```

Options and parameters

-h

Displays help/usage and version information.

-q

Terminates the interactive mode and exits the command.

-z or -zx

Makes the **raidar** command enter the interactive mode. The `-zx` option guards performing of the HORCM in the interactive mode. When this option detects a HORCM shut down, interactive mode terminates.

OpenVMS cannot use the `-zx` option.

-I[H][M] [instance#]or -I[TC][SI] [instance#]

Specifies the command as [HORC]/[HOMRCF] or [TrueCopy]/[ShadowImage], and used when specifying the CCI instance number.

-p <port#> <targ> <lun> [mun]....

Monitors one or more (up to 16) devices at a time.

- `<port#>`: Specifies the port to be reported: CL1-A to CL1-R and CL2-A to CL2-R (excluding CL1-I, CL1-O, CL2-I, CL2-O). In case of the expanded port, specify from following: CL3-A to CL3-R (excluding CL3-I and CL3-O), or CLG-A to CLG-R (excluding CLG-I and CLG-O).

  The port is not case sensitive (for example, CL1-A= cl1-a= CL1-a= cl1-A, CL3-a= CL3-A= cl3-a= cl3-A).

- `<targ>`: Specifies the SCSI TID of the specified port.

- `<lun>`: Specifies the LUN on the specified TID.

- `[mun]`: Specifies the MU number of the specified LUN within the range of 0 to 63 (ShadowImage or Copy-on-Write Snapshot only).

-pd[g] <raw_device> [mun][-n]

Allows designation of an LDEV by raw device file name. The `-pdg` option is used to show a LUN on the host view by finding a host group.

This -n option is specified to use a device file for NVMe-oF. Specifying this option enables the use of both device files for NVMe-oF and SCSI.

If the HORCC_NVME environment variable is specified, both device files for NVMe-oF and SCSI can be used without the -n option specified.

To specify this option, a command device must be enabled in a namespace on the storage system. For more information about whether a command device can be enabled in a namespace on the storage system, see the *Provisioning Guide* for your storage system.

-s [interval] or -sm [interval]

Designates the time interval in seconds.

- -s: Interprets the time interval as seconds.

- -sm: Interprets the time interval as minutes.

- [interval]: Designates the time interval value (1 to 60). If not specified, the default interval (3) is used.

- [count]: Designates the number of repeats. When this parameter omitted, this command repeats until CNTL-C.

Returned values

None

Error codes

None

Examples

The following shows an example of the **raidar** command and its output.

```
# raidar -p cl1-a  15 6   -p cl1-b  14 5   -p cl1-a  12 3    -s 3
```

```
TIME[03] PORT   T L VOL    STATUS IOPS HIT(%)W(%)IOCNT
13:45:25 -      - - -       -    -    -     -    -
13:45:28 CL1-A 15 6 SMPL   -     200.0 80.0  40.0  600
         CL1-B 14 5 P-VOL PAIR  133.3 35.0  13.4  400
         CL1-A 12 3 P-VOL PSUS  200.0 35.0  40.6  600
```

Description of the raidar command output:

IOPS

# of I/Os (read/write) per second (total I/O rate).

HIT(%)

Hit rate for read I/Os (read hit rate).

W(%)

Ratio of write I/Os to total I/Os (percent writes).

IOCNT

number of write and read I/Os.

raidqry*

The **raidqry** command (RAID query) displays the configuration of the connected host and RAID storage system.

Syntax

```
raidqry { -h | -q | -z[x] | -I[H][M][instance#] or
     -I[TC][SI][instance#] | -l[m] | -r <group> | [ -f ] | -g | -c}
```

Options and parameters

-h

Displays help/usage and version information.

-q

Terminates the interactive mode and exits the command.

-z or -zx

Makes the **raidqry** command enter the interactive mode. The -zx option guards performing of the HORCM in the interactive mode. When this option detects a HORCM shut down, interactive mode terminates.

OpenVMS cannot use the -zx option.

-I[H][M] [instance#]or -I[TC][SI] [instance#]

Specifies the command as [HORC]/[HOMRCF] or [TrueCopy]/[ShadowImage], and used when specifying the CCI instance number.

-I[m]

Displays the configuration information for the local host and the local RAID storage system.

- -l option: Display DKCMAIN microcode version in CCI format.

- -lm option: Display DKCMAIN microcode version in Device Manager - Storage Navigator format.

-r <group>

Displays the configuration information for the remote host and the remote storage system which contains the specified group.

-f

Displays the host name (ip_address) as specified in the configuration definition file. Use this option if 'floatable IP address' is used for the host name (ip_address) in the configuration file.

-g

This option is used when displaying the lists of group name (dev_group) which described in the configuration file of a local host (instance).

| -c

Displays the information about the copy group in an MxN configuration managed by CCI instances and the instances that perform the processing to maintain the data consistency in the copy group. In a Universal Replicator MxN configurations with redundant CCI, one of the CCI instances performs the processing, such as issuing a time stamp (CTQ-Marker) to maintain the data consistency in the consistency group. The other instance is in the standby status in case of failure. If you issue the **pairsplit** command from this standby instance, an error (EX_INVVOL) might occur. Use this option to check which instance can issue the **pairsplit** command in an MxN configuration with redundant CCI.

Returned values

None

Error codes

None

Example 1

The following shows the example of the **raidqry** command and its output.

```
# raidqry -l
```

```
No Group Hostname    HORCM_ver  Uid  Serial#   Micro_ver Cache(MB)
 1   ---     HOSTA  01-22-03/06    0    30053 50-04-00/00   256
 1   ---     HOSTA  01-22-03/06    1    30054 50-04-00/00   256
```

```
# raidqry -lm
```

```
No Group    Hostname    HORCM_ver   Uid  Serial#    Micro_version
 1  ---     JSSA1012  01-37-03/04    0   302614    80-04-20-00/00
```

```
# raidqry -r oradb
```

```
No Group Hostname    HORCM_ver  Uid  Serial#   Micro_ver Cache(MB)
 1 oradb    HOSTA  01-22-03/06    0    30053 50-04-00/00   256
 2 oradb    HOSTB  01-22-03/06    0    30053 50-04-00/00   256
```

Chapter 2: Data management commands

```
 1 oradb    HOSTA  01-22-03/06    1     30054 50-04-00/00   256
 2 oradb    HOSTB  01-22-03/06    1     30054 50-04-00/00   256
```

```
# raidqry -l  -f
```

```
No Group Floatable Host  HORCM_ver Uid Serial# Micro_ver Cache(MB)
 1   ---          FH001 01-22-03/06   0  30053 50-04-00/00 256
```

Output of the `raidqry` command (Example 1):

No

> This column shows the order when the group name (dev_group) which is described in the configuration definition file has multiple remote hosts.

Group

> When the `-r` option is used, this column shows the group name (dev_group) which is described in the configuration definition file.

Hostname

> When using -l[m] option, this column shows the host name of local host. When using `-r` option, this column shows the host name of remote host which is included the group name (dev_group) described in a configuration definition file. Over 30 characters long of the host name is not displayed.

Floatable Host

> When the `-f` option is used, this column shows the host name (ip_address) which is described in the configuration definition file. Up to 30 host names can be displayed. The `-f` option interprets the host name as utilizing floatable IP for a host.

HORCM_ver

> This column shows the version of CCI on the local or remote host. The `-l[m]` option specifies local host. The `-r` option specifies remote host.
>
> **HORCM_ver**: This column shows the version of CCI on the local or remote host. The `-l[m]` option specifies local host. The `-r` option specifies remote host.

Uid Serial# Micro_ver (Micro_version)

> This column shows unitID, serial number, and (DKCMAIN) microcode version of the storage system that is connected to the local or remote host. The `-l[m]` option specifies local host. The `-lm` option displays Micro_version. The `-r` option specifies remote host.

Cache(MB)

> Shows the logical cache capacity (in MB) of the storage system connected to the local or remote host. The `-l` option specifies local host, and `-r` specifies remote host.

Example 2

```
# raidqry -g
```

```
GNo  Group          RAID_type   IV/H  IV/M  MUN/H  MUN/M
  1  ora            HTC_RAID      12     9      4     64
  2  orb            XP_RAID       12     9      4     64
  3  orc            HTC_DF         8     6      1      1
```

Output of the `raidqry` command (Example 2):

GNo

> The order of the group name (dev_group) described in the configuration definition file.

Group

> The group name (dev_group) described in the configuration definition file.

RAID_type

> The type of RAID configured in the group.

IV/H

> The interface version for TrueCopy/TrueCopy Async/Universal Replicator/global-active device in a group, and this is used for maintenance.

IV/M

> The interface version for ShadowImage/Copy-on-Write Snapshot/Volume Migration in a group, and this is used for the maintenance.

MUN/H

> The number of maximum MUs for Universal Replicator in a group.

MUN/M

> The number of maximum MUs for ShadowImage/Copy-on-Write Snapshot in a group.

Example 3

```
# raidqry -c
GNo  Group  S
  1  ora    N
  2  orb    N
  3  orc    Y
```

Output of the `raidqry` command (Example 3):

GNo

> The order of the group name (dev_group) described in the configuration definition file.

Group

> The group name (dev_group) described in the configuration definition file.

Chapter 2: Data management commands

S

> Displays the information about the copy group in an MxN configuration managed by CCI instances and the instances that perform the processing to maintain the data consistency in the copy group.
>
> - **Y**: This instance performs the processing to maintain the data consistency in the copy group, or it does not support a redundant configuration. You can run the **pairsplit** command from this instance. For more information about the requirements for the redundancy support, see the Command Control Interface User and Reference Guide.
>
> - **N**: Instances other than this instance perform the processing to maintain the data consistency in the copy group, or the instances are in the pair status that cannot run the **pairsplit** command. You cannot run the **pairsplit** command from this instance.

raidvchkset*

The **raidvchkset** command sets the parameters for validation checking of the specified volumes, and can also be used to turn off all validation checking without specifying [type]. Unit of checking for the validation is based on the group of CCI configuration definition file.

This command is controlled as protection facility. This command is rejected with EX_ERPERM by connectivity checking between CCI and the RAID storage system.

Syntax

```
raidvchkset { -h | -q | -z[x] | -I[H][M][instance#] or
  -I[TC][SI][instance#]
  | -g <group> | -d <pair Vol> | -d[g] <raw_device> [MU#][-n]
  | -d[g] <seq#> <LDEV#> [MU#] | -nomsg | -vg  [type] [rtime]
  | -vext <size> }
```

Options and parameters

-h

> Displays help/usage and version information.

-q

> Terminates the interactive mode and exits the command.

-z or -zx

> Makes the **raidvchkset** command enter the interactive mode. The -zx option guards performing of the HORCM in the interactive mode. When this option detects a HORCM shut down, interactive mode terminates.
>
> OpenVMS cannot use the -zx option.

-I[H][M] [instance#]or -I[TC][SI] [instance#]

Specifies the command as [HORC]/[HOMRCF] or [TrueCopy]/[ShadowImage], and used when specifying the CCI instance number.

-g <group>

Specifies a group name written in the configuration definition file.

-d <pair Vol>

Specifies paired logical volume name defined in the configuration definition file. When this option is specified, the command is executed for the specified paired logical volume.

-d[g] <raw_device> [MU#] [-n]

Searches whether the specified raw_device is included in a group on the configuration definition file (local instance). If it is in the group, the target volume is executed as the paired logical volume (-d) or group (-dg). This option is effective without specification of -g <group> option. If the specified raw_device is contained in two or more groups, the command is executed for the first group.

This -n option is specified to use a device file for NVMe-oF. Specifying this option enables the use of both device files for NVMe-oF and SCSI.

If the HORCC_NVME environment variable is specified, both device files for NVMe-oF and SCSI can be used without the -n option specified.

To specify this option, a command device must be enabled in a namespace on the storage system. For more information about whether a command device can be enabled in a namespace on the storage system, see the *Provisioning Guide* for your storage system.

-d[g] <seq#> <LDEV#> [MU#]

Searches whether the specified LDEV is included in a group on the configuration definition file (local instance). If it is in the group, the target volume is executed as the paired logical volume (-d) or group (-dg). This option is effective without specification of '-g <group>' option. If the specified LDEV is contained in two or more groups, the command is executed on the first group. The <seq#> <LDEV#> values can be specified in hexadecimal (add the '0x' prefix) or decimal notation.

- When specifying the serial number for VSP 5000 series, add a "5" at the beginning of the serial number. For example, for serial number 12345, enter 512345.

- When specifying the serial number for VSP G1x00 and VSP F1500, add a "3" at the beginning of the serial number. For example, for serial number 12345, enter 312345.

-nomsg

Suppresses messages to be displayed when this command is executed. It is used to execute this command from a user program. This option must be specified at the beginning of a command argument. The command execution log is not affected by this option.

-vg [type] [rtime]

Specifies the following guard type to the target volumes for Data Retention Utility. If [type] is not specified, this option releases all of the guarding.

The following values are available to specify on `[type]`.

- inv: The target volumes are concealed from SCSI Inquiry command by responding 'unpopulated volume'.

- sz0: The target volumes replies with 'SIZE 0' through SCSI Read capacity command.

- rwd: The target volumes are prohibited from reading and writing.

- wtd: The target volumes are prohibited from writing.

- svd: If the target volume is SMPL, it is protected from paircreate (from becoming an S-VOL). If the target volume is P-VOL, it is protected from pairresync restore or pairresync swaps(p). If the target volume is SVOL_PSUS(SSUS), it is protected from pairresync synchronous copy.

  svd option can be used with the other option (inv, sz0, rwd, wtd) in parallel. For example, if you want to protect the absolute volume from the writing executed by the copy series software product and the host access, set the both wtd and svd options. The only setting of wtd option cannot protect the absolute volume from the writing by the copy processing of the copy series software product.

`[rtime]`: Specifies the retention time in days. If [rtime] is not specified, the default value 0 (zero) defined by the storage system is used. This option is ignored in 9900V microcode versions 21-06-*xx* or 21-07-*xx* (default = infinite).

- This option sets each four flags for guarding type as follows:

typeINQRCAPREADWRITE inv1111 Sz00111 rwd0011 wtd0001

-vext <size>

Used when extending the LUN capacity of a Dynamic Provisioning volume.

The increment size of capacity can be specified in bytes or blocks. When specifying in bytes, the unit is t/T (terabyte), g/G (gigabyte), m/M (megabyte), or k/K (kilobyte). If this unit is omitted, block (512 bytes) is used.

Examples for extending 1GB (gigabyte) are: -vext 1G, -vext 1g, -vext 1024M, -vext 1024m, -vext 1048576K, -vext 1048576k, -vext 2097152

Examples for extending 1KB (kilobyte) are: -vext 1K, -vext 1k, -vext 2

LUN capacity and usage rate for Dynamic Provisioning volume can be verified by referring 'LU_CAP' of the 'raidvchkdsp -v aou' or 'raidvchkdsp -v aoub' command. "Aou" (allocation on use) refers to Dynamic Provisioning.

> 📄 **Note:**
>
> When a group operation is specified, a warning message appears, and this command enters the interactive mode.

Returned values

The **raidvchkset** command sets the following returned values during exit allowing you to check the execution results.

- Normal termination: 0

- Abnormal termination: Specific error codes (see Error codes below) and generic error codes.

Error codes

The **raidvchkset -vg** option command returns the following error code as well as generic error codes. See the table below.

Category	Error Code	Error Message	Recommended Action	Value
Volume Status (Unrecoverable)	EX_EPRORT	Mode changes denied due to retention time	Confirm the retention time for a target volume by using **raidvchkscan -v gflag** command.	208

Examples

Disables writing to volumes in oralog group:

```
raidvchkset -g oralog -vg wtd
```

Disables writing and sets retention time to volumes in oralog group:

```
raidvchkset -g oralog -vg wtd 365
```

Releases all guarding to volumes in oralog group:

```
raidvchkset -g oralog -vg
```

raidvchkdsp*

The **raidvchkdsp** command displays the parameters for validation checking of the specified volumes. Unit of checking for the validation is based on the group of CCI configuration definition file.

Syntax

```
raidvchkdsp { -h | -q | -z[x] | -I[H][M][instance#] or
    -I[TC][SI][instance#] | -g <group> | -d <pair Vol>
    | -d[g] <raw_device> [MU#][-n] | -d[g] <seq#> <LDEV#> [MU#]
    | -f[xde] | -v <op> | -c }
```

Options and parameters

-h

>Displays help/usage and version information.

-q

>Terminates the interactive mode and exits the command.

-z or -zx

>Makes the `raidvchkdsp` command enter the interactive mode. The `-zx` option guards performing of the HORCM in the interactive mode. When this option detects a HORCM shut down, interactive mode terminates.

>OpenVMS cannot use the `-zx` option.

-I[H][M] [instance#]or -I[TC][SI] [instance#]

>Specifies the command as [HORC]/[HOMRCF] or [TrueCopy]/[ShadowImage], and used when specifying the CCI instance number.

-g <group>

>Specifies a group name written in the configuration definition file.

-d <pair Vol>

>Specifies paired logical volume name defined in the configuration definition file. When this option is specified, the command is executed for the specified paired logical volume.

-d[g] <raw_device> [MU#] [-n]

>Searches whether the specified raw_device is included in a group on the configuration definition file (local instance). If it is in the group, the target volume is executed as the paired logical volume (-d) or group (-dg). This option is effective without specification of `-g <group>` option. If the specified raw_device is contained in two or more groups, the command is executed for the first group.

>This `-n` option is specified to use a device file for NVMe-oF. Specifying this option enables the use of both device files for NVMe-oF and SCSI.

>If the HORCC_NVME environment variable is specified, both device files for NVMe-oF and SCSI can be used without the `-n` option specified.

>To specify this option, a command device must be enabled in a namespace on the storage system. For more information about whether a command device can be enabled in a namespace on the storage system, see the *Provisioning Guide* for your storage system.

-d[g] <seq#> <LDEV#> [MU#]

Searches whether the specified LDEV is included in a group on the configuration definition file (local instance). If it is in the group, the target volume is executed as the paired logical volume (-d) or group (-dg). This option is effective without specification of '-g <group>' option. If the specified LDEV is contained in two or more groups, the command is executed on the first group. The <seq#> <LDEV#> values can be specified in hexadecimal (add the '0x' prefix) or decimal notation.

- When specifying the serial number for VSP 5000 series, add a "5" at the beginning of the serial number. For example, for serial number 12345, enter 512345.

- When specifying the serial number for VSP G1x00 and VSP F1500, add a "3" at the beginning of the serial number. For example, for serial number 12345, enter 312345.

-fx

Displays the LDEV number in hexadecimal.

-fd

Displays the relationship between the Device_File and the paired volumes, based on the group (as defined in the local instance configuration definition file). If Device_File column shows 'Unknown' to host (instance), then the volume is not recognized on own host, and **raidvchkdsp** command is rejected in protection mode. Non-permitted volume is shown without LDEV# information (LDEV# '-').

Specify ENABLE for HORCM_NVME in the configuration file. If ENABLE is not specified, a device file name for NVMe-oF is not displayed.

-fe

Displays the serial# and LDEV# of the external LUNs mapped to the LDEV for the target volume by adding to last column (ignores the format of 80 column).

-c

When CCI starts, HORCM_DEV in horcm.conf is translated from port/target/lun numbers to the CU:Ldev information, on one hand HORCM_LDEV in horcm.conf is translated from the CU:Ldev information to port/target/lun numbers, because RAID needs to specify 'Port#, Targ#, Lun#' and 'LDEV' for specifying the target device, and then HORCM keeps this information as internal database for the configuration.

If a storage administrator changes the LDEV to LUN/port mapping, such as

- a new/different LDEV is mapped to a previously used port/LUN, or

- an LDEV is mapped to a different/new port

then pair operations might be rejected because the new mapping is different from the mapping information of the database in the running CCI instance. A **pairdisplay** command shows the real LDEV mapping at the time of the command execution and hence shows different information than what is stored in the internal database of the CCI instance.

The `-c` option for **`raidvchkdsp`** allows you to see if there is a difference between the current running CCI instance information and the real mapping. This indication should be used to find such issues which indicate that:

- the CCI instance should be restarted to discover and use the new mapping information, or

- a configuration change occurred without changing the affected configuration files of the CCI instance.

Example change from LDEV#785 to LDEV#786:

raidvchkdsp -g VG000 -c

Group PairVol Port# TID LU Seq# LDEV# LDEV#(conf) -change-> LDEV#

VG000 vg0001 CL4-E-0 0 17 63528 786 785(conf) -change-> 786

raidvchkdsp -g VG000 -c -fx

Group PairVol Port# TID LU Seq# LDEV# LDEV#(conf) -change-> LDEV#

VG000 vg0001 CL4-E-0 0 17 63528 312 311(conf) -change-> 312

Example remove LDEV#785 from a port:

raidvchkdsp -g VG000 -c

Group PairVol Port# TID LU Seq# LDEV# LDEV#(conf) -change-> LDEV#

VG000 vg0001 CL4-E-0 0 17 63528 - 785(conf) -change-> NO LDEV

raidvchkdsp -g VG000 -c -fx

Group PairVol Port# TID LU Seq# LDEV# LDEV#(conf) -change-> LDEV#

VG000 vg0001 CL4-E-0 0 17 63528 - 311(conf) -change-> NO LDEV

> **📄 Note:**
>
> If there have not been any changes to the specified volumes, this option displays nothing.

-v [op]

Specifies the following operation that displays each parameter for validation checking:

- gflag: Displays the parameter for guarding on the specified target volumes.

- pool : This option displays the information about the Dynamic Provisioning pool to which the LDEV belongs. If the LDEV does not belong to a Dynamic Provisioning pool, information about the Thin Image or Copy-on-Write Snapshot pool is displayed. See Example 2 below.

- aou[b] : Displays the LUN capacity and usage rate for Dynamic Provisioning volume corresponding to the group of RM configuration file, and displays the ID of the pool to which LDEV belongs.

- aoub : Displays the LUN capacity in blocks (1 block = 512 bytes).

Returned values

None

Error codes

None

Example 1

Example of the **raidvchkdsp** command with the -v gflag option and its output:

```
# raidvchkdsp -g vg01 -fd -v gflag
```

```
Group PairVol Device_File Seq# LDEV# GI-C-R-W-S PI-C-R-W-S R-Time
vg01   oradb1  c4t0d2      2332  2     E E D D E   E E D D E   365
vg01   oradb2  c4t0d3      2332  3     E E D D E   E E D D E    -
```

Description of the raidvchkdsp command output with the –v gflag option:

GI-C-R-W-S

Displays the flags for guarding as for the target volume.

- I=E: Enabled for Inquiry command.
- I=D: Disabled for Inquiry command.
- C=E: Enabled for Read Capacity command.
- C=D: Disabled for Read Capacity command.
- R=E: Enabled for Read command.
- R=D: Disabled for Read command.
- W=E: Enabled for Write command.
- W=D: Disabled for Write command.
- S=E: Enabled for becoming the S-VOL.
- S=D: Disabled for becoming the S-VOL.

PI-C-R-W-S

Displays the permission flags that show whether each mode flag can be changed to enable or not.

- I=E: 'I' flag can be changed to enable.
- I=D: 'I' flag cannot be changed to enable.
- C=E: 'C' flag can be changed to enable.
- C=D: 'C' flag cannot be changed to enable.
- R=E: 'R' flag can be changed to enable.
- R=D: 'R' flag cannot be changed to enable.
- W=E: 'W' flag can be changed to enable.
- W=D: 'W' flag cannot be changed to enable.

- ▪ S=E: 'S' flag can be changed to enable.
- ▪ S=D: 'S' flag cannot be changed to enable.

R-Time

Displays the retention time for write protect in days. The hyphen (-) shows that the retention time is infinite. The application knows whether the target volume is denied to change to writing enable by referring 'R-Time'.

R-time (Rtime) is identical to rtime and both of them indicate Retention Time. This setting value can normally be identified as a value of R-time that is output by **raidychkdsp** (the logging format is Rtime=xxxx). However, R-time (Rtime) is indicated as the value of 'Retention Time + 1000000' when the expiration lock is enabled. The setting of **raidvchkset** command in this status is denied.

Audit lock status is shown as the retention time plus 1000000. 'R-Time + 1000000' shows the retention time with Audit lock status.

Example 2

Example of the **raidvchkdsp** command with the `-v pool` option and its output:

```
raidvchkdsp -g vg01 -v pool
```

```
Group PairVol Port# TID LU Seq# LDEV# Bsize Available   Capacity
Vg01   oradb1 CL2-D   2  7 62500  167  2048    100000  1000000000
Vg01   oradb2 CL2-D   2 10 62500  170  2048    100000  1000000000
```

Description of the `raidvchkdsp` command output with the `-v pool` option:

Bsize

Displays the data block size of the pool in blocks (512 bytes).

Available(Bsize)

Displays the available capacity for the volume data on the Copy-on-Write Snapshot pool in units of Bsize.

Capacity(Bsize)

Displays the total capacity in the Copy-on-Write Snapshot pool in units of Bsize.

Example 3

Example of the **raidvchkdsp** command with the `-v aou` option and its output ("aou" (allocation on use) refers to Dynamic Provisioning):

```
# raidvchkdsp -v aou -g AOU
```

```
Group  PairVol Port# TID LU Seq# LDEV# Used(MB) LU_CAP(MB) U(%)
```

Chapter 2: Data management commands

```
T(%) PID
AOU    AOU_001 CL2-D  2   7 62500  167   20050  1100000   10
70   1
AOU     AOU_002 CL2-D  2  10 62500  170 110000  1100000   10
70   1
```

Description of the `raidvchkdsp` command output with the `-v aou` option:

Used(MB)

Displays the usage size of the allocated block on this LUN. Range: 0 ≤ Used (MB) < LU_CAP(MB) + 42 MB

LU_CAP(MB)

Displays the LUN capacity responded to the 'Readcapacity' command as SCSI interface.

LU_CAP(BLK)

Displays the LUN capacity (in block/512 bytes) responded to the 'Readcapacity' command as SCSI interface.

U(%)

Displays the usage rate of the allocated block on the Dynamic Provisioning pool containing this LU.

T(%)

Displays the threshold rate being set to the Dynamic Provisioning pool as the high water mark.

PID

Displays the Dynamic Provisioning pool ID assigned to this Dynamic Provisioning volume.

raidvchkscan*

The `raidvchkscan` command displays the port of the storage system (9900V or later), target ID, LDEV mapped for LUN# and MU#, and status of LDEV, regardless of the configuration definition file.

📄 **Note:**

This command is rejected with EX_ERPERM by connectivity checking between CCI and the RAID storage system.

Syntax

```
raidvchkscan { -h | -q | -z[x] | -I[H][M][instance#] or
    -I[TC][SI][instance#] | -p <port#> [hgrp] | -pd[g] <raw_device>[-n]
    | -s <seq#> | -t <target> | -l <lun> | [ -f[x] ] | -v <op> }
```

Options and parameters

-h

Displays help/usage and version information.

-q

Terminates the interactive mode and exits the command.

-z or -zx

Makes the **`raidvchkscan`** command enter the interactive mode. The `-zx` option guards performing of the HORCM in the interactive mode. When this option detects a HORCM shut down, interactive mode terminates.

OpenVMS cannot use the `-zx` option.

-I[H][M] [instance#]or -I[TC][SI] [instance#]

Specifies the command as [HORC]/[HOMRCF] or [TrueCopy]/[ShadowImage], and used when specifying the CCI instance number.

-p <port#> [hgrp]

Specifies the port ID of the port to be scanned. Valid ports are CL1-A to CL1-R and CL2-A to CL2-R (excluding CL1-I, CL1-O, CL2-I, CL2-O). In addition, it is able to specify from CL3-a to CL3-r (except CL3-i, CL3-o), or CL4-a to CL4-r (except CL4-i, CL4-o) for the expanded port.

The port is not case sensitive (for example, CL1-A= cl1-a= CL1-a= cl1-A, CL3-a= CL3-A= cl3-a= cl3-A). This option must be specified if '-find' or '-pd <raw_device>'option is not specified.

Specify [hgrp] to display only the LDEVs mapped to a host group on a port.

-pd[g] <raw_device> [-n]

Specifies the raw device name. This option finds Seq# and port_name of the storage system to which the specified device can be connected, and scans the port of the storage system which corresponds with the unit ID that searches the unit ID from Seq#. This option must be specified if the '-find' option is not specified. If this option is specified, the following -s <seq#> option is invalid.

-pdg: Shows a LUN on the host view by finding a host group.

This `-n` option is specified to use a device file for NVMe-oF. Specifying this option enables the use of both device files for NVMe-oF and SCSI.

If the HORCC_NVME environment variable is specified, both device files for NVMe-oF and SCSI can be used without the `-n` option specified.

To specify this option, a command device must be enabled in a namespace on the storage system. For more information about whether a command device can be enabled in a namespace on the storage system, see the *Provisioning Guide* for your storage system.

-s <seq#>

Used to specify the seq# (serial#) of the storage system when this option cannot specify the unit ID which is contained for '-p <port>' option. This option scans the port specified by '-p <port>' option of the storage system which corresponds with the unit ID that searches the unit ID from seq#. If this option is specified, then the unit ID which is contained in '-p <port>' option is invalid.

- When specifying the serial number for VSP 5000 series, add a "5" at the beginning of the serial number. For example, for serial number 12345, enter 512345.

- When specifying the serial number for VSP G1x00 and VSP F1500, add a "3" at the beginning of the serial number. For example, for serial number 12345, enter 312345.

-t <target>

Specifies a target ID (0 to 15) of the specified port. If this option is not specified, the command applies to all target IDs.

-l <lun>

Specifies a LUN (0 to 2047) of the specified target ID. For VSP 5000 series, the LUN that can be specified is (0 to 4095). If this option is not specified, the command applies to all LUNs. If this option is specified, the TID must also be specified.

-fx

Displays the LDEV number in hexadecimal notation.

-v [op]

Specifies the following operation that displays each parameter for validation checking:

gflag: Displays the parameter for guarding on the specified target volumes.

pool: This option displays the information about the Dynamic Provisioning pool to which the LDEV belongs. If the LDEV does not belong to a Dynamic Provisioning pool, information about the Thin Image or Copy-on-Write Snapshot pool is displayed. See Example 2 below.

aou[b]: Displays the LUN capacity and usage rate for only Dynamic Provisioning volume mapped to the specified port, and displays the ID of the pool to which LDEV belongs.

aoub: Displays the LUN capacity in blocks (512 bytes).

Returned values

None

Error codes

None

Example 1

Example of the **raidvchkscan** command with the `-v gflag` option

```
# raidvchkscan -p CL1-A -v gflag
```

```
PORT# /ALPA/C TID# LU#    Seq#  Num LDEV#   GI-C-R-W-S  PI-C-R-W-S R-Time
CL1-A / ef/ 0    0   0   2332   1    0    E E D D E   E E D D E    365
CL1-A / ef/ 0    0   1   2332   1    1    E E D D E   E E D D E     -
CL1-A / ef/ 0    0   2   2332   1    2    E E D D E   E E D D E     0
```

Description of the raidvchkscan command output with the `-v gflag` option:

GI-C-R-W-S

Displays the flags for guarding as for the target volume.

- I=E: Enabled for Inquiry command.

- I=D: Disabled for Inquiry command.

- C=E: Enabled for Read Capacity command.

- C=D: Disabled for Read Capacity command.

- R=E: Enabled for Read command.

- R=D: Disabled for Read command.

- W=E: Enabled for Write command.

- W=D: Disabled for Write command.

- S=E: Enabled for becoming the S-VOL.

- S=D: Disabled for becoming the S-VOL.

PI-C-R-W-S

Displays the permission flags that show whether each mode flag can be changed to enable or not.

- I=E: 'I' flag can be changed to enable.

- I=D: 'I' flag cannot be changed to enable.

- C=E: 'C' flag can be changed to enable.

- C=D: 'C' flag cannot be changed to enable.

- R=E: 'R' flag can be changed to enable.

- R=D: 'R' flag cannot be changed to enable.

- W=E: 'W' flag can be changed to enable.

- W=D: 'W' flag cannot be changed to enable.

- S=E: 'S' flag can be changed to enable.

- S=D: 'S' flag cannot be changed to enable.

Chapter 2: Data management commands

R-Time

> Displays the retention time for write protect in days. The hyphen (-) shows that the retention time is infinite. The application knows whether the target volume is denied to change to writing enable by referring 'R-Time'.

📄 **Note:** Audit lock status is shown as the retention time plus 1000000. 'R-Time + 1000000' shows the retention time with Audit lock status.

Example 2

Example of the **raidvchkscan** command with the `-v pool` option:

```
# raidvchkscan -v pool -p CL2-d-0
```

```
PORT# /ALPA/C TID# LU#   Seq#  Num LDEV#   Bsize    Available     Capacity
CL2-D-0 /e4/ 0   2   0  62500   1   160    2048        100000   1000000000
CL2-D-0 /e4/ 0   2   1  62500   1   161    2048        100000   1000000000
```

Description of the raidvchkscan command output with the –v pool option:

Bsize

> Displays the data block size of the pool in blocks (512 bytes).

Available(Bsize)

> Displays the available capacity for the volume data on the Copy-on-Write Snapshot pool in units of Bsize.

Capacity(Bsize)

> Displays the total capacity in the Copy-on-Write Snapshot pool in units of Bsize.

Example 3

Example of the **raidvchkscan** command with the `-v aou` option ("aou" (allocation on use) refers to Dynamic Provisioning)

```
# raidvchkscan -v aou -p CL2-d-0
```

```
PORT# /ALPA/C TID# LU#   Seq#  Num LDEV#  Used(MB)  LU_CAP(MB)  U(%) T(%) PID
CL2-D-0 /e4/ 0   2   0  62500   1   160    20050     1100000     1   60   1
CL2-D-0 /e4/ 0   2   1  62500   1   161   200500     1100000    18   60   2
```

Description of the raidvchkscan command output with the –v aou option:

Used(MB)

> Displays the usage size the allocated block on this LUN. Range: 0 ≤ Used (MB) < LU_CAP(MB) + 42MB

Chapter 2: Data management commands

LU_CAP(MB)

Displays the LUN capacity responded to the 'Readcapacity' command as SCSI interface.

LU_CAP(BLK)

Displays the LUN capacity (in block/512 bytes) responded to the 'Readcapacity' command as SCSI interface.

U(%)

Displays the usage rate of the allocated block on the Dynamic Provisioning pool containing this LU.

T(%)

Displays the threshold rate being set to the Dynamic Provisioning pool as the high water mark.

PID

Displays the Dynamic Provisioning pool ID assigned to this Dynamic Provisioning volume.

raidvchkscan for Universal Replicator*

The **raidvchkscan** command supports the (-v jnl [t] [unit#]) option to find the journal volume list. It also displays any information for the journal volume.

Syntax

```
raidvchkscan { -h | -q | -z[x] | -I[H][M][instance#] or -I[TC][SI][instance#]| -v jnl
[t] [unit#] | [ -s <seq#>] | [ -f[x] ] }
```

Options and parameters

-h

Displays help/usage and version information.

-q

Terminates the interactive mode and exits the command.

-z or -zx

Makes the **raidvchkscan** command enter the interactive mode. The `-zx` option guards performing of the HORCM in the interactive mode. When this option detects a HORCM shut down, interactive mode terminates.

OpenVMS cannot use the `-zx` option.

-I[H][M] [instance#]or -I[TC][SI] [instance#]

Specifies the command as [HORC]/[HOMRCF] or [TrueCopy]/[ShadowImage], and used when specifying the CCI instance number.

-s <seq#>

> Used to specify the seq# (serial#) of the storage system when this option cannot specify unitID which is contained for '-v jnl' option. If this option is specified, the unitID which is contained in '-v jnl' is invalid.
>
> - When specifying the serial number for VSP 5000 series, add a "5" at the beginning of the serial number. For example, for serial number 12345, enter 512345.
>
> - When specifying the serial number for VSP G1x00 and VSP F1500, add a "3" at the beginning of the serial number. For example, for serial number 12345, enter 312345.

-fx

> Displays the LDEV number in hexadecimal notation.

-v jn

> Displays information for the journal volume.

-v jnlt

> Displays the DOW , DPW , and APW time-out values for controlling the journal.

Returned values

None

Error codes

Non

Example 1

Example of the **raidvchkscan** command with the -v jnl 0 option

```
# raidvchkscan -v jnl 0
```

```
JID MU CTG  JNLS  AP  U(%)  Q-Marker   Q-CNT   D-SZ(BLK)   Seq# Nnm LDEV#
001 0   1   PJNN  4   21    43216fde      30      512345  62500  2   265
002 1   2   PJNF  4   95    3459fd43   52000      512345  62500  3   270
002 2   2   SJNS  4   95    3459fd43   52000      512345  62500  3   270
003 0   3   PJSN  4    0           -       -      512345  62500  1   275
004 0   4   PJSF  4   45    1234f432      78      512345  62500  1   276
005 0   5   PJSE  0    0           -       -      512345  62500  1   277
006 -   -   SMPL  -    -           -       -      512345  62500  1   278
007 0   6   SMPL  4    5    345678ef      66      512345  62500  1   278
```

Description of the raidvchkscan command output with the -v jnl 0 option:

JID

> Journal ID

MU

Mirror descriptions on Universal Replicator.

CTG

Consistency group ID

JNLS

Status of the journal:

- SMPL: This means the journal volume which does not have a pair, or in the state of deleting.

- P(S)JNN: This means 'P(S) Journal Volume is in Normal status'.

- P(S)JNS: This means 'P(S) Journal Volume is suspended in Normal status' created with `-nocsus` option.

- P(S)JSN: This means 'P(S) Journal Volume is suspended in Normal status'.

- P(S)JNF: This means 'P(S) Journal Volume is in Full status'.

- P(S)JSF: This means 'P(S) Journal Volume is suspended in Full status'.

- P(S)JSE: This means 'P(S) Journal Volume is suspended by an error (including link failures)'.

- P(S)JES: This means 'P(S) Journal Volume is suspended by an error' created with `-nocsus` option.

AP

Displays the following two conditions (status) according to the pair status.

Shows the number of active paths on the initiator port in Universal Replicator links. 'Unknown' is shown as '-'.
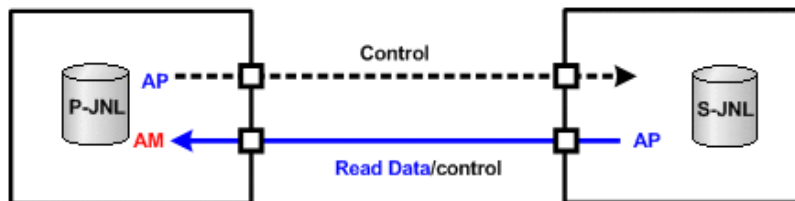
AM

The activity monitor that detects whether or not there is a request for data from the initiator at regular intervals. If AM detects a time-out, the P-JNL state is changed from P-JNN to PJSE.

📄 **Note:** The same path information is used for AP for three commands (`pairvolchk`, `pairdisplay`, `raidvchkscan`). The differential is that pairvolchk and pairdisplay are to show a special meaning with SSUS(SSWS) state.

Q-Marker

Displays the sequence # of the journal ID, called the Q-marker. For P-JNL, Q-Marker shows the latest sequence # on the P-JNL volume. For S-JNL, the Q-Marker shows the latest sequence # of the cache(DFW).

Q-CNT

Displays the number of remaining Q-Markers within each journal volume.

The following figure shows an example of Q-Marker and Q-CNT

U(%)
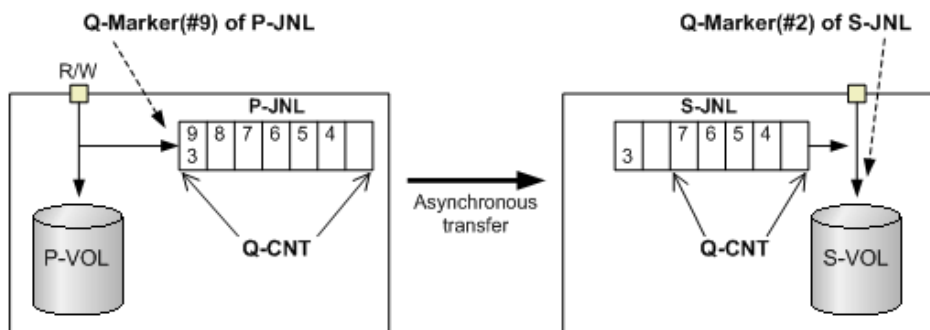
Displays the usage rate of the journal data.

D-SZ

Displays the capacity for the journal data on the journal volume. For details about the displayed capacity, see the *Hitachi Universal Replicator User Guide*.

Seq#

Displays the serial number of the RAID storage system.

- The serial number for VSP 5000 series is displayed with a "5" at the beginning ("512345" = serial number 12345).

- The serial number for VSP G1x00 and VSP F1500 is displayed with a "3" at the beginning ("312345" = serial number 12345).

Num

Displays the number of LDEVs composing the journal volume.

LDEV#

Displays the following LDEV numbers of journal volumes for storage systems that support reserve journal volumes:

- LDEV number of the master journal volume for the master journal

- LDEV number of the restore journal volume for the restore journal. Whether to support reserve journal volumes depends on the model of your storage system. For the support status, see the *Hitachi Universal Replicator User Guide* for your model.

The following table lists information about the different journal volume statuses. QCNT=0 indicates that the number of remaining Q-Markers is '0'. The letter 'N' indicates a non-zero.

JNLS		Other Information		
P-JNL	**S-JNL**	**QCNT**	**AP**	**Description**
SMPL		0	-	Configured as journal volume, but NOT pair

JNLS		Other Information		Description
P-JNL	S-JNL	QCNT	AP	
		N	-	Deleting the journal volume
PJNN (PJNS)	SJNN (SJNS)	0	-	Normal state of the journal volume without data
PJNN (PJNS)	-	N	-	Normal state of the journal volume with data
-	SJNN (SJNS)	N	N	Normal state of the journal volume with data
			0	Still normal state of the journal volume at Link failure
PJSN	SJSN	0	-	Suspended journal volume via operation
		N	-	Suspending the journal volume
PJNF	-	N	-	High water mark state
PJSF	SJSF	0	-	Suspended journal volume due to full journal
		N	-	Suspending the journal volume due to full journal
PJSE	-	0	-	Suspended journal volume due to failure/Link failure
		N	-	Suspending the journal volume due to failure/Link failure
-	SJSE	0	N	Suspended journal volume due to failure
			0	Suspended journal volume due to Link failure
		N	N	Suspending the journal volume due to failure
			0	Suspending the journal volume due to Link failure

Example 2

Example of the **raidvchkscan** command with the -v jnlt option

```
# raidvchkscan -v jnlt
```

```
JID MU CTG  JNLS  AP  U(%)   Q-Marker   Q-CNT   D-SZ(BLK)   Seq# DOW PBW APW
001  0  1   PJNN   4   21    43216fde      30      512345  63528  20 300  40
```

Chapter 2: Data management commands

```
002  1   2   PJNF   4   95      3459fd43    52000      512345  63528  20 300  40
003  0   3   PJSN   4   0          -          -        512345  63528  20 300  40
```

Description of the `raidvchkscan` command output with the `-v jnlt` option:

DOW

Data Overflow Watch' timer (in seconds) setting per the Journal.

PBW:

Path Blockade Watch timer setting (in seconds) per the Journal. If the setting is more than 3600 seconds, it displays 6000 seconds.

APW

This shows 'Active Path Watch' timer (in seconds) for detecting Link failure.

raidvchkscan for Thin Image, Copy-on-Write Snapshot, and HDP pools*

The `raidvchkscan` command supports the option (`-v pid[a] [unit#]`) to find the Thin Image, Copy-on-Write Snapshot, or Dynamic Provisioning pool settings via SVP, and displays information for the Thin Image, Copy-on-Write Snapshot, or Dynamic Provisioning pool.

Syntax

```
raidvchkscan {  -h | -q | -z[x] | -I[H][M][instance#] or
    -I[TC][SI][instance#] | -v  pid[a][s][b][-fp] [unit#]
    | [ -s <seq#>] |  [ -f[x] | }
```

Options and parameters

-h

Displays help/usage and version information.

-q

Terminates the interactive mode and exits the command.

-z or -zx

Makes the raidvchkdsp command enter the interactive mode. The `-zx` option guards performing of the HORCM in the interactive mode. When this option detects a HORCM shutdown, interactive mode terminates.

OpenVMS cannot use the `-zx` option.

-I[H][M] [instance#]or -I[TC][SI] [instance#]

Specifies the command as [HORC]/[HOMRCF] or [TrueCopy]/[ShadowImage], and used when specifying the CCI instance number.

-s <seq#>

Used to specify the Seq# (serial#) of the storage system when this option cannot specify unitID which is contained for '-v jnl[a]' option. If this option is specified, the unitID which is contained in '-v jnl[a]' is invalid.

- When specifying the serial number for VSP 5000 series, add a "5" at the beginning of the serial number. For example, for serial number 12345, enter 512345.

- When specifying the serial number for VSP G1x00 and VSP F1500, add a "3" at the beginning of the serial number. For example, for serial number 12345, enter 312345.

-fx

Displays the LDEV number in hexadecimal notation.

-v pidb -fp

Displays the actual capacity of the pool.

The following GUI displays the actual capacity information as the capacity of the accelerated compression enabled pool. Specify this option to confirm the capacity equivalent to the capacity displayed in the GUI.

- Device Manager - Storage Navigator (VSP 5000 series, VSP E series, VSP G130, G/F350, G/F370, G/F700, G/F900)

- Storage Advisor Embedded

If the microcode version does not support this option, nothing is displayed when this option is specified.

-v pid[s]

Displays information for the Thin Image/Copy-on-Write Snapshot/Dynamic Provisioning pool.

-v pid[a]

Displays information for the Dynamic Provisioning pool.

-v pidb

Supported storage systems:

- VSP 5000 series

- VSP G1x00 and VSP F1500

- VSP E series

- VSP G130, G/F350, G/F370, G/F700, G/F900

- VSP G200, G400, G600, G800 and VSP F400, F600, F800

Displays basic information for the pool.

If the microcode version does not support this option, nothing is displayed when this option is specified.

Returned values

None

Chapter 2: Data management commands

Error codes

None

Example 1

Example of the **raidvchkscan** command with the `-v pid` option:

```
# raidvchkscan -v pid 0
```

```
PID  POLS   U(%)  SSCNT  Available(MB)  Capacity(MB)    Seq#  Num  LDEV#  H(%)
FMT_CAP(MB)
001  POLN    10    330     10000000      1000000000    62500   2    265   80   100
002  POLF    95   9900       100000      1000000000    62500   3    270   70   100
003  POLS   100  10000          100      1000000000    62500   1    275   70   100
004  POLE     0      0            0               0    62500   0      0   80   100
```

Description of the `raidvchkscan` command output with the `-v pid` option:

PID

> Displays the Thin Image/Copy-on-Write Snapshot/Dynamic Provisioning pool ID.

POLS

> Displays the status of the Thin Image/Copy-on-Write Snapshot/Dynamic Provisioning pool.

- POLN: "Pool Normal" (The pool is in the normal status.)

- POLF: "Pool Full" (The pool is in the overflow status exceeding the threshold.)

- POLS: "Pool Suspend" (The pool is in the overflow status exceeding the threshold and is blocked.)

- POLE: "Pool failure" (The pool is suspended in the failure status.) In this status, the pool information cannot be displayed.

U(%)

> Displays the usage rate of the Thin Image/Copy-on-Write Snapshot/Dynamic Provisioning pool.

SSCNT

> Displays the number of Thin Image/Copy-on-Write Snapshot volumes in the Thin Image/Copy-on-Write Snapshot pool or the total number of Dynamic Provisioning volumes mapped in this Dynamic Provisioning pool.

Available(MB)

> Displays the available capacity for the volume data on the Thin Image/Copy-on-Write Snapshot/Dynamic Provisioning pool.

Capacity(MB)

> Displays the total capacity in the Thin Image/Copy-on-Write Snapshot/Dynamic Provisioning pool.

Chapter 2: Data management commands

Seq#

Displays the serial number of the RAID storage system.

- The serial number for VSP 5000 series is displayed with a "5" at the beginning ("512345" = serial number 12345).

- The serial number for VSP G1x00 and VSP F1500 is displayed with a "3" at the beginning ("312345" = serial number 12345).

Num

Displays the number of LDEVs composing the Thin Image/Copy-on-Write Snapshot/ Dynamic Provisioning pool.

LDEV#

Displays the first number of LDEV composing the Thin Image/Copy-on-Write Snapshot/Dynamic Provisioning pool.

H(%)

Displays the threshold rate being set to the Thin Image/Copy-on-Write Snapshot/ Dynamic Provisioning pool as the high water mark. 'Unknown' is shown as '-'.

FMT_CAP(MB)

Displays the formatted pool capacity. If there is no valid information for the pool, a hyphen (-) is displayed.

Example 2

Example of the **raidvchkscan** command with the `-v pida` option

```
# raidvchkscan -v pida 0
```

```
PID  POLS U(%)   AV_CAP(MB)   TP_CAP(MB)  W(%)  H(%) Num LDEV#   LCNT   TL_CAP(MB)
001  POLN  10     45000000     50000000    50    80   2   265     33     65000000
002  POLF  95        10000    100000000    50    80   3   270    900    100000000
004  POLN   0     10000000    100000000    80    90   2   280      0            0
```

Description of the output of raidvchkscan command with the `-v pida` option:

PID

Displays the Dynamic Provisioning pool ID.

POLS

Displays the status of the Dynamic Provisioning pool.

- POLN: "Pool Normal" (The pool is in the normal status.)

- POLF: "Pool Full" (The pool is in the overflow status exceeding the threshold.)

- POLS: "Pool Suspend" (The pool is in the overflow status exceeding the threshold and is blocked.)

- POLE: "Pool failure" (The pool is suspended in the failure status.) In this status, the pool information cannot be displayed.

U(%)

> Displays the usage rate of the Dynamic Provisioning pool.

AV_CAP(MB)

> Displays the available capacity for the Dynamic Provisioning volumes mapped to this pool.

TP_CAP(MB)

> Displays the total capacity of the Dynamic Provisioning pool.

W(%)

> Displays the threshold value for 'WARNING' set for this Dynamic Provisioning pool.

H(%)

> Displays the threshold rate set for the Dynamic Provisioning pool as the high water mark.

Num

> Displays the number of LDEVs composing the Dynamic Provisioning pool.

LDEV#

> Displays the first number of LDEV composing the Dynamic Provisioning pool.

LCNT

> Displays the total number of Dynamic Provisioning volumes mapped to this Dynamic Provisioning pool.

TL_CAP(MB)

> Displays the total capacity of all Dynamic Provisioning volumes mapped to this Dynamic Provisioning pool.

Example 3

Example of the **raidvchkscan** command with the `-v pidb` option:

```
# raidvchkscan -v pidb 0
```

```
PID POLS U(%) LCNT SSCNT Available(MB) Capacity(MB) Snap_Used(MB)
TL_CAP(MB) BM TR_CAP(MB) RCNT Seq# Num LDEV# W(%) H(%) STIP VCAP(%)
TYPE PM PT POOL_NAME
001 POLN 0 11001 11001 46998 46998 0
2432398 NB 0 0 300050 1 0 70 80 YES - OPEN
N HDP dp_ti_pool
001 POLN 0 11001 11001 46998 46998 0
2432398 NB 0 0 300050 1 0 70 80 YES - OPEN
N HDP dp_ti_pool
```

Example of the **raidvchkscan** command with the `-v pidb -fp` option:

```
# raidvchkscan -v pidb 0 -fp
PID POLS U(%) LCNT SSCNT Available(MB) Capacity(MB) Snap_Used(MB)
TL_CAP(MB) BM TR_CAP(MB) RCNT Seq# Num LDEV# W(%) H(%) STIP VCAP(%)
```

Chapter 2: Data management commands

```
TYPE PM PT POOL_NAME
001 POLN 0 11001 11001 46998 46998 -
2432398 NB 0 0 300050 1 0 70 80 YES -
OPEN N HDP dp_ti_pool
001 POLN 0 11001 11001 46998 46998 -
2432398 NB 0 0 300050 1 0 70 80 YES -
OPEN N HDP dp_ti_pool
```

Description of the output of the `raidvchkscan` command with the `-v pidb` or `-v pidb -fp` option:

PID

> Displays the pool ID.

POLS

> Displays the status of the pool:
>
> - POLN: "Pool Normal" (The pool is in the normal status.)
>
> - POLF: "Pool Full" (The pool is in the overflow status exceeding the threshold.)
>
> - POLS: "Pool Suspend" (The pool is in the overflow status exceeding the threshold and is blocked.)
>
> - POLE: "Pool failure" (The pool is blocked due to a failure.) In this status, the pool information cannot be displayed.

U(%)

> Displays the usage rate of the pool. The actual usage rate is displayed if the `-fp` option is specified.

LCNT

> Displays the total number of Dynamic Provisioning virtual volumes mapped to the pool. A hyphen (-) is displayed if the information is not available for this pool.

SSCNT

> Displays the total number of snapshot data items mapped to the pool. A hyphen (-) is displayed if the information is not available for this pool.

Available(MB)

> Displays the available capacity for the volume data in the pool. Displays the actual pool capacity when the -fp option is specified.

Capacity(MB)

> Displays the total capacity of the pool. Displays the actual pool capacity when the -fp option is specified.

Snap_Used(MB)

> Displays the capacity used for Thin Image data in megabytes. If the used capacity is less than 1 MB, the value is rounded up. A hyphen (-) is displayed if the information is not available for this pool.
>
> If you entered the -fp option, a hyphen (-) is displayed.

Chapter 2: Data management commands

TL_CAP(MB)

Displays the total capacity of all Dynamic Provisioning virtual volumes and Thin Image pairs mapped to the pool. A hyphen (-) is displayed if the information is not available for this pool.

BM

Displays the I/O Blocking Mode of the pool:

- PF (Pool Full): If the pool is full, you cannot read from or write to the target DP-VOL. If the pool VOL is blocked, you can read from or write to the target DP-VOL.

- PB (Pool vol Blockade): If the pool VOL is blocked, you cannot read from or write to the target DP-VOL. If the pool is full, you can read from or write to the target DP-VOL.

- FB (Full or Blockade): If the pool is full or pool VOL is blocked, you cannot read from or write to the target DP-VOL.

- NB (No Blocking): If the pool is full or pool VOL is blocked, you can read from or write to the target DP-VOL.

- - (Not supported): The configuration does not support the I/O Blocking Mode.

TR_CAP(MB)

Displays the sum of the pool capacities reserved for the volumes for which Full Allocation or Proprietary Anchor is enabled. For a configuration that does not support Full Allocation, a hyphen (-) is displayed.

RCNT

Displays the number of volumes for which Full Allocation is enabled that are mapped to a pool. For a configuration that does not support Full Allocation, a hyphen (-) is displayed.

Seq#

Displays the serial number.

- The serial number for VSP 5000 series is displayed with a "5" at the beginning ("512345" = serial number 12345).

- The serial number for VSP G1x00 and VSP F1500 is displayed with a "3" at the beginning ("312345" = serial number 12345).

Num

Displays the number of LDEVs belonging to the pool.

LDEV#

Displays the number of the first LDEV in the pool. "65535 (ffff)" is shown while the pool is being created.

W(%)

Displays the threshold value for WARNING set for the pool. A hyphen (-) is displayed if the information is not available for this pool.

H(%)

Displays the threshold value set for the pool as the high water mark.

Chapter 2: Data management commands

STIP

Displays the setting for Thin Image pairs when the high water mark threshold is exceeded.

- YES: Thin Image pairs are suspended.

- NO: Thin Image pairs are not suspended.

- - (hyphen): The information is not available for this pool.

VCAP(%)

Displays the percentage of the subscription limit of V-VOL and Thin Image pair to the pool capacity.

- UNLIMITED: Unlimited.

- - (hyphen): The information is not available for this pool.

For VSP G130, G/F350, G/F370, G/F700, G/F900, a hyphen (-) indicating invalid is displayed.

TYPE

Displays the platform type of the pool:

- OPEN: Pool for open systems

- M/F: Pool for mainframe systems

PM

Displays the pool status:

- N: Normal status.

- S: Shrinking or rebalancing.

- NT: The pool for Thin Image is in the normal status.

- ST: The pool for Thin Image is shrinking or rebalancing.

PT

Displays the pool type:

- HDP: Pool for Dynamic Provisioning

- HDT: Pool for Dynamic Tiering

- RT: Pool for active flash

- TI: Pool for Thin Image

- CW: Pool for Copy-on-Write Snapshot

- DM: Pool for Dynamic Provisioning that has the data direct mapping attribute

POOL_NAME

Displays the pool name.

horcmstart

The **horcmstart** command is a script that starts HORCM. This script also sets the environment variables for HORCM as needed (for example, HORCM_CONF, HORCM_LOG, HORCM_LOGS).

When setting the instance number by specifying the HORCMINST environment variable, make sure that the setting value of the environment variable is in the range of 0 to 2047, and then execute **horcmstart**.

Syntax

```
horcmstart.sh { inst ... }      (UNIX systems)
horcmstart.exe { inst ... }     (Windows systems)
```

Options and parameters

inst

Specifies the HORCM instance number (range= from 0 to 2047). When this option is specified, the **horcmstart** shell script sets the environment variables (HORCMINST, HORCM_CONF, HORCM_LOG, HORCM_LOGS) corresponding to the instance number, and starts the specified HORCM instance. (Environment variables set by the user become invalid.) When this option is not specified, the **horcmstart** shell script starts 1 HORCM and uses the environment variables set by the user. If you have designated full environment variables, use horcmstart.sh without any arguments. If you did not designate environment variables (HORCM_CONF, HORCM_LOG, HORCM_LOGS), then this shell script sets the environment variables as follows:

For UNIX-based platforms:

- If HORCMINST is specified:

  - HORCM_CONF = /etc/horcm*.conf (* is instance number)

  - HORCM_LOG = /HORCM/log*/curlog

  - HORCM_LOGS = /HORCM/log*/tmplog

- If HORCMINST is not specified:

  - HORCM_CONF = /etc/horcm.conf

  - HORCM_LOG = /HORCM/log/curlog

  - HORCM_LOGS = /HORCM/log/tmplog

For Windows platform:

- If HORCMINST is specified:

  - HORCM_CONF = %windir%\horcm*.conf (* is the instance number)

  - HORCM_LOG = \HORCM\log*\curlog

  - HORCM_LOGS = \HORCM\log*\tmplog

- If HORCMINST is not specified:

  - HORCM_CONF = %windir%\horcm.conf

  - HORCM_LOG = \HORCM\log\curlog

  - HORCM_LOGS = \HORCM\log\tmplog

[environmental variable]: The HORCM_LOGS environment variable is used when specifying the log file directory for automatic storing. When HORCM starts up, the log files created in the operation are stored automatically in the HORCM_LOGS directory. This log directory must give an equality class with HORCM_LOG

HORCMSTART_WAIT (for waiting the CCI instance with start-up). Horcmgr does fork/exec() horcmd_XX as daemon process, and verifies/waits until HORCM become ready state. The timeout is used for only avoiding infinite loop, currently the default time is 200 sec in consideration of maximum LDEV. However, it may be needed to change the default timeout value for starting HORCM under high-loading of the server, or the remote command device. In such a case, this environmental variable is used when changing a timeout value (in seconds) from the current default value (200 sec), this value must be specified more than 5 seconds and multiple of 5 seconds. For example, setting 500 sec:

```
HORCMSTART_WAIT=500 Export HORCMSTART_WAIT
```

For OpenVMS platform: OpenVMS needs to make the Detached LOGINOUT.EXE Process as a JOB in the background by using the 'RUN /DETACHED' command. For details see **Requirements and restrictions for OpenVMS** (item **4**) in the *Installation and Configuration Guide*.

horcmshutdown

The `horcmshutdown` command is a script for stopping HORCM.

Syntax

```
horcmshutdown.sh {inst...}      (UNIX systems)
horcmshutdown.exe {inst...}     (Windows systems)
```

Options and parameters

inst

Specifies the HORCM (CCI) instance number (range= from 0 to 2047). When this option is specified, the command stops the specified HORCM instance. When this option is not specified, the command refers to the instance (environment variable HORCMINST) of the execution environment of this shell script and stops the following the HORCM instance.

- When HORCMINST is specified, this command stops the HORCM instance of the execution environment of this shell script.

- When HORCMINST is not specified, this command stops the HORCM having no instance setting.

> ⚠️ **Caution:** After direction of stopping HORCM instance, this command returns a response just before stopping HORCM instance. Thus to return the response of this command, it does not mean HORCM instance disappeared.

horcctl*

The HORCM software have logs that identify the cause of software and/or hardware errors as well as a tracing function for investigating such errors. The location of the log files depends on the user's command execution environment and the HORCM execution environment. The command trace file and core file reside together under the directory specified in the HORCM execution environment.

The `horcctl` command can be used for both maintenance and troubleshooting. The `horcctl` command allows you to change and display the internal trace control parameters (for example, level, type, buffer size) of the HORC Manager (CCI) software. If a new value for a parameter is not specified, the current trace control parameter is displayed.

> ⚠️ **Caution:**
>
> Do not change the trace level unless directed to do so by a Hitachi Vantara
> representative. Level 4 is the normal trace level setting. Levels 0-3 are for
> troubleshooting. Setting a trace level other than 4 may impact problem resolution.
> If you request a change of the trace level using the **horcctl -l <level>**
> command, a warning message appears, and this command enters interactive
> mode.

Syntax

```
horcctl { -h | -q | -z[x] | -I[H][M][instance#] or
    -I[TC][SI][instance#] | -d | -c | -l <level> | -b <y/n>
    | -s <size(KB)> | -t <type> | -S | -D[I] | -C
    | [-u <-unitid> | -ND | -NC | -g <group> }
```

Options and parameters

-h

Displays help/usage and version information.

-q

Terminates the interactive mode and exits the command.

-z or -zx

Makes the **horcctl** command enter the interactive mode. The -zx option guards
performing of the HORCM in the interactive mode. When this option detects a HORCM
shut down, interactive mode terminates.

OpenVMS cannot use the -zx option.

-I[H][M] [instance#]or -I[TC][SI] [instance#]

Specifies the command as [HORC]/[HOMRCF] or [TrueCopy]/[ShadowImage], and
used when specifying the CCI instance number.

-d

Interprets the control options following this option (-l <level>, -b <y/n>, -s <size(KB)>,
and -t <type>) as the parameters of the CCI commands.

-c

Interprets the control options following this option (-l <level>, -b <y/n> and -t <type>)
as the parameters of the HORC Manager (HORCM).

-l <level>

Sets the trace level (range = 0 to 15). If a negative value is specified, the trace mode is
canceled. A negative value 'n' must be specified as '--n'.

> ⚠ **Caution:**
>
> Do not change the trace level unless directed to do so by a Hitachi Vantara representative. Level 4 is the normal trace level setting. Levels 0-3 are for troubleshooting. Setting a trace level other than 4 may impact problem resolution. If you request a change of the trace level using the `horcctl -l <level>` command, a warning message appears, and this command enters interactive mode.

-b <y/n>

    Sets the trace writing mode: Y = buffer mode, N = synchronous mode.

-t <type>

    Sets the trace type (range = 0 to 511). When this option is used, only traces of the specified type are output. One or more values can be specified.

-s <size(KB)>

    Sets the trace buffer size in increments of 1024 bytes (default = 1 MB).

-S

    Shuts down HORCM.

-D

    Displays the command device name currently used by HORCM. If the command device is blocked due to online maintenance (microcode replacement) of the storage system, you can check the command device name in advance using this option.

-C

    Changes the command device name being used by HORCM and displays the new command device name. If the command device is blocked due to online maintenance (microcode replacement) of the storage system, you can change the command device in advance using this option.

    The horcctl -D -C command designates a protection mode command device by adding an asterisk (*) to the device file name as shown in the following example.

    HP-UX example with command device security:

    `# horcctl -D`

    Current control device = /dev/rdsk/c0t0d0*

    The `horcctl -DI` command shows the number of CCI instances of when HORCM has been started as shown in the following example.:

    HP-UX example without command device security:

    `# horcctl -DI`

    Current control device = /dev/rdsk/c0t0d0 AI = 14 TI = 0 CI = 1, where

-     AI: NUM of actual instances in use

-     TI: NUM of temporary instances in the storage system

-     CI: NUM of instances using current (own) instance

-u \<unitid\>

> Used when specifying the unit ID of a command device as the target. This option is effective when the -D or -C option is specified. If this option is not specified, the unit ID is 0.

-ND -g \<group\>

> Displays the network address and port name being used by HORCM. The $-g$ \<group\> option is used when specifying the group name defined in the configuration definition file.

-NC -g \<group\>

> Changes the network address and port name being used by HORCM and displays the new network address name. The $-g$ \<group\> option specifies the group name defined in the configuration definition file.

horctakeoff

This is a scripted command for executing several HORC operation commands combined. It checks the volume attribute (optionally specified) and decides a takeover action. The **horctakeoff** operation is defined to change from 3DC multi-target to 3DC multi-hop with the state of running APP, after that the **horctakeover** command is able to configure 3DC multi-target on the remote site without stopping the APP. The granularity of either a logical volume or volume group can be specified with this command.

Syntax

```
horctakeoff | -h | -q | -z[x] | -I[H][M][instance#] or
     -I[TC][SI][instance#] | -g[s] <group> | -d[s] <pair Vol>
     | -d[g][s] <raw_device> [MU#] | -d[g][s] <seq#> <LDEV#> [MU#]
     | -jp <id> | -js <id> | [-t  <timeout> ]| -nomsg }
```

Options and parameters

-h

> Displays help/usage and version information.

-q

> Terminates the interactive mode and exits the command.

-z or -zx

> Makes the **raidvchkdsp** command enter the interactive mode. The $-zx$ option guards performing of the HORCM in the interactive mode. When this option detects a HORCM shutdown, interactive mode terminates.

> OpenVMS cannot use the $-zx$ option.

-I[H][M] [instance#]or -I[TC][SI] [instance#]

> Specifies the command as [HORC]/[HOMRCF] or [TrueCopy]/[ShadowImage], and used when specifying the CCI instance number.

-g[s] <group>

> Specifies a group name (defined in the configuration definition file). The command is executed for the specified group unless the `-d <pair Vol>` option shown below is specified.

-d[s] <pair Vol>

> Specifies a logical (named) volume (defined in the configuration definition file). When this option is specified, the command is executed for the specified paired logical volume.

-d[g][s] <raw_device> [MU#]

> Searches the configuration definition file (local instance) for a volume that matches the specified raw device. If a volume is found, the command is executed on the paired volume (-d) or group (-dg). This option is effective without specification of the `-g <group>` option. If the specified raw_device is listed in multiple device groups, this applies to the first one encountered.

-d[g][s] <seq#> <LDEV#> [MU#]

> Searches the configuration definition file (local instance) for a volume that matches the specified sequence # and LDEV. If a volume is found, the command is executed on the paired logical volume (-d) or group (-dg). This option is effective without specification of the `-g <group>` option. If the specified LDEV is listed in multiple device groups, this applies to the first one encountered. <seq#> <LDEV#> can be specified in a hexadecimal (by addition of '0x') or decimal.

> 📄 **Note:**
>
> - When specifying the serial number for VSP 5000 series, add a "5" at the beginning of the serial number. For example, for serial number 12345, enter 512345.
>
> - When specifying the serial number for VSP G1x00 and VSP F1500, add a "3" at the beginning of the serial number. For example, for serial number 12345, enter 312345.

-jp <id> (HORC/UR only)

> The **horctakeoff** command can change 3DC configuration from 3DC multi-target to 3DC multi-hop. To create a 3DC multi-hop (TC_Sync to TC_Sync/UR P-VOL to UR), you must specify a journal ID for UR P-VOL. So this option is used for that purpose. If this option is not specified, a journal ID for UR P-VOL used for 3DC multi-target is inherited automatically.

-js <id>(HORC/UR only)

> The **horctakeoff** command can be changed 3DC configuration from 3DC multi-target to 3DC multi-hop. To create a 3DC multi-hop (TC_Sync to TC_Sync/UR to UR S-VOL), you must specify a journal ID for UR S-VOL. So this option is used for that purpose. If this option is not specified, a journal ID for UR S-VOL used with 3DC multi-target is inherited automatically. The CTG ID is also inherited automatically for the internal **paircreate** command.

-t <timeout>

> Specifies the maximum time to wait for the Sync P-VOL to Sync S-VOL delta data resynchronizing operation. Used for the internal **pairresync** command with the time-out period in seconds. If this option is not specified, the default timeout value (7200 sec) is used.

-noms

> Suppresses messages when this command is executed from a user program. This option must be specified at the beginning of the command arguments.

Returned values

The **horctakeoff** command sets the following returned values during exit allowing you to check the execution results.

▪ **Normal termination:**

- **0:**

▪ **Abnormal termination:**

- **other than 0:** The **horctakeoff** command returns the following error codes as well as generic error.

Error codes

Unrecoverable error should have been done without re-execute by handling of an error code. If the command failed, then the detailed status is logged in the CCI command log ($HORCC_LOG), even though the user script has no error handling.

Category	Error Code	Error Message	Value
Volume status (Unrecoverable)	EX_ENQVOL	Unmatched volume status within the group	236
	EX_INCSTG	Inconsistent status in group	229
	EX_EVOLCE	Pair Volume combination error	235
	EX_VOLCRE	Local and Remote Volume currency error	223
Timer (Recoverable)	EX_EWSTOT	Timeout waiting for specified status	233

Example 1
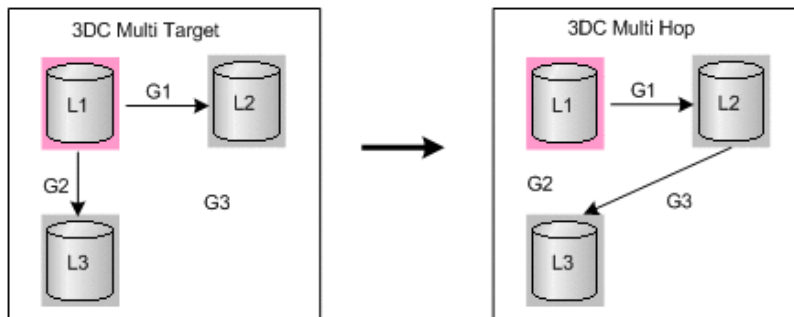
horctakeoff command on L1 local site

```
# horctakeoff -g G1 -gs G2
```

```
horctakeoff : 'pairsplit -g G1 -S -FHORC 2' is in progress
horctakeoff : 'pairsplit -g G1' is in progress
horctakeoff : 'pairsplit -g G2 -S' is in progress
horctakeoff : 'paircreate -g G1 -gs G2 -FHORC 2 -nocopy -f async
-jp 0 -js 1' is in progress
horctakeoff : 'pairsplit -g G1 -FHORC 2' is in progress
horctakeoff : 'pairresync -g G1' is in progress
horctakeoff : 'pairresync -g G1 -FHORC 2' is in progress
horctakeoff : horctakeoff done
```

Example 2
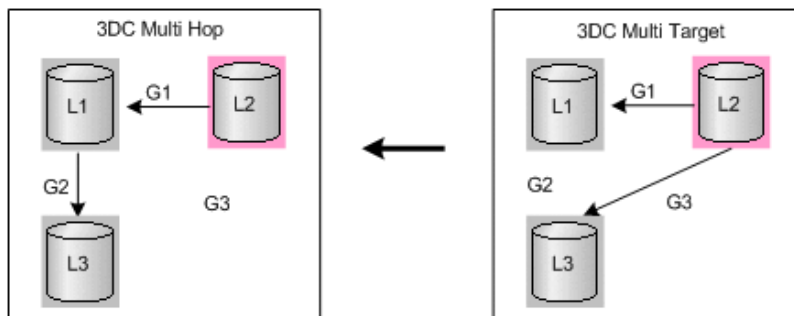
horctakeoff command on L2 local site

```
# horctakeoff -g G1 -gs G3
```

```
horctakeoff : 'pairsplit -g G1 -S -FHORC 1' is in progress.
horctakeoff : 'pairsplit -g G1' is in progress.
horctakeoff : 'pairsplit -g G3 -S' is in progress.
horctakeoff : 'paircreate -g G1 -gs G3 -FHORC 1 -nocopy -f async
-jp 0 -js 1' is in progress.
horctakeoff : 'pairsplit -g G1 -FHORC 1' is in progress.
```

Chapter 2: Data management commands

```
horctakeoff : 'pairresync -g G1' is in progress.
horctakeoff : 'pairresync -g G1 -FHORC 1' is in progress.
horctakeoff : horctakeoff done.
```

Example 3
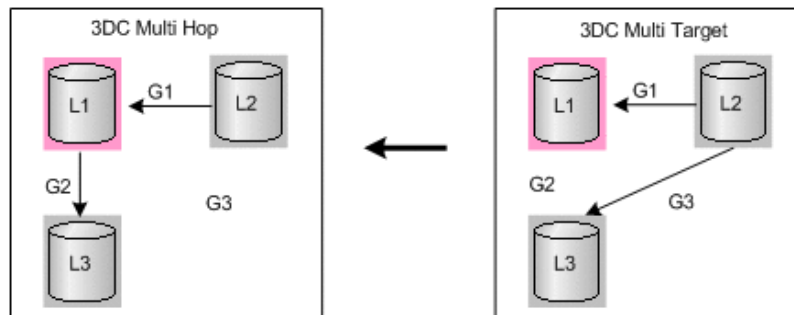
horctakeoff command on L1 remote site

```
# horctakeoff -g G1 -gs G2
```

```
horctakeoff : 'pairsplit -g G2 -S' is in progress.
horctakeoff : 'pairsplit -g G1' is in progress.
horctakeoff : 'pairsplit -g G1 -FHORC 2 -S' is in progress.
horctakeoff : 'paircreate -g G2 -vl -nocopy -f async -jp 0
-js 1' is in progress.
horctakeoff : 'pairsplit -g G2' is in progress.
horctakeoff : 'pairresync -g G1' is in progress.
horctakeoff : 'pairresync -g G2' is in progress.
horctakeoff : horctakeoff done.
```

Example 4
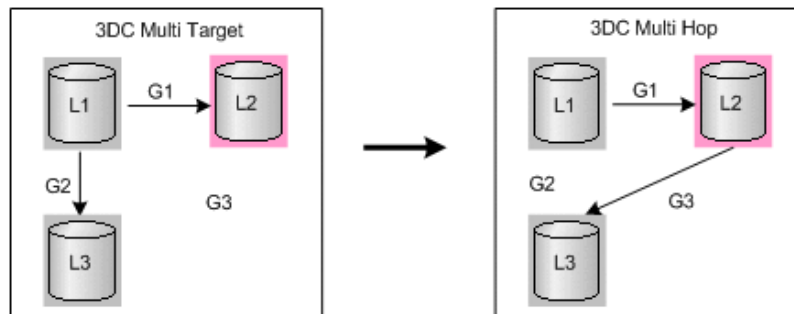
horctakeoff command on L2 remote site

```
# horctakeoff -g G1 -gs G3
```

```
horctakeoff : 'pairsplit -g G3 -S' is in progress.
horctakeoff : 'pairsplit -g G1' is in progress.
```

Chapter 2: Data management commands

```
horctakeoff : 'pairsplit -g G1 -FHORC 1 -S' is in progress.
horctakeoff : 'paircreate -g G3 -vl -nocopy -f async -jp 0
-js 1' is in progress.
horctakeoff : 'pairsplit -g G3' is in progress.
horctakeoff : 'pairresync -g G1' is in progress.
horctakeoff : 'pairresync -g G3' is in progress.
horctakeoff : horctakeoff done.
```

Chapter 3: Subcommands

CCI provides subcommands for the Windows platform and environment variable subcommands.

Windows subcommands

The CCI subcommands for Windows are executed as an option of another command. When you specify a subcommand as the only option of a command, you do not need to start HORCM. If another option of the command and the subcommand are specified on the same command line, place the other option after the subcommand.

The Windows subcommands are `findcmddev`, `drivescan`, `portscan`, `sync(d)`, `mount`, and `umount(d)`.

findcmddev

The `findcmddev` subcommand searches for command devices within the specified range of disk drive numbers. If it is found, the command device appears in the same format as in the configuration definition file. This subcommand is used when the command device name is not known and when the HORCM is not started.

> ⚠️ **Caution:** The `findcmddev` subcommand must be used when HORCM is running.

> 📄 **Note:** The `findcmddev` subcommand searches for the physical and logical drives associated with the command device. If the command device is indicated as a logical drive in addition to a physical drive, then a drive letter is assigned to the command device. You must delete the drive letter assigned to the command device to prevent utilization by general users.

The 'Volume{GUID}' must be made by setting a partition using the disk management. Do not format it (no file system). In a SAN environment, the physical drive number might be changed on every reboot. For this case, use the Volume (GUID) that keeps as the same name.

Syntax

```
-x findcmddev drive#(0-N)
```

Argument

drive#(0-N)

Specifies the range of disk drive numbers on the Windows system.

Example

The following shows an example of the **findcmddev** subcommand used as an option of the **raidscan** command and its output. This example searches for command devices in the range of disk drive numbers 0 through 20.

```
D:\HORCM\etc> raidscan  -x  findcmddev   hdisk0, 20
cmddev of Ser# 62496 = \\.\PhysicalDrive0
cmddev of Ser# 62496 = \\.\E:
cmddev of Ser# 62496 = \\.\Volume{b9b31c79-240a-11d5-a37f-00c00d003b1e}
```

drivescan

The **drivescan** subcommand displays the relationship between the disk numbers assigned by the Windows system and the LDEVs on the RAID storage system, and also displays attribute and status information for each LDEV.

Syntax

```
-x drivescan drive#(0-N)
```

Argument

drive#(0-N)

Specifies the range of data drive numbers on the Windows system.

Example

The following shows an example of the **drivescan** subcommand used as an option of the **raidscan** command and its output. This example displays the devices for the range of disk drive numbers from 0 to 20.

```
raidscan -x drivescan  harddisk0,20
```

```
Harddisk 0..Port[1] PhId[0] TId[0] Lun[0] [HITACHI] [DK328H-43WS]
Harddisk 1..Port[2] PhId[4] TId[29] Lun[0] [HITACHI] [OPEN-3]
           Port[CL1-J] Ser#[   30053] LDEV#[    9(0x009)]
           HORC=P-VOL HOMRCF[MU#0 = SMPL MU#1 = SMPL MU#2 = SMPL]
           RAID5[Group  2- 1] SSID = 0x0008 CTGID = 3
Harddisk 2..Port[ 2] PhId[ 4] TId[29] Lun[ 1] [HITACHI] [OPEN-3]
           Port[CL1-J] Ser#[ 30053] LDEV#[  10(0x00A)]
           HORC=S-VOL HOMRCF[MU#0 = SMPL MU#1 = SMPL MU#2 = SMPL]
           RAID5[Group  2- 1] SSID = 0x0004 CTGID = 3
```

Chapter 3: Subcommands

```
Harddisk 3..Port[2] PhId[4] TId[29] Lun[ 6] [HITACHI] [OPEN-3-CM]
          Port[CL1-J] Ser#[ 30053] LDEV#[  15(0x00F)]
```

Description of the `drivescan` subcommand output:

Harddisk #

Shows the data drive recognized by the Windows system.

Port

Shows the port number on the device adapter recognized by the Windows system.

PhId

Shows the bus number on the device adapter port recognized by Windows system.

TId

Shows the target ID of the data drive(s) connected to device adapter port. For the detail of fibre-to-SCSI address conversion, see the *Command Control Interface Installation and Configuration Guide*.

Lun

Shows the LU number of the data drive connected to device adapter port.

Port[CLX-Y]

Shows the port number on the storage system.

Ser#

Shows the production number (serial number) of the storage system.

- The serial number for VSP 5000 series is displayed with a "5" at the beginning ("512345" = serial number 12345).

- The serial number for VSP G1x00 and VSP F1500 is displayed with a "3" at the beginning ("312345" = serial number 12345).

LDEV#

Shows the LDEV ID (hexadecimal) of the specified volume.

HORC

Shows the TrueCopy/global-active device attribute (P-VOL, S-VOL, SMPL) of the specified volume.

HOMRCF

Shows the ShadowImage or Copy-on-Write Snapshot attribute (P-VOL, S-VOL, or SMPL) and MU number (0-2) of the specified volume.

RAIDX[Group]

Shows the physical location (frame number-parity group number) of the specified volume and the RAID level of this parity group.

SSID

Shows the SSID (hexadecimal) of the specified volume.

> 📄 **Note:** SSID is a parameter used by enterprise storage systems. Although SSID is not used by VSP E series, VSP G130, G/F350, G/F370, G/F700, G/F900, VSP G200, G400, G600, G800 and VSP F400, F600, F800, or HUS VM, the set value is displayed.

CTGID (TrueCopy Async, Universal Replicator, global-active device only)
Shows the consistency group ID of specified volume.

portscan

The `portscan` subcommand displays the devices on the specified port(s).

Syntax

```
-x portscan port#(0-N)
```

Argument

port#(0-N)
Specifies the range of port numbers on the Windows system.

Example

The following shows an example of the `portscan` subcommand used as an option of the `raidscan` command and its output. This example displays the devices for the range of ports from 0 to 20.

```
raidscan  -x  portscan  port0,20
```

```
PORT[ 0] IID [ 7] SCSI Devices
     PhId[ 0] TId[ 3] Lun[ 0] [MATSHIT] [CD-ROM CR-508]...Claimed
     PhId[ 0] TId[ 4] Lun[ 0]  [HP    ] [C1537A       ]...Claimed
PORT[ 1] IID [ 7] SCSI Devices
     PhId[ 0] TId[ 0] Lun[ 0] [HITACHI] [DK328H-43WS  ]...Claimed
PORT[ 2] IID [ 7] SCSI Devices
     PhId[ 0] TId[ 5] Lun[ 0]  [HITACHI] [OPEN-3       ]...Claimed
     PhId[ 0] TId[ 5] Lun[ 1]  [HITACHI] [OPEN-3       ]...Claimed
     PhId[ 0] TId[ 5] Lun[ 2]  [HITACHI] [OPEN-3       ]...Claimed
     PhId[ 0] TId[ 6] Lun[ 0]  [HITACHI] [3390-3A      ]...Claimed
```

> 📄 **Note:** This example displays the devices for the range of ports from 0 to 20.

Description of the `portscan` subcommand output:

Port
Shows the port number on the device adapter recognized by the Windows system

IID

Shows the initiator ID on the specified device adapter port

Phid

Shows the BUS number on the specified device adapter port

Tid

Shows the target ID of the data drive(s) connected to device adapter port.

LUN

Shows the LU number of each data drive connected to device adapter port. This item shows LDEV# of the partner who becomes a pair in or among the RAID storage systems.

sync, syncd

The `sync` (synchronization) subcommand sends unwritten data remaining on the Windows server to the specified device(s) to synchronize the pair(s) before the CCI command is executed. The `syncd` (synchronization delay) subcommand waits for the delayed IO for dismount after issued '`sync`'.

Syntax

```
-x sync[d] A: B: C:  ...
-x sync[d] all
-x sync[d] drive#(0-N)
-x sync[d] Volume#(0-N) ...
-x sync[d] D:\Directory or \Directory pattern...
```

Arguments

A: B: C: [\directory or \Directory pattern]

Specifies the logical drive that you want to synchronize. Data is flushed into the specified logical drive and the physical drive corresponding to the logical drive. If the specified logical drive has the directory mount volumes then SYNC is executed to all of the volumes on the logical drive as shown below:

```
pairsplit -x sync D:
[SYNC] D: HarddiskVolume2
[SYNC] D:\hd1 HarddiskVolume8
[SYNC] D:\hd2 HarddiskVolume9
```

[\directory or \Directory pattern] is used to find the directory mount point on the logical drive. If the directory is specified, then SYNC does execute to a directory mounted volume only.

```
pairsplit -x sync  D:\hd1
[SYNC] D:\hd1 HarddiskVolume8
```

Chapter 3: Subcommands

If the directory pattern is specified, then SYNC does execute to any directory mounted volumes identified to '\directory pattern'.

```
pairsplit  -x  sync  D:\h
[SYNC] D:\hd1 HarddiskVolume8
[SYNC] D:\hd2 HarddiskVolume9
```

all

Synchronizes all logical drives and the physical drives corresponding to the logical drives assuming that they are on the data drives. The logical drive on which the CCI software is installed and the logical drive containing the Windows directory are excluded. If the logical drive has the directory mount volumes then SYNC is executed to all volumes on the logical drive as shown below:

```
pairsplit  -x  sync  all
[SYNC] C: HarddiskVolume1
[SYNC] D:\hd1 HarddiskVolume8
[SYNC] D:\hd2 HarddiskVolume9
[SYNC] G: HarddiskVolume10
```

drive#(0-N)

Specifies the physical drives to be flushed.

Volume#(0-N)

Specifies the LDM volumes to be flushed. Volume# must be specified '\Vol#,\Dms#,\Dmt#, \Dmr# or Volume{...}' as LDM volume for Windows systems.

To flush HarddiskVolumeX:

```
-x sync  \VolX
```

For information on '\Vol#, \Dms#,\Dmt#,\Dmr# or Volume{...}' for LDM volumes, see **Volume Discovery Function** in the *User and Reference Guide*.

Examples

The following examples show the `sync` subcommand used as an option of the `pairsplit` command.

`sync` subcommand with `pairsplit`

For the following example, the data remaining on logical drives C: and D: is written to disk, all pairs in the specified group are split (status = PSUS), and read/write access is enabled for all S-VOLs in the specified group.

```
pairsplit -x sync  C: D:  -g  oradb  -rw
```

`sync` subcommand with `pairsplit -S`

For the following example, the data remaining on physical devices harddisk2 and harddisk3 is written to disk, all pairs in the specified group are deleted (status = SMPL), which enables read/write access for all secondary volumes.

```
pairsplit -x sync  hdisk2  hdisk3 -g oradb -S
```

> 📄 **Note:** The `sync` subcommand has the following behavior on any conditions:
>
> - If the logical drives designated as the objects of the `sync` subcommand are not opened to any applications, then `sync` flushes the system buffer to a drive and makes the dismount state for this drive.
>
> - If the logical drives designated as the objects of the `sync` subcommand are already opened to any applications, then `sync` only flushes the system buffer to a drive.

This flushes the system buffer before `pairsplit` without unmounting the P-VOL (opening state), and indicates as [WARNING] below:

```
pairsplit -x sync  C:
WARNING: Only flushed to [\\.\C:] drive due to be opening
[SYNC] C: HarddiskVolume3
```

> 📄 **Note:** `syncd` has the following behavior as well:
>
> - If the logical drives designated as the objects of the `sync` subcommand are not opened to any applications, then `syncd` flushes the system buffer to a drive and waits (30 sec) the delayed (paging) IO for dismount after made the dismount state about the drive. If the logical drives are opened to applications, the `syncd` subcommand waits (30 sec) after the flush of system buffer.
>
>   This avoids a problem that NTFS on P-VOL is split on inconsistent state because Windows 2003 delays the IO for dismounting.

> 📄 **Note:** If `sync` has failed, you need to confirm the following conditions:
>
> - The logical and physical drives designated as the objects of the `sync` subcommand are not opened to any applications. For example, confirm that Explore is not pointed on the target drive. If Explore is pointed on the target drive, the target drive is opening.
>
> - The `sync` subcommand does not ignore the detected error on the NT file system, so `sync` executes successfully in normal case (NO ERROR case) only on NT file system. For example, confirm the target drive has no failure on the system for Event Viewer. In this case, you must reboot the system or delete the partition and reconfigure the target drive.

mount

The `mount` subcommand mounts the specified drive (indicates HDD, SSD, SCM, and FMD) to the specified partition on the specified data drive using the drive letter. When the `mount` subcommand is executed without an argument, all currently mounted drives (including directory mounted volumes) are displayed, and logical drive has been mounting an LDM volume then displays Harddisk#[n] configured an LDM volume.

Syntax

```
-x mount
-x mount drive: hdisk# [partition#] (for Windows)
-x mount drive: Volume#(0-N) (for Windows)
-x mount drive: [\directory] Volume#(0-N) (for Windows)
```

Arguments

drive: hdisk# [partition #]

Specifies the logical drive, data drive (number), and partition to be mounted.

drive: [\directory] Volume#

Specifies the logical drive and LDM volume name and number to be mounted. Volume# must be specified '\Vol# or \Dms# or \Dmt# or \Dmr# ' as LDM volume for Windows.

To mount HarddiskVolumeX: -x mount C: hdX or -x mount C: \VolX

For information on '\Vol# or \Dms# or \Dmt# or \Dmr#' for LDM volumes, see **Volume Discovery Function** in the *User and Reference Guide*.

[\directory]: Specifies the directory mount point on the logical drive.

pairsplit -x mount D:\hd1 \Vol8 D:\hd1 <+> HarddiskVolume8 pairsplit -x mount D:\hd2 \Vol9 D:\hd2 <+> HarddiskVolume9

> ⚠️ **Caution:** The partition on the specified data drive must be recognized on the Windows system.
>
> [\directory] for the mount must be specified a mount point without embedded space character.
>
> If [\directory] is detected as mount point with embedded space (that is, aaabbb), then the directory is shown by adding '…' to first strings as below.
>
> pairsplit -x mount
>
> Drive FS_name VOL_name Device Partition ... Port PathID Targ Lun
>
> D: NTFS Null Harddiskvolume3 ... Harddisk2
>
> D:\aaa... NTFS Null Harddiskvolume4 ... Harddisk3
>
> The same method is used for the '`inqraid $LETALL`' and '`raidscan -pi $LETALL -find`' commands.

Chapter 3: Subcommands

Example

The following example shows the **mount** subcommand used as an option of the **pairsplit** command and its outputs.

Examples for Windows

This example executes the **mount** subcommand from a subcommand option of **pairsplit**. It mounts the "F:" drive to the harddiskvolume2, and then displays mounted devices.

```
pairsplit -x mount  F: hdisk2
pairsplit  -x mount
```

```
Drive  FS_name VOL_name Device Partition ... Port PathID Targ Lun
C      NTFS    Null     Harddiskvolume1  ...    Harddisk0
F:     NTFS    Null     Harddiskvolume2  ...    Harddisk1
D:     NTFS    Null     Harddiskvolume3  ...    Harddisk2
D:\hd1 NTFS    Null     Harddiskvolume4  ...    Harddisk3
D:\hd2 NTFS    Null     Harddiskvolume5  ...    Harddisk4
G:     NTFS    Null     HarddiskDmVolumes\…\Volume1...Harddisk5[3]
```

This example executes **mount** from command option of the **pairsplit** command and then displays the mounted devices:

Description of the mount subcommand output:

Drive

> Shows the logical drive recognized by the Windows system

FS_name

> Shows the name of the file system formatted on the logical drive

VOL_name

> Shows the volume label name for the logical drive

Device, Partition

> Shows the device name and partition for the mounted logical drive

Port,PathID,Targ,Lun

> Shows the port number, path group ID (bus), target ID, and LUN for the device adapter of the mounted logical drive.

umount, umountd

The **umount** subcommand unmounts the specified logical drive and deletes the drive letter. Before deleting the drive letter, this subcommand executes sync internally for the specified logical drive and flushes unwritten data. The **umountd** subcommand unmounts the logical drive after waiting the delayed IO for dismount.

Syntax

```
-x umount[d] drive: [time]


-x umount[d] drive:[\directory] [time] (for Windows)
```

Argument

drive

Specifies the mounted logical drive.

[\directory] [time]

Specifies the directory mount point on the logical drive. This command option calls 'mountvol /P' internally, if 'USE_MOUNTVOL_P' environment variable is specified. In case of Windows, it is required to specify 'USE_MOUNTVOL_P' variable to avoid a problem of mount.

pairsplit -x umount D:\hd1 D:\hd1 <-> HarddiskVolume8 set USE_MOUNTVOL_P=1
pairsplit -x umount D:\hd2 D:\hd2 <-> HarddiskVolume9

Example for waiting 45 sec:

pairsplit -x umount D: 45

D: <-> HarddiskVolume8

⚠️ **Caution: Umountd** has above restriction and it will prompt the delayed IO for dismount. Wait 30 seconds until the completion and release the mount point after making dismount status of the logical drive. This avoids a problem (Windows 2003 only) that the delayed writing for dismount as Event ID51, 57.

Example

The following example shows the **umount** subcommand used as an option of the **pairsplit** command.

```
pairsplit -x umount F:  -x umount G: -g oradb -rw
pairsplit -x mount
```

```
Drive FS_name VOL_name Device    Partition... Port PathID Targ Lun
C:    FAT     Null     Harddisk0 Partition1...   1     0     0   0
Z:    Unknown Unknown  CdRom0              ...  Unknown
```

This example unmounts the F: and G: drives, splits all pairs in the specified group (status = PSUS), enables read/write access to all secondary volumes in the specified group, and then displays all mounted drives.

Description of the umount subcommand output:

Drive:
> Shows the logical drive recognized by the Windows system

FS_name
> Shows the name of the file system formatted on the logical drive

VOL_name
> Shows the volume label name for the logical drive

Device, Partition
> Shows the device name and partition for the mounted logical drive

Port,Phid,Targ,Lun
> Shows the port number, path group ID (bus), target ID, and LUN for the device adapter of mounted logical drive.

> **📄 Note:**
>
> The `umount` subcommand flushes (sync) the system buffer of the associated drive before deleting the drive letter. If `umount` fails, confirm the following conditions:
>
> - The logical and physical drives designated as the objects of the `umount` subcommand are not opened to any applications. For example, confirm that Explore is not pointed on the target drive. If it is, then the target drive is opening.
>
> - The `umount` subcommand does not ignore the detected error on the NT file system, so that `umount` is successful in a normal case (NO ERROR case) only on NT file system. For example, confirm the target drive has no failure on the system for Event Viewer. If so, you must reboot the system or delete the partition and reconfigure the target drive.

> **📄 Note:**
>
> The `umountd` subcommand has the following behavior as well.
>
> Unmount the logical drive after waiting (30 sec) the delayed (paging) IO for dismount after flushed the system buffer to a drive.
>
> This avoids a problem (Windows 2003 only) that NTFS on P-VOL is split on inconsistent state because Windows 2003 (SP1) delays the IO for dismounting. This also avoids a problem that the delayed (paging) IO for dismounting is written on SVOL_PAIR(Writing Disable) state by rescan, and logged as Windows event (that is, ID51,57).
>
> These problems do not occur on Windows 2008 and later systems.

Chapter 3: Subcommands

Environment variable subcommands

If no environment variables are set in the execution environment, the following environment variable subcommands set or cancel an environment variable within the CCI command.

- **setenv**: The **setenv** subcommand sets the specified environment variable(s).

- **usetenv**: The **usetenv** subcommand deletes the specified environment variable(s).

- **env**: The **env** subcommand displays the environment variable(s).

- **sleep**: The **sleep** subcommand causes CCI to wait for the specified time.

Syntax

```
-x setenv vaname value
-x usetenv vaname
-x env
-x sleep time
```

Arguments

Varname

Specifies the environment variable to be set or canceled.

Value

Specifies the value or character string of the environment variable to be set.

Time

Specifies the sleep time in seconds.

> ⚠️ **Caution:** The environment variables must be set before connecting to HORCM. And it must be specified during interactive mode ($-z[x]$ option). If specified with other than interactive mode, all specified environment variables are not enabled. Changing an environment variable after a CCI command execution error is invalid.

Example

The following examples show the **setenv** and **usetenv** subcommands used as an option of the **raidscan** command. This example changes from 'HORC' to 'HOMRCF' an execution environment of the **raidscan** command that makes a dialog mode, because of establishing 'HORCC_MRCF' as an environment variable.

```
raidscan[HORC]:  -x setenv  HORCC_MRCF 1
raidscan[HOMRCF]:
```

```
raidscan[HOMRCF]:  -x usetenv  HORCC_MRCF
raidscan[HORC]:
```

Chapter 4: Command tools

CCI provides the following command tools:

- inqraid

- mkconf

- rmawk

inqraid

The **inqraid** command is a tool used to confirm the drive connection between the storage system and host system. The **inqraid** command displays the relationship between special file(s) on the host system and actual physical drive of the RAID storage system.

Syntax

```
/HORCM/usr/bin/inqraid [-h | quit | -inqdump
| -fx[p][l][g][c][h][n] | -find[c] | <special file>
| -CLI[W][P][N][B] | -sort [-CM ] | -export][-n]

/HORCM/etc/inqraid [-h | quit | -inqdump | -fx[p][l][g][c][h][n]
| -find[c] | <special file> |-CLI[W][P][N][B] | -sort [-CM ]
| -gvinf | -svinf | -gplba | -fv| -export]
```

Options and parameters

-h

Displays Help/Usage.

quit

Terminates from waiting STDIN and exits this command.

-inqdump

Displays information for standard inquiry with Dump Image of hexadecimal.

-fx

Displays the LDEV number (hexadecimal).

-find[c]

Finds the appropriate group within the configuration file using a special file provided by STDIN.

- `-find`: Searches a group on the configuration definition file (local instance) from <special file> of STDIN by using **pairdisplay** command, and uses the following options of the **pairdisplay** command to display its state. This option must be specified HORCMINST as command execution environment.

  For ShadowImage/Copy-on-Write Snapshot:

  ```
  pairdisplay -d <Seq#> <LDEV#> 0 1 2 -l [-fx] [-CLI] 2>/dev/null
  ```

  For TrueCopy/Universal Replicator/global-active device:

  ```
  pairdisplay -d <Seq#>  <LDEV#>  -l  [-fx]  [-CLI]  2>/dev/null
  ```

  > **📄 Note:**
  >
  > <Seq#> and <LDEV#> are included using SCSI Inquiry command.
  >
  > - When specifying the serial number for VSP 5000 series, add a "5" at the beginning of the serial number. For example, for serial number 12345, enter 512345.
  >
  > - When specifying the serial number for VSP G1x00 and VSP F1500, add a "3" at the beginning of the serial number. For example, for serial number 12345, enter 312345.

- `<special file>`: This option is used to specify the special file name as argument of command. If no argument, this command makes mode that waits for STDIN. The input from the special file has to be waited.

- ▪ `-findc`: Uses the following options of the **pairdisplay** command, and displays with CLI format by editing an output of **pairdisplay** command. This option must be specified HORCMINST as command execution environment.

  For ShadowImage/Copy-on-Write Snapshot:

  ```
  pairdisplay -d <Seq#>  <LDEV#>  <MU#>  -fd   -CLI  2>/dev/null
  ```

  For TrueCopy/Universal Replicator/global-active device:

  ```
  pairdisplay -d <Seq#>  <LDEV#>  -fd   -CLI   2>/dev/null
  ```

  - • When specifying the serial number for VSP 5000 series, add a "5" at the beginning of the serial number. For example, for serial number 12345, enter 512345.

  - • When specifying the serial number for VSP G1x00 and VSP F1500, add a "3" at the beginning of the serial number. For example, for serial number 12345, enter 312345.

- ▪ `<special file>`: Specifies a special file name as the argument of a command. If no argument, this command makes mode that waits for STDIN. The input from the special file has to be waited.

-CLI

Specifies the display of structured column output for command line interface (CLI) parsing. Also used for '-find' option. The delimiters between columns can be spaces and/or dashes (-).

-CLIWP, -CLIWN

Displays the WWN (world wide name for host adapter) with command line interface (CLI) format, also used for '-find' option.

-sort[CM]

Sorts the target devices by Serial#,LDEV# order.

- ▪ The serial number for VSP 5000 series is displayed with a "5" at the beginning ("512345" = serial number 12345).

- ▪ The serial number for VSP G1x00 and VSP F1500 is displayed with a "3" at the beginning ("312345" = serial number 12345).

[CM] Searches command device from the specified special file (STDIN or argument) and displays the command device only in structure definition file image. This option is valid within '-sort' option.

-gvinf or -gvinfex

Windows systems only. The -gvinfex option is for GPT disk only.

Gets the signature and volume layout information of a raw device file provided via STDIN or arguments, and saves this information to the system disk with the following format:

\WindowsDirectory\VOLssss_llll.ini

where ssss = serial#, llll = LDEV#

- The serial number for VSP 5000 series is displayed with a "5" at the beginning ("512345" = serial number 12345).

- The serial number for VSP G1x00 and VSP F1500 is displayed with a "3" at the beginning ("312345" = serial number 12345).

Normally this option is used by the Windows Disk Management after setting the signature and volume layout information for S-VOL. You do not need to be aware of this file.

-svinf[=PTN] or -svinfex[=PTN]

Windows systems only. The -svinfex[=PTN] option is for GPT disk only.

Reconfigure the signature and volume layout information that was saved to the system disk to a device provided by STDIN or arguments. Gets the serial# and LDEV# for the target device issuing SCSI Inquiry, and sets the signature and volume layout information into VOLssss_llll.ini file to the readout device. This option will set correctly because the signature and volume layout information is managed by the serial# and LDEV# without depend on Harddisk#, even if Harddisk# is changed by the configuration changes.

ssss = serial#

- The serial number for VSP 5000 series is displayed with a "5" at the beginning ("512345" = serial number 12345).

- The serial number for VSP G1x00 and VSP F1500 is displayed with a "3" at the beginning ("312345" = serial number 12345).

[=PTN]: Specifies a strings pattern to interpret the strings provided by STDIN or argument as a raw device.

\Device\HarddiskVolume#(number) is made in a sequential order executed -svinf to Harddisk, and its number will remain the same as long as the system configuration is not changed. If you want to make \Device\HarddiskVolume#(number) more absolutely, then make \Device\HarddiskVolume# in serial# and LDEV# order by using the '-sort' option as shown below:

[VOL61459_451_5296A763] -> Harddisk3 [OPEN-3]

[VOL61459_452_5296A760] -> Harddisk4 [OPEN-3]

[VOL61459_453_5296A761] -> Harddisk5 [OPEN-3]

-gplba or -gplbaex

Windows systems only. The `-gplbaex` option is for GPT disk only.

Displays usable LBA on a physical drive in units of 512 bytes, and specifies `[slba]` `[elba]` options for **raidvchkset** command.

Example:

```
C:\HORCM\etc>inqraid $Phys -CLI -gplba -sort
Harddisk11 : SLBA = 0x00003f00 ELBA = 0x000620d9 PCNT = 7 [OPEN-3-CVS    ]
Harddisk12 : SLBA = 0x00003f00 ELBA = 0x00042ad1 PCNT = 4 [OPEN-3-CVS    ]
Harddisk13 : SLBA = 0x0000003f ELBA = 0x000620d9 PCNT = 1 [OPEN-3-CVS    ]
```

SLBA: Displays usable starting LBA in units of 512 bytes. ELBA: Displays usable ending LBA (ELBA -1) in units of 512 bytes. PCNT: Displays the number of partitions.

Example for setting of Harddisk11: `C:\HORCM\etc>raidvchkset -d hd11 -vs 16 0x00003f00 0x000620d9`

-fv

Windows systems only.

Displays the Volume{GUID} via $Volume with wide format. The serial number (serial#) is SERIAL.

- The serial number for VSP 5000 series is displayed with a "5" at the beginning ("512345" = serial number 12345).

- The serial number for VSP G1x00 and VSP F1500 is displayed with a "3" at the beginning ("312345" = serial number 12345).

Example:

```
C:\HORCM\etc>inqraid -CLI $Vol -fv
```

```
DEVICE_FILE      PORT   SERIAL  LDEV CTG  H/M/12  SSID R:Group PRODUCT_ID
Volume{cec25efe-d3b8-11d4-aead-00c00d003b1e}\Vol3\Dsk0  CL2-D  62496  256
-       -      -        - OPEN-3-CVS-CM
```

-fl

If the target device file is set as a protection volume, the device file name is shown by appending '*'. It is valid when specified with "-CLI" option. If the `-fl` option is specified, the data protection volume is a Data Retention Utility volume.

- The serial number for VSP 5000 series is displayed with a "5" at the beginning ("512345" = serial number 12345).

- The serial number for VSP G1x00 and VSP F1500 is displayed with a "3" at the beginning ("312345" = serial number 12345).

-fg

Shows a LUN on the host view by finding a host group.

-CLIB -sort [-fh | -fc]

Displays the number of the tables of the differential bitmap which are required on the shared memory. This option is used to find the number of pairs that can be created to a paired volume in the storage system. Specify this option with the -sort option to sort the specified special files (the standard input or the argument) in order of the following priority, Serial#, and then LDEV#.

- When you specify `-fh` option: the number of tables of the differential bitmap for TC, TCz, UR, and GAD is displayed.

- When you specify `-fc` option: the number of tables of the differential bitmap in a cylinder size for TC, TCz, UR, and GAD is displayed.

- When you omit the `-fh` option and `-fc` option: the number of tables of the differential bitmap for SI and SIz is displayed.

Example to display the number of tables of the differential bitmap for SI and SIz:

```
# ls /dev/rdsk/* | inqraid -CLIB -sort
DEVICE_FILE PORT   SERIAL LDEV   SL CL +SI/SI UNUSED PRODUCT_ID
c1t0d0      CL1-E 63516 0       0  0    -     -      OPEN-9-CM
c1t0d1      CL1-E 63516 12288   0  0    1     30718  OPEN-3
c1t0d2      CL1-E 63516 12403   0  0    4     30718  OPEN-9
c1t0d3      CL1-E 63516 12405   0  0    9     30718  OPEN-E
c1t0d4      CL1-E 63516 12800   0  0    12    30718  OPEN-8
c1t0d5      CL1-E 63516 12801   0  0    18    30718  OPEN-8*2
c1t0d6      CL1-E 63516 13057   0  0    31    30718  OPEN-L
c2t0d6      CL2-E 63516 13057   0  0    31    30718  OPEN-L
```

Example to display the number of tables of the differential bitmap for TC, TCz, UR, and GAD:

```
# ls /dev/rdsk/* | inqraid -CLIB -sort -fh
DEVICE_FILE PORT   SERIAL LDEV   SL CL +TC/UR UNUSED PRODUCT_ID
c1t0d0      CL1-E 63516 0       0  0    -     -      OPEN-9-CM
c1t0d1      CL1-E 63516 12288   0  0    1     11605  OPEN-3
c1t0d2      CL1-E 63516 12403   0  0    3     11605  OPEN-9
c1t0d3      CL1-E 63516 12405   0  0    10    11605  OPEN-E
c1t0d4      CL1-E 63516 12800   0  0    11    11605  OPEN-8
c1t0d5      CL1-E 63516 12801   0  0    13    11605  OPEN-8*2
c1t0d6      CL1-E 63516 13057   0  0    21    11605  OPEN-L
c2t0d6      CL2-E 63516 13057   0  0    21    11605  OPEN-L
```

- SERIAL: Serial number

- SL: The SLPR number of LDEV.

- CL: The CLPR ID of LDEV.

- +SI/SI: Shows the accumulated total number of tables of the differential bitmap for SI and SIz. The increased number shows necessary tables of the differential bitmap for one SI or SIz pair .

> 📄 **Note:**
>
> The number of tables of the differential bitmap for the following LDEV are excepted from the accumulated total.
>
> - LDEVs whose Serial# and LDEV# are same (they are calculated as one LDEV if the multiple ports share the LDEV).
>
> - An LDEV which does not manage differential bitmap on a shared memory.
>
> - An LDEV which is used as a command device.

- +TC/UR: Shows the accumulated total number of tables of the differential bitmap for TC, TCz, UR, and GAD. The increased number shows necessary tables of the differential bitmap for TC, TCz, UR, and GAD. The '-fc' option displays the number of tables of differential bitmap in a cylinder size.

> 📄 **Note:**
>
> The number of tables of the differential bitmap for the following LDEVs are not included in the accumulated total:
>
> - LDEVs whose Serial# and LDEV# are same (they are calculated as one LDEV if the multiple ports share the LDEV).
>
> - An LDEV which does not manage differential bitmap on a shared memory.
>
> - An LDEV which is used as a command device.

- UNUSED: Shows the number of tables of unused differential bitmap for SI, SIz, TC, TCz, UR, and GAD.

-CLI -fn

Displays the LDEV nickname in the PRODUCT ID. If the storage system does not support LDEV nickname, '-' is displayed. The serial number (serial#) is SERIAL.

- The serial number for VSP 5000 series is displayed with a "5" at the beginning ("512345" = serial number 12345).

- The serial number for VSP G1x00 and VSP F1500 is displayed with a "3" at the beginning ("312345" = serial number 12345).

Example:

```
# ls /dev/rdsk/c57t4* | ./inqraid -CLI -fn
DEVICE_FILE PORT  SERIAL LDEV CTG C/B/12 SSID R:Group LDEV_NIC_NAME
c57t4d0     CL1-D 62496  32    -   s/P/ss 0004 5:01-03 my_volume_1
c57t4d3     CL1-D 62496  35    -   s/P/ss 0004 5:01-03 my_volume_2
```

```
c57t4d4     CL1-D 62496  36   -   s/P/ss 0004 5:01-01 -
c57t4d5     CL1-D 62496  37   -   s/P/ss 0004 5:01-02 -
```

-CLI -export

This option outputs as export format after obtaining the storage system device information from the specified special file (STDIN or argument). The device information includes 'Keyword, Serial#, Ldev#, Device file name, ...'. This export formatted file can import as STDIN of the raidscan -find verify or the raidscan -find inst. If the application server and CCI server are operated on different hosts, volume discovery can be operated with using this option among the hosts.

- The serial number for VSP 5000 series is displayed with a "5" at the beginning ("512345" = serial number 12345).

- The serial number for VSP G1x00 and VSP F1500 is displayed with a "3" at the beginning ("312345" = serial number 12345).

Example:

```
# ls /dev/rdsk/c1t* | inqraid -CLI -export
INQRAID:@CL4-G@64015@0@124@OPEN-V-CM@/dev/rdsk/c1t0d0s2
INQRAID:@CL4-G@64015@1@124@OPEN-V-CM@/dev/rdsk/c1t0d1s2
INQRAID:@CL4-G@64015@2@95@OPEN-V@/dev/rdsk/c1t0d2s2
INQRAID:@CL4-G@64015@3@95@OPEN-V@/dev/rdsk/c1t0d3s2
INQRAID:@CL4-G@64015@4@95@OPEN-V@/dev/rdsk/c1t0d4s2
INQRAID:@CL4-G@64015@5@95@OPEN-V@/dev/rdsk/c1t0d5s2
INQRAID:@CL4-G@64015@7@95@OPEN-V@/dev/rdsk/c1t0d7s2
```

Restrictions on device naming

STDIN or special files are specified as follows:

- HP-UX:
  - /dev/rdsk/*
  - /dev/rdisk/disk*

- Solaris:
  - /dev/rdsk/*s2
  - c*s2

- Linux:
  - /dev/sd...
  - /dev/rd...
  - /dev/raw/raw*

- ▪ zLinux:
  - • /dev/sd...
  - • /dev/dasd…
  - • /dev/rd...
  - • /dev/raw/raw*
- ▪ AIX:
  - • /dev/rhdisk*
  - • /dev/hdisk*
  - • hdisk*
- ▪ DIGITAL or Tru64:
  - • /dev/rrz*c
  - • /dev/rdisk/dsk*c
  - • /dev/cport/scp*
- ▪ DYNIX:
  - • /dev/rdsk/sd*
  - • sd* for only unpartitioned raw device
- ▪ IRIX64:
  - • /dev/rdsk/*vol
  - • /dev/rdsk/node_wwn/*vol/*
  - • /dev/dsk/*vol
  - • /dev/dsk/node_wwn/*vol/*
- ▪ Windows: For information about LDM volumes for Windows systems, see **Remote Volume Discovery** in the *Command Control Interface User and Reference Guide*.
  - • hdX-Y
  - • $LETALL
  - • $Volume
  - • $Phys
  - • D:\Vol(Dms,Dmt,Dmr)X\DskY
  - • \Vol(Dms,Dmt,Dmr)X\DskY
  - • D:\DskX\pY
  - • \DskX\pY

Lines starting with '#' via STDIN are interpreted as comments.

[-n]

This -n option is specified to use a device file for NVMe-oF. Specifying this option enables the use of both device files for NVMe-oF and SCSI.

Chapter 4: Command tools

If the HORCC_NVME environment variable is specified, both device files for NVMe-oF and SCSI can be used without the `-n` option specified.

To specify this option, a command device must be enabled in a namespace on the storage system. For more information about whether a command device can be enabled in a namespace on the storage system, see the *Provisioning Guide* for your storage system.

Example 1: using `inqraid` and `system` command to display the connection between STDIN special file and actual physical drive of storage system

HP-UX system:

```
# ioscan -fun | grep rdsk | ./inqraid
/dev/rdsk/c0t2d1 -> [HP] CL2-D Ser = 30053 LDEV = 9 [HP  ]
          [ OPEN-3  ] HORC = P-VOL  HOMRCF[MU#0 = SMPL MU#1 = SMPL
         MU#2 = SMPL] RAID5[Group  2- 1] SSID = 0x0008 CTGID = 3
/dev/rdsk/c0t4d0 -> [HP] CL2-D Ser = 30053 LDEV = 14 [HP ]
          [OPEN-3-CM    ] RAID5[Group  2- 1] SSID = 0x0008
```

Linux and zLinux system:

```
# ls /dev/sd* | ./inqraid
/dev/sdh -> CHNO =   0 TID =   1 LUN =   7
            [HP] CL2-B Ser =   30053 LDEV =  23 [HP    ]
            [OPEN-3    ] HORC = P-VOL  HOMRCF[MU#0 = SMPL
            MU#1 = SMPL MU#2 = SMPL] RAID5[Group  1- 2]
            SSID = 0x0004 CTGID = 2
/dev/sdi -> CHNO =   0 TID =   4 LUN =   0
            [HP] CL2-B Ser =   30053 LDEV =  14 [HP    ]
            [OPEN-3-CM    ] RAID5[Group  1- 2] SSID = 0x0004
```

Linux system (namespace for NVMe-oF):

Namespace = Y is displayed if a namespace is used.

```
[root@localhost ~]# ls /dev/nvme4n1 | /HORCM/usr/bin/inqraid -n
/dev/nvme4n1 -> [ST] CL1-D Ser =  500054 LDEV =1536 [HITACHI ] [OPEN-V-CM      ]
                A-LUN[PoolID 0002] SSID = 0x0011 Namespace = Y
```

Solaris system:

```
# ls /dev/rdsk/* | ./inqraid
/dev/rdsk/c0t2d1 -> [HP] CL2-D Ser =   30053 LDEV =   9 [HP     ]
            [OPEN-3    ] CA = P-VOL  BC[MU#0 = SMPL MU#1 = SMPL
            MU#2 = SMPL] RAID5[Group  2- 1] SSID = 0x0008
            CTGID = 3
/dev/rdsk/c0t4d0 -> [HP] CL2-D Ser =   30053 LDEV =  14 [HP     ]
            [OPEN-3-CM    ] RAID5[Group  2- 1] SSID = 0x0008
```

AIX system:

```
# lsdev -C -c disk | grep hdisk | ./inqraid
hdisk1 -> [SQ] CL2-D Ser =  30053 LDEV =  9 [HITACHI ] [OPEN-3  ]
          HORC = P-VOL  HOMRCF[MU#0 = SMPL MU#1 = SMPL
          MU#2 = SMPL] RAID5[Group  2- 1] SSID = 0x0008 CTGID = 3
hdisk2 -> [SQ] CL2-D Ser = 30053 LDEV = 14 [HITACHI ] [OPEN-3-CM ]
          RAID5[Group  2- 1] SSID = 0x0008
```

Windows system:

```
C:\HORCM\etc> echo hd1-2 | inqraid ( or inqraid hd1-2 )
Harddisk 1 -> [SQ] CL2-D Ser = 30053 LDEV = 9 [HITACHI ] [OPEN-3 ]
              HORC = P-VOL  HOMRCF[MU#0 = SMPL MU#1 = SMPL
              MU#2 = SMPL] RAID5[Group  2- 1] SSID = 0x0008
              CTGID = 3
Harddisk 2 -> [SQ] CL2-D Ser = 30053
              LDEV = 14 [HITACHI ] [OPEN-3-CM ] RAID5[Group  2- 1]
              SSID = 0x0008
```

Tru64 UNIX system:

```
# ls /dev/rdisk/dsk* | ./inqraid
/dev/rdisk/dsk10c -> [SQ] CL2-D Ser = 30053 LDEV = 9 [HITACHI ]
                     [OPEN-3 ] HORC = P-VOL  HOMRCF[MU#0 = SMPL
                     MU#1 = SMPL MU#2 = SMPL] RAID5[Group  2- 1]
                     SSID = 0x0008 CTGID = 3
/dev/rdisk/dsk11c -> [SQ] CL2-D Ser = 30053 LDEV = 14 [HITACHI ]
                     [OPEN-3-CM  ] RAID5[Group 2- 1] SSID = 0x0008
```

IRIX system with FC_AL:

```
# ls /dev/rdsk/*vol | ./inqraid
/dev/rdsk/dks1d6vol -> [SQ] CL2-D Ser = 30053 LDEV = 9 [HITACHI ]
                       [OPEN-3 ] HORC = P-VOL
                       HOMRCF[MU#0 = SMPL MU#1 = SMPL MU#2 = SMPL]
                       RAID5[Group  2- 1] SSID = 0x0008 CTGID = 3
/dev/rdsk/dks1d7vol -> [SQ] CL2-D Ser = 30053 LDEV = 14 [HITACHI ]
                       [OPEN-3-CM  ]
                       RAID5[Group  2- 1] SSID = 0x0008
```

IRIX system with fabric:

```
# ls /dev/rdsk/*/*vol/* | ./inqraid
/dev/rdsk/50060e8000100262/lun3vol/c8p0 -> [SQ] CL2-D Ser = 30053
                 LDEV = 9 [HITACHI] [OPEN-3 ] HORC = P-VOL
                 HOMRCF[MU#0 = SMPL MU#1 = SMPL MU#2 = SMPL]
                 RAID5[Group  2- 1] SSID = 0x0008 CTGID = 3
/dev/rdsk/50060e8000100262/lun4vol/c8p0 -> [SQ] CL2-D Ser = 30053
```

```
                    LDEV = 14 [HITACHI] [OPEN-3-CM]
                    RAID5[Group  2- 1] SSID = 0x0008
```

OpenVMS system:

```
$ inqraid dka145-146
DKA145 -> [ST] CL2-D Ser = 30053 LDEV = 9 [HITACHI ] [OPEN-3 ]
          HORC = P-VOL  HOMRCF[MU#0 = SMPL MU#1 = SMPL
          MU#2 = SMPL] RAID5[Group  2- 1] SSID = 0x0008 CTGID = 3
DKA146 -> [ST] CL2-D Ser = 30053 LDEV = 14 [HITACHI ] [OPEN-3-CM ]
          RAID5[Group  2- 1] SSID = 0x0008
```

Description of the `inqraid` command tool output for example 1:

CLX -Y:
> Displays the port number on the RAID storage system.

Ser
> Displays the production (serial#) number on the RAID storage system.
>
> - The serial number for VSP 5000 series is displayed with a "5" at the beginning ("512345" = serial number 12345).
>
> - The serial number for VSP G1x00 and VSP F1500 is displayed with a "3" at the beginning ("312345" = serial number 12345).

LDEV
> Displays the LDEV# in the RAID storage system.

HORC
> Displays the attribute ('P-VOL/S-VOL/SMPL') of a volume in the RAID storage system.

HOMRCF
> Displays the attribute ('P-VOL/S-VOL/SMPL') of a volume as MU#0-2 of ShadowImage or Copy-on-Write Snapshot in the RAID storage system.

Group
> Displays the relationship of the physical volume mapped to LDEV.

LDEV Mapping	Display Formats
RAID Group	RAID1[Group Group number - Sub number] RAID5[Group Group number - Sub number] RAID6[Group Group number - Sub number]
Copy-on-Write Snapshot S-VOL	SNAPS[PoolID poolID number]
Unmapped	UNMAP[Group 00000]
External LUN	E-LUN[Group External Group number]

Chapter 4: Command tools

LDEV Mapping	Display Formats
Dynamic Provisioning volume	A-LUN[PoolID poolID number]

SSID

Displays the storage subsystem ID (hexadecimal) of the LDEV in the RAID storage system.

> **Note:** SSID is a parameter used by enterprise storage systems. Although SSID is not used by VSP E series, VSP G130, G/F350, G/F370, G/F700, G/F900, VSP G200, G400, G600, G800 and VSP F400, F600, F800, or HUS VM, the set value is displayed.

CTGID

Displays the consistency group ID of TrueCopy Async and Universal Replicator when the LDEV has been defined as the P-VOL or S-VOL of the TrueCopy Async or Universal Replicator.

CHNO

Displays the channel number on the device adapter that recognizes on the Linux host. Displayed only for Linux systems.

TID

Displays target ID of the data drive that connects on the device adapter port. Displayed only for Linux systems.

LUN

Displays logical unit number of the data drive that connects on the device adapter port. Displayed only for Linux systems.

> **Note:**
>
> The display of group, SSID, and CTG ID depends on the storage system microcode level. The **CHNO**, **TID**, and **LUN** items are displayed only for Linux systems.

Example 2: `inqraid` with `-find` option (Linux shown)

```
ls /dev/sd* | inqraid -find
/dev/sdb -> No such on the group
Group  PairVol (L/R) (Port#,TID,LU),Seq#,LDEV#.P/S, Status,Fence,
Seq#,P-LDEV# M
oradb  oradev2 (L)   (CL2-N , 3, 2) 8071   22..SMPL ----- ------,
----- ---- -
->/dev/sdc
```

Example 3: `inqraid` with `-find` option (HP-UX shown)

```
# echo /dev/rdsk/c23t0d0 /dev/rdsk/c23t2d3 | ./inqraid -find
Group   PairVol(L/R) (Port#,TID,LU-M),Seq#,LDEV#.P/S,Status,
Seq#,P-LDEV# M
horc1   dev00(L)      (CL2-J, 0, 0-0) 61456  192..S-VOL SSUS,
-----   193  -
->/dev/rdsk/c23t0d0
Group   PairVol(L/R) (Port#,TID,LU-M),Seq#,LDEV#.P/S,Status,
Seq#,P-LDEV# M
horc1   dev10(L)      (CL2-J , 2,  3-0) 61456 209..S-VOL SSUS,
-----   206  -
->/dev/rdsk/c23t2d3
```

Example 4: `inqraid` with `-findc` option (HP-UX shown)

```
# echo /dev/rdsk/c23t0d0 /dev/rdsk/c23t2d3 | ./inqraid -findc
DEVICE_FILE M Group PairVol P/S   Stat R_DEVICE M P/S  Stat LK
c23t0d0     0 horc1   dev00 S-VOL SSUS  c23t0d1 0 P-VOL PSUS OK
/dev/rdsk/c23t0d0[1] -> No such on the group
/dev/rdsk/c23t0d0[2] -> No such on the group
DEVICE_FILE M Group  PairVol P/S  Stat  R_DEVICE M P/S   Stat LK
c23t2d3     0 horc1    dev10 S-VOL SSUS  c23t2d2 0 P-VOL PSUS OK
/dev/rdsk/c23t2d3[1] -> No such on the group
/dev/rdsk/c23t2d3[2] -> No such on the group

# echo /dev/rdsk/c23t0d0 /dev/rdsk/c23t2d3 | ./inqraid -findc -CLI
DEVICE_FILE  M Group PairVol P/S Stat  R_DEVICE  M P/S  Stat LK
c23t0d0      0 horc1   dev00 S-VOL SSUS  c23t0d1  0 P-VOL PSUS OK
c23t2d3      0 horc1   dev10 S-VOL SSUS  c23t2d2  0 P-VOL PSUS OK
```

Description of the `inqraid` tool output for example 2 to 4:

DEVICE_FILE

Device file name.

M:

MU# of local and remote.

Group

Group name (dev_group) defined in the configuration file.

PairVol

Paired vol. name (dev_name) within the group defined in the configuration file.

P/S

Volume attribute (P-VOL or S-VOL or simplex).

Stat

Status of the paired volume.

inqraid

R_DEVICE

> Device file name of remote site.

LK

> Check result of the paired volume connection path.

Example 5: `inqraid` with `-cli` option (Linux shown)

```
# ls /dev/sd* | ./inqraid -CLI
DEVICE_FILE PORT  SERIAL LDEV CTG  H/M/12 SSID R:Group PRODUCT_ID
sdh         CL2-B 30053  23   2    S/P/ss 0004 5:02-01 OPEN-3
sdi         CL1-A 64015  14   -    -      0004 E:00002 OPEN-3-CM
sdj         -     -      -    -    -      -    -       -
```

Description of the `inqraid` with `-cli` option (Linux)

DEVICE_FILE:

> Displays the device file name only.

PORT

> Displays the RAID storage system port number.

SERIAL

> Displays the production (serial#) number of the storage system.
>
> - The serial number for VSP 5000 series is displayed with a "5" at the beginning ("512345" = serial number 12345).
>
> - The serial number for VSP G1x00 and VSP F1500 is displayed with a "3" at the beginning ("312345" = serial number 12345).

LDEV

> Displays the LDEV# within the storage system.

CTG

> Displays the consistency group ID of TrueCopy Async and Universal Replicator when the LDEV has been defined as a TrueCopy Async or Universal Replicator P-VOL or S-VOL.

H/M/12

> Displays the attribute (P = P-VOL, S = S-VOL, s = SMPL) of a remote copy volume (TC/UR/GAD), local copy volume (SI/HTI/Copy-on-Write Snapshot), or local copy MU#1,2 volume.
>
> - H: Status of MU#0 for remote copy
> - M: Status of MU#0 for local copy
> - 1: Status of MU#1 for local copy
> - 2: Status of MU#2 for local copy

SSID

> Displays the storage subsystem ID of an LDEV in the storage system.

Chapter 4: Command tools

Command Control Interface Command Reference   205

> **Note:** SSID is a parameter used by enterprise storage systems. Although SSID is not used by VSP E series, VSP G130, G/F350, G/F370, G/F700, G/F900, VSP G200, G400, G600, G800 and VSP F400, F600, F800, or HUS VM, the set value is displayed.

R:Group

Displays the physical position of an LDEV according to mapping of LDEV in the storage system.

LDEV Mapping	R:	Group
RAID Group	RAID Level 1: RAID1 5: RAID5 6: RAID6	RAID Group number - Sub number
Copy-on-Write Snapshot S-VOL	S	Pool ID number
Unmapped	U	00000
External LUN	E	External group number
Dynamic Provisioning volume	A	Pool ID number

PRODUCT_ID

Displays product-id field in the STD inquiry page.

> **Note:**
>
> For a command device, PORT/SERIAL/LDEV/PRODUCT_ID is the SCSI Inquiry information for the external command device, if the command device is mapped as ELUN(**R:** =E).

Example 6: `inqraid` with `-cliwp` and `-cliwn` options (HP-UX shown)

```
# echo /dev/rdsk/c23t0d0 /dev/rdsk/c23t0d1 | ./inqraid -CLIWP
DEVICE_FILE PWWN             AL PORT  LUN  SERIAL LDEV PRODUCT_ID
c23t0d0     500060e802f01018  - CL2-J   -   61456  192     OPEN-3
c23t0d1     500060e802f01018  - CL2-J   -   61456  193     OPEN-3


# echo /dev/rdsk/c0t2d3 | ./inqraid -CLIWN
DEVICE_FILE  NWWN            AL PORT  LUN  SERIAL LDEV PRODUCT_ID
c0t2d3       5000E000E0005000  - CL1-A   -   30015  2054 OPEN3-CVS
```

Description of the `inqraid` output with `-cliwp` and `-cliwn` options (HP-UX shown)

DEVICE_FILE

Displays the device file name only.

WWN

CLIWP option displays Port_WWN of the host adapter included in the STD inquiry page. CLIWN option displays Node_WWN of host adapter included in STD inquiry page.

AL

Always displays '-'.

PORT

Displays the RAID storage system port number.

LUN

Always displays '-'.

SERIAL

Displays the production (serial#) number of the storage system.

- The serial number for VSP 5000 series is displayed with a "5" at the beginning ("512345" = serial number 12345).

- The serial number for VSP G1x00 and VSP F1500 is displayed with a "3" at the beginning ("312345" = serial number 12345).

LDEV

Displays the LDEV# within the storage system.

PRODUCT_ID

Displays product-id field in the STD inquiry page.

Example 7: `inqraid` with `-sort[cm]` option (HP-UX shown)

```
#ioscan -fun | grep rdsk | ./inqraid -sort -CM -CLI
HORCM_CMD
#dev_name               dev_name               dev_name
#UnitID 0 (Serial# 30012)
/dev/rdsk/c0t3d0        /dev/rdsk/c1t2d1
#UnitID 1 (Serial# 30013)
/dev/rdsk/c2t3d0
```

Example 8: `inqraid` with `-gvinf` option

```
D:\HORCM\etc>inqraid $Phys -gvinf -CLI
\\.\PhysicalDrive0:
# Harddisk0      -> [VOL61459_448_DA7C0D91] [OPEN-3    ]
\\.\PhysicalDrive1:
# Harddisk1      -> [VOL61459_449_DA7C0D92] [OPEN-3    ]
\\.\PhysicalDrive2:
```

```
# Harddisk2          -> [VOL61459_450_DA7C0D93] [OPEN-3    ]
                          -> S/N  LDEV  Signature
```

Example 9: `inqraid` with `-svinf[=PTN]` option

```
D:\HORCM\etc>pairdisplay -l -fd -g URA
Group PairVol(L/R) Device_File M ,Seq#,LDEV#.P/S,Status, Seq#,
   P-LDEV# M
URA   URA_000(L)   Harddisk3   0 61459  451..S-VOL SSUS,-----
   448     -
URA   URA_001(L)    Harddisk4  0 61459  452..S-VOL SSUS,-----
   449     -
URA   URA_002(L)    Harddisk5  0 61459  453..S-VOL SSUS,-----
   450     -


D:\HORCM\etc>pairdisplay -l -fd -g URA | inqraid -svinf=Harddisk
[VOL61459_451_5296A763] -> Harddisk3        [OPEN-3      ]
[VOL61459_452_5296A760] -> Harddisk4        [OPEN-3      ]
[VOL61459_453_5296A761] -> Harddisk5        [OPEN-3      ]
```

> ⚠️ **Caution:** If the S-VOL is created with '`Noread`' option (ShadowImage only) and the system is rebooted, the system cannot create a Device object (\Device \HarddiskVolume#) and Volume{guid} for S-VOL, but a Device object (\Device \HarddiskVolume#) and Volume{guid} is created by using -svinf option after splits the S-VOL.

mkconf

The **`mkconf`** command is a CCI command tool used to make a configuration file from a special file list (raw device file list) provided via STDIN.

Syntax

```
/HORCM/usr/bin/mkconf.sh    (UNIX systems)
\HORCM\Tool\mkconf.exe      (Windows and OpenVMS systems)
mkconf.sh  [ -g[g] <group> [-m <mu#>] [-i <inst#>] [-s <service>]
[-a] ][-n]
mkconf.exe [ -g[g] <group> [-m <mu#>] [-i <inst#>] [-s <service>]
[-a] [-c <drive#>] ]
```

Options and parameters

No argument

No option displays help/usage and version information.

-g <group>

Specifies the 'dev_group' name for a configuration file. If not specified, 'VG' is used as default.

-gg

Shows a LUN on the host view by finding a host group.

-m <mu#>

Specifies the mirror descriptor for ShadowImage and Copy-on-Write Snapshot volumes. TrueCopy volume does not specify the mirror descriptor.

-i <inst#>

Specifies the instance number for HORCM.

-s <service>

Specifies the service name (port name) for a configuration file. If not specified, '52323' is used as default.

-a

Specifies an addition of the group to a configuration file.

-c <drive#>

Windows systems only.

Specifies the range of drive numbers that should be examined in order to discover the command devices. If not specified, '$PhysicalDrive' is used as default.

-c <DKA#-#>

OpenVMS systems only.

Specifies the range of drive numbers that should be examined to discover the command devices. If not specified, '1DGA0-10000 DKA0-10000 DGA0-10000' is used as default.

[-n]

This `-n` option is specified to use a device file for NVMe-oF. Specifying this option enables the use of both device files for NVMe-oF and SCSI.

If the HORCC_NVME environment variable is specified, both device files for NVMe-oF and SCSI can be used without the `-n` option specified.

To specify this option, a command device must be enabled in a namespace on the storage system. For more information about whether a command device can be enabled in a namespace on the storage system, see the *Provisioning Guide* for your storage system.

Example 1: command example for special file list creation

The following are command examples that output the special file list provided via STDIN by using the **mkconf** command.

The special file list must contain the special file name and the hint path.

- Windows: `c:\HORCM\etc\inqraid $Phys -CLI`

- LINUX: `ls /dev/sd*`

- unix: `ls /<device-file-storage-directory-path>/*`

Example 2: `mkconf` command tool (HP-UX shown)

In this example, the configuration file is created as 'horcm*.conf' in the current directory. The log directory of HORCM is specified as 'log*' in the current directory. You must modify the 'ip_address & service' of an existing configuration file as needed.

```
# cd /tmp/test
# ls /dev/rdsk/* > /etc/devicelist.txt
# cat /etc/devicelist.txt | /HORCM/usr/bin/mkconf.sh -g ORA -i 9
-m 0
starting HORCM inst 9
HORCM inst 9 starts successfully.
HORCM Shutdown inst 9 !!!
A CONFIG file was successfully completed.
starting HORCM inst 9
HORCM inst 9 starts successfully.
DEVICE_FILE        Group PairVol  PORT   TARG  LUN M  SERIAL  LDEV
/dev/rdsk/c23t0d0  ORA   ORA_000  CL2-J   0     0 0   61456   192
/dev/rdsk/c23t0d1  ORA   ORA_001  CL2-J   0     1 0   61456   193
/dev/rdsk/c23t0d2  ORA   ORA_002  CL2-J   0     2 0   61456   194
/dev/rdsk/c23t0d3  ORA   ORA_003  CL2-J   0     3 0   61456   195
/dev/rdsk/c23t0d4  ORA   ORA_004  CL2-J   0     4 0   61456   256
/dev/rdsk/c23t0d5  ORA   ORA_005  CL2-J   0     5 0   61456   257
/dev/rdsk/c23t0d6  ORA   ORA_006  CL2-J   0     6 0   61456   258
/dev/rdsk/c23t0d7  -     -        -       -     - 0   61456   259
HORCM Shutdown inst 9 !!!
Please check '/tmp/test/horcm9.conf','/tmp/test/log9/curlog/
horcm_*.log', and modify 'ip_address & service'.

# ls                     <=Verify configuration and log files.
horcm9.conf  log9
# vi *.conf      <=Verify config file, check ip address & service.

# Created by mkconf.sh on Mon Jan 22 17:59:11 JST 2001

HORCM_MON
#ip_address       service         poll(10ms)      timeout(10ms)
127.0.0.1         52323             1000               3000

HORCM_CMD
#dev_name              dev_name                    dev_name
#UnitID 0 (Serial# 61456)
/dev/rdsk/c23t3d0
```

Example 3: `mkconf` command tool (HP-UX shown)

```
HORCM_DEV
#dev_group       dev_name     port#      TargetID     LU#      MU#
# /dev/rdsk/c23t0d0 SER = 61456 LDEV = 192 [FIBRE FCTBL = 4] ORA
ORA_000  CL2-J  0         0         0
# /dev/rdsk/c23t0d1 SER = 61456 LDEV = 193 [FIBRE FCTBL = 4] ORA
ORA_001  CL2-J  0         1         0
# /dev/rdsk/c23t0d2 SER = 61456 LDEV = 194 [FIBRE FCTBL = 4] ORA
ORA_002  CL2-J  0         2         0
# /dev/rdsk/c23t0d3 SER = 61456 LDEV = 195 [FIBRE FCTBL = 4] ORA
ORA_003  CL2-J  0         3         0
# /dev/rdsk/c23t0d4 SER = 61456 LDEV = 256 [FIBRE FCTBL = 4] ORA
ORA_004  CL2-J  0         4         0
# /dev/rdsk/c23t0d5 SER = 61456 LDEV = 257 [FIBRE FCTBL = 4] ORA
ORA_005  CL2-J  0         5         0
# /dev/rdsk/c23t0d6 SER = 61456 LDEV = 258 [FIBRE FCTBL = 4] ORA
ORA_006  CL2-J  0         6         0
# ERROR [CMDDEV] /dev/rdsk/c23t0d7 SER = 61456  LDEV = 259
[OPEN-3-CM]                                          <=See Notes.

HORCM_INST
#dev_group    ip_address    service
ORA           127.0.0.1     52323    <=Check and update as needed.
```

> 📄 **Note:**
>
> - A unitID is added to the Serial# order. If two or more command devices exist in the storage system, select the device file that is shared among the storage system ports on a priority basis, and treat as an alternate command device.
>
> - If the STDIN device includes the command device, the target device is commented out as shown below:
>
>   ```
>   # ERROR [CMDDEV] /dev/rdsk/c23t0d7   SER =  61456  LDEV = 259 [ OPEN-3-
>   CM    ]
>   ```
>
> - If the STDIN device is shared among multiple device files and already displayed as a target device, the target device is commented out as shown below:
>
>   ```
>   # ERROR [LDEV LINK] /dev/rdsk/c24t0d3  SER =  61456  LDEV = 195
>   [FIBREFCTBL = 4]
>   ```
>
> - If the STDIN device does not have appropriate mirror description (MU#), the target device is commented out as shown below:
>
>   ```
>   # ERROR [INVALID MUN (2 < 1)] /dev/rdsk/c24t0d3  SER = 61456 LDEV = 195
>   [ OPEN-3 ]
>   ```
>
> - If the STDIN device is mixed among the storage systems of differential mirror control, the target device is commented out as shown below:
>
>   ```
>   # ERROR [MIXING RAID TYPE] /dev/rdsk/c24t0d3   SER = 61456  LDEV = 195
>   [ OPEN-3]
>   ```

rmawk

The `rmawk` command is a scriptable command for associating with the pair operation commands and raidcom commands.

This command provides basically the following three functions:

- Outputs by filtering STDIN with the specified conditions.

```
Command | rmawk @1-eq:PAIR -a @2-ne:0
```

- Performs the command with the variable parameter specified by interpreting STDIN with the specified conditions.

```
Command | rmawk @3-eq:TAR exe="Command line @1"
```

- Tests by interpreting the output command with the specified conditions or waits until the conditions become TRUE, by performing the specified command.

```
rmawk exe="Command line" @2-eq:PAIR -o @2-eq:PSUS
rmawk exe="Command line" @2-eq:COPY timeout=300
```

Syntax

```
/HORCM/usr/bin/rmawk  [-h | exe="…" | sys="…" | timeout=value | interval=value | -BL
| -AT |  -EC[VAL] | @variable | operators ]
```

```
\HORCM\etc\rmawk  [-h | exe="…" | sys="…" | timeout=value | interval=value | -BL | -
AT | -EC[VAL] | @variable | operators ]
```

Options and parameters

No argument or -h

No option displays help/usage and version information.

exe="command_line"

Specifies the command line to be performed when the result of the specified formula and the conditions are TRUE. The testing/waiting specifies the command to be its target.

"@variable" with the field variables of STDIN can be included in the command line.

Example:

```
exe="raidcfg -a qry -o hgrp -pport @1"
exe="raidcfg -a reg -o hgrp -pport @1-@2 -pname @3"
exe="type map.txt | rmawk @@L-eq:@L exe=\"raidcfg -a map -o
snap -pname @1 -pldev @5 @@1\""
```

exe="print …. "

Specifies printing (with Line Feed) when the result of the specified formula and the conditions are TRUE.

"@variable" with the field variables of STDIN can be included in the command line. When "@variable" is specified as "@C#/C#*", the variable is interpreted as character string. When "@variable" is specified as "@C#?/C#+/L/C/R", the variable is interpreted as integer.

The printing range is from "print" to the end including spaces.

Example:

```
exe="print PORT=@1 LDEV=@3"
exe="print PORT=@1 LDEV=@3 Total CAP = @4+"
```

exe="printn…. " or exe="prints…. "

Specifies printing (without Line Feed) when the result of the specified formula and the conditions are TRUE.

"@variable" with the field variables of STDIN can be included in the command line. When "@variable" is specified as "@C#/C#*", the variable is interpreted as character string. When "@variable" is specified as "@C#?/C#+/L/C/R", the variable is interpreted as integer.

The printing range is from "print" to the end including spaces.

Example:

```
exe="printn PORT=@1 LDEV=@3"
```

> 📄 **Note:** "prints..." is used to print to the strings buffer of '@0*'.

exe=exit

Specifies to exit this command when the result of the specified formula and the conditions are TRUE. The return value of exit is "@R" variable.

sys="command_line"

Specifies the command line to be performed in the transparent mode when the result of the specified formula and the conditions are TRUE.

"@variable" and "@expression" in the command line are not interpreted as variables but are filtered out as variables to path them to the specified command.

Example for UNIX:

```
sys='rmawk exe="pairdisplay -g G1 -CLI -l" @6-eq:PAIR'
```

Example for Windows:

```
sys="rmawk exe=\"pairdisplay -g G1 -CLI -l\" @6-eq:PAIR"
```

Example for openVMS:

```
sys="rmawk exe=""pairdisplay -g G1 -CLI -l"" @6-eq:PAIR"
```

Exception: When "exe="..."" is specified in the command line and there is "@variable" as an argument in this command, it is interpreted as "@variable" that is including the field variables of STDIN.

Example for UNIX:

```
sys='rmawk exe="pairdisplay -d @3 @5 -CLI -l" @6-eq:PAIR'
```

Example for Windows:

```
sys="rmawk exe=\"pairdisplay -d @3 @5 -CLI -l\" @6-eq:PAIR"
```

Example for OpenVMS:

```
sys="rmawk exe=""pairdisplay -d @3 @5 -CLI -l"" @6-eq:PAIR"
```

timeout=value

Specifies the timeout time in seconds for waiting by interpreting this output command with the specified conditions when the specified command is performed.

interval=value

Specifies the interval time in seconds for waiting by interpreting this output command with the specified conditions when the specified command is performed.

-BL

Specifies to include blank lines from STDIN. This is used to detect the blank lines when there are in the STDIN.

-AT

- When "-AT" is specified (the default setting does not specify this option)

  Specifies to testing or waiting until the conditions become TRUE in all the command lines, by interpreting the output of the specified command with the specified conditions.

  Testing ends as an ERROR (1) if there is even one FALSE in the specified condition. And waiting waits until all the lines become TRUE by retrying the specified command until becoming timeout if there is even one FALSE in the lines of the specified conditions. If it becomes timeout, the command ends as an ERROR (1).

- When "-AT" is not specified

  It tests or waits until at least one line becomes TRUE, by interpreting the output of the specified command with the specified conditions.

  If there is one TRUE at least in the specified conditions, the testing ends as NORMAL (0). If there is one TRUE at least in the specified conditions, the waiting ends as NORMAL (0). If all the lines are FALSE, it waits until at least one line become TRUE by retrying the specified command until becoming timeout. If it becomes timeout, the command ends as an ERROR (1).

-EC[VAL]

Performs only once with ending of lines or "exe=exit" as Line#0. If you need to perform something special at end of lines, specify this option that is called as Line#0. This is enabled in the run mode.

If VAL(integer) is specified, @C#? variable for all are initialized to the specified VAL before starting. If VAL is invalid such as -ECM, then @C#? variable for all are initialized to the MAX value with 64-bit integer.

@variable

Specifies the following variables. The variables can be specified in both the left and right side of the formula. The variables can also be specified as arguments of the command to be performed or of printing.

- @C#

  Field strings corresponding to the column number (Column#) from STDIN. If "C#" is specified in hexadecimal ("0x..."), then its field is treated as hexadecimal strings.

- @C#*

  Strings from the field to the end corresponding to the column number from STDIN.

- @C#?

  64 bit variables corresponding to the column number from STDIN (initial value 0).

- @C#+

  64 bit variables adding field variables corresponding to the column number from STDIN (initial value 0).

- @0

  A variable that indicates one line (This is used when printing or searching within the line.).

- @L

  A variable that indicates the line number from STDIN.

- @C

  A variable that indicates the number of column in each line of STDIN.

- @R

  A return value of performing command.

- @0*

  A string variable that can be memorized one line.

- null

  A special variable that tests if the character string is null or not.

- @C#?t

  A variable that can be specified within exe="…" to print the C#? variable as TOD (time of day). If the C#? variable is zero, the current time is printed.

- @C#?x

  A variable that can be specified within exe="…" to print the C#? variable as hexadecimal strings by adding 0x.

- @@

  A variable that can be specified within exe="…" to print the @ character.

operators (-operator:)

Specifies the following comparison operation. Variables in the field are operated in 64-bit integers.

The numerical values in the field strings are "1234..." or "0x1234..." format without octal code.

- -eq:

  Treats the target as character strings or integers by comparing with "==". When "@C#/C#" is specified in the left side of an equation, this is treated as character strings. When "@C#?/C#+/L/C/R" is specified in it, this is treated as integers.

- -ne:

  Treats the target as character strings or integers by comparing with "!=". When "@C#/C#" is specified in the left side of an equation, this is treated as character strings. When "@C#?/C#+/L/C/R" is specified in it, this is treated as integers.

- -gt:

  Treats the target as integers by comparing with ">".

- -ge:

  Treats the target as integers by comparing with ">=".

- -lt:

  Treats the target as integers by comparing with "<".

- -le:

  Treats the target as integers by comparing with "<=".

operators (=operator:)

Specifies the following arithmetical operation. Variables in the field are operated in 64-bit integers.

The numerical values in the field strings are "1234..." or "0x1234..." format without octal code.

- =ad:

  Treats the target as integers by operating with "+ (Addition)".

- =sb:

  Treats the target as integers by operating with "- (Subtraction)".

- =ml:

  Treats the target as integers by operating with "* (Multiplication)".

- =dv:

    Treats the target as integers by operating with "/ (Division)".

- =st:

    Treats the target as integers by substituting with "= (Setting)".

operators (-operator)

Specifies the following logical operation for the result of the comparison operation.

- -o

    Judges the result of the comparison operation as the "Logical OR".

- -a

    Judges the result of the comparison operation as the "Logical AND".

- -n

    Inverts the result of the comparison operation. (TRUE becomes FALSE, FALSE becomes TRUE.)

Returned values

The **rmawk** command sets the following returned values during exit allowing you to check the execution results.

- **Normal termination:**

  - **0:** The command ends normally with the specified condition.

  - **1:** Testing: The specified condition is FALSE ; Waiting: The specified condition is Timeout.

- **Abnormal termination:**

  - **125:** The command ends with a syntax error.

  - **126:** The command ends with a system error.

Examples of comparison expression

- @20-eq:PAIR: Compares if the character string in Column #20 from STDIN matches "PAIR".

- @20-eq:PSU*: Compares if "PSU" is included in the character string in Column #20 from STDIN.

- @0-eq:PSU*: Compares if "PSU" is included in one line from STDIN. This is equivalent to "grep PSU".

- @20-eq:@21: Compares if the character strings in Column #20 and Column #21 match.

- @20-ge:50: Compares if it is "value >= 50" of Column #20 from STDIN.

- @L-ge:20: Compares if it is the current "number of lines >= 20" from STDIN.

- @C-ge:20: Compares if it is "number of columns >= 20" of the current lines from the STDIN.

- @R-gt:0: Compares if it is "return value > 0" of the specified command.

Examples of arithmetic expression

- @8?=ad:@8: Performs "@8? = @8? + @8". This is equivalent to "@8+".

- @8=ad:@5: Performs "@8? = @8 + @5".

- @8=ad:@5?: Performs "@8? = @8 + @5?".

- @8=ad:30: Performs "@8? = @8 + 30?".

- @8=st:30: Performs "@8? = 30".

- @5=st:@5: Performs "@5? = @5" for converting to the integer.

- :

- :

- :

Example:

```
Command | rmawk @8?=ad:@8 exe="print Total = @8?"
Command | rmawk exe="print Total = @8+"
```

Example of the command option format

Performs the operation by interpreting the specified option format and classifying them in the following three function formats.

- Performs the operation as a filter if there is no "exe="...""" in the specified option.

```
Command | rmawk [@expression1] Logical operator [@expression2] …
```

Example:

```
Command | rmawk @1-eq:PAIR -a @2-ne:0
```

- Performs the operation as a testing if there is one "exe="..."" but no "@variable" in the specified option.

```
rmawk exe="Command line" [@expression1] Logical operator [@expression2] …
```

Example:

```
rmawk exe="Command line" @2-eq:PAIR -o @2-eq:PSUS
```

And performs as an waiting command if "timeout=value" is specified as an option.

```
rmawk exe="Command line" [@expression1] Logical operator [@expression2] timeout=6
```

Example:

```
rmawk exe="Command line" @2-eq:COPY timeout=300
```

- Other than the above, it performs the multiple "exe="Command line"" commands that include contexts written in multiple formulas and the variable parameters by interpreting the STDIN.

```
Command | rmawk [@expression1] Logical operator [@expression2] exe="Command line"
```

Example:

```
Command | rmawk @3-eq:TAR exe="Command line @1"
Command | rmawk @3-eq:TAR exe="Command line(true) @1" -n exe="Command line(false)
@1"
```

Because the command runs with the result of TRUE, multiple command can be performed if you keep writing lines ("exe="command line2"", "exe="command line3"", and so on).

```
Command | rmawk [@expression1] Logical operator [@expression2] exe="Command line1"
exe="Command line2" exe="Command line3"
```

Example:

```
Command | rmawk @3-eq:TAR exe="Command line @1" exe="Command line2 @1"
Command | rmawk @3-eq:TAR exe="Command line(true) @1" exe="Command line2(true) @1"
-n exe="Command line(false) @1"
```

Chapter 5: Configuration setting commands

The CCI `raidcom` command enables you to perform configuration and provisioning operations on the Hitachi RAID storage systems.

📄 **Note:** If a failure occurs in a processor on an MP blade or controller, some commands cannot be executed on any MP blades or controllers (even if no failure occurs on those MP blades or controllers). For details, see Command execution with MP blade or controller failure (on page 228).

raidcom

Specifies a configuration change.

Syntax

Displaying help

```
raidcom { [-h]
```

Log-in and log-out

```
raidcom {-login [<user_name> <password>] | -logout] [-s <seq#>
      | -u <unit#>] [-I[H][M] <instance#>] [-I[TC][SI] <instance#>]
```

Configuration change: line-by-line mode

```
raidcom <action> <object> [<param> <value>…] [-fx] [-checkmode <check mode>] [-store
<filename>] [-nomsg] [-login [<user_name> <password>] [-resource <resource_grp_id>…]
| -logout] [-s <seq#> | -u <unit#>] [-I[H][M] <instance#>] [-I[TC][SI] <instance#>]
```

Configuration change: transaction mode

```
raidcom -zt <filename> [-load <work_filename>]
      [-checkmode <check mode>]
```

Options and parameters

[-h]
    Displays help of raidcom.

[-login [<user_name> <password>]]

Specifies user authentication for the storage system. Specifies user name and password.

If no user authentication is done yet and the -login option is omitted, the input of the user name and the password is required. And, in case of omitting the user name and the password by specifying -login option, the input of the user name and the password is required.

The maximum number of the user who can login at same time is 512.

[-logout]

Deletes the cache of the session control table in the storage system and logs out from the storage system (command device). Also deletes all the authentication files corresponding to the storage system.

At the time of next login, the CCI command requires the user name and the password. If an application that uses the command device exists in the host, the application also requires the user name and the password. If the same user is set for multiple hosts, the session control table of each host is managed so that the logout is applied only to the relevant host.

[-s <seq#>]

Specifies the serial number.

- When specifying the serial number for VSP 5000 series, add a "5" at the beginning of the serial number. For example, for serial number 12345, enter 512345.

- When specifying the serial number for VSP G1x00 and VSP F1500, add a "3" at the beginning of the serial number. For example, for serial number 12345, enter 312345.

[-u <Unit#>]

Specifies the unit ID of command device as HORCM_CMD of configuration file.

[-I[H][M] <instance#>] or [-I[TC][SI] <instance#>]

Specifies CCI instance number by specifying the command as HORC/HOMRCF or TrueCopy/ShadowImage. If you specify only -I, the command refers to the setting of HORCC_MRCF environment variable.

<action> <object> [<param> <value>…]

Specifies each command name that is expressed in `raidcom add copy_grp` and later section and the parameters and values that are required at each operation.

[-fx]

Displays the LDEV number in hexadecimal notation.

[-checkmode <check mode>]

Specifies when executing Precheck function (execute checking commands only). It specifies the following value.

`Precheck`: When this option is specified, the actual process is not executed for the storage system even if the command is executed.

This option is available to specify the setting of $HORCC_NO_EXEC environment variable and $HORCC_CTX_CHK environment variable.

[-store <file name>]

Specifies the file name of the configuration file to be created for implementation check.

[-nomsg]

Prevents displaying messages. This option is required to be defined on the beginning of the command parameter.

[-zt <file name>]

Specifies the script file.

[-load <file name>]

Specifies the file name of the file (Configuration file) that is created for implementation check.

[-resource <resource_grp_id>…]

Limits the operations coverage to the specified resource groups. When this option is specified in reference commands, objects that are in the specified resource groups are displayed. When the resource group specified by this option does not contain the resource specified for executing the command, the command is not executed because of the EX_EGPERM error.

Returned values

Unless otherwise stated, the `raidcom` commands set the following returned values:

- **0:** Normal termination.

- **other than 0:** Abnormal termination

For details, see the section describing command error messages in *Command Control Interface User and Reference Guide*.

Examples

Performing user authentication (login) by the user ID: USER01 and the password: PASS01.

```
# raidcom -login USER01 PASS01
```

Performing log-out.

```
# raidcom -logout
```

Performing user authentication (login) to instance 99 by the user ID: USER01 and the password: PASS01.

```
# raidcom -login USER01 PASS01 -I99
```

Performing syntax check and the context check of the script file (the actual processing is not executed).

```
# raidcom -zt <script file> -checkmode precheck
```

Performing syntax check, the context check, and the implementation check of the script file (the actual processing is not executed).

```
# raidcom -zt <script file> -load <work file> -checkmode precheck
```

> **Note:** If the user who does not have the write permission for the HORCM directory and the directories under the directory performs a raidcom command, the following WARNING might be output. If the WARNING is output, confirm whether you have the write permission for the directory, and then retry the command. If only the WARNING occurs, the returned value will be 0 (normal termination).
>
> ```
> raidcom: [WARNING] Couldn't create User file for Serial#(<XXXXX>).
> ```
>
> Serial#: Displays the product serial number.

Method for specifying LDEV number

The methods for specifying LDEV numbers are shown in the following table.

Specification method	Example
Decimal number	-ldev_id 300
Hexadecimal number	-ldev_id 0x12C
	-ldev_id 01:2C

Methods for specifying multiple LDEVs

Methods for specifying multiple LDEVs by LDEV IDs

The methods for specifying multiple LDEV IDs in a single command are shown below. Some methods do not apply to some commands.

```
-ldev_id 300-305
-ldev_id 0x12C-0x131
-ldev_id 01:2C-01:31
-ldev_id 300 -cnt 6
-ldev_id 300 301 302 303 304 305
```

Method for specifying LDEVs by device group

When you configure multiple LDEVs for pools or journals (for example, changing the resource group), you must issue the command to each LDEV. However, if you define LDEVs that configure pools or journals, you can issue the command to all LDEVs defined as a device group by specifying each device group at a time.

If you issue the command to LDEVs by specifying a device group, the command is executed in each LDEV registered in the device group. If an error occurs in an LDEV while the command is being executed, the execution of command stops at the LDEV where the error occurred. The command is not executed in the remaining LDEVs. If an error occurs, solve the error. Then restore the LDEV and issue the command to all remaining LDEVs.

- Creating a pool by specifying a device group:

```
# raidcom add device_grp -device_grp_name dg_pool1 data1
-ldev_id 512 513 514 515
```

```
# raidcom add dp_pool -pool_id 2 -grp_opt ldev -device_grp_name dg_pool1
```

- Creating a journal by specifying a device group:

```
# raidcom add device_grp -device_grp_name dg_jnl1 data1
-ldev_id 512 513 514 515
```

```
# raidcom add journal -journal_id 2 -grp_opt ldev
-device_grp_name dg_jnl1
```

Operations where multiple LDEVs can be specified

You can use "-ldev_id <ldev#>" to specify multiple LDEVs at the same time in the following operations only:

(a) Displaying LDEV information

```
# raidcom get ldev -ldev_id 100-103
```

```
# raidcom get ldev -ldev_id 100 -cnt 4
```

> 📄 **Note:**
>
> Specifying multiple LDEV as follows cannot be performed: `#raidcom get ldev -ldev_id 100 101 103`

(b) Creating a journal

```
# raidcom add journal -journal_id 1 -ldev_id 265 266
```

```
# raidcom add journal -journal_id 1 -ldev_id 265-266
```

```
# raidcom add journal -journal_id 1 -ldev_id 265 -cnt 2
```

(c) Creating a pool

Creating a pool for Copy-on-Write Snapshot:

```
# raidcom add snap_pool -pool 1 -ldev_id 365 366 367
```

```
# raidcom add snap_pool -pool 1 -ldev_id 365-367
```

```
# raidcom add snap_pool -pool 1 -ldev_id 365 -cnt 3
```

Creating a pool for Dynamic Provisioning or Dynamic Provisioning for Mainframe:

```
# raidcom add dp_pool -pool 1 -ldev_id 465 466 467
```

```
# raidcom add dp_pool -pool 1 -ldev_id 465-470
```

```
# raidcom add dp_pool -pool 1 -ldev_id 465 -cnt 5
```

(d) Creating a device group

```
# raidcom add device_grp -device_grp_name DevG2 dev101 -ldev_id 101 105 201
```

```
# raidcom add device_grp -device_grp_name DevG2 dev101 -ldev_id 101-105
```

```
# raidcom add device_grp -device_grp_name DevG2 dev101 -ldev_id 101 -cnt 5
```

(e) Deleting a device group

```
# raidcom delete device_grp -device_grp_name DevG3 -ldev_id 200 201 204
```

```
# raidcom delete device_grp -device_grp_name DevG3 -ldev_id 200-204
```

```
# raidcom delete device_grp -device_grp_name DevG3 -ldev_id 200 -cnt 5
```

Specifying and displaying storage system serial numbers

Specifying or displaying the serial number of the storage system varies depending on storage systems.

- When specifying the serial number for VSP 5000 series, add a "5" at the beginning of the serial number. For example, for serial number 12345, enter 512345.

- When specifying the serial number for VSP G1x00 and VSP F1500, add a "3" at the beginning of the serial number. For example, for serial number 12345, enter 312345.

- The serial number for VSP 5000 series is displayed with a "5" at the beginning ("512345" = serial number 12345).

- The serial number for VSP G1x00 and VSP F1500 is displayed with a "3" at the beginning ("312345" = serial number 12345).

> **Note:** When you specify the serial number for the storage system in the GUI (for example, HDvM - SN, Hitachi Ops Center Administrator, or Hitachi Command Suite), use only the five-digit serial number. Do not add "3" or "5" in front of the serial number.

Maximum number of acceptable asynchronous commands

The following list shows the maximum numbers of asynchronous commands that the storage system can accept. When you continuously run a large number of asynchronous commands, make sure not to exceed these maximum numbers. For details, see the topic regarding asynchronous commands in the *Command Control Interface User and Reference Guide*. If the following number is exceeded, split the commands so that the asynchronous commands exceeding the maximum number of acceptable asynchronous commands are not run on the storage system.

- VSP 5000 series: 2,560

- VSP G1x00 and VSP F1500: 1,280

- VSP E series: 2,560

- VSP G130, G/F350, G/F370, G/F700, G/F900: 2,560

- VSP G200, G400, G600, G800 and VSP F400, F600, F800: 2,560

Command execution with MP blade or controller failure

If a failure occurs in a processor on an MP blade or controller, some commands cannot be executed on any MP blades or controllers, even if no failure occurs on those MP blades or controllers. The following table lists the commands.

Command group	Commands
raidcom <action> clpr	- add - delete - modify
raidcom <action> user_system_opt	modify
raidcom <action> copy_grp	- add - delete
raidcom <action> device_grp	- add - delete
raidcom <action> drive	modify
raidcom <action> external_grp	- add - check_ext_storage - delete - disconnect - modify
raidcom <action> external_storage	discover
raidcom <action> external_iscsi_name	- add - delete
raidcom <action> external_chap_user	modify
raidcom <action> initiator_chap_user	modify
raidcom <action> journal	- add - delete - modify
raidcom <action> ldev	- add - delete - extend[1]

Command group	Commands
	▪ initialize ▪ modify
raidcom <action> license	▪ add ▪ delete ▪ modify
raidcom <action> quorum	▪ add ▪ delete ▪ modify ▪ replace
raidcom <action> local_replica_opt	▪ modify
raidcom <action> remote_replica_opt	▪ modify
raidcom <action> lun	▪ add[2]
raidcom <action> path	▪ add ▪ check_ext_storage ▪ delete ▪ disconnect ▪ modify
raidcom <action> pool	▪ delete ▪ initialize ▪ modify ▪ monitor ▪ reallocate ▪ rename
raidcom <action> system	▪ modify
raidcom <action> port	▪ modify[3]
raidcom <action> parity_grp	▪ add ▪ delete ▪ initialize ▪ modify

Chapter 5: Configuration setting commands

Command group	Commands
raidcom <action> rcu	- add - delete - modify
raidcom <action> rcu_iscsi_port	- add - delete
raidcom <action> rcu_path	- add - delete
raidcom <action> ssid	- add - delete
raidcom <action> resource	- add - delete - map - modify - unmap
raidcom <action> snap_pool	add
raidcom <action> dp_pool	add
raidcom <action> server	- add - delete - modify
raidcom <action> qos_grp	- add - delete - modify
raidcom <action> nvm_subsystem	- add - delete - modify
raidcom <action> nvm_subsystem_port	- add - delete
raidcom <action> host_nqn	- add - delete - modify

Chapter 5: Configuration setting commands

Command group	Commands
raidcom <action> namespace_path	▪ add ▪ delete
raidcom <action> namespace	▪ add ▪ delete ▪ modify
1. For the **raidcom extend ldev** command, this restriction applies only when you run the command with the following options specified: `raidcom extend ldev -request_id auto` 2. For the **raidcom add lun** command, this restriction applies only when you run the command with the following options specified: `raidcom add lun -request_id auto` 3. For the **raidcom modify port** command, this restriction applies only when you run the command with the following options specified: `raidcom modify port {-t10pi {enable|disable}	-iscsi_virtual_port_mode {enable|disable}	-port_mode {scsi	nvme} -request_id auto	-delete_login_host_nqn}`	

Resource group operation

When you have multiple resource group authorities, use the `-resource` option to see the resource group information.

Examples

```
# raidcom get resource
```

```
RS_GROUP        RGID    stat      Lock_owner   Lock_host    Serial#
meta_resource     0    Unlocked  -            -              64568
RSG_CLI1          1    Unlocked  -            -              64568
RSG_CLI2          2    Unlocked  -            -              64568
```

```
# raidcom get port
```

```
PORT   TYPE   ATTR SPD LPID FAB CONN SSW SL Serial# WWN       PHY_PORT
CL1-A FIBRE ELUN AUT   EF N   FCAL N   0  64568 50060e8006fc3800 -
CL1-B FIBRE TAR  AUT   EF N   FCAL N   0  64568 50060e8006fc3801 -
```

```
CL1-C FICON TAR    -  -  -   -    -   0  64568 50060e8006fc3802 -
CL1-D FICON TAR    -  -  -   -    -   0  64568 50060e8006fc3803 -
```

```
# raidcom get port -resource 1
```

```
PORT  TYPE  ATTR SPD LPID FAB CONN SSW SL Serial# WWN      PHY_PORT
CL1-B FIBRE TAR  AUT   EF N   FCAL N   0  64568 50060e8006fc3801 -
CL1-C FICON TAR    -   -  -   -    -   0  64568 50060e8006fc3802 -
```

```
# raidcom get port -resource 2
```

```
PORT  TYPE  ATTR SPD LPID FAB CONN SSW SL Serial# WWN      PHY_PORT
CL1-A FIBRE ELUN AUT   EF N   FCAL N   0  64568 50060e8006fc3800 -
CL1-D FICON TAR    -   -  -   -    -   0  64568 50060e8006fc3803 -
```

Resource lock operation

When multiple users perform operations for a single resource, lock the resource group to which the target resource is allocated before performing operations. This applies to both the setting and viewing operations.

The following examples show locking resource groups, executing commands, and then unlocking resource groups.

Examples

- Performing a user authentication by User ID: USER01, Password: PASS01.

```
# raidcom -login USER01 PASS01
```

- Locking a resource group: rsg001.

```
# raidcom lock resource -resource_name rsg001
```

- Creating LDEV#100 and #101.

```
# raidcom add ldev -parity_grp_id 5-2 -ldev_id 100 -capacity 10g
```

```
# raidcom add ldev -parity_grp_id 5-2 -ldev_id 101 -capacity 10g
```

- Unlocking resource group: rsg001

```
# raidcom unlock resource -resource_name rsg001
```

Chapter 5: Configuration setting commands

Resource locking and CCI commands

If you execute a CCI command when the specified resource is locked, the specified resource groups cannot be used by other users. Commands can be executed when the specified resources are not locked. However, if another user locks the resource, CCI commands will result in error. The following tables show the relationship between CCI commands and resources that need to be locked, except for the required options for the commands. In the following tables:

- Res. group: resource group

- Lock/auth: Resource locking and resource authority are required.

- Auth: Only resource authority check is specified.

- (VSP only) Lock by user: If another user locks the applicable resource, the command execution fails. To execute the command, the applicable resource authority is not required.

Relationship between commands and resources that need to be locked

Res. group	Port	Host group	Parity group	External volume/ VDEV	LDEV	Command	Option
-	-	-	-	-	Lock/auth	raidcom add copy_grp	-
-	-	-	-	-	Lock/auth	raidcom delete copy_grp	-
-	-	-	-	-	Auth	raidcom get copy_grp	-
-	-	-	-	-	Lock/auth	raidcom add device_grp	-
-	-	-	-	-	Lock/auth	raidcom delete device_grp	-
-	-	-	-	-	Auth	raidcom get device_grp	-
-	Lock/auth	-	-	-	Lock/auth	raidcom add external_grp	-
-	-	-	-	Lock/auth	-	raidcom check_ext_storage external_grp	-
-	-	-	-	Lock/auth	Auth	raidcom check_ext_storage external_grp	-ldev_id <ldev#>
-	-	-	-	Lock/auth	-	raidcom delete external_grp	-

Res. group	Port	Host group	Parity group	External volume/ VDEV	LDEV	Command	Option
-	-	-	-	Lock/auth	-	raidcom disconnect external_grp	-
-	-	-	-	Lock/auth	Auth	raidcom disconnect external_grp	-ldev_id <ldev#>
-	-	-	-	Auth	-	raidcom get external_grp	-
-	-	-	-	Lock/auth	-	raidcom modify external_grp	-
-	Auth	-	-	-	-	raidcom discover external_storage	-
-	-	Lock/ auth	-	-	-	raidcom add host_grp	-
-	-	Lock/ auth	-	-	-	raidcom delete host_grp	-
-	Auth	Auth	-	-	-	raidcom get host_grp	-
-	-	Lock/ auth	-	-	-	raidcom modify host_grp	-
-	-	Lock/ auth	-	-	-	raidcom add hba_wwn	-
-	-	Lock/ auth	-	-	-	raidcom delete hba_wwn	-
-	Auth	Auth	-	-	-	raidcom get hba_wwn	-
-	-	-	-	-	Lock/auth	raidcom add journal	-
-	-	-	-	-	Lock/auth	raidcom delete journal	-
-	-	-	-	-	Auth	raidcom get journal	-
-	-	-	-	-	Lock/auth	raidcom modify journal	-
-	-	-	Lock/ auth	-	Lock/auth	raidcom add ldev	-parity_ grp_id <gno-sgno>
-	-	-	-	Lock/auth	Lock/auth	raidcom add ldev	-external_ grp_id <gno-sgno>
-	-	-	-	-	Lock/auth	raidcom add ldev	-

Res. group	Port	Host group	Parity group	External volume/ VDEV	LDEV	Command	Option
-	-	-	-	-	Lock/auth	raidcom delete ldev	-
-	-	-	-	-	Lock/auth	raidcom extend ldev	-
-	-	-	-	-	Auth	raidcom get ldev	-
-	-	-	-	-	Lock/auth	raidcom initialize ldev	-
-	-	-	-	-	Lock/auth	raidcom modify ldev	-
-	-	Lock/auth	-	-	Lock/auth	raidcom add lun	-
-	-	Lock/auth	-	-	Lock/auth	raidcom delete lun	-
-	Auth	-	-	-	-	raidcom discover lun	-
-	-	Auth	-	-	-	raidcom get lun	-
-	-	Lock/auth	-	-	Lock/auth	raidcom modify lun	-
-	Lock/auth	-	-	-	-	raidcom add path	-
-	Lock/auth	-	-	-	-	raidcom check_ext_storage path	-
-	Lock/auth	-	-	-	-	raidcom delete path	-
-	Lock/auth	-	-	-	-	raidcom disconnect path	-
-	-	-	-	-	Auth	raidcom get path	-
-	-	-	-	-	Lock/auth	raidcom delete pool	-
-	-	-	-	-	Auth	raidcom get pool	-
-	-	-	-	-	Lock/auth	raidcom modify pool	-
-	-	-	-	-	Lock/auth	raidcom rename pool	-
-	Auth	-	-	-	-	raidcom get port	-
-	Lock/auth	-	-	-	-	raidcom modify port	-
-	-	-	Auth	-	-	raidcom get parity_grp	-
-	Lock/auth	-	-	-	-	raidcom add rcu	-

Res. group	Port	Host group	Parity group	External volume/ VDEV	LDEV	Command	Option
-	Lock/auth	-	-	-	-	raidcom delete rcu	-
-	Auth	-	-	-	-	raidcom get rcu	-
-	Lock/auth	-	-	-	-	raidcom modify rcu	-
-	Lock/auth	-	-	-	-	raidcom add rcu_path	-
-	Lock/auth	-	-	-	-	raidcom delete rcu_path	-
Auth	-	-	-	-	-	raidcom get resource	-
Auth	-	-	-	-	-	raidcom lock resource	-
Auth	-	-	-	-	-	raidcom unlock resource	-
-	-	-	-	-	Lock/auth	raidcom add snap_pool	-
-	-	-	-	-	Auth	raidcom get snap_pool	-
-	*	-	-	-	-	raidcom add ssid	-
-	*	-	-	-	-	raidcom delete ssid	-
-	-	-	-	-	Lock/auth	raidcom add dp_pool	-
-	-	-	-	-	Auth	raidcom get dp_pool	-
-	-	Lock/ auth	-	-	-	raidcom set hba_wwn	-
-	-	Lock/ auth	-	-	-	raidcom reset hba_wwn	-
-	-	-	-	-	Lock/auth	raidcom monitor pool	-
-	-	-	-	-	Lock/auth	raidcom reallocate pool	-
-	-	-	-	-	-	raidcom get command_status	-
-	-	-	-	-	-	raidcom reset command_status	-
-	-	-	-	-	-	raidcom add resource	-
Lock/ auth	-	-	-	-	Lock/auth	raidcom add resource	-ldev_id <ldev#>

Res. group	Port	Host group	Parity group	External volume/ VDEV	LDEV	Command	Option
Lock/ auth	Lock/auth	-	-	-	-	raidcom add resource	-port <port#>
Lock/ auth	-	Lock/ auth	-	-	-	raidcom add resource	-port <port#> <host group name>
Lock/ auth	-	-	Lock/ auth	-	-	raidcom add resource	-parity_grp_id <gno-sgno>
Lock/ auth	-	-	-	Lock/auth	-	raidcom add resource	-external_ grp_id <gno-sgno>
Auth	-	-	-	-	-	raidcom delete resource	-
Lock/ auth	-	-	-	-	Lock/auth	raidcom delete resource	-ldev_id <ldev#>
Lock/ auth	Lock/auth	-	-	-	-	raidcom delete resource	-port <port#>
Lock/ auth	-	Lock/ auth	-	-	-	raidcom delete resource	-port <port#> <host group name>
Lock/ auth	-	-	Lock/ auth	-	-	raidcom delete resource	-parity_grp_id <gnosgno>
Lock/ auth	-	-	-	Lock/auth	-	raidcom delete resource	-external_ grp_id <gno-sgno>
Lock/ auth	-	-	-	-	-	raidcom modify resource	-
-	Lock/auth	-	-	-	Lock/auth	raidcom map resource	-
-	Lock/auth	-	-	-	Lock/auth	raidcom unmap resource	-
-	-	-	-	-	-	raidcom get error_message	-
-	Auth	-	-	-	-	raidcom add spm_wwn	-port <port#>
-	Auth	-	-	-	-	raidcom add spm_group	-port <port#>

Res. group	Port	Host group	Parity group	External volume/ VDEV	LDEV	Command	Option
-	*	Auth	-	-	-	raidcom add spm_group	-port <port#> [<host group name>] -spm_host_grp
-	Auth	-	-	-	-	raidcom delete spm_wwn	-port <port#>
-	Auth	-	-	-	-	raidcom delete spm_group	-port <port#>
-	*	Auth	-	-	-	raidcom delete spm_group	-port <port#> [<host group name>] -spm_host_grp
-	Auth	-	-	-	-	raidcom modify spm_wwn	-port <port#>
-	Auth	-	-	-	-	raidcom modify spm_group	-port <port#>
-	*	Auth	-	-	-	raidcom modify spm_group	-port <port#> [<host group name>] -spm_host_grp
-	Auth	-	-	-	-	raidcom get spm_wwn	-port <port#>
-	Auth	-	-	-	-	raidcom get spm_group	-port <port#>
-	*	Auth	-	-	-	raidcom get spm_group	-port <port#> [<host group name>] -spm_host_grp
-	Auth	-	-	-	-	raidcom monitor spm_wwn	-
-	Auth	-	-	-	-	raidcom monitor spm_group	-
-	-	Lock/ auth	-	-	-	raidcom add hba_iscsi	-port <port#> [<host group name>]

Res. group	Port	Host group	Parity group	External volume/ VDEV	LDEV	Command	Option
-	-	Lock/ auth	-	-	-	raidcom delete hba_iscsi	-port <port#> [<host group name>]
-	-	Lock/ auth	-	-	-	raidcom set hba_iscsi	-port <port#> [<host group name>]
-	-	Lock/ auth	-	-	-	raidcom reset hba_iscsi	-port <port#> [<host group name>]
-	-	Auth	-	-	-	raidcom get hba_iscsi	-port <port#> [<host group name>]
-	-	Lock/ auth	-	-	-	raidcom add chap_user	-port <port#> [<host group name>]
-	-	Lock/ auth	-	-	-	raidcom delete chap_user	-port <port#> [<host group name>]
-	-	Lock/ auth	-	-	-	raidcom set chap_user	-port <port#> [<host group name>]
-	-	Lock/ auth	-	-	-	raidcom reset chap_user	-port <port#> [<host group name>]
-	-	Auth	-	-	-	raidcom get chap_user	-port <port#> [<host group name>]
-	Auth	-	-	-	-	raidcom send ping	-port <port#>
-	Lock/auth	-	-	-	-	raidcom add external_iscsi_name	-
-	Lock/auth	-	-	-	-	raidcom delete external_iscsi_name	-
-	Lock/auth	-	-	-	-	raidcom modify external_chap_user	-
-	Lock/auth	-	-	-	-	raidcom modify initiator_chap_user	-

Chapter 5: Configuration setting commands

Res. group	Port	Host group	Parity group	External volume/ VDEV	LDEV	Command	Option
-	Auth	-	-	-	-	raidcom get external_iscsi_name	-
-	Auth	-	-	-	-	raidcom get initiator_iscsi_name	-
-	Auth	-	-	-	-	raidcom discover external_iscsi_name	-
-	Auth	-	-	-	-	raidcom check external_iscsi_name	-
-	Lock/auth	-	-	-	-	raidcom add rcu_iscsi_port	-
-	Lock/auth	-	-	-	-	raidcom delete rcu_iscsi_port	-
-	Auth	-	-	-	-	raidcom get rcu_iscsi_port	-
-	-	-	Lock/ auth	-	-	raidcom modify parity_grp	-
-	-	-	-	-	-	raidcom modify local_replica_opt	-
-	-	-	-	-	-	raidcom get local_replica_opt	-
-	-	-	-	-	-	raidcom get license	-
-	-	-	-	-	Lock/auth	raidcom modify quorum	-quorum_id <quorum id>
-	-	-	-	-	Auth	raidcom get quorum	-quorum_id <quorum id>
-	-	-	-	-	Lock/auth	raidcom replace quorum	-quorum_id <quorum id> -ldev_id <ldev#>
-	-	-	Lock/ auth	Lock/auth	Lock/auth	raidcom modify clpr	-ldev_id <ldev#> | -parity_grp_id <gno-sgno> | -external_grp_i d <gno-sgno>

Res. group	Port	Host group	Parity group	External volume/ VDEV	LDEV	Command	Option
-	-	-	Auth	Auth	Auth	raicom get clpr	-key pg, -key vvol
-	-	-	-	-	Auth	raidcom add snapshot	-
-	-	-	-	-	Auth	raidcom delete snapshot	Except: -range tree
-	-	-	-	-	Lock/auth	raidcom delete snapshot	-range tree
-	-	-	-	-	Auth	raidcom modify snapshot	-
-	-	-	-	-	Auth	raidcom map snapshot	-
-	-	-	-	-	Auth	raidcom unmap snapshot	-
-	-	-	-	-	Auth	raidcom get snapshot	-
-	-	-	-	-	Auth	raidcom replace snapshot	-
-	-	-	-	-	Auth	raidcom get system_opt	-key destage -cu <CU#>
-	-	-	-	-	Lock/auth	raidcom initialize pool	-ppid {<pool ID#> |<pool naming>}
-	-	-	-	-	-	raidcom get system	-
-	-	-	-	-	-	raidcom get ssid	-
-	-	-	-	-	-	raidcom add server	-
-	-	-	-	-	-	raidcom delete server	-
-	-*	Lock/ auth	-	-	-*	raidcom modify server	-
-	-	-	-	-	-	raidcom get server	-
-	-	-	-	-	Lock/auth	raidcom add quorum	-ldev_id <ldev#>
-	-	-	-	-	Lock/auth	raidcom delete quorum	-quorum_id <quorum id>
-	-	-	-	-	-	raidcom get apn	-

Res. group	Port	Host group	Parity group	External volume/ VDEV	LDEV	Command	Option
* If another user locks the applicable resource, the command execution fails. To execute the command, the applicable resource authority is not required.							

Relationship between commands supported only by VSP 5000 series and resources that need to be locked

Res. group	Port	Host group	NVM subsystem	Parity group	External volume/ VDEV	LDEV	Command	Option
-	-	-	-	-	-	Lock/ auth	raidcom add qos_grp	-ldev_id <ldev#>
-	-	-	-	-	-	Lock/ auth	raidcom delete qos_grp	-ldev_id <ldev#>
-	-	-	-	-	-	Lock/ auth	raidcom modify qos_grp	-
-	-	-	-	-	-	Auth	raidcom get qos_grp	-key resource
-	-	-	-	-	-	-	raidcom monitor resource	-

Relationship between commands supported only by VSP 5000 series, VSP E1090, and resources that need to be locked

Res. group	Port	Host group	NVM subsystem	Parity group	External volume/ VDEV	LDEV	Command	Option
-	-	-	Lock/auth	-	-	-	raidcom add nvm_subsystem	-
-	-	-	Lock/auth	-	-	-	raidcom modify nvm_subsystem	-
-	-	-	Lock/auth	-	-	-	raidcom delete nvm_subsystem	-

Res. group	Port	Host group	NVM subsystem	Parity group	External volume/ VDEV	LDEV	Command	Option
-	Auth	-	Auth	-	-	-	raidcom get nvm_subsystem	-key namesp ace -key port
-	Lock/ auth	-	Lock/auth	-	-	-	raidcom add nvm_subsystem_ port	-
-	Lock/ auth	-	Lock/auth	-	-	-	raidcom delete nvm_subsystem_ port	-
-	-	-	Auth	-	-	-	raidcom get nvm_subsystem_ port	-
-	-	-	Lock/auth	-	-	-	raidcom add host_nqn	- nvm_su bsystem _id <nvm subsyste m id>
-	-	-	-	-	-	-	raidcom modify host_nqn	-
-	-	-	Lock/auth	-	-	-	raidcom delete host_nqn	-
-	-	-	Auth	-	-	-	raidcom get host_nqn	-
-	-	-	Lock/auth	-	-	Lock/ auth	raidcom add namespace_path	-
-	-	-	Lock/auth	-	-	Lock/ auth	raidcom delete namespace_path	-
-	-	-	Auth	-	-	Auth	raidcom get namespace_path	-ldev_id <ldev#>
-	-	-	Lock/auth	-	-	Lock/ auth	raidcom add namespace	-

Chapter 5: Configuration setting commands

Res. group	Port	Host group	NVM subsystem	Parity group	External volume/ VDEV	LDEV	Command	Option
-	-	-	Lock/auth	-	-	Lock/ auth	raidcom delete namespace	-
-	-	-	Lock/auth	-	-	Lock/ auth	raidcom modify namespace	-
-	-	-	Auth	-	-	Auth	raidcom get namespace	-

Relationship between commands supported only by VSP E series, VSP G130, G/F350, G/F370, G/F700, G/F900, VSP G200, G400, G600, G800 and VSP F400, F600, F800 and resources that need to be locked

Res. group	Port	Host group	Parity group	External volume/ VDEV	LDEV	Command	Option
-	-	-	-	-	-	raidcom add license	-
-	-	-	-	-	-	raidcom delete license	-
-	-	-	-	-	-	raidcom modify license	-
-	-	-	Lock/ auth	-	-	raidcom initialize parity_grp	-
-	-	-	Auth	-	-	raidcom get drive	-
-	-	-	-	-	-	raidcom modify drive*	-
-	-	-	-	-	-	raidcom add partiy_grp	-
-	-	-	Lock/ auth	-	Lock/ auth	raidcom delete parity_grp	-
* If another user locks a resource, the command fails.								

Relationship between commands supported only by VSP and resources that need to be locked

Res. group	Port	Host group	Parity group	External volume/ VDEV	LDEV	Command	Option
-	Locked by user	Lock/ auth	-	-	-	raidcom add spm_group	-port <port#> [<host group name>] - spm_host_grp
-	Locked by user	Lock/ auth	-	-	-	raidcom delete spm_group	-port <port#> [<host group name>] - spm_host_grp
-	Locked by user	Lock/ auth	-	-	-	raidcom modify spm_group	-port <port#> [<hostgroup name>] - spm_host_grp
-	-	Auth	-	-	-	raidcom get spm_group	-port <port#> [<host group name>] - spm_host_grp

Request ID function

Request ID is an identification number that manages the execution results of asynchronous commands that can specify the -request_id auto option. If you specify the -request_id auto option for an asynchronous command, a request ID is automatically assigned. For details about the asynchronous commands that can specify the -request_id auto option, see each asynchronous command.

Even if the execution of the asynchronous command fails, a request ID might be assigned and the request ID might be output. Each time the request ID is output, confirm the execution result of the command using the **raidcom get command_status** command.

Up to 65,280 request IDs can be assigned per storage system. If no request ID is available, the command fails with EX_IDEXHA.

If EX_IDEXHA fails, release unnecessary request IDs in the following procedure, and then retry the command. Or, release the request ID by another user using the request ID, and then retry the command.

1. Specify the request ID and confirm the execution result using the `raidcom get command_status` command.

2. Release the request ID whose execution result is confirmed by using the `raidcom reset command_status -request_id <request #>` command or the `raidcom reset command_status -request_id all` command.

> 📄 **Note:**
>
> - Use the `raidcom reset command_status -request_id <request #>` command to release the specific request ID.
>
> - Use the `raidcom reset command_status -request_id all` command to release all the request IDs.

If you specify `-request_id all` for the `raidcom reset command_status` command, the request ID of the command whose execution result is not confirmed is also deleted.

raidcom add clpr

Supported storage systems:

- VSP 5000 series

- VSP E series

- VSP G/F350, G/F370, G/F700, G/F900

Adds CLPR. If no option is specified, no operation is performed.

This command is executed asynchronously with the command input. Check the completion of this process using the `raidcom get command_status` command.

> 📄 **Note:** If a failure occurs in a processor on an MP blade or a controller, this command cannot be run even on an MP blade or a controller on which no failure occurs.

Syntax

```
raidcom add clpr -clpr <clpr#> [-clpr_name <clpr name>] -cache_size <Cache Size(MB)>
```

Options and parameters

-clpr <clpr#>

Specifies the CLPR number (1 to 31) to add. You can specify a number that is not assigned. An error occurs for the following:

- 0 or a number greater than 31 is specified as a CLPR number.

- The specified CLPR number is already used.

-clpr_name <clpr name>

Sets a CLPR name. An error occurs if you specify any of the following:

- A character other than alphanumeric characters

- A name consisting of more than 16 characters

- A reserved CLPR name

- A CRPR name that is already used

-cache_size <Cache Size(MB)>

Specifies the cache size. An error occurs for the following:

- 0 to 4095 MB is specified.

- The size that is larger than 4095 MB is specified, but it is not a multiple of 2048 MB.

- Capacity exceeding the maximum capacity is specified.

Example

Creating a CLPR of CLPR number: 2, CLPR name: TEST02, and cache size: 8192 MB.

```
#raidcom add clpr -clpr 2 -clpr_name TEST02 -cache_size 8192
```

raidcom delete clpr

Supported storage systems:

- VSP 5000 series

- VSP E series

- VSP G/F350, G/F370, G/F700, G/F900

Deletes CLPR. If no option is specified, no operation is performed.

This command is executed asynchronously with the command input. Check the completion of this process on the raidcom get command_status command.

📄 **Note:** If a failure occurs in a processor on an MP blade or a controller, this command cannot be run even on an MP blade or a controller on which no failure occurs.

Syntax

```
raidcom delete clpr -clpr <clpr#>
```

Options and parameters

-clpr <clpr#>

Specifies the CLPR number (1 to 31) to delete. An error occurs if you specify either of the following:

- 0 or a number greater than 31 is specified as a CLPR number.

- A CLPR for which the parity group, LDEV, or external group is registered is specified.

Example

Deleting CLPR of CLPR number: 2.

```
#raidcom delete clpr -clpr 2
```

raidcom get clpr

Displays the CLPR information for the storage system.

Syntax

```
raidcom get clpr
```

Options and parameters

None

Examples

```
# raidcom get clpr
```

```
CLPR CLPR_NAME     TC_CAP(MB) TU_CAP(MB) WP_CAP(MB) SF_CAP(MB) U(%)  W(%)  S(%)
001  Oracle_DB_PROD   20000     10000       2000          0 50 10 0
003  Oracle_DB_BACK   10000      5000        500          0 50  5 0
```

Description of the `raidcom get clpr` output:

CLPR

CLPR ID (decimal)

CLPR_NAME

Nickname of the CLPR

TC_CAP(MB)

Capacity of cache memory of the CLPR

TU_CAP(MB)

Used capacity of cache memory of the CLPR

WP_CAP(MB)

Capacity of write pending data of the CLPR

SF_CAP(MB)

Capacity of sidefiles of the CLPR

U(%)

Usage rate of cache memory of the CLPR

W(%)

Rate of write pending data of the CLPR

S(%)

Usage rate of sidefiles of the CLPR

raidcom modify clpr

Modifies a CLPR.

📄 **Note:** If a failure occurs in a processor on an MP blade or a controller, this command cannot be run even on an MP blade or a controller on which no failure occurs.

Syntax

```
raidcom modify clpr -clpr <clpr#> { -ldev_id <ldev#> | -parity_grp_id <gno-sgno> | -
external_grp_id <gno-sgno>} | {[-clpr_name <clpr name>] [-cache_size <Cache
Size(MB)>]}
```

Options and parameters

-clpr <clpr#>

Specifies a CLPR number (0-31).

For example:

- -clpr 2

-ldev_id <ldev#>

Specifies an LDEV number (0-65279).

For example:

- -ldev_id 200

You cannot use this option to specify an LDEV in a LUSE volume.

-parity_grp_id <gno-sgno>

Specifies a parity group number (gno: 1-52, sgno: 1-32).

For example:

- 3-1

-external_grp_id <gno-sgno>

Specifies an external volume group number (gno: 1-16384, sgno: 1-4096).

For example:

- 52-11

-clpr_name <clpr name>

Supported storage systems:

- VSP 5000 series

- VSP E series

- VSP G/F350, G/F370, G/F700, G/F900

Specify this option to change the CLPR name. This option cannot be specified in the same command line as the `-ldev_id`, `-parity_grp_id`, and `-external_grp_id` options. An error occurs if you specify any of the following:

- A character other than alphanumeric characters

- A name consisting of more than 16 characters

- A reserved CLPR name. However, you can restore the default value of each CLPR.

- A CLPR name which is already used

-cache_size <Cache Size(MB)>

Supported storage systems:

- VSP 5000 series

- VSP E series

- VSP G/F350, G/F370, G/F700, G/F900

Specify this option to change the cache size. This option cannot be specified in the same command line as the `-ldev_id`, `-parity_grp_id`, and `-external_grp_id` options. An error occurs if you specify any of the following:

- CLPR0

- 0 to 4095 MB

- Larger than 4095 MB, but not a multiple of 2048 MB

- Capacity exceeding the maximum capacity

Examples

Moving the LDEV 02:00 to the CLPR ID 2.

```
# raidcom modify clpr -clpr 2 -ldev_id 0x0200
```

Chapter 5: Configuration setting commands

Moving the parity group 5-2 to the CLPR ID 2

```
# raidcom modify clpr -clpr 2 -parity_grp_id 5-2
```

Moving the external volume group 1-1 to the CLPR ID 2.

```
# raidcom modify clpr -clpr 2 -external_grp_id 01-01
```

Changing the CLPR name of CLPR ID 2 to TEST02.

```
# raidcom modify clpr -clpr 2 -clpr_name TEST02
```

Changing the cache size of CLPR ID 2 to 10 GB.

```
# raidcom modify clpr -clpr 2 -cache_size 10240
```

raidcom get system_opt

Supported storage systems:

- VSP 5000 series

- VSP E series

- VSP G130, G/F350, G/F370, G/F700, G/F900

- VSP G200, G400, G600, G800 and VSP F400, F600, F800

Views system options.

Syntax

```
raidcom get system_opt [-key {destage -cu <CU#> | mode -lpr <system|clpr#>}]
```

Options and parameters

-key <value>

Displays the setting status of a system option.

- destage -cu <CU#>: Displays the setting status of the write through operation. Nothing is displayed if the CU which is not installed is specified.

- mode -lpr <system|clpr#>: Displays the setting status of a system option mode. Nothing is displayed if a CLPR number which is not assigned is specified.

Examples

Displaying the basic information of the system option (without the -key option).

```
#raidcom get system_opt
Serial# : 563528
SDR : interleave
```

```
DCP : medium
CC : Enable
DS : Enable
LFT : 255
```

Output of the raidcom get system_opt command:

Serial#: Product serial number.

SDR: Spare disk recover. Displays the speed of copying data to a spare drive.

- interleave: I/Os from the host are prioritized. Copies data for slots, and then stops the copy processing for a certain period of time.

- fullspeed: Copy processing is prioritized.

DCP: Disk copy pace. When interleave is specified for SDR, displays the speed of copying data to a spare drive.

- faster: Copying data to the drive takes precedence over jobs from the host.

- medium: Optimal mode. The copying time is determined by the load of host I/Os

- slower: Jobs from the host takes precedence over the drive copy.

CC: Correction copy. Specifies the operation to be performed when the drive is blocked.

- Enable: When the drive is blocked, correction copy to the spare drive is performed.

- Disable: When the drive is blocked, correction copy to the spare drive is not performed.

DS: Dynamic sparing. Displays the behavior when a drive failure occurred because the threshold was exceeded.

- Enable: If a drive failure occurs because the threshold was exceeded, data is copied to the spare drive automatically.

- Disable: If a drive failure occurs because the threshold was exceeded, data is not copied to the spare drive automatically.

LFT: Link failure threshold. Displays the threshold (0 to 255) for notifying a link failure.

Displaying the setting status of the write through operation for each LDEV of CU#: 0.

```
#raidcom get system_opt -key destage -cu 0
Serial# LDEV# DESTAGE
563528 100 E
563528 101 E
563528 102 E
563528 103 E
```

Output of the raidcom get system_opt command:

DESTAGE: Displays the setting status of the write through operation.

- E: The write through operation setting is disabled (the write through suppression setting is enabled).

- D: The write through operation setting is enabled (the write through suppression setting is disabled).

Chapter 5: Configuration setting commands

Specifying " system " and displaying the setting status of system option mode.

```
#raidcom get system_opt -key mode -lpr system
Serial# : 563528
LPR : system
CACHE_T : level2
COM_CTL : 0xffffffff
MODE_ID : 117 256 2047
```

Output of the raidcom get system_opt command:

LPR: Displays the specified " system " or the CLPR number.

CACHE_T: CACHE tuning. Displays the setting status (level1 to level5) of the cache tuning level. A hyphen (-) is displayed if LPR is not " system " .

COM_CTL: Command control. Displays the setting status (0x00000000 to 0xffffffff) of information for switching read ahead conditions. A hyphen (-) is displayed if LPR is not " system " .

MODE_ID: Lists the numbers of the system option modes set to ON.

raidcom modify user_system_opt

Supported storage systems:

▪ VSP E590, VSP E790 , and VSP E1090

Changes the advanced system settings.

> **Note:** If a failure occurs in a processor on an MP blade or a controller, this command cannot be run even on an MP blade or a controller on which no failure occurs.

Syntax

```
raidcom modify user_system_opt -opt_id <Option ID> -opt <enable|disable>
```

Options and parameters

-opt_id <Option ID>

Specifies the option ID of the advanced system settings in decimal.

▪ 5: If you enable this option, the differential data is maintained in a pool with which a DP-VOL that has the user capacity (up to 4,194,304 MB) is linked when a new TC, UR, or GAD pair is created using the DP-VOL, or resynchronized (differential data management). Note that differential data is maintained, regardless of this setting, in a pool linked with the DP-VOL that has the user capacity (greater than 4,194,304 MB) when a new TC, UR, or GAD pair using the DP-VOL is created.

▪ 6: If you enable this option, the differential data is maintained in a pool with which a DP-VOL that has the user capacity (up to 4,194,304 MB) is linked when a new TC, UR, or GAD pair is created using the DP-VOL (differential data management). Note that differential data is maintained, regardless of this setting, in a pool linked with the DP-VOL that has the user capacity (greater than 4,194,304 MB) when a new TC, UR, or GAD pair using the DP-VOL is created.

▪ 9: If you enable this option, when tier relocation is suspended by the system, an alert is issued to users. For details about an alert (SIM) to be issued, see the *Provisioning Guide*.

> 📄 **Note:** If both the option IDs 5 and 6 of the advanced system settings are enabled, the setting of the option ID 5 is ignored.

--opt <enable|disable>

Specifies the settings to enable or disable use of the advanced system settings option specified with -opt_id.

▪ enable: Enables the specified option of the advanced system settings.

▪ disable: Disables the specified option of the advanced system settings.

Examples

Enables the option ID 5 of the advanced system settings.

```
# raidcom modify user_system_opt -opt_id 5 -opt enable
```

raidcom get user_system_opt

Supported storage systems:

▪ VSP E590, VSP E790 , and VSP E1090

Views the enabled option of the advanced system settings.

Syntax

```
raidcom get user_system_opt
```

Options and parameters

None

Examples

Viewing the enabled option of the advanced system settings.

```
# raidcom get user_system_opt
OPT_ID : 5 6
```

Description of each column in output example:

OPT_ID

Displays the enabled option ID of the advanced system settings.

raidcom get command_status

It displays error information of the configuration setting command (asynchronous command) to be executed asynchronously.

When an error occurs with the execution of an asynchronous command, the total number of errors or error information such as error codes (SSB1 and SSB2) are stored in the storage system at the first occurrence. After executing asynchronous command, check the error information by executing this command.

The error information is deleted if you execute the `raidcom reset command_status` command, or log out from the storage system. For details about `raidcom reset command_status`, see . When you log out from the system, all error information for the session you used is deleted.

The error codes of the asynchronous commands are stored in different areas depending on request ID is specified or not. The options at the time of command execution and the status for each storage area are as follows.

Asynchronous command type	Storage area per login session[1]	Storage area per request ID[2]
-request_id auto is specified	Stored	Stored
-request_id auto is not specified	Stored	Not stored
Notes:		

1. This area stores the error information prepared in the storage system for each login session. This area stores the error of the first occurrence during the login session. However, for an error of the second occurrence, error codes of SSB1 and SSB2 are not stored. To be able to refer the error code when an error occurs, reset the error information that is stored in the storage area for each login session by executing `raidcom reset command_status` before executing asynchronous commands.

Asynchronous command type	Storage area per login session[1]	Storage area per request ID[2]
2. This area stores the error information prepared for each request ID. The stored error code can be checked individually by specifying the request ID.		

Syntax

```
raidcom get command_status [-time <time(sec)>] [-request_id <request#>]
```

Options and parameters

[-time <time(sec)>]

Specifies waiting time to complete the process of asynchronous command.

If this option is omitted, the default waiting time (CMD_DEF_TMOUT: 7200 sec.) is set.

[-request_id <request#>]

Specifies the request ID. Specifying this option displays error information of the command specified by the request ID. If you specify the request ID of the command issued by another user, this command ends without displaying the information.

This command interprets <request#> as a hexadecimal number. If the specified <request#> satisfies either of the following conditions, EX_INVARG is returned:

- <request#> contains characters other than numbers and letters (a to z, A to Z).

- <request#> contains nine or more characters.

Returned Values

Either of the following returned values is returned to exit (), which allows users to check the execution results using a user program or a script.

- **0:** Normal Termination

- **1:** One or more errors occurred. Abnormal termination.

Examples

Displaying error information of the asynchronous command:

```
# raidcom get command_status
```

```
HANDLE SSB1 SSB2 ERR_CNT Serial# Description
7E30   2E20 6000      4  64034   The pool ID is not installed
```

Specifying request ID: 1, and displaying error information of the asynchronous command:

```
# raidcom get command_status -request_id 1
REQID R SSB1 SSB2 Serial# ID Description
00000001 - 2E20 6000 64034 3 The pool ID is not installed
```

Description of the `raidcom get command_status` output:

HANDLE

> Handle number that uniquely identifies the user

SSB1

> SSB1 error code. For details about error codes, see the *Command Control Interface User and Reference Guide*.

SSB2

> SSB2 error code. For details about error codes, see the *Command Control Interface User and Reference Guide*.

ERR_CNT

> Total number of errors in this user handle

Serial#

> Serial number.

REQID

> Request ID of the command.

R

> Displays if the error is caused by a failure of another command.

> - 0: The error occurred due to an error caused by a failure of another command. SSB1, SSB2, and the description show error codes and error information of the failure of another command.

> - T: The command failed due to the command specified by the `-request_id` option. SSB1, SSB2, and the description show error codes and error information of the command specified by the option.

> - - (hyphen): This information is not available for this error.

Chapter 5: Configuration setting commands

ID

Displays ID of the object related to the request ID. A hyphen (-) is displayed if the ID information is not available.

- An LDEV number is displayed if the request ID that is output in the following cases specified for the `-request_id` option.

    - When the `-ldev_id auto -request_id auto` option is specified in the **raidcom add ldev** command.

    - When any of the following options are specified in the **raidcom modify ldev** command.

        - `-upper_throughput_io <upper throughput io> -request_id auto`

        - `-upper_data_trans_mb <upper data trans mb> -request_id auto`

        - `-upper_alert_time <upper alert time> -request_id auto`

        - `-lower_throughput_io <lower throughput io> -request_id auto`

        - `-lower_data_trans_mb <lower data trans mb> -request_id auto`

        - `-lower_alert_time <lower alert time> -request_id auto`

        - `-response_priority <#priority> -request_id auto`

        - `-response_alert_time <response alert time> -request_id auto`

    - A hyphen (-) is displayed if the request ID that is output when the `-request_id auto` option is specified for the **raidcom map resource** command or the **raidcom unmap resource** command is specified for the `-request_id` option.

- An LU number is displayed if the request ID that is output when the lun_id auto -request_id auto option is specified for the raidcom add lun command is specified for the -request_id option.

- A server ID is displayed if the request ID that is output when executing the **raidcom add server** command, **raidcom delete server** command, or **raidcom modify server** command is specified for the `-request_id` option.

- A QoS group ID is displayed if the request ID that is output when executing the **raidcom add qos_grp** command, **raidcom delete qos_grp** command, or **raidcom modify qos_grp** command is specified for the `-request_id` option.

Chapter 5: Configuration setting commands

- An NVM subsystem ID is displayed if the request ID that is output when running the **raidcom add nvm_subsystem** command, **raidcom modify nvm_subsystem** command, **raidcom delete nvm_subsystem** command, **raidcom add nvm_subsystem_port** command, **raidcom delete nvm_subsystem_port** command, **raidcom add host_nqn** command, **raidcom modify host_nqn** command, **raidcom delete host_nqn** command, **raidcom add namespace_path** command, or **raidcom delete namespace_path** command is specified for the -request_id option.

- A namespace ID is displayed if the request ID that is output when running the **raidcom add namespace** command, **raidcom modify namespace** command, or **raidcom delete namespace** command is specified for the -request_id option.

Description

Error information. If no error occurred, a hyphen (-) is displayed.

EX_EWSTOT

Timeout while waiting the result of the command execution

For details, see the section describing **Command error messages** in the *Command Control Interface User and Reference Guide*.

raidcom reset command_status

Resets the error information of the configuration setting command that is stored in the storage system and executed asynchronously (Asynchronous command).

The command options and the information to be reset are as follows:

Option type	Storage area per login session	Storage area per request ID
No option	Deleted	-
-request_id <request#>	-	Deleted[1]
-request_id all	-	Deleted[2]

Notes:

1. Deletes only the error information of the specified <request #> used in the login session in which the command is executed. If you specify <request #> used in other login session, the deletion is not be performed.

2. Deletes all error information for each request ID used in the login session in which the command is executed.

Syntax

```
raidcom reset command_status [-request_id <request#> | -request_id all]
```

Chapter 5: Configuration setting commands

Options and parameters

[-request_id <request#> | -request_id all]

Specifies the request ID of a command whose error information is cleared.

- -request_id <request#>: Specifies the request ID of a command whose error information is cleared.

- -request_id all: Clears all error information of the raidcom add ldev command which is specified with the -ldev_id auto option and executed by the user who executes the raidcom reset command_status command.

For <request#>, type the request ID output by the `raidcom add ldev` command. This command interprets <request#> as a hexadecimal number. If the specified <request#> satisfies any of the following conditions, EX_INVARG or EX_CMDRJE is returned:

- <request#> contains characters other than numbers and letters (a to z, A to Z) (EX_INVARG is returned).

- <request#> contains nine or more characters (EX_INVARG is returned).

- <request#> is a hexadecimal number, and it is interpreted as 0x00000000 (EX_CMDRJE is returned).

- <request#> is a hexadecimal number, and it is interpreted as a value between 0x0000ff01 and 0xffffffff (EX_CMDRJE is returned).

Examples

Resetting the error information of the asynchronous command.

```
# raidcom reset command_status
```

Resetting the error information of the command for request ID: 1.

```
# raidcom reset command_status -request_id 1
```

Resetting the error information of the raidcom add ldev command with the -ldev_id auto option executed by the user.

```
# raidcom reset command_status -request_id all
```

raidcom add copy_grp

Creates a copy group.

> **Note:** If a failure occurs in a processor on an MP blade or a controller, this command cannot be run even on an MP blade or a controller on which no failure occurs.

Syntax

```
raidcom add copy_grp -copy_grp_name <copy group name>
     <device group name> [<device group name>] [-mirror_id <mu#>]
     [-journal_id <journal ID#>]
```

Options and parameters

-copy_grp_name <copy group name><device group name>[<device group name>]
Specifies the device group (maximum 32 characters) configuring a copy group (maximum 32 characters).

You can specify up to two device group names. If you specify more than two, the option is ignored.

For a copy group for an ShadowImage pair, specify two device groups.

For a copy group for a TrueCopy pair, specify only one device group for the relevant storage system side (primary/main or secondary/remote).

[-mirror_id <mu#>]
Specifies the mirror ID.

If this option is omitted (by **raidcom get copy_grp**), "-" is displayed.

> 📄 **Note:**
>
> The mirror ID to be registered is as follows depending on whether the environment variable is set or not and whether the option is specified or not.
>
> | Option | HORCC_MRCF environment variable is not set | HORCC_MRCF environment variable is set |
> |---|---|---|
> | -IM and -IH options are not specified | hx | x |
> | -IM option is specified | x | x |
> | -IH option is specified | hx | hx |
>
> x: The mirror ID specified in <mu#>

[-journal_id <journal ID#>]
Specifies the journal number (0-255).

If this option is omitted (by raidcom get copy_grp), "-" is displayed.

Example

Creating a copy group (ora) by device groups (grp1, grp2).

```
# raidcom add copy_grp -copy_grp_name ora grp1 grp2
```

Chapter 5: Configuration setting commands

raidcom delete copy_grp

Deletes the specified copy group.

📄 **Note:** If a failure occurs in a processor on an MP blade or a controller, this command cannot be run even on an MP blade or a controller on which no failure occurs.

Syntax

```
raidcom delete copy_grp -copy_grp_name <copy group name>
```

Options and parameters

-copy_grp_name <copy group name>
Specifies the name of the copy group (maximum 32 characters).

Example

Deleting the copy group: ora.

```
# raidcom delete copy_grp -copy_grp_name ora
```

raidcom get copy_grp

Displays the information of the specified copy group.

Syntax

```
raidcom get copy_grp
```

Options and parameters

None.

Example

Displaying copy group information.

```
# raidcom get copy_grp
```

```
COPY_GROUP LDEV_GROUP MU# JID# Serial#
ora grp1            0   -       64034
ora grp2            0   -       64034
```

Description of the raidcom get copy_grp output:

COPY_GROUP
> Copy group name

LDEV_GROUP
> Device group name that composes copy group

MU#
> Mirror ID to which the device group belongs. If -mirror_id is not specified at the creation, "-" is displayed.

JID#
> Journal number to which device group belongs. If -journal_id is not specified at the creation, "-" is displayed.

Serial#
> Product serial number. The display of the product serial number varies depending on storage systems.
>
> - The serial number for VSP 5000 series is displayed with a "5" at the beginning ("512345" = serial number 12345).
>
> - The serial number for VSP G1x00 and VSP F1500 is displayed with a "3" at the beginning ("312345" = serial number 12345).

raidcom add device_grp

Assigns a device name to the specified LDEV, and creates a device group. If the group already exists, the LDEV is added to the group. If the LDEV also already exists in the specified device group, the specified LDEV name is set.

> 📄 **Note:** If a failure occurs in a processor on an MP blade or a controller, this command cannot be run even on an MP blade or a controller on which no failure occurs.

Syntax

```
raidcom add device_grp -device_grp_name <device group name>
     <device name> -ldev_id <ldev#>… [-cnt <count>]
```

Options and parameters

device_grp_name <device group name><device name>
> Specifies the device group name (maximum 32 characters) and the device name in the device group (maximum 32 characters). If multiple LDEVs are specified, the same device name is set for all of them.

-ldev_id <ldev#> …
> Specifies the LDEV number (0-65279).

When you specify an LDEV that is part of a LUSE volume, all LDEVs in the LUSE volume have the same name. For example:

- -ldev_id 200

- -ldev_id 100-110

- -ldev_id 100 -cnt 10

Up to 64 of LDEVs can be specified.

When you specify an LDEV in a LUSE volume, the number of LDEVs in the LUSE volume is included.

[-cnt <count>]
Specifies the count (2-64).

If this option is omitted, the count is set to one.

Example

Assigning a device name: data1 to an LDEV: 400 and adding it to the device group: grp1.

```
# raidcom add device_grp -device_grp_name grp1 data1 -ldev_id 400
```

raidcom delete device_grp

Deletes the specified LDEV from the specified group. When the last LDEV is deleted, the device group is also deleted.

📄 **Note:** If a failure occurs in a processor on an MP blade or a controller, this command cannot be run even on an MP blade or a controller on which no failure occurs.

Syntax

```
raidcom delete device_grp -device_grp_name <device group name>
     -ldev_id <ldev#>… [-cnt <count>]
```

Options and parameters

-device_grp_name <device group name>
Specifies the device group name (maximum 32 characters)

-ldev_id <ldev#> …
Specifies the LDEV number (0-65279). For example:

- -ldev_id 200

- -ldev_id 100-110

- -ldev_id 100 -cnt 10

Up to 64 of LDEVs can be specified. When you specify an LDEV in a LUSE volume, the number of LDEVs in the LUSE volume is included.

[-cnt <count>]

Specifies the count (2-64).

If this option is omitted, the count is set to one.

Example

Deleting the LDEV400 from the device group: grp1.

```
# raidcom delete device_grp -device_grp_name grp1 -ldev_id 400
```

raidcom get device_grp

Displays the LDEV information for the specified device group, or lists all device groups.

Syntax

```
raidcom get device_grp [-device_grp_name <device group name>]
```

Options and parameters

[-device_grp_name <device group name>]

Displays the device (LDEV) information for the specified device group (maximum 32 characters).

If this option is omitted, the list of the registered device groups is displayed.

Example for displaying device group information

```
# raidcom get device_grp
```

```
LDEV_GROUP Serial#
grp1 64034
grp2 64034
grp3 64034
```

Example for displaying device group information: grp1

```
# raidcom get device_grp -device_grp_name grp1
```

```
LDEV_GROUP  LDEV_NAME  LDEV#  Serial#
grp1        data1      400    64034
grp1        data2      401    64034
```

Chapter 5: Configuration setting commands

Description of the `raidcom get device_grp` output:

LDEV_GROUP
> Device group name

LDEV_NAME
> Device name in the device group

LDEV#
> LDEV number

Serial#
> Product serial number.

raidcom get drive

Supported storage systems:

- VSP 5000 series
- VSP E series
- VSP G130, G/F350, G/F370, G/F700, G/F900
- VSP G200, G400, G600, G800 and VSP F400, F600, F800

Displays drive information.

Syntax

```
raidcom get drive [-parity_grp_id <gno-sgno> | -usage <usage>]
```

Options and parameters

-parity_grp_id <gno-sgno>
> Specifies the parity group number (gno: 1 to 52, sgno: 1 to 32)

> Example:

> 3-1

-usage <usage>
> Specifies the drive usage.

> Specify the following character strings for <usage>:

> - data: Data drive
> - spare: Spare drive
> - free: Unused drive

Example

Displaying drive information.

```
#raidcom get drive
```

```
LOCATION    TYPE    RPM    TOTAL_CAP(GB)    CODE         USAGE    STS    GROUP
1-1         SAS    15000          300     DKS5C-K300SS   DATA     NML    1-1
```

Description of the `raidcom get drive` output:

LOCATION

Displays the location of the drive in the xx-yy format. In Device Manager - Storage Navigator, the drive location is displayed in the HDDxx-yy format.

TYPE

Displays the drive type.

(VSP 5000 series) SSD or SCM is displayed as a drive type of a drive whose drive type code (CODE) is SPx5x-YxxxNC, where x is any character string or number. The displayed character string depends on the DKCMAIN microcode version.

RPM

Displays the number of revolutions of a drive in rpm. For SSD, a hyphen (-) is displayed.

TOTAL_CAP(GB)

Displays the capacity of a drive in gigabytes.

CODE

Displays the drive type code.

USAGE

Displays the drive usage.

- DATA: Data drive
- SPARE: Spare drive
- FREE: Unused drive

STS

Displays the drive status.

- NML: Normal.
- WAR: A blocked part exists.
- CPY: Drive copy is in process.
- CPI: Copy is incomplete.
- RSV: Spare disk is unusable.
- FAI: Blocked due to a failure.

▪ BLK: Blocked due to maintenance.

▪ UNK: The status is unknown.

GROUP

Displays the parity group number of the drive if it is contained in a parity group. If the drive is not contained in a parity group, a hyphen (-) is displayed.

raidcom modify drive

Supported storage systems:

▪ VSP E series

▪ VSP G130, G/F350, G/F370, G/F700, G/F900

▪ VSP G200, G400, G600, G800 and VSP F400, F600, F800

Configures spare drives or cancels the settings.

This command is executed asynchronously with the command input. Use the `raidcom get command_status` command to check if the command is completed.

📄 **Note:** If a failure occurs in a processor on an MP blade or a controller, this command cannot be run even on an MP blade or a controller on which no failure occurs.

Syntax

```
raidcom modify drive -drive_location <drive location> -spare {enable|disable}
```

Options and parameters

-drive_location <drive location>

Specifies the drive location.

Example:

When the drive location is HDD00-01:

0-1

-spare {enable|disable}

Sets a spare drive or cancels the setting.

▪ enable: Sets a drive as a spare drive.

▪ disable: Cancels the spare drive setting.

Example

Setting a drive whose drive location is HDD00-01 as a spare drive.

```
# raidcom modify drive -drive_location 0-1 -spare enable
```

raidcom get error_message

Displays the error message for the specified error code.

Syntax

```
raidcom get error_message -ssb <ssb1> <ssb2>
```

Options and parameters

-ssb <ssb1> <ssb2>

Specifies the error code as a hexadecimal number (add the 0x prefix).

- ▪ <ssb1>: Specifies SSB1 of the error code.

- ▪ <ssb2>: Specifies SSB2 of the error code.

Example

Displaying the error message for the error code whose SSB1 is 0x2E00 and SSB2 is 0x0023.

```
# raidcom get error_message -ssb 0x2E00 0x0023
CAUSE : Volume capacity is too small.
```

raidcom add external_grp

To use an external volume, map the volume of the external storage system to the external volume group on the local storage system. Only one external volume can be mapped in a single operation.

Adds an external volume to the specified external volume group, and connects to an external LUN on the specified external port/wwn. Only one external VOL is added in each operation. If the external volume group already exists, the external volume is added to the external volume group.

This command is executed asynchronously with the command input. Check the completion of this process using the **raidcom get command_status** command.

This command is rejected by EX_ENOOBJ in the following cases:

- ▪ The specified iSCSI port cannot be found.

- ▪ The iSCSI virtual port mode is enabled, but the specified iSCSI virtual port ID is not correct.

To fix this error, specify the correct port and iSCSI virtual port ID.

> **Note:** If a failure occurs in a processor on an MP blade or a controller, this command cannot be run even on an MP blade or a controller on which no failure occurs.

Chapter 5: Configuration setting commands

📄 **Note:** If the local storage system port is a Fiber Channel Bidirectional port, it is recommended that you always execute the command by specifying the `-safety_check enable` option unless instructed by the Command Control Interface User and Reference Guide.

If you do not specify the `-safety_check enable` option, the following problems might occur. For the details, see the Hitachi Universal Volume Manager User Guide.

- If you use a non-Hitachi external storage system:

  The I/O path from the external storage system to the local storage system that uses the route between the specified local storage system port and the external storage system port might be disconnected.

- If you use a Hitachi external storage system:

  The external or remote path connection that uses the route between the specified local storage system port and the external storage system port might be temporarily disconnected (the external or remote path will be reconnected immediately and no blockage occurs).

Syntax

```
raidcom add external_grp -path_grp <path group#>
  -external_grp_id <gno-sgno> -port <port#> {-external_wwn
  <wwn strings> | -external_iscsi_name <external iscsi name>
  -external_address <IP address> [-iscsi_virtual_port_id
  <iSCSI virtual port ID>]} -lun_id <lun#>
  [-emulation <emulation type>] [-clpr <clpr#>]
  [-external_attribute migration][-data_direct_mapping enable]
  [-command_device y -ldev_id <ldev#>][-safety_check enable]
```

Options and parameters

-path_grp <path group#>
　　　Specifies the external VOL path group number (0-63231).

-external_grp_id <gno-sgno>
　　　Specifies the external volume group number (gno:1-16384, sgno:1-4096). For example:

- 52-11

-port <port#>
　　　Specifies the port number. Specifies the port number whose attribute is External. For example:

- CL1-A

Displays an external port.

-external_wwn <wwn strings>

Specifies the WWN (16-digit hexadecimal value) of the external storage system. The 17th digit and later are ignored. The value can be split in units of 4 bytes by "," (comma). For example:

- 210000e08b0256f8

- 210000e0,8b0256f8

-external_iscsi_name <external iscsi name>

Specifies the iSCSI name of the iSCSI target on the external storage system in an iqn format or an eui format.

- iqn format: `iqn.` and the subsequent maximum 219 characters.

- eui format: `eui.` and the subsequent 16 characters in hexadecimal notation.

> 📄 **Note:**
>
> You can specify the pseudo WWN of the iSCSI target to the `-external_wwn` option instead of specifying the iSCSI target by using the `-external_iscsi_name` option and the `-external_address` option. The pseudo WWN corresponds to the iSCSI name and IP address of the iSCSI target on the external storage system. By executing the **raidcom get external_iscsi_name** command, you can confirm the suitability between the pseudo WWN and the iSCSI name and IP address of the iSCSI target on the external storage system.

[-iscsi_virtual_port_id <iSCSI virtual port ID>]

Specifies the iSCSI virtual port ID when the specified port is iSCSI and the iSCSI virtual port mode is enabled. If this option is omitted, 0 is assumed.

An error occurs if you specify this option in any of the following cases:

- The virtual port of the specified iSCSI virtual port ID is not found.

- The iSCSI virtual port mode is disabled.

- The device does not support the iSCSI virtual port mode.

-external_address <IP address>

Specifies the IP address of the iSCSI target on the external storage system. An IPv4 address or an IPv6 address can be specified.

Note that the following addresses cannot be specified when you specify the IPv4 address:

- Network address (for example: 192.168.10.0, 0.120.10.1)

- Broadcast address (for example: 255.255.255.255, 10.1.255.255)

- Loop back address (for example: 127.0.0.1)

Chapter 5: Configuration setting commands

Note that the following addresses cannot be specified when you specify the IPv6 address:

- Unspecified address (for example: ::)

- Multicast address (for example: ff:1024:1215::01)

- Loop back address (for example: ::1)

> 📄 **Note:**
>
> You can specify the pseudo WWN of the iSCSI target to the –
> `external_wwn` option instead of specifying the iSCSI target by using the
> `-external_iscsi_name` option and the `-external_address` option.
> The pseudo WWN corresponds to the iSCSI name and IP address of the
> iSCSI target on the external storage system. By executing the **raidcom
> get external_iscsi_name** command, you can confirm the suitability
> between the pseudo WWN and the iSCSI name and IP address of the
> iSCSI target on the external storage system.

-lun_id <lun#>
Displays LUN (0-4095) of the external storage system port.

[-emulation <emulation type>]
Specifies the emulation type. If this option is omitted, OPEN-V is used.

The valid values for <emulation type> are:

- OPEN-3, OPEN-8, OPEN-9, OPEN-E, OPEN-K, OPEN-L, OPEN-V

- 3390-1, 3390-2, 3390-3, 3390-A, 3390-3A, 3390-3B, 3390-3C, 3390-3R, 3390-9, 3390-9A, 3390-9B, 3390-9C, 3390-L, 3390-LA, 3390-LB, 3390-LC, 3390-M, 3390-MA, 3390-MB, 3390-MC, 3390-V

- 3380-3, 3380-3A, 3380-3B, 3380-3C

Some emulation types cannot be specified according to the type of device.

> ⚠️ **Caution:**
>
> You can specify 3390-3 or 3390-3R as the emulation type, but these are
> unable to be mixed. You can specify the 3380 series or the 3390 series as
> the emulation type for each parity group, but they are unable to be mixed
> for each 32 address boundary because of the OS restriction.

[-clpr <clpr#>]
Specifies the CLPR number.

[-external_attribute migration]
Specifies if an attribute of NDM functions is set.

[-data_direct_mapping enable]
Specifies to set the data direct mapping attribute. The data direct mapping attribute is automatically set to an LDEV which is created in the external volume group having the data direct mapping attribute.

Chapter 5: Configuration setting commands

[-command_device y -ldev_id <ldev#>]

Supported storage systems:

- VSP 5000 series
- VSP G1x00 and VSP F1500
- VSP E series
- VSP G130, G/F350, G/F370, G/F700, G/F900

Specifies to map an external volume as a remote command device. The LDEV number specified by -ldev_id <ldev#> is set for the remote command device.

[-safety_check enable]

Specify the Fiber Channel bidirectional port for the `-port` option. If you specify any other port, this option is ignored. When this option is specified, processing that might cause the disconnection between the Bidirectional port specified by the `-port` option and the Fiber Channel port of the external storage system specified by the `-external_wwn` option is suppressed. When processing is suppressed, the external volume will not be added.

Examples

Mapping an LU: 0 defined to the external storage system port: 50060e80,05fa0f36 connected to the port: CL1-A (External port) by the External Volume Group#1-1 and the path group#1.

```
# raidcom add external_grp -path_grp 1 -external_grp_id 1-1
-port CL1-A -external_wwn 50060e80,05fa0f36 -lun_id 0 -safety_check enable
```

Mapping an LU:0 defined to the external storage system iSCSI target (iSCSI name: iqn.z2, IP address: 158.214.135.100) which is connected to the port: CL1-A (iSCSI port) of the local storage system by the External Volume Group#1-1 and the path group#1.

```
# raidcom add external_grp -path_grp 1 -external_grp_id 1-1 -port CL1-A -
external_iscsi_name iqn.z2 -external_address 158.214.135.100 -lun_id 0
```

Mapping an LU:0 defined to the external storage system port: 50060e80,05fa0f36 connected to the port: CL1-A (External port) by the External Volume Group#1-1 and the path group#1, and setting the attribute for the NDM function and the data direct mapping attribute.

```
# raidcom add external_grp -path_grp 1 -external_grp_id 1-1 -port CL1-A -external_wwn
50060e80,05fa0f36 -lun_id 0 -external_attribute migration -data_direct_mapping enable
-safety_check enable
```

Mapping an LU:0 defined to the external storage system iSCSI target (iSCSI name: iqn.z2, IP address: 158.214.135.100) connected to the iSCSI port: CL1-A, iSCSI virtual port ID: 2 of the local storage system by the External Volume Group#1-1 and the path group#1.

```
# raidcom add external_grp -path_grp 1 -external_grp_id 1-1 -port CL1-A
-iscsi_virtual_port_id 2 -external_iscsi_name iqn.z2 -external_address
158.214.135.100 -lun_id 0
```

Mapping an LU: 0 defined to the external storage system port: 50060e80, 05fa0f36 connected to the port: CL1-A (External port) by the External Volume Group#1-1 and the path group#1, and setting the LDEV number: 1.

```
# raidcom add external_grp -path_grp 1 -external_grp_id 1-1 -port CL1-A
-external_wwn 50060e80,05fa0f36 -lun_id 0 -command_device y -ldev_id 1 -safety_check
enable
```

raidcom check_ext_storage external_grp

Specifies the external volume group, check the connection for the external VOL, and then restart using. Only one external VOL is operated in each operation.

An LDEV or device group can be specified instead of an external volume group:

- If an LDEV is specified, CCI finds the external volume groups to which the specified LDEV belongs and displays the result.

- If a device group is specified, CCI finds the external volume groups to which the specified device group belongs and displays the result.

If no LDEV exists in the external volume, the command is rejected with EX_ENOOBJ.

> **Note:** If a failure occurs in a processor on an MP blade or a controller, this command cannot be run even on an MP blade or a controller on which no failure occurs.

Syntax

```
raidcom check_ext_storage external_grp {-external_grp_id
   <gno-sgno> | -ldev_id <ldev#>} | -grp_opt <group option>
   -device_grp_name <device group name> [<device name>]}
```

Options and parameters

-external_grp_id <gno-sgno>
Specifies the external volume group number (gno:1-16384, sgno:1-4096). For example:

- 52-11

-ldev_id <ldev#>
Specifies the LDEV number (0-65279). For example:

- -ldev_id 200

-grp_opt <group option>
Specifies the device information extracted from the LDEV belonging to the device group. Specify "ldev" (fixed). The information about LDEV belonging to the device group is used.

Chapter 5: Configuration setting commands

-device_grp_name <device group name> [<device name>]

> Specifies the name of device group (maximum 32 characters) to be operated.
>
> To operate an LDEV in the device group, specify the device name of LDEV (maximum 32 characters) within the device group.
>
> If the device name is omitted, all of the LDEVs in the device group are operated.

Examples

Executing the confirmation of existence and the LDEV recovery for the external volume group #1-1.

```
# raidcom check_ext_storage external_grp -external_grp_id 1-1
```

Executing the confirmation of connection and the LDEV recovery for the external volume group including the external volume (LDEV:200).

```
# raidcom check_ext_storage external_grp  -ldev_id 200
```

Executing the confirmation of connection and the LDEV recovery for the external volume group including the LDEV belonging to the device group: grp1.

```
# raidcom check_ext_storage external_grp -grp_opt ldev -device_grp_name grp1
```

raidcom delete external_grp

Releases the mapping of the external volume to delete the registered external VOLs from the configuration. Only one external VOL is deleted in each operation. When the last external volume is deleted, the path group is also deleted.

This command is executed asynchronously with the command input. Check the completion of this process using the **raidcom get command_status** command.

> 📄 **Note:** If a failure occurs in a processor on an MP blade or a controller, this command cannot be run even on an MP blade or a controller on which no failure occurs.

Syntax

```
raidcom delete external_grp -external_grp_id <gno-sgno> [-forcible]
```

Options and parameters

-external_grp_id <gno-sgno>

> Specifies the external volume group number (gno:1 to 16384, sgno:1 to 4096). For example:
>
> - 52-11

[-forcible]

Supported storage systems:

- VSP 5000 series

- VSP E series

- VSP G/F350, G/F370, G/F700, G/F900

- VSP G200, G400, G600, G800 and VSP F400, F600, F800

Specify this option to delete external volumes whose connection is not disconnected. Specifying this option deletes an external volume without writing data on cache to the external volume. If your microcode version does not support this option, this option is ignored even if it is specified.

Examples

Deleting the external volume group #1-1.

```
# raidcom delete external_grp -external_grp_id 1-1
```

Deleting external volume 1-1 whose connection is not disconnected without writing data on cache to the external volume.

```
# raidcom delete external_grp -external_grp_id 1-1 -forcible
```

raidcom disconnect external_grp

Disconnects the connection to the external volumes. Only one external VOL is operated in each operation.

You can specify an LDEV defined for the external volume group or a device group to which the LDEV in the external volume group belongs instead of the external volume group.

- If an LDEV is specified, CCI finds the external volume groups to which the specified LDEV belongs and displays the result.

- If a device group is specified, CCI finds the external volume groups to which the LDEV in the specified device group belongs and displays the result.

If no LDEV exists in the external volume group, the command is rejected with EX_ENOOBJ.

Before finishing the write processing from the cache to the external volume, the processing of **raidcom disconnect external_grp** command ends. Check the status (STS) using the **raidcom get path** command, and confirm the finishing of the write processing (destaging). (destaging). The following are the statuses (STS) after executing the **raidcom disconnect external_grp** command.

- NML: It means the previous status of receiving the request by the **raidcom disconnect external_grp** command.

- SYN: Write processing (destaging) is in process.

- DSC: Write processing (destaging) has finished.
- BLK: Write processing (destaging) has failed.

For details, see **raidcom get path**.

📄 **Note:** If a failure occurs in a processor on an MP blade or a controller, this command cannot be run even on an MP blade or a controller on which no failure occurs.

Syntax

```
raidcom disconnect external_grp {-external_grp_id <gno-sgno>
    | -ldev_id <ldev#>} | -grp_opt <group option>
    -device_grp_name <device group name> [<device name>]}
```

Options and parameters

-external_grp_id <gno-sgno>

Specifies the external volume group number (gno:1-16384, sgno:1-4096). For example:

- 52-11

-ldev_id <ldev#>

Specifies the LDEV number (0-65279). For example:

- -ldev_id 200

-grp_opt <group option>

Specifies the device information extracted from the LDEV belonging to the device group.

Specify "ldev" (fixed). The information of LDEV belonging to the device group is used.

-device_grp_name <device group name> [<device name>]

Specifies the name of device group (maximum 32 characters) to be operated.

To operate an LDEV in the device group, specify the device name of LDEV (maximum 32 characters) within the device group.

If the device name is omitted, all of the LDEVs belonging in the device group are operated.

Examples

Disconnect the connection to the external volume group #1-1.

```
# raidcom disconnect external_grp -external_grp_id 1-1
```

Disconnect the connection to the external volume group including the external volume (LDEV:200) to "blocked".

```
# raidcom disconnect external_grp -ldev_id 200
```

Disconnect the connection to the external volume group including the LDEV belonging to the device group: grp1.

```
# raidcom disconnect external_grp -grp_opt ldev -device_grp_name grp1
```

raidcom get external_grp

Displays the information of registered external volumes.

Syntax

```
raidcom get external_grp [-external_grp_id <gno-sgno>]
```

Options and parameters

[-external_grp_id <gno-sgno>]

Specifies the external volume group number (gno:1-16384, sgno:1-4096).

If this option is omitted, the list of the registered external volumes is displayed.

If this option is specified. the LDEV information defined for the specified external volume group is displayed. For example:

52-11

Example

Displaying external volume information (The display might not be in ascending order).

```
#raidcom get external_grp
```

```
T GROUP Num_LDEV U(%) AV_CAP(GB) R_LVL  E_TYPE  SL  CL  DRIVE_TYPE  M
E 1-1   0        0      100       -      OPEN-V  0   0   OPEN-V  N
E 1-2   0        0       30       -      OPEN-V  0   0   OPEN-V  N
```

Description of the each column in output example:

T

Type of the volume group

R

Parity group, E: External volume group

GROUP

External volume group number

Num_LDEV

Number of LDEV assigned to the external volume group

U(%)
> Usage rate of the external volume group

AV_CAP(GB)
> Available capacity (free space) for the external volume group

R_LVL
> RAID level of the parity group. As the external volume group is not relevant, "-" (bar) is displayed.

E_TYPE
> Base emulation type of the external volume group

SL
> SLPR to which the external volume group belongs (always displays 0)

CL
> CLPR to which the external volume group belongs

DRIVE_TYPE
> Product ID included in the SCSI Inquiry command of the external volume group

M
> Displays whether the allocation of external volume group to CLPR is changed.
> - Y: The allocation of external volume group is being changed.
> - N: The allocation of external volume group is not changed.
> - - (hyphen): The information is not available (not supported).

Displaying the external volume information by specifying the external volume group:

```
# raidcom get external_grp -external_grp_id 01-01
```

```
T GROUP P_NO LDEV# STS LOC_LBA SIZE_LBA Serial# SP
E 1-1 0 - NML 0x000000000000 0x000000003f00 64034 -
E 1-1 1 200 NML 0x000000003f00 0x000000010000 64034 R
E 1-1 2 201 REG 0x000000013f00 0x000000010000 64034 V
E 1-1 3 - DEL 0x000000023f00 0x0000f0000000 64034 -
```

Description of each column in the above output example:

T
> Type of the volume group

R
> Parity group, E: External volume group

GROUP
> External volume group number.

P_NO

Partition number in this external volume group.

LDEV#

LDEV number assigned to this external volume group.

STS

Displays the following status.

- ▪ NML: an LDEV is installed.

- ▪ REG: an LDEV is being created.

- ▪ DEL: an LDEV is being deleted.

LOC_LBA

Starting point of LBA for this partition on this external volume group, in blocks (512 bytes).

SIZE_LBA:

Partition size of this external volume group, in blocks (512 bytes).

Serial#

Product serial number.

SP

Displays whether the LDEV uses the expanded space of the parity group.

- ▪ V: LDEV uses the expanded space.

- ▪ R: LDEV does not use the expanded space.

- ▪ - (hyphen): LDEV is not mounted.

raidcom modify external_grp

Changes the attribute of external volume options (cache mode, cache inflow control mode, and MP blade ID setting).

📄 **Note:** If a failure occurs in a processor on an MP blade or a controller, this command cannot be run even on an MP blade or a controller on which no failure occurs.

Syntax

```
raidcom modify external_grp -external_grp_id <gno-sgno>
     {-cache_mode {y|n|through|sync} | -cache_inflow {y|n}
     | -mp_blade_id <mp#>} | -load_balance <mode>
   |-alua_switch {y|n}}
```

Options and parameters

-external_grp_id <gno-sgno>
> Specifies the external volume group number (gno:1-16384, sgno:1-4096). For example:
>
> - 52-11

-cache_mode {y|n|through|sync}
> Specifies whether to use the cache mode for an external volume. If the relevant external volume group is split into the multiple LDEVs and these LDEVs are allocated to the pools, you cannot change the setting for the cache mode.
>
> - y: Write cache enabled (E)
>
> - n: Write cache disabled (D)
>
> - through: Cache through (T). Available only when the attribute of the relevant external volume is NDM.
>
> - sync: Write Sync mode (S). Available only when the attribute of the relevant external volume is NDM.

-cache_inflow {y|n}
> Specifies whether to use the Cache Inflow Control mode for an external volume. If the external volume group consists of multiple LDEVs and these LDEVs are allocated to the pool, you cannot change this parameter.
>
> - y: Cache Inflow Control mode enabled (E)
>
> - n: Cache Inflow Control mode disabled (D)

-mp_blade_id <mp#>
> Specifies the MP blade ID (0-15). For example:
>
> -mp_blade_id 2

-load_balance <mode>
> Specifies load distribution mode of the alternate paths.
>
> - normal: normal round robin
>
> - extended: extended round robin
>
> - disable: disables the alternate paths

-alua_switch {y|n}
> Specifies whether the ALUA mode is used.

Examples

Turning the cache mode of the external volume group #01-01 ON.

```
# raidcom modify external_grp -external_grp_id 01-01 -cache_mode y
```

Enabling the Cache Inflow Control mode of the external volume group #01-01 ON.

```
# raidcom modify external_grp -external_grp_id 01-01 -cache_inflow y
```

Chapter 5: Configuration setting commands

Changing the MP blade ID of the external volume group #01-01 to "3".

```
# raidcom modify external_grp -external_grp_id 01-01 -mp_blade_id 3
```

raidcom discover external_storage

Searches the port information on the external storage system connected to the external port.

If you execute the **raidcom discover external_storage** command while the **raidcom discover lun** command or another **raidcom discover external_storage** command is being executed, the external storage system might not be displayed. If this happens, confirm that the storage system in which the command is being executed and the external storage system are connected correctly, and the LU of the external storage system is configured correctly.

If both the storage system and the external storage system have no problem, wait a while, and then execute one command at a time for a storage system.

> 📄 **Note:** If a failure occurs in a processor on an MP blade or a controller, this command cannot be run even on an MP blade or a controller on which no failure occurs.

> 📄 **Note:**
>
> If the local storage system port is a Fiber Channel Bidirectional port, it is recommended that you always execute the command by specifying the `-safety_check enable` option unless instructed by the user guide.
>
> If you do not specify the `-safety_check enable` option, the following problems might occur. For the details, see the Hitachi Universal Volume Manager User Guide.
>
> - If you use a non-Hitachi external storage system:
>
>   The I/O path from the external storage system to the local storage system that uses the route between the specified local storage system port and the searched external storage system port might be disconnected.
>
> - If you use the Hitachi external storage system:
>
>   The external or remote path connection that uses the route between the specified local storage system port and the searched external storage system port might be temporarily disconnected (the external or remote path will be reconnected immediately and no blockage occurs).

Syntax

```
raidcom discover external_storage -port <port#> [-safety_check enable] [-
discovery_external_wwn <wwn strings>]
```

Chapter 5: Configuration setting commands

Options and parameters

-port <port#>

Specifies the port number. Specifies the port number where the attribute is "ELUN(External)". For example:

- CL1-A

Displays an external port.

[-safety_check enable]

Specify the Fiber Channel bidirectional port for the `-port` option. If you specify any other port, this option is ignored.

When this option is specified, processing that might cause a disconnection between the bidirectional port specified by the `-port` option and the Fiber Channel port of the external storage system is suppressed. When processing is suppressed, the following information is displayed for Serial#, VENDOR_ID, and PRODUCT_ID.

- Serial#: "-" （hyphen）

- VENDOR_ID: "Unknown"

- PRODUCT_ID: "Unknown"

[-discovery_external_wwn <wwn strings>]

If you specify this option, information only for the specified external storage system port is displayed. Specify the external storage system WWN (Hexadecimal) in 8 bytes. You can separate 4 bytes using "," (comma).

(Example)

- 21000e08b0256f8

- 21000e0,8b0256f8

Use this option in one of the following methods:

1. Specify the `-discovery_external_wwn <wwn strings>` option without the `-safety_check enable` option.

   Use this method to check the Serial #, VENDOR_ID, and PRODUCT_ID information of the port specified by the `-port` option as a bidirectional port and the external storage system port.

   When using this method, make sure that the route is not applicable to the precautions described in the Hitachi Universal Volume Manager User Guide before executing the command. If applicable, this method might temporarily disconnect the external or remote path that uses the route between the external storage system port specified by `-discovery_external_wwn` and the bidirectional port specified by the `-port` option. If the external storage system is made by a third party, the I/O path from the external storage system to the local storage system might be disconnected.

2. Specify the `-discovery_external_wwn <wwn strings>` option with the `-safety_check enable` option.

   Use this method to check whether the external storage port specified by the `-discovery_external_wwn` option is connected when the route is applicable to the route mentioned above.

Example

Displaying the external storage system ports from the port: CL1-A.

```
# raidcom discover external_storage -port CL1-A
```

```
PORT WWN PM USED Serial# VENDOR_ID PRODUCT_ID
CL1-A 50060e8005fa0f36 M YES 60010 HITACHI VSP
CL1-A 50060e8005fa0f38 M YES 60010 HITACHI VSP
```

Displaying the external storage system ports from the port: CL1-A by specifying the `-safety_check enable` option.

```
# raidcom discover external_storage -port CL1-A -safety_check enable
PORT WWN PM USED Serial# VENDOR_ID PRODUCT_ID
CL1-A 50060e8005fa0f36 M No   -       Unknown Unknown
CL1-A 50060e8005fa0f38 M YES 60010 HITACHI VSP
```

Checking the Serial#, VENDOR_ID, and PRODUCT_ID of the external storage system port 50060e8005fa0f36 from the port: CL1-A.

```
# raidcom discover external_storage -port CL1-A -discovery_external_wwn
50060e8005fa0f36
PORT   WWN             PM USED Serial# VENDOR_ID PRODUCT_ID
CL1-A 50060e8005fa0f36 M YES   60010   HITACHI   VSP
```

Checking the connection with the external storage system port 50060e8005fa0f36 from the port: CL1-A.

```
# raidcom discover external_storage -port CL1-A -safety_check enable -
discovery_external_wwn 50060e8005fa0f36
PORT   WWN              PM USED Serial# VENDOR_ID PRODUCT_ID
CL1-A 50060e8005fa0f36 M No    -        Unknown   Unknown
```

Description of each column in output example:

PORT

   Displays the external port number of the storage system.

WWN

   Displays the WWN which can be referred to from the port.

Chapter 5: Configuration setting commands

PM:

Displays the path mode for external path.

- ▪ M: Multi

- ▪ S: Single

- ▪ A: APLB

USED

Displays whether this target WWN is used.

- ▪ YES: Used

- ▪ NO: Not used

Serial#

Product serial number of the external storage system.

If the `-safety_check enable` option is specified, "-" (hyphen) is displayed when processing that might cause disconnection is suppressed.

VENDOR_ID

Displays the vendor name of the external storage system. "OTHER" is displayed if an unsupported external storage system is connected.

If the `-safety_check enable` option is specified, "Unknown" is displayed when processing that might cause disconnection is suppressed.

PRODUCT_ID

Displays the system name of the external storage system. "OTHER" is displayed if an unsupported external storage system is connected.

If the `-safety_check enable` option is specified, "Unknown" is displayed when processing that might cause disconnection is suppressed.

raidcom add host_grp

Creates a host group or an iSCSI target on the specified port.

- ▪ If the port type is FIBRE or FCoE (Fibre Channel over Ethernet), this command creates a host group.

- ▪ If the port type is iSCSI, this command creates the iSCSI target (equivalent of a host group) and the iSCSI name.

If the specified port does not exist, the command is rejected with EX_ENOOBJ.

If the specified port/host group already exists, the existing host group name is changed to the specified host group name.

If the specified port, iSCSI target, and iSCSI name already exist, the existing information is changed to the specified information.

The specified host group name must be unique in a port.

Syntax

```
raidcom add host_grp -port <port#> -host_grp_name <host group name> [-iscsi_name
<target iscsi name>]
```

Options and parameters

-port <port#>

Specifies the port number and the host group ID. The host group ID can be omitted. For example:

- CL1-A-g

    A is the port number, and *g* is the host group ID (0-254).

- CL1-A

    If the host group ID is not specified, an unused host group ID is automatically assigned. However, if multiple commands without host group ID are executed simultaneously for the same port, the same host group ID might be assigned to multiple newly created host groups.

    To avoid this behavior, use the **raidcom lock resource** command in advance to lock the resource group to which the unused host group ID belongs. If you lock the resource group, commands executed by other users cannot take out unused host group IDs from the resource group. If you set multiple host groups for the same port, execute the next **raidcom add host_grp** command after the execution of the current **raidcom add host_grp** command is complete.

    After the execution of all the **raidcom add host_grp** commands are complete, use the **raidcom unlock resource** command to unlock the resource group.

-host_grp_name <host group name>

If the port type is FIBRE or FCoE, specifies the HOST group name. Up to 64 characters can be set by CCI. If more than 64 characters are set, commands that specify host group name by CCI cannot be executed.

If the port type is iSCSI, specifies the iSCSI target name. Up to 32 English one-byte characters can be set by CCI.

-iscsi_name <target iscsi name>

Specifies the iSCSI name by using either one of two formats, iqn or eui.

- iqn format: "iqn." followed by up to 219 English one-byte characters. The permitted characters are:

    - Alphabet

    - Number (0-9)

    - Period (.)

    - Hyphen (-)

    - Colon (:)

- eui format: "eui" (case sensitive) followed by a 16-digit hexadecimal value.

Chapter 5: Configuration setting commands

If this option is omitted, the default settings are specified. The default value depends on the serial number, the port number, or the target ID.

> 📄 **Note:**
>
> - You cannot register multiple iSCSI names to the same port if the only difference among them is the case (lowercase or uppercase). For example, if iqn.win2k8.example.of.iqn.form is already registered as an iSCSI name, you cannot register iqn.win2k8.example.of.iqn.FORM as a different iSCSI name.
>
> - If you input an iSCSI name in iqn format, use lowercase letters only.
>
> - If you input an iSCSI name in eui format, use lowercase letters for "eui.", and use uppercase letters for the following hexadecimal value.

Examples

Creating a host group ID: 3, the host group name: a host group of Win_export, to the port: CL4-E.

```
# raidcom add host_grp -port CL4-E-3 -host_grp_name Win_export
```

Creating a host group ID: allocated automatically, the host group name: a host group of Win_export, to the port: CL4-E.

```
# raidcom add host_grp -port CL4-E -host_grp_name Win_export
```

Creating an iSCSI name: iqn.2014-04.jp.co.hitachi:xxx.h70.i.62510.1A.FF, iSCSI target name: Target00, to the port: CL4-E.

```
# raidcom add host_grp -port CL4-E -host_grp_name Target00
-iscsi_name iqn.2014-04.jp.co.hitachi:xxx.h70.i.62510.1A.FF
```

raidcom delete host_grp

Deletes the specified host group or iSCSI target. This command also deletes the WWN/IQN or LUNs settings of the host registered for the host group or iSCSI target.

If the port type of the port with the specified host group is FIBRE or FCoE (Fibre Channel over Ethernet), this command deletes the host group, and WWN and LUNs settings of the host registered to the host group.

If the port type of the port with the specified iSCSI target is iSCSI, this command deletes the iSCSI target, and WWN and LUNs settings of the host (initiator) registered to the iSCSI target. However, if the host group ID of the host group or the target ID of the iSCSI target is 0, this command changes the settings back to the default. If the LUSE configuration is defined, this command releases the LUSE.

When the specified host group is associated with an SPM group, the following information is deleted after the host group is deleted:

- Association between the SPM group and the host group

- SPM settings of the WWN registered in the host group

- Registration of the WWN in the SPM group

If no WWN is registered in the SPM group, the SPM group itself is also deleted.

If the specified port does not exist, the command is rejected with EX_ENOOBJ.

If a specific host mode option is set for the specified host group, the command might be rejected with EX_CMDRJE. For details about specific host mode options, see the *Provisioning Guide* for your storage system.

Syntax

```
raidcom delete host_grp -port <port#> [<host group name>]
```

Options and parameters

-port <port#> [<host group name>]
Specifies the port number, host group ID, and host group name (iSCSI target name for iSCSI). If the host group name or the iSCSI target name is more than 64 characters, use the host group ID or the iSCSI target ID. For example:

- CL1-A-g (g is from 0 to 254)

- CL1-A Linux_X86

- CL1-A Target00

Examples

Deleting the host group of port CL4-E, host group #7.

```
# raidcom delete host_grp -port CL4-E-7
```

Deleting the host group of port CL4-E, host group name: host group of Win_export.

```
# raidcom delete host_grp -port CL4-E Win_export
```

Deleting the target of port CL4-E, iSCSI target name: Target01.

```
# raidcom delete host_grp -port CL4-E Target01
```

raidcom get host_grp

Displays the information about all host groups or an iSCSI target that are defined on the specified port.

Syntax

```
raidcom get host_grp {-port <port#> [<host group name>]| -allports} [-key <keyword>]
```

Options and parameters

-port <port#> [<host group name>]

Specifies the port number, host group ID, and host group name (iSCSI target name for iSCSI). If the host group name or the iSCSI target name is more than 64 characters, use the host group ID or the iSCSI target ID.

For example:

- CLI-A

- CL1-A-g (g is from 0 to 254)

- CL1-A Linux_X86

- CL1-A Target00

If you specified the above settings, the information about all host groups that are defined on the specified port is displayed regardless of the examples.

-allports

Displays the information about the host groups or the iSCSI target that is set to all ports.

[-key <keyword>]

Specifies the information to be displayed. Shows display keywords can be specified in *<keyword>* as follows:

- host_grp: Specifies the host group IDs that are not installed.

- server: Specifies the information about the server managed by Storage Advisor Embedded on which the host group or the iSCSI target is registered.

   This parameter can be specified only for VSP E series and VSP G130, G/F350, G/F370, G/F700, G/F900.

- detail: Specifies the option to display the resource group ID of the host groups or iSCSI targets. -key detail and -allports cannot be specified at the same time. When the -key detail is specified, host group information about both implemented and unimplemented host groups is displayed.

Example

Displaying the host group information set for the port: CL4-E (in case of FIBRE, FCoE (Fibre Channel over Ethernet), HNASS, or HNASU).

```
# raidcom get host_grp -port CL4-E
```

```
PORT   GID  GROUP_NAME    Serial#    HMD         HMO_BITs
CL4-E   0   Linux_x86     63528      LINUX/IRIX   2 13
CL4-E   1   Solaris       63528      SOLARIS      2 22
```

Chapter 5: Configuration setting commands

```
CL4-E    2   HP-UX           63528       HP-UX           40
CL4-E    3   Win_export      63528       WIN_EX          7
CL4-E    5   DEC             63528       TRU64           14
CL4-E    6   OpenVMS         63528       OVMS
CL4-E  254   RMTEST          63528       LINUX           7
```

```
# raidcom get host_grp -port CL4-E -key host_grp
```

```
PORT   GID  GROUP_NAME      Serial#     HMD         HMO_BITs
CL4-E  0    Linux_x86       63528       LINUX/IRIX  2 13
CL4-E  1    Solaris         63528       SOLARIS     2 22
CL4-E  2    HP-UX           63528       HP-UX       40
CL4-E  3    Win_export      63528       WIN_EX      7
CL4-E  5    DEC             63528       TRU64       14
CL4-E  6    OpenVMS         63528       OVMS
CL4-E  10   -               63528       -
CL4-E  11   -               63528       -
CL4-E  12   -               63528       -
CL4-E  13   -               63528       -
CL4-E  254  RMTEST          63528       LINUX       7
```

Displaying the iSCSI target information set for the port: CL4-E (in case of iSCSI)

```
# raidcom get host_grp -port CL4-E
```

```
PORT   GID  GROUP_NAME IQN       AMD  D Serial# HMD         HMO_BITs
CL4-E   0   Linux_x86  iqn.z1… CHAP S 63528   LINUX/IRIX 2 13
CL4-E   1   Solaris    iqn.z2… CHAP S 63528   SOLARIS    2 22
CL4-E   2   HP-UX      iqn.z3… CHAP S 63528   HP-UX      40
```

Displaying the information about the host groups or the iSCSI targets that are set for the port: CL4-E (in case of iSCSI and Fibre Channel)

```
# raidcom get host_grp -port CL4-E -key server
```

```
PORT      GID    GROUP_NAME  Serial#     SRVID        SRV_NAME
CL4-E     0      Linux_x86   63528         N              -
CL4-E     1      Solaris     63528         1           "server01"
CL4-E     2      HP-UX       63528         N              -
CL4-E     3      Win_export  63528         N              -
CL4-E     5      DEC         63528         N              -
CL4-E     6      OpenVMS     63528         N              -
CL4-E     254 RMTEST         63528         N              -
```

Chapter 5: Configuration setting commands

Displaying the resource group information about the host groups or the iSCSI targets that are set for the port: CL4-E (in case of iSCSI or Fibre Channel)

```
# raidcom get host_grp -port CL4-E -key detail
PORT GID RGID GROUP_NAME Serial# HMD HMO_BITs
CL4-E 0 0 "Linux_x86" 63528 LINUX/IRIX 2:13
CL4-E 1 0 "Solaris" 63528 SOLARIS 2:22
CL4-E 2 1 "HP-UX" 63528 HP-UX 40
CL4-E 3 0 "Win_export" 63528 WIN_EX 7
CL4-E 5 0 "DEC" 63528 TRU64 14
CL4-E 6 0 "OpenVMS" 63528 OVMS -
CL4-E 10 0 - 63528 - -
CL4-E 11 0 - 63528 - -
CL4-E 12 0 - 63528 - -
CL4-E 13 0 - 63528 - -
CL4-E 254 0 "RMTEST" 63528 LINUX -
```

Displaying the information about the host groups or the iSCSI targets that are set to all ports

```
# raidcom get host_grp -allports
```

```
PORT   GID    GROUP_NAME     Serial# HMD
CL1-A  0      Linux_x86       63528 LINUX/IRIX
CL1-A  1      Solaris         63528 SOLARIS
CL1-A  2      HP-UX           63528 HP-UX
CL1-A  3      Win_export      63528 WIN_EX
CL1-A  5      DEC             63528 TRU64
CL1-A  6      OpenVMS         63528 OVMS
CL1-A  254    RMTEST          63528 LINUX
CL1-B  0      Linux_x86       63528 LINUX/IRIX
CL1-B  1      Solaris         63528 SOLARIS
CL1-B  2      HP-UX           63528 HP-UX
CL1-B  3      Win_export      63528 WIN_EX
CL1-B  5      DEC             63528 TRU64
CL1-B  6      OpenVMS         63528 OVMS
CL1-B  254    RMTEST          63528 LINUX
```

```
# raidcom get host_grp -allports -key host_grp
```

```
PORT   GID    GROUP_NAME     Serial#    HMD
CL1-A  0      Linux_x86       63528     LINUX/IRIX
CL1-A  1      Solaris         63528     SOLARIS
CL1-A  2      HP-UX           63528     HP-UX
CL1-A  3      Win_export      63528     WIN_EX
CL1-A  5      DEC             63528     TRU64
```

Chapter 5: Configuration setting commands

```
CL1-A  6     OpenVMS         63528   OVMS
CL1-A  10    -               63528   -
CL1-A  11    -               63528   -
CL1-A  12    -               63528   -
CL1-A  13    -               63528   -
CL1-A  254   RMTEST          63528   LINUX
CL1-B  0     Linux_x86       63528   LINUX/IRIX
CL1-B  1     Solaris         63528   SOLARIS
CL1-B  2     HP-UX           63528   HP-UX
CL1-B  3     Win_export      63528   WIN_EX
CL1-B  5     DEC             63528   TRU64
CL1-B  6     OpenVMS         63528   OVMS
CL1-B  10    -               63528   -
CL1-B  11    -               63528   -
CL1-B  12    -               63528   -
CL1-B  13    -               63528   -
CL1-B  254   RMTEST          63528   LINUX
```

```
# raidcom get host_grp -allports -key server
```

```
PORT      GID      GROUP_NAME   Serial#     SRVID     SRV_NAME
CL1-A     0     Linux_x86      63528       N                -
CL1-A     1       Solaris        63528        1        "server01"
CL1-A     2       HP-UX          63528       N       -
CL1-A     3       Win_export     63528       N      -
CL1-A     5       DEC            63528       N       -
CL1-A     6       OpenVMS        63528       N      -
CL1-A     254   RMTEST         63528       N      -
CL1-B     0     Linux_x86      63528       N      -
CL1-B     1       Solaris        63528       N      -
CL1-B     2       HP-UX          63528       N       -
CL1-B     3       Win_export     63528       N       -
CL1-B     5       DEC            63528       N
-
```

Displaying only the host group IDs that are allocated to available resource groups for users.

Description of each column in output example:

PORT
> Displays the port number.

GID
> Displays the host group ID of a port.

GROUP_NAME
> Displays the host group name of a port or target alias.

Chapter 5: Configuration setting commands

For unimplemented host groups or target aliases, a hyphen (-) is displayed. If the `-key` option is specified, a double quotation (") is added to the beginning and end of the name of the implemented host groups or target aliases.

IQN

Displays the iSCSI Qualified Name of the port.

AMD

Displays the authentication mode of the iSCSI target.

- CHAP: CHAP authentication is enabled.

- NONE: Authentication is disabled.

- BOTH: Both CHAP authentication and connection by no-authentication are enabled.

D

Displays the direction of the authentication mode of the iSCSI target.

- S: Unidirectional (An initiator is recognized from the target side)

- D: Bidirectional (An initiator is recognized from the target side, and a target is recognized from the initiator)

Serial#

Product serial number.

- The serial number for VSP 5000 series is displayed with a "5" at the beginning ("512345" = serial number 12345).

- The serial number for VSP G1x00 and VSP F1500 is displayed with a "3" at the beginning ("312345" = serial number 12345).

HMD

Displays the following HOST MODE for the host adapter setting on host group.

- HP-UX, SOLARIS, AIX, WIN, LINUX/IRIX, TRU64, DYNIX, OVMS, NETWARE, HI-UX

- VMWARE, HP-XP, VMWARE_EX, WIN_EX, UVM

HMO_BITs

Displays the host mode options of the host group. For details, see the *Provisioning Guide* for the storage system.

If the host mode options are two or more, when case the `-key detail` option is specified, the host mode options which are set are displayed separately by a colon (:). If the `-key detail` option is specified, when there is no host mode option that is set, a hyphen (-) is displayed.

SRVID

Displays the server ID for the Storage Advisor Embedded on which the host groups or the iSCSI targets are registered in decimal. Displays "N" if they are not registered on the server. Displays a hyphen (-) if the microcode version does not support the server ID indication.

Chapter 5: Configuration setting commands

SRV_NAME
Displays the nickname of the server identified by SRVID. A double quotation (") is added at the beginning and end of the server nickname. A hyphen (-) is displayed if SRVID is" N "or hyphen (-).

RSGID
Displays the resource group ID of the host groups or the iSCSI targets.

raidcom modify host_grp

Sets a host mode to the host group or an iSCSI target on the specified port.

If the specified host group does not exist, the command is ignored.

In case of iSCSI, set the CHAP authentication (enable/disable, or unidirectional/bidirectional).

Syntax

```
raidcom modify host_grp -port <port#> [<host group name>]
   -host_mode <host mode> [-host_mode_opt <host mode option>... |
   -set_host_mode_opt <host mode option>... | -reset_host_mode_opt]
   [-authmethod {CHAP|NONE|BOTH}] [-mutual {enable|disable}]
```

Options and parameters

-port <port#>[<host group name>]
Specifies the port number, host group ID, and host group name (iSCSI target name for iSCSI). If the host group name or the iSCSI target name is more than 64 characters, use the host group ID or the iSCSI target name.

- CL1-A-g (g is from 0 to 254)
- CL1-A Linux_X86
- CL1-A Target00

-host_mode <host mode>
Specifies the host mode using the following strings and values. For the supported host modes, see the *Provisioning Guide* for the storage system.

- LINUX or IRIX (0x00)
- VMWARE (0x01)
- HP-UX (0x03)
- OVMS (0x05)
- TRU64 (0x07)
- SOLARIS (0x09)
- NETWARE (0x0a)

- WIN (0x0c)

- AIX (0x0f)

- VMWARE_EX (0x21)

- WIN_EX (0x2c)

- UVM (0x4c)

[-host_mode_opt <host mode option>...]

This parameter remains for the compatibility with the old version. Use the –
`set_host_mode_opt` option and `-reset_host_mode_opt` option.

[-set_host_mode_opt <host mode option> ...]

Specifies the host mode option. The other host mode options which you do not specify
is cleared. For details about the host mode option, see the *Provisioning Guide* for the
storage system.

[-reset_host_mode_opt]

Resets all host mode options. For details about the host mode option, see the
Provisioning Guide for the storage system.

[-authmethod {CHAP|NONE|BOTH}]

Specifies the CHAP authentication mode. Even if the CHAP user name has not been
specified yet, the CHAP authentication mode can be specified. This option must be
specified in parallel with specifying the host mode.

- CHAP: CHAP authentication is enabled.

- NONE: Authentication is disabled.

- BOTH: Both CHAP authentication and connection by no-authentication are
  enabled.

[-mutual {enable|disable}]

Specifies the CHAP authentication mode: unidirectional authentication, or bidirectional
authentication. Even if the CHAP authentication mode is specified to NONE, CHAP
authentication can be specified (when the authentication mode will be changed to
CHAP, or BOTH, the specified mode will be enabled). This option must be specified in
parallel with specifying the host mode.

- enable: specifies bidirectional CHAP authentication (an initiator is recognized from
  the target side, and a target is recognized from the initiator).

- disable: specifies unidirectional CHAP authentication (an initiator is recognized
  from the target).

Examples

Setting the host mode: HP-UX for the port: CL4-E, the host group: #2.

```
# raidcom modify host_grp -port CL4-E-2 -host_mode HP-UX
```

Setting the host mode: HP-UX and the host mode option: 2, 13 for the port: CL4-E, the host group: #2.

```
# raidcom modify host_grp -port CL4-E-2 -host_mode HP-UX -set_host_mode_opt 2 13
```

Clearing the host mode options of the host mode: HP-UX of the port: CL4-E, the host group: #2.

```
# raidcom modify host_grp -port CL4-E-2 -host_mode HP-UX -reset_host_mode_opt
```

Setting the host mode: HP-UX and the bidirectional CHAP authentication enabled for the port: CL4-E, the host group: #2.

```
# raidcom modify host_grp -port CL4-E-2 -host_mode HP-UX -authmethod CHAP -mutual
enable
```

raidcom add chap_user

This command sets the CHAP user name for the specified iSCSI target. Also this command registers the CHAP user name of the host on the initiator set in the specified iSCSI target. If the specified CHAP user name of the host on the initiator already exists, the registration is ignored.

Syntax

```
raidcom add chap_user -port <port#> [<host group name>]
    {-target_chap_user <user name>|-initiator_chap_user
    <user name>}
```

Options and parameters

-port <port#>[<host group name>]
> Specifies the port number, target ID, or iSCSI target name. For example:
>
> ▪ CL1-A-g (g is from 0 to 254)
> ▪ CL1-A Target00

-target_chap_user <user name>
> Specifies the CHAP user name for the iSCSI target. You can specify up to 223 characters. The maximum number of CHAP user names of the iSCSI target is 1 for each iSCSI target.
>
> For example: storage01

-initiator_chap_user <user name>
> Specifies the CHAP user name on the initiator that is set as the iSCSI target. You can specify up to 223 characters.
>
> For example: Linux-abc

Chapter 5: Configuration setting commands

Examples

To set the CHAP user name "storage01" to the iSCSI target whose port is CL4-E and the host group ID is 0:

```
# raidcom add chap_user -port CL4-E-0 -target_chap_user storage01
```

To set the CHAP user name "storage02" to the host whose iSCSI target name is target00 and port is CL4-E:

```
# raidcom add chap_user -port CL4-E Target00 -target_chap_user storage02
```

To register the CHAP user name "Linux-abc" to the host on the initiator whose iSCSI target name is target00 and port is CL4-E:

```
# raidcom add chap_user -port CL4-E Target00 -initiator_chap_user Linux-abc
```

raidcom delete chap_user

This command deletes the CHAP user name from the specified iSCSI target. Also this command deletes the CHAP user name from the host on the initiator set in the specified iSCSI target. If the specified CHAP user name does not exist, the command is ignored.

Syntax

```
raidcom delete chap_user -port <port#> [<host group name>]
    {-target_chap_user <user name>|-initiator_chap_user
    <user name>}
```

Options and parameters

-port <port#>[<host group name>]
    Specifies the port number, target ID, or iSCSI target name. For example:

- CL1-A-g (g is from 0 to 254)
- CL1-A Target00

-target_chap_user <user name>
    Specifies the CHAP user name of the iSCSI target.

    For example: storage01

-initiator_chap_user <user name>
    Specifies the CHAP user name of the host on the initiator.

    For example: Linux-abc

Examples

To delete the CHAP user name "storage01" from the iSCSI target whose port is CL4-E and the host group ID is 0:

```
# raidcom delete chap_user -port CL4-E-0 -target_chap_user storage01
```

To delete the CHAP user name "storage02" from the host whose iSCSI target name is target00 and port is CL4-E:

```
# raidcom delete chap_user -port CL4-E Target00 -target_chap_user storage02
```

To delete the CHAP user name "Linux-abc" from the host on the initiator whose iSCSI target name is target00 and port is CL4-E:

```
# raidcom delete chap_user -port CL4-E Target00 -initiator_chap_user Linux-abc
```

raidcom set chap_user

This command sets the password, called "secret", for the specified CHAP user. Both the host on the initiator and the host on the target are set by this command. To avoid a secret is given as an argument directly, the proper prompt is displayed to enter the secret. If the secret already exists for the specified CHAP user, the secret is overwritten.

Syntax

```
raidcom set chap_user -port <port#> [<host group name>] {-target_chap_user <user
name> -secret |-initiator_chap_user <user name> -secret}
```

Options and parameters

-port <port#>[<host group name>]
> Specifies the port number, target ID, or iSCSI target name. For example:

>> ▪ CL1-A-g (g is from 0 to 254)

>> ▪ CL1-A Target00

-target_chap_user <user name>
> Specifies the CHAP user name of the iSCSI target.

> For example: storage01

-initiator_chap_user <user name>
> Specifies the CHAP user name of the host on the initiator.

> For example: Linux-abc

-secret
> Displays the prompt for entering a secret.

Input characters as "secret", within the range of 12 characters to 32 characters, or an error occurs.

Examples

To set the "iSCSI-secret" for the "secret" to the user whose CHAP user name is storage01, port is CL4-E, and the target ID of the iSCSI target is 0:

```
# raidcom set chap_user -port CL4-E-0 -target_chap_user storage01 -secret
Enter Secret :
```

(Enter "iSCSI-secret" after the "Enter secret: " above. The entered characters are not displayed on the screen.)

To register the secret: Linux-secret for the CHAP user name "Linux-abc" from the host on the initiator whose iSCSI target name is target00 and port is CL4-E:

```
# raidcom set chap_user -port CL4-E Target00 -initiator_chap_user Linux-abc -secret
Enter Secret :
```

(Enter "Linux-secret" after the "Enter secret: " above. The entered characters are not displayed on the screen.)

raidcom reset chap_user

This command removes the secret from the specified CHAP user. Both the host on the initiator and the host on the target are set by this command.

Syntax

```
raidcom reset chap_user -port <port#> [<host group name>]
    {-target_chap_user <user name>|-initiator_chap_user
    <user name>}
```

Options and parameters

-port <port#>[<host group name>]
Specifies the port number, target ID, or iSCSI target name. For example:

- CL1-A-g (g is from 0 to 254)
- CL1-A Target00

-target_chap_user <user name>
Specifies the CHAP user name of the iSCSI target.

For example: storage01

**-initiator_chap_user

Specifies the CHAP user name of the host on the initiator.

For example: Linux-abc

Examples

To delete the secret for the CHAP user name "storage01" from the iSCSI target whose port is CL4-E and the target ID is 0:

```
# raidcom reset chap_user -port CL4-E-0 -target_chap_user storage01
```

To delete the secret for the CHAP user name "Linux-abc" from the initiator host whose port is CL4-E:

```
# raidcom reset chap_user -port CL4-E Target00 -initiator_chap_user Linux-abc
```

raidcom get chap_user

This command indicates the CHAP user name of the iSCSI target on the specified port and the CHAP user name of the host bus adapter on the initiator that is registered in the iSCSI target.

Syntax

```
raidcom get chap_user -port <port#> [<host group name>]
```

Options and parameters

-port <port#>[<host group name>]

Specifies the port number, target ID, or iSCSI target name. For example:

- CL1-A-g (g is from 0 to 254)
- CL1-A Target00

Example

To display the CHAP user name whose port is CL4-E and whose host group ID is 0:

```
# raidcom get chap_user -port CL4-E-0
```

```
PORT   GID GROUP_NAME      CHAP_user           Serial# WAY Sec
CL4-E   0  Linux_x86       raidmanager         63528   INI *
CL4-E   0  Linux_x86       raidmanager1        63528   INI *
CL4-E   0  Linux_x86       raidmanager2        63528   INI *
CL4-E   0  Linux_x86       oracle              63528   TAR *
```

Description of each column in the output example:

PORT
> Displays the port.

GID
> Displays the host group ID of the port.

GROUP_NAME
> Displays the iSCSI target name of the port.

CHAP_user
> It indicates the CHAP user name of the iSCSI target and the CHAP user name of the host bath adapter which is registered in the iSCSI target.

Serial#
> Displays the Seq#.

WAY
> Indicates whether the CHAP user name on the iSCSI target or the CHAP user name on the host bus adapter (initiator).
>
> - TAR: iSCSI target side
>
> - INI: Host bus adapter (initiator) side

Sec
> An asterisk (*) is always displayed.

raidcom add hba_wwn

Registers the WWN of the host adapter to the host group on the specified port to add a host.

If the specified WWN already exists, this command is ignored. (VSP and VSP G1x00 and VSP F1500 only) If the specified host group is associated with a n SPM group, the WWN of the host adapter is registered. After that, the WWN is registered in the SPM group, and then the SPM information is set for the WWN.

Syntax

```
raidcom add hba_wwn -port <port#> [<host group name>]
     -hba_wwn <WWN strings>
```

Options and parameters

-port <port#>[<host group name>]
> Specifies the port number, host group ID, and host group name. If the host group name is more than 64 characters, use the host group ID. For example:
>
> - CL1-A-g (g is from 0 to 254)
>
> - CL1-A Linux_X86

-hba_wwn <WWN strings>

>Specifies the WWN (16-digit hexadecimal value) of the host adapter. The 17th digit and later are ignored. The value can be split in units of 4 bytes by "," (comma). For example:

>- 210000e08b0256f8
>- 210000e0,8b0256f8

Example

Setting the WWN of host adapter: 210000e0,8b0256f8 to the port: CL4-E, the host group #0.

```
# raidcom add hba_wwn -port CL4-E-0 -hba_wwn 210000e0,8b0256f8
```

raidcom delete hba_wwn

Deletes the specified host (WWN) from the host group.

If the specified WWN does not exist, this command is ignored. (VSP and VSP G1x00 and VSP F1500 only) If the specified host group is associated with a n SPM group, the WWN of the host adapter is deleted, and then the SPM settings of the WWN and the registration of the WWN in the SPM group are deleted.

Syntax

```
raidcom delete hba_wwn -port <port#> [<host group name>]
      -hba_wwn <WWN strings>
```

Options and parameters

-port <port#>[<host group name>]

>Specifies the port number, host group ID, and host group name. If the host group name is more than 64 characters, use the host group ID. For example:

>- CL1-A-g (g is from 0 to 254)
>- CL1-A Linux_X86

-hba_wwn <WWN strings>

>Specifies the WWN (16-digit hexadecimal value) of the host adapter. The 17th digit and later are ignored. The value can be split in units of 4 bytes by "," (comma). For example:

>- 210000e08b0256f8
>- 210000e0,8b0256f8

Example

Deleting the WWN of host adapter: 210000e0,8b039800 set for the port: CL4-E, the host group #0.

```
# raidcom delete hba_wwn -port CL4-E-0 -hba_wwn 210000e0,8b039800
```

raidcom get hba_wwn

Displays the WWN of the HBA registered to the host group.

Syntax

```
raidcom get hba_wwn -port <port#> [<host group name>]
```

Options and parameters

-port <port#>[<host group name>]

Specifies s a port number, a host group ID, and a host group name. It cannot be specified when more than 64 characters is set for the host group name. Use the host group ID. If setting the host group ID or the host group name is omitted, the information about host group ID 0 is displayed. For example,

- CLI-A
- CL1-A-g (g is from 0 to 254)
- CL1-A Linux_X86

Example

Displaying the WWN of host adapter set for the port: CL4-E, the host group 0.

```
# raidcom get hba_wwn -port CL4-E-0
```

```
PORT  GID GROUP_NAME            HWWN Serial#      NICK_NAME
CL4-E   0  Linux_x86 210000e08b0256f8   63528 ORA_NODE0_CTL_0
CL4-E   0  Linux_x86 210000e08b039c15   63528 ORA_NODE1_CTL_0
```

Description of each column in output example:

PORT

Displays the port number.

GID

Displays the host group ID of a port.

GROUP_NAME

Displays the host group name of a port.

Chapter 5: Configuration setting commands

HWWN

Displays the WWN of registered host adapter.

Serial#

Displays the storage system serial number:

- The serial number for VSP 5000 series is displayed with a "5" at the beginning ("512345" = serial number 12345).

- The serial number for VSP G1x00 and VSP F1500 is displayed with a "3" at the beginning ("312345" = serial number 12345).

NICK_NAME

Displays the Nick Name of the WWN of host adapter.

raidcom reset hba_wwn

Deletes the nickname from the specified WWN on the specified port.

If there is no specified port, the command is rejected with EX_ENOOBJ.

Syntax

```
raidcom reset hba_wwn -port <port#>[<host group name>]
     -hba_wwn <WWN strings>
```

Options and parameters

-port <port#>[<host group name>]

Specifies the port number. For example:

- CL1-A-g (g: 0-254)

- CL1-A Linux_X86

-hba_wwn <WWN strings>

Specifies the WWN (16-digit hexadecimal value) of the host adapter. The 17th digit and later are ignored. The value can be split in units of 4 bytes by "," (comma). For example:

- 210000e08b0256f8

- 210000e0,8b0256f8

Example

Deleting the nickname that is given to the connection host "WWN:210000e0,8b0256f8" to which the port: CL4-E and host group #0 are set.

```
# raidcom reset hba_wwn -port CL4-E-0 -hba_wwn 210000e0,8b0256f8
```

raidcom set hba_wwn

Sets a nickname (maximum 64 characters) to the specified WWN on the specified port.

If the specified port does not exist, the command is rejected with EX_ENOOBJ.

If a nickname exists for the specified WWN, it is changed as NEW.

Syntax

```
raidcom set hba_wwn -port <port#>[<host group name>]
     -hba_wwn <WWN strings> -wwn_nickname <WWN Nickname>
```

Options and parameters

-port <port#>[<host group name>]
> Specifies the port number. For example:
>
> - CL1-A-g (g: 0-254)
> - CL1-A Linux_X86

-hba_wwn <WWN strings>
> Specifies the WWN (16-digit hexadecimal value) of the host adapter. The 17th digit and later are ignored. The value can be split in units of 4 bytes by "," (comma). For example:
>
> - 210000e08b0256f8
> - 210000e0,8b0256f8

-wwn_nickname <WWN Nickname>
> Specifies the nickname (maximum 64 characters) to be assigned to the WWN of a specific port.
>
> You cannot specify the same nickname to another WWN in the same port.

Example

Giving the nickname of "ORA_NODE0_CTL_0" to the connection host "WWN: 210000e0,8b0256f8" to which the port: CL4-E and host group #0 are set.

```
# raidcom set hba_wwn -port CL4-E-0 -hba_wwn 210000e0,8b0256f8
-wwn_nickname ORA_NODE0_CTL_0
```

raidcom add hba_iscsi

This command registers the iSCSI name (on the initiator) of the host bus adapter on the iSCSI target of the specified port in order to add hosts. If the specified iSCSI name already exists, the command is ignored.

Syntax

```
raidcom add hba_iscsi -port <port#> [<host group name>]
     -hba_iscsi_name <initiator iscsi name>
```

Options and parameters

-port <port#>[<host group name>]

Specifies the port number, target ID, or iSCSI target name. For example:

- CL1-A-g (g is from 0 to 254)

- CL1-A Target00

-hba_iscsi_name <initiator iscsi name>

Specifies the iSCSI name of the host bus adapter (initiator). You can specify within 223 characters. The maximum number of the host bus adapter is 255 for each port. For example:

- iqn.win2k8.example.of.iqn.form

- eui.0123456789ABCDEF

Notes:

- You cannot register multiple iSCSI names to the same port if the only difference among them is the case (lowercase or uppercase). For example, if iqn.win2k8.example.of.iqn.form is already registered as an iSCSI name, you cannot register iqn.win2k8.example.of.iqn.FORM as a different iSCSI name.

- If you input an iSCSI name in iqn format, use lowercase letters only.

- If you input an iSCSI name in eui format, use lowercase letter s for "eui.", and use uppercase letters for the following hexadecimal value.

Examples

To set the iSCSI name "iqn.win2k8.example.of.iqn.form" of the host bus adapter whose port is CL4-E and whose target ID is 0:

```
# raidcom add hba_iscsi -port CL4-E-0 -hba_iscsi_name iqn.win2k8.example.of.iqn.form
```

To set the iSCSI name "eui.0123456789ABCDEF" of the host bus adapter whose port is CL4-E and whose iSCSI target name is Target00:

```
# raidcom add hba_iscsi -port CL4-E Target00 -hba_iscsi_name eui.0123456789ABCDEF
```

raidcom delete hba_iscsi

Supported storage systems:

- VSP 5000 series

- VSP G1x00 and VSP F1500

- VSP E series

- VSP G130, G/F350, G/F370, G/F700, G/F900

- VSP G200, G400, G600, G800 and VSP F400, F600, F800

This command deletes the host (initiator iSCSI name) from the host group. If the specified initiator iSCSI name does not exist, the command is ignored.

Syntax

```
raidcom delete hba_iscsi -port <port#> [<host group name>]
     -hba_iscsi_name <initiator iscsi name>
```

Options and parameters

-port <port#>[<host group name>]
        Specifies the port number, target ID, or iSCSI target name. For example:

- CL1-A-g (g is from 0 to 254)

- CL1-A Target00

-hba_iscsi_name <initiator iscsi name>
        Specifies the iSCSI name of the host bus adapter (initiator). For example:

- iqn.win2k8.example.of.iqn.form

- eui.0123456789ABCDEF

Examples

To delete the iSCSI name "iqn.win2k8.example.of.iqn.form" of the host bus adapter from the host group whose port is CL4-E and the target ID is 0:

```
# raidcom delete hba_iscsi -port CL4-E-0 -hba_iscsi_name
iqn.win2k8.example.of.iqn.form
```

To delete the iSCSI name "eui.0123456789ABCDEF" of the host bus adapter from the host group whose port is CL4-E and the iSCSI target name is Target00:

```
# raidcom delete hba_iscsi -port CL4-E Target00 -hba_iscsi_name eui.0123456789ABCDEF
```

raidcom set hba_iscsi

Supported storage systems:

- VSP 5000 series

- VSP G1x00 and VSP F1500

- VSP E series

- VSP G130, G/F350, G/F370, G/F700, G/F900

- VSP G200, G400, G600, G800 and VSP F400, F600, F800

This command sets a nickname for the iSCSI name of the initiator on the specified port. If the specified port does not exist, the command is rejected with EX_ENOOBJ. If the nickname already exists for the specified initiator iSCSI name, the existing nickname is overwritten.

Syntax

```
raidcom set hba_iscsi -port <port#>[<host group name>]
     -hba_iscsi _name <initiator iscsi name>
     -iscsi_nickname <initiator iscsi Nickname>
```

Options and parameters

-port <port#>[<host group name>]

Specifies the port number, target ID, or iSCSI target name. For example:

- CL1-A-g g is from 0 to 254)

- CL1-A Target00

-hba_iscsi_name <initiator iscsi name>

Specifies the iSCSI name of the host bus adapter (initiator). For example:

- iqn.win2k8.example.of.iqn.form

- eui.0123456789ABCDEF

-iscsi_nickname <initiator iscsi Nickname>

Specifies the nickname given to the iSCSI name of the initiator. You can specify up to 32 characters.

Example

This command sets the nickname: ORA_NODE0_CTL_0 for the connection host iSCSI name: iqn.win2k8.example.of.iqn.form whose port is CL4-E and the target ID is 0.

```
# raidcom set hba_iscsi -port CL4-E-0 -hba_iscsi_name iqn.win2k8.example.of.iqn.form -
iscsi_nickname ORA_NODE0_CTL_0
```

raidcom reset hba_iscsi

Supported storage systems:

- VSP 5000 series

- VSP G1x00 and VSP F1500

- VSP E series

- VSP G130, G/F350, G/F370, G/F700, G/F900

- VSP G200, G400, G600, G800 and VSP F400, F600, F800

This command removes the nickname from the iSCSI name of the initiator on the specified port. When the specified port does not exist, the command will be rejected with EX_ENOOBJ.

Syntax

```
raidcom reset hba_iscsi -port <port#> [<host group name>] -hba_iscsi_name <initiator iscsi name>
```

Options and parameters

-port <port#>[<host group name>]

Specifies the port number, target ID, or iSCSI target name. For example:

- CL1-A-g g is from 0 to 254)

- CL1-A Linux_X86

-hba_iscsi_name <initiator iscsi name>

Specifies the iSCSI name of the host bus adapter (initiator). For example:

- iqn.win2k8.example.of.iqn.form

- eui.0123456789ABCDEF

Examples

This command removes the nickname from the connection host iSCSI name: iqn.win2k8.example.of.iqn.form whose port ID is CL4-E and the host group is 0.

```
# raidcom reset hba_iscsi -port CL4-E-0 -hba_iscsi_name iqn.win2k8.example.of.iqn.form
```

raidcom get hba_iscsi

Supported storage systems:

- VSP 5000 series
- VSP G1x00 and VSP F1500
- VSP E series
- VSP G130, G/F350, G/F370, G/F700, G/F900
- VSP G200, G400, G600, G800 and VSP F400, F600, F800

This command displays the iSCSI name of the host bus adapter on the initiator for each iSCSI target, which is registered as an iSCSI target.

Syntax

```
raidcom get hba_iscsi -port <port#> [<host group name>]
```

Options and parameters

-port <port#>[<host group name>]

Specifies the port number, target ID, or iSCSI target name. For example:

- CL1-A-g (g is from 0 to 254)
- CL1-A Target00

Examples

To display the iSCSI name of the host bus adapter whose port is CL4-E and the iSCSI target ID is 0:

```
# raidcom get hba_iscsi -port CL4-E-0
```

```
PORT   GID GROUP_NAME      IQN            Serial# NICK_NAME
CL4-E    0  Linux_x86      iqn.z1…          63528  ORA_NODE0_CTL_0
CL4-E    0  Linux_x86      iqn.z2…          63528  ORA_NODE1_CTL_0
```

Description of each column in the output example:

PORT

Displays the port.

GID

Displays the target ID of the port.

GROUP_NAME

Displays the iSCSI target name of the port.

IQN
> Displays the iSCSI name of the registered host bus adapter.

Serial#
> Displays the product serial number.

NICK_NAME
> Displays the nickname of the iSCSI name for the host bus adapter.

raidcom add external_iscsi_name

Supported storage systems:

- VSP 5000 series

- VSP G1x00 and VSP F1500

- VSP E series

- VSP G/F350, G/F370, G/F700, G/F900

- VSP G200, G400, G600, G800 and VSP F400, F600, F800

This command registers the iSCSI name of the iSCSI target on the external storage system to the iSCSI port of the local storage system.

When the iSCSI name has been registered in the iSCSI port of the specified local storage system, this command sets the CHAP authentication mode, and enables or disables the mutual CHAP authentication mode. When the iSCSI name has been registered in the port other than the iSCSI port of the specified local storage system, the command registers the iSCSI name to the iSCSI port of the specified local storage system. In this case, the information about the CHAP authentication mode and the mutual CHAP authentication which have been set to the iSCSI target is used by the local storage system.

If the specified iSCSI port does not exist, the command is rejected with EX_ENOOBJ.

> **Note:** If a failure occurs in a processor on an MP blade or a controller, this command cannot be run even on an MP blade or a controller on which no failure occurs.

Syntax

```
raidcom add external_iscsi_name -port <port#> -iscsi_name
   <external iscsi name> -address <external IP address>
   [-authmethod {CHAP|NONE}] [-mutual {enable|disable}]
   [-tcp_port <value>]
   [-iscsi_virtual_port_id <iSCSI virtual port ID>]
```

Options and parameters

-port <port#>
> Specifies the port number of the local storage system. For example:
>
> ▪ CL1-A

-iscsi_name <external iscsi name>
> Specifies the iSCSI name of the iSCSI target on the external storage system in an iqn format or an eui format.
>
> ▪ iqn format: `iqn.` and the subsequent maximum 219 characters.
>
> ▪ eui format: `eui.` and the subsequent 16 characters in hexadecimal notation.

-address <external IP address>
> Specifies the IP address of the iSCSI target on the external storage system. An IPv4 address or an IPv6 address can be specified.
>
> Note that the following addresses cannot be specified when you specify the IPv4 address:
>
> ▪ Network address (for example: 192.168.10.0, 0.120.10.1)
>
> ▪ Broadcast address (for example: 255.255.255.255, 10.1.255.255)
>
> ▪ Loop back address (for example: 127.0.0.1)
>
> Note that the following addresses cannot be specified when you specify the IPv6 address:
>
> ▪ Unspecified address (for example: ::)
>
> ▪ Multicast address (for example: ff:1024:1215::01)
>
> ▪ Loop back address (for example: ::1)

[-authmethod {CHAP|NONE}]
> Can configure the CHAP authentication mode. Even if the CHAP user name has not been specified yet, the CHAP authentication mode can be configured.
>
> ▪ CHAP: CHAP authentication mode is enabled.
>
> ▪ NONE: The authentication mode is not configured.
>
> If this option is omitted, the current setting value is maintained. The initial value is "NONE".

[-mutual {enable|disable}]
> Enables or disables the mutual CHAP authentication mode. Even if -authmethod option is configured to NONE, the mutual CHAP authentication mode can be enabled or disabled. When -authmethod option changes to CHAP from NONE, this option setting becomes to be enabled.
>
> ▪ enable: The mutual CHAP authentication is enabled. An iSCSI target recognizes the iSCSI initiator and vice versa.
>
> ▪ disable: The mutual CHAP authentication is disabled. An iSCSI target recognizes the iSCSI target.

Chapter 5: Configuration setting commands

If this option is omitted, the current setting value is maintained. The initial value is "disable".

[-tcp_port <value>]

Specifies the TCP port number of the iSCSI target on the external storage system. If this option is omitted, the TCP port number of the iSCSI target port which is specified with -port option is configured.

[-iscsi_virtual_port_id <iSCSI virtual port ID>]

Specifies the iSCSI virtual port ID when the iSCSI virtual port mode is enabled. If this option is omitted, 0 is assumed.

An error occurs if you specify this option in any of the following cases:

- The virtual port of the specified iSCSI virtual port ID is not found.

- The iSCSI virtual port mode is disabled.

- The device does not support the iSCSI virtual port mode.

Examples

When you register the iSCSI name of the iSCSI target (iSCSI name: iqn.z1, IP address: 158.214.135.100) on the external storage system to the iSCSI port: CL4-E of the local storage system:

```
# raidcom add external_iscsi_name -port CL4-E -iscsi_name iqn.z1 -address
158.214.135.100
```

When you register the iSCSI name of the iSCSI target (iSCSI name: iqn.z1, IP address: 158.214.135.100) on the external storage system to the iSCSI port: CL4-E, virtual port ID: 1 of the local storage system:

```
#raidcom add external_iscsi_name -port CL4-E -iscsi_name iqn.z1 -address
158.214.135.100 -iscsi_virtual_port_id 1
```

When you change the CHAP authentication mode of the iSCSI target (iSCSI name: iqn.z1, IP address: 158.214.135.100) on the external storage system which is registered in the iSCSI port: CL4-E of the local storage system to "CHAP" and the mutual CHAP authentication mode to "enable":

```
# raidcom add external_iscsi_name -port CL4-E -iscsi_name iqn.z1 -address
158.214.135.100 -authmethod CHAP -mutual enable
```

When you register all iSCSI names of the iSCSI target on the external storage system which is registered in the port: CL2-E of the local storage system to the port: CL4-E:

```
# raidcom get external_iscsi_name | rmawk @1-eq:CL2-E exe="raidcom add
external_iscsi_name -port CL4-E -address @3 -iscsi_name @4"
```

When you register all iSCSI names of the iSCSI target on the external storage system which is registered in the port: CL2-E of the local storage system to the port: CL4-E. At this time, configure whether to enable the CHAP authentication mode and the mutual CHAP authentication mode:

```
# raidcom get external_iscsi_name | rmawk @1-eq:CL2-E | @7-eq:D exe="raidcom add
external_iscsi_name -port CL4-E -address @3 -iscsi_name @4 -authmethod @6!u -mutual
enable" -n exe="raidcom add external_iscsi_name -port @1 -address @3 -iscsi_name @4 -
authmethod @6!u -mutual disable"
```

When you search the iSCSI name of the iSCSI target which exists in the iSCSI port (IP address: 10.213.60.111) on the external storage system, and register the detected iSCSI name to the iSCSI port: CL4-E of the local storage system.

```
# raidcom discover external_iscsi_name -port CL4-E -address 10.213.60.111 | rmawk @5-
eq:N exe="raidcom add external_iscsi_name -port @1 -address @3 -iscsi_name @6"
```

raidcom delete external_iscsi_name

Supported storage systems:

- VSP 5000 series
- VSP G1x00 and VSP F1500
- VSP E series
- VSP G/F350, G/F370, G/F700, G/F900
- VSP G200, G400, G600, G800 and VSP F400, F600, F800

This command deletes the iSCSI name of the iSCSI target on the external storage system which is registered in the specified iSCSI port of the local storage system.

If the specified iSCSI port of the local storage system does not exist, the command is rejected with EX_ENOOBJ. If the iSCSI name of the iSCSI target on the specified external storage system does not registered in the iSCSI port of the specified local storage system, this command is ignored.

📄 **Note:** If a failure occurs in a processor on an MP blade or a controller, this command cannot be run even on an MP blade or a controller on which no failure occurs.

Syntax

```
raidcom delete external_iscsi_name -port <port#> -iscsi_name
   <external iscsi name> -address <external IP address>
   [-iscsi_virtual_port_id <iSCSI virtual port ID>]
```

Options and parameters

-port <port#>

Specifies the port number of the local storage system. For example:

- CL1-A

-iscsi_name <external iscsi name>

Specifies the iSCSI name of the iSCSI target on the external storage system in an iqn format or an eui format.

- iqn format: `iqn.` and the subsequent maximum 219 characters.

- eui format: `eui.` and the subsequent 16 characters in hexadecimal notation.

-address <external IP address>

Specifies the IP address of the iSCSI target on the external storage system. An IPv4 address or an IPv6 address can be specified.

Note that the following addresses cannot be specified when you specify the IPv4 address:

- Network address (for example: 192.168.10.0, 0.120.10.1)

- Broadcast address (for example: 255.255.255.255, 10.1.255.255)

- Loop back address (for example: 127.0.0.1)

Note that the following addresses cannot be specified when you specify the IPv6 address:

- Unspecified address (for example: ::)

- Multicast address (for example: ff:1024:1215::01)

- Loop back address (for example: ::1)

[-iscsi_virtual_port_id <iSCSI virtual port ID>]

Specifies the iSCSI virtual port ID when the iSCSI virtual port mode is enabled. If this option is omitted, 0 is assumed.

An error occurs if you specify this option in any of the following cases:

- The virtual port of the specified iSCSI virtual port ID is not found.

- The iSCSI virtual port mode is disabled.

- The device does not support the iSCSI virtual port mode.

Examples

When you delete the iSCSI name: iqn.z1 of the iSCSI target (IP address: 158.214.135.100) on the external storage system from the iSCSI port: CL4-E of the local storage system:

```
# raidcom delete external_iscsi_name -port CL4-E -iscsi_name iqn.z1 -address
158.214.135.100
```

When you delete the iSCSI name: iqn.z1 of the iSCSI target (IP address: 158.214.135.100) on the external storage system from the iSCSI port: CL4-E, virtual port ID: 1 of the local storage system:

```
#raidcom delete external_iscsi_name -port CL4-E -iscsi_name iqn.z1 -address
158.214.135.100 -iscsi_virtual_port_id 1
```

When you delete all iSCSI names of the iSCSI target on the external storage system which are registered in the iSCSI port: CL4-E of the local storage system:

```
# raidcom get external_iscsi_name | rmawk @1-eq:CL4-E exe="raidcom delete
external_iscsi_name -port @1 -address @3 -iscsi_name @4"
```

When you delete all iSCSI names of the iSCSI target which exists on the iSCSI port (IP address: 158.214.135.100) of the external storage system from the iSCSI port of the local storage system:

```
# raidcom get external_iscsi_name | rmawk @3-eq:158.214.135.100 exe="raidcom delete
external_iscsi_name -port @1 -address @3 -iscsi_name @4"
```

raidcom get external_iscsi_name

Supported storage systems:

- VSP 5000 series
- VSP G1x00 and VSP F1500
- VSP E series
- VSP G/F350, G/F370, G/F700, G/F900
- VSP G200, G400, G600, G800 and VSP F400, F600, F800

This command displays the iSCSI name of the iSCSI target on the external storage system which is registered in the iSCSI port of the specified local storage system.

If the iSCSI port does not exist in the specified local storage system, the command is rejected with EX_ENOOBJ.

Only the iSCSI names registered in the port to which the user who executes the command can refer are output. For details about the port to which the user can refer, see the descriptions of the relationship of the resource group and the command operation in the *Command Control Interface User and Reference Guide*.

Syntax

```
raidcom get external_iscsi_name [-port <port#> [-iscsi_virtual_port_id <iSCSI virtual
port ID>]]
```

Options and parameters

[-port <port#>]

> Specifies the port number of the local storage system. For example:

> ▪ CL1-A

[-iscsi_virtual_port_id <iSCSI virtual port ID>]

> Specifies the iSCSI virtual port ID when the iSCSI virtual port mode of the specified port is enabled. Only the iSCSI initiator of the port that can be referenced by the user who executes the command is output. If the microcode or firmware version does not support this option, nothing is output.

Examples

Displays all iSCSI targets of the external storage system which are registered in all iSCSI ports of the local storage system.

```
# raidcom get external_iscsi_name
```

```
PORT   Serial# IP_ADDR        IQN       WWN(pseudo)      AMD  D CHAP_user    Sec
ISCSI_VP_ID IP_PORT#
CL4-E  63528 158.214.135.100 iqn.z1    50060e80070a3640 CHAP D Win_SQL_EX
*           -        3260
CL2-E  63528 158.214.135.100 iqn.z2    50060e80070a3641 CHAP S -
-           -        3260
CL2-E  63528 158.214.135.102 iqn.z3    50060e80070a3642 CHAP S -
-           -        3260
CL4-E  63528 158.214.135.100 iqn.z2    50060e80070a3643 CHAP S -
-           -        3260
CL4-E  63528 158.214.135.102 iqn.z3    50060e80070a3644 CHAP S -
-           -        3260
CL4-E  63528 158.214.135.102 iqn.z4    50060e80070a3645 NONE S -
-           -        3260
CL4-E  63528 158.214.135.102 iqn.z5    50060e80070a3646 NONE S -
-           -        3260
```

Displays all iSCSI names of the iSCSI target on the external storage system which are registered in iSCSI port: CL4-E of the local storage system.

```
# raidcom get external_iscsi_name -port CL4-E
```

```
PORT   Serial# IP_ADDR        IQN       WWN(pseudo)      AMD  D CHAP_user    Sec
ISCSI_VP_ID IP_PORT#
CL4-E  63528 158.214.135.100 iqn.z1    50060e80070a3640 CHAP D Win_SQL_EX
*           -        3260
CL4-E  63528 158.214.135.100 iqn.z2    50060e80070a3643 CHAP S -
-           -        3260
CL4-E  63528 158.214.135.102 iqn.z3    50060e80070a3644 CHAP S -
```

Chapter 5: Configuration setting commands

```
-            -      3260
CL4-E  63528 158.214.135.102 iqn.z4    50060e80070a3645 NONE S -
-            -      3260
CL4-E  63528 158.214.135.102 iqn.z5    50060e80070a3646 NONE S -
-            -      3260
```

Displays all iSCSI names of the iSCSI target of the external storage system registered in the iSCSI port: CL2-E, virtual port ID: 1 of the local storage system.

```
#raidcom get external_iscsi_name -port CL2-E -iscsi_virtual_port_id 1
```

```
PORT Serial# IP_ADDR IQN WWN(pseudo) AMD D CHAP_user Sec ISCSI_VP_ID IP_PORT#
CL2-E 63528 158.214.135.100 iqn.z2 50060e80070a3641 CHAP S - - 1     3260
CL2-E 63528 158.214.135.102 iqn.z3 50060e80070a3642 CHAP S - - 1     3260
```

Description of each column in output example:

PORT

Displays the port number.

Serial#

Displays the product serial number.

IP_ADDR

Displays the IP address of the iSCSI target on the external storage system.

IQN

Displays the iSCSI name of the iSCSI target on the external storage system.

WWN(pseudo)

Displays the pseudo WWN of the iSCSI target on the external storage system. The pseudo WWN matches the iSCSI name of the iSCSI target and the IP address on the external storage system. The pseudo WWN is managed by each storage system. Therefore, if the iSCSI target on an external storage is shared within multiple storage systems, the pseudo WWN which is corresponded with an iSCSI target depends on the storage system.

AMD

Displays the authentication mode of the iSCSI target on the external storage system.

- CHAP: CHAP authentication is enabled.

- NONE: The authentication mode is not configured.

D

Displays the direction of the authentication mode of the iSCSI target.

- S: single directional (An iSCSI target recognizes the iSCSI target.)

- D: dual-directional (An iSCSI target recognizes the iSCSI initiator and vice versa.)

Chapter 5: Configuration setting commands

CHAP_user

Displays the CHAP user name of the iSCSI target on the external storage system. If the CHAP user name is not set, a hyphen (-) is displayed.

Sec

Displays an asterisk (*) when the secret is set to the iSCSI target on the external storage system. Otherwise, a hyphen (-) is displayed.

ISCSI_VP_ID

Displays the virtual port number when the virtual port mode is enabled. A hyphen (-) is displayed when the virtual port mode is disabled.

IP_PORT#

Displays the TCP port number of the iSCSI target on the external storage system.

raidcom get initiator_iscsi_name

Supported storage systems:

- VSP 5000 series

- VSP G1x00 and VSP F1500

- VSP E series

- VSP G/F350, G/F370, G/F700, G/F900

- VSP G200, G400, G600, G800 and VSP F400, F600, F800

This command displays the iSCSI initiator of the iSCSI port on the specified local storage system.

If the iSCSI port does not exist in the specified local storage system, the command is rejected with EX_ENOOBJ.

Only iSCSI initiators of the port to which the user who executes the command can refer are output. For details about the port to which the user can refer, see the descriptions of the relationship of the resource group and the command operation in the *Command Control Interface User and Reference Guide*.

Syntax

```
raidcom get initiator_iscsi_name -port <port#> [-iscsi_virtual_port_id <iSCSI virtual
port ID>]
```

Options and parameters

-port <port#>

Specifies the port number of the local storage system. For example:

- CL1-A

Chapter 5: Configuration setting commands

[-iscsi_virtual_port_id <iSCSI virtual port ID>]
> Specifies the iSCSI virtual port ID when the iSCSI virtual port mode of the specified port is enabled. Only the iSCSI initiator of the port that can be referenced by the user who executes the command is output. If the microcode version does not support this option, nothing is output.

Examples

Displays the iSCSI initiator of the iSCSI port: CL4-E on the local storage system.

```
# raidcom get initiator_iscsi_name -port CL4-E
```

```
PORT    Serial# IP_ADDR         IQN       IP_PORT#  CHAP_user    Sec  ISCSI_VP_ID
CL4-E     63528 158.214.197.100 iqn.z1       3260  Elun_INI_4E  *            -
```

Displays the iSCSI initiator of the iSCSI port: CL2-E on the local storage system.

```
# raidcom get initiator_iscsi_name -port CL2-E
```

```
PORT    Serial# IP_ADDR         IQN       IP_PORT#  CHAP_user    Sec  ISCSI_VP_ID
CL2-E     63528 158.214.197.101 iqn.zx       3260  Elun_INI_2E  *            -
```

Displays the iSCSI initiator of the iSCSI port: CL2-E, virtual port: 1 on the local storage system.

```
# raidcom get initiator_iscsi_name -port CL2-E -iscsi_virtual_port_id 1
```

```
PORT    Serial# IP_ADDR         IQN       IP_PORT#  CHAP_user    Sec  ISCSI_VP_ID
CL2-E     63528 158.214.197.102 iqn.zy       3260 Elun_INI_2E *            1
```

Description of each column in output example:

PORT
> Displays the port number.

Serial#
> Displays the product serial number.

IP_ADDR

Displays the IP address which is set to the iSCSI initiator of the iSCSI port on the local storage system. See the table below for more information. To display the details of the IP address, use the `-key opt` option of the **raidcom get port** command.

IPv6 mode status	Address acquisition mode	Address status	IP address to display
Invalid	-	-	IP address of IPv4
Valid	Manual acquisition (MM)	-	Link local address
	Automatic acquisition (AM)	Valid (VAL)	
		Invalid (INV)	-(hyphen)
		Acquiring (ACQ)	
		Duplication (DUP)	

IQN

Displays the iSCSI name which is set to the iSCSI initiator of the iSCSI port on the local storage system.

IP_PORT#

Displays the TCP port number of the iSCSI target which is registered in the iSCSI port on the local storage system.

CHAP_user

Displays the CHAP user name which is set in the iSCSI initiator of the iSCSI port on the local storage system. If the CHAP user name is not set, a hyphen (-) is displayed.

Sec

Displays an asterisk (*) when the secret is set to the iSCSI initiator of the iSCSI port on the local storage system. Otherwise, a hyphen (-) is displayed.

ISCSI_VP_ID

Displays the virtual port number when the virtual port mode is enabled. A hyphen (-) is displayed when the virtual port mode is disabled.

raidcom discover external_iscsi_name

Supported storage systems:

- VSP 5000 series

- VSP G1x00 and VSP F1500

- VSP E series

- VSP G/F350, G/F370, G/F700, G/F900

- VSP G200, G400, G600, G800 and VSP F400, F600, F800

From the iSCSI port of the local storage system, this command searches the iSCSI targets which is registered in the port of the external storage system, and then displays the iSCSI name of the iSCSI target.

To execute this command, the host which executes the command has to support IPv6. If the host does not support an IPv6, the command is rejected with EX_ENOSUP.

Syntax

```
raidcom discover external_iscsi_name -port <port#>
   -address <external IP address> [-tcp_port <value>]
   [-iscsi_virtual_port_id <iSCSI virtual port ID>]
```

Options and parameters

-port <port#>

Specifies the port number of the local storage system. For example:

- CL1-A

-address <external IP address>

Specifies the IP address of the iSCSI target on the external storage system. An IPv4 address or an IPv6 address can be specified.

Note that the following addresses cannot be specified when you specify the IPv4 address:

- Network address (for example: 192.168.10.0, 0.120.10.1)

- Broadcast address (for example: 255.255.255.255, 10.1.255.255)

- Loop back address (for example: 127.0.0.1)

Note that the following addresses cannot be specified when you specify the IPv6 address:

- Unspecified address (for example: ::)

- Multicast address (for example: ff:1024:1215::01)

- Loop back address (for example: ::1)

Chapter 5: Configuration setting commands

[-tcp_port <value>]

Specifies the TCP port number of the iSCSI target on the external storage system. If this option is omitted, the TCP port number of the iSCSI target port which is specified with -port option is configured.

[-iscsi_virtual_port_id <iSCSI virtual port ID>]

Specifies the iSCSI virtual port ID when the iSCSI virtual port mode is enabled. If this option is omitted, 0 is assumed.

An error occurs if you specify this option in any of the following cases:

- The virtual port of the specified iSCSI virtual port ID is not found.

- The iSCSI virtual port mode is disabled.

- The device does not support the iSCSI virtual port mode.

Examples

Searches the iSCSI name of the iSCSI target which is registered in the iSCSI port (IP address: 10.213.60.111) on the external storage system from the iSCSI port: CL4-E of the local storage system, and then displays the iSCSI name:

```
# raidcom discover external_iscsi_name -port CL4-E -address 10.213.60.111
```

```
PORT    Serial# IP_ADDR         IP_PORT# R  IQN  ISCSI_VP_ID
CL4-E    63528 10.213.60.111    3260  N  iqn.z1         -
CL4-E    63528 10.213.60.111    3260  N  iqn.z2         -
```

Searches the iSCSI name of the iSCSI target which is registered in the iSCSI port (IP address: 10.213.60.112) on the external storage system from the iSCSI port: CL4-E, iSCSI virtual port ID: 1 of the local storage system, and then displays the iSCSI name:

```
#raidcom discover external_iscsi_name -port CL4-E -address 10.213.60.112 -
iscsi_virtual_port_id 1
```

```
PORT Serial# IP_ADDR IP_PORT# R IQN ISCSI_VP_ID
CL4-E 63528 10.213.60.112 3260 N iqn.z1 1
CL4-E 63528 10.213.60.112 3260 N iqn.z2 1
```

Description of each column in output example:

PORT

Displays the port number.

Serial#

Displays the product serial number.

IP_ADDR

Displays the IP address of the iSCSI target on the external storage system.

IP_PORT#

Displays the TCP port number of the iSCSI target on the external storage system.

R

Displays whether or not the searched iSCSI target of the external storage system has been registered in the iSCSI port of the local storage system.

- Y: The iSCSI target has been registered in the iSCSI port.

- N: The iSCSI target has not been registered in the iSCSI port.

IQN

Displays the searched iSCSI name of the iSCSI target on the external storage system.

ISCSI_VP_ID

Displays the virtual port number when the virtual port mode is enabled. A hyphen (-) is displayed when the virtual port mode is disabled.

raidcom check external_iscsi_name

Supported storage systems:

- VSP 5000 series

- VSP G1x00 and VSP F1500

- VSP E series

- VSP G/F350, G/F370, G/F700, G/F900

- VSP G200, G400, G600, G800 and VSP F400, F600, F800

This command attempts to login to the iSCSI target on the external storage system which has been registered in the local storage system, and then displays the result of the login.

To execute this command, the host which executes the command has to support IPv6. If the host does not support an IPv6, the command is rejected with EX_ENOSUP.

Syntax

```
raidcom check external_iscsi_name [-port <port#> [-iscsi_name <external iscsi name> -
address <external IP address>]
   [-iscsi_virtual_port_id <iSCSI virtual port ID>]]
```

Options and parameters

[-port <port#>]

Specifies the port number of the local storage system. For example:

- CL1-A

[-iscsi_name <external iscsi name> -address <external IP address>]

Specify this option to display the login result when attempting to log in to the iSCSI target of the specified external storage system.

For -iscsi_name <external iscsi name>, specify the iSCSI name of the iSCSI target on the external storage system in the iqn or eui format:

- iqn format: iqn. followed by up to 219 characters.

- eui format: eui. followed by 16 characters in hexadecimal notation.

For -address <external IP address>, specify the IP address of the iSCSI target on the external storage system. An IPv4 address or an IPv6 address can be specified. If the specified iSCSI target is not registered in the specified port or iSCSI virtual port, the command is rejected with EX_ENOOBJ.

[-iscsi_virtual_port_id <iSCSI virtual port ID>]

Specifies the iSCSI virtual port ID when the iSCSI virtual port mode is enabled. When the iSCSI virtual port mode is enabled, if both this option and the -iscsi_name <external iscsi name> -address <external IP address> option are not specified, information including all iSCSI virtual port IDs is displayed.

An error occurs if you specify this option in any of the following cases:

- The virtual port of the specified iSCSI virtual port ID is not found.

- The iSCSI virtual port mode is disabled.

- The device does not support the iSCSI virtual port mode.

Examples

Attempts to log in to the iSCSI target of all external storage systems which are registered in the iSCSI port of the local storage system, and then displays the result of login:

```
# raidcom check external_iscsi_name
```

```
PORT    Serial# IP_ADDR         IQN         IP_PORT# AMD  D CHAP_user    Sec LOGIN
ISCSI_VP_ID
CL4-E    63528 158.214.135.100 iqn.z1        3260 CHAP D Win_SQL_EX   *   OK
-
CL2-E    63528 158.214.135.100 iqn.z2        3260 CHAP S -            -   OK
-
CL2-E    63528 158.214.135.102 iqn.z3        3260 CHAP S -            -   OK
-
CL4-E    63528 158.214.135.100 iqn.z2        3260 CHAP S -            -   OK
-
CL4-E    63528 158.214.135.102 iqn.z3        3260 CHAP S -            -   OK
-
CL4-E    63528 158.214.135.102 iqn.z4        3260 NONE S -            -   NG
-
CL4-E    63528 158.214.135.102 iqn.z5        3260 NONE S -            -   NG
-
```

Chapter 5: Configuration setting commands

Attempts to log in to the iSCSI target of the external storage systems which is registered in the iSCSI port: CL4-E of the local storage system, and then displays the result of login:

```
# raidcom check external_iscsi_name -port CL4-E
```

```
PORT    Serial# IP_ADDR         IQN        IP_PORT# AMD  D CHAP_user     Sec LOGIN
ISCSI_VP_ID
CL4-E    63528 158.214.135.100 iqn.z1       3260 CHAP D Win_SQL_EX   *   OK
-
CL4-E    63528 158.214.135.100 iqn.z2       3260 CHAP S -           -   OK
-
CL4-E    63528 158.214.135.102 iqn.z3       3260 CHAP S -           -   OK
-
CL4-E    63528 158.214.135.102 iqn.z4       3260 NONE S -           -   NG
-
CL4-E    63528 158.214.135.102 iqn.z5       3260 NONE S -           -   NG
-
```

Attempts to log in to the iSCSI target (iSCSI Name: iqn.z1, IP Address: 158.214.135.100) on the external storage system which is registered in the iSCSI port: CL4-E on the local storage system, and then displays the result of login:

```
# raidcom check external_iscsi_name -port CL4-E -iscsi_name iqn.z1 -address
158.214.135.100
```

```
PORT Serial# IP_ADDR IQN IP_PORT# AMD D CHAP_user Sec LOGIN ISCSI_VP_ID
CL4-E 63528 158.214.135.100 iqn.z1 - - - - - OK
```

Attempts to log in to the iSCSI target of the external storage system which is registered in the iSCSI port: CL2-E, iSCSI virtual port ID: 0 of the local storage system, and then displays the result of login:

```
# raidcom check external_iscsi_name -port CL2-E -iscsi_virtual_port_id 0
```

```
PORT    Serial# IP_ADDR         IQN        IP_PORT# AMD  D CHAP_user     Sec LOGIN
ISCSI_VP_ID
CL2-E    63528 158.214.135.100 iqn.z2       3260 CHAP S -           -   OK
0
CL2-E    63528 158.214.135.102 iqn.z3       3260 CHAP S -           -   OK
0
```

Description of each column in output example:

PORT

Displays the port number.

Serial#

Displays the product serial number.

IP_ADDR

Displays the IP address of the iSCSI target on the external storage system.

IQN

Displays the iSCSI name of the iSCSI target on the external storage system.

IP_PORT#

Displays the TCP port number of the iSCSI target on the external storage system.

AMD

Displays the authentication mode of the iSCSI target on the external storage system.

- CHAP: CHAP authentication is enabled.

- NONE: The authentication mode is not configured.

D

Displays the direction of the authentication mode of the iSCSI target.

- S: single directional (An iSCSI target recognizes the iSCSI target.)

- D: dual-directional (An iSCSI target recognizes the iSCSI initiator and vice versa.)

CHAP_user

Displays the CHAP user name of the iSCSI target on the external storage system. If the CHAP user name is not set, a hyphen (-) is displayed.

Sec

Displays an asterisk (*) when the secret is set to the iSCSI target of the external storage system. Otherwise, a hyphen (-) is displayed.

LOGIN

Displays the result of the login.

ISCSI_VP_ID

Displays the iSCSI virtual port ID. A hyphen (-) is displayed when the iSCSI virtual port is disabled.

A hyphen (-) is also displayed if the -iscsi_name <external iscsi name> -address <external IP address> option is specified and the -iscsi_virtual_port_id <value> option is omitted for the enabled iSCSI virtual port.

raidcom modify external_chap_user

This command sets the CHAP user name and the secret (password) to the iSCSI target of the specified external storage system.

When you omit both the CHAP user name and secret, the CHAP user name and the secret which are set in the iSCSI target are deleted. If the iSCSI target of the specified external storage system has been registered in multiple iSCSI ports of the local storage system, the settings are applied to all iSCSI ports of the local storage system.

> 📄 **Note:** If a failure occurs in a processor on an MP blade or a controller, this command cannot be run even on an MP blade or a controller on which no failure occurs.

Syntax

```
raidcom modify external_chap_user -port <port#> -iscsi_name
   <external iscsi name> -address <external IP address>
   [-iscsi_virtual_port_id <iSCSI virtual port ID>]
   [-chap_user <user name> ] [-secret]
```

Options and parameters

-port <port#>
> Specifies the port number of the local storage system. For example:
>
> ▪ CL1-A

-iscsi_name <external iscsi name>
> Specifies the iSCSI name of the iSCSI target on the external storage system in an iqn format or an eui format.
>
> ▪ iqn format: `iqn.` and the subsequent maximum 219 characters.
>
> ▪ eui format: `eui.` and the subsequent 16 characters in hexadecimal notation.

-address <external IP address>
> Specifies the IP address of the iSCSI target on the external storage system. An IPv4 address or an IPv6 address can be specified.
>
> Note that the following addresses cannot be specified when you specify the IPv4 address:
>
> ▪ Network address (for example: 192.168.10.0, 0.120.10.1)
>
> ▪ Broadcast address (for example: 255.255.255.255, 10.1.255.255)
>
> ▪ Loop back address (for example: 127.0.0.1)
>
> Note that the following addresses cannot be specified when you specify the IPv6 address:
>
> ▪ Unspecified address (for example: ::)
>
> ▪ Multicast address (for example: ff:1024:1215::01)
>
> ▪ Loop back address (for example: ::1)

[-iscsi_virtual_port_id <iSCSI virtual port ID>]
> Specifies the iSCSI virtual port ID when the specified port is iSCSI and the iSCSI virtual port mode is enabled. If this option is omitted, 0 is assumed.

Chapter 5: Configuration setting commands

An error occurs if you specify this option in any of the following cases:

- The virtual port of the specified iSCSI virtual port ID is not found.

- The iSCSI virtual port mode is disabled.

- The device does not support the iSCSI virtual port mode.

[-chap_user <user name>]

Specifies the CHAP user name of the iSCSI target on the external storage system. You can input up to 223 characters for the name.

For example: storage01

[-secret]

Displays the prompt for entering a secret. You can specify the secret within the range of 12 to 32 characters.

If the characters you input for the secret are less than 12 characters, an error occurs. If the characters you input for the secret are more than 32 characters, the first 32 characters are valid.

For the details of available characters for the secret when you use CCI, see the supported characters topic. Note that the space and back slash cannot be used.

Examples

When you set the CHAP user name: Elun_TAR_4E and the secret to the iSCSI target (iSCSI name: iqn.z1, IP address: 158.214.135.100) on the external storage system which is registered in the port: CL1-A of the local storage system:

```
# raidcom modify external_chap_user -port CL1-A -iscsi_name iqn.z1 -address
158.214.135.100 -chap_user Elun_TAR_4E -secret
Enter Secret :
```

When you set the secret to the iSCSI target (iSCSI name: iqn.z1, IP address: 158.214.135.100) on the external storage system which is registered in the port: CL1-A of the local storage system:

```
# raidcom modify external_chap_user -port CL1-A -iscsi_name iqn.z1 -address
158.214.135.100 -secret
Enter Secret :
```

When you delete the registered CHAP user name and the secret from the iSCSI target (iSCSI name: iqn.z1, IP address: 158.214.135.100) on the external storage system which is registered in the port: CL1-A of the local storage system:

```
# raidcom modify external_chap_user -port CL1-A -iscsi_name iqn.z1 -address
158.214.135.100
```

When you set the CHAP user name and the secret for the iSCSI target (iSCSI name: iqn.z1, IP address: 158.214.135.100) on the external storage system which is registered in the port: CL1-A, iSCSI virtual port ID: 2 of the local storage system:

```
# raidcom modify external_chap_user -port CL1-A -iscsi_virtual_port_id 2 -iscsi_name
iqn.z1 -address 158.214.135.100 -chap_user Elun_TAR_4E -secret
Enter Secret :
```

raidcom modify initiator_chap_user

This command sets the CHAP user name and the secret to the iSCSI initiator of the specified local storage system. When you omit both the CHAP user name and secret, the CHAP user name and the secret which are set in the iSCSI initiator are deleted.

> 📄 **Note:** If a failure occurs in a processor on an MP blade or a controller, this command cannot be run even on an MP blade or a controller on which no failure occurs.

Syntax

```
raidcom modify initiator_chap_user -port <port#>
   [-chap_user  <user name> ] [-secret]
```

Options and parameters

-port <port#>

Specifies the port number of the local storage system. For example:

- CL1-A

[-chap_user <user name>]

Specifies the CHAP user name of the iSCSI initiator on the external storage system. You can input up to 223 characters for the name. For example:

storage01

[-secret]

Displays the prompt for entering a secret. You can specify the secret within the range of 12 to 32 characters.

If the characters you input for the secret are less than 12 characters, an error occurs. If the characters you input for the secret are more than 32 characters, the first 32 characters are valid.

For the details of available characters for the secret when you use CCI, see the topic Supported characters. Note that the space and back slash cannot be used.

Chapter 5: Configuration setting commands

Examples

When you set the CHAP user name: Elun_INI_4E of the iSCSI initiator and the secret to the iSCSI port: CL4-E of the local storage system:

```
# raidcom modify initiator_chap_user -port CL4-E -chap_user Elun_INI_4E -secret
Enter Secret :
```

When you set the secret of the iSCSI initiator to the iSCSI port: CL4-E of the local storage system:

```
# raidcom modify initiator_chap_user -port CL4-E -secret
Enter Secret :
```

When you delete the CHAP user name and the secret of the iSCSI initiator from the iSCSI port: CL4-E of the local storage system:

```
# raidcom modify initiator_chap_user -port CL4-E
```

raidcom add journal

Registers (adds) a volume to a journal.

This makes a journal for open systems (Open System) or mainframe (M/F System) according to the specified LDEVs.

If the journal already exists, the specified LDEV is added to the journal.

If the journal does not exist, you must create it first, and then add an LDEV.

If the -timer_type option is specified, this command makes the journal as "M/F System". If not, this command makes the journal as "Open System".

A device group can also be specified instead of an LDEV.

This command is executed asynchronously with the command input. Check the completion of this process using the **raidcom get command_status** command.

> **Note:** If a failure occurs in a processor on an MP blade or a controller, this command cannot be run even on an MP blade or a controller on which no failure occurs.
>
> If cache memory is blocked, the operation cannot be performed.

Syntax

```
raidcom add journal -journal_id <journal ID#> {-ldev_id <ldev#> …
    [-cnt <count>] | -grp_opt <group option>
    -device_grp_name <device group name> [<device name>]}
    [-mp_blade_id <mp#> | -timer_type <timer type> ]
```

Options and parameters

-journal_id <journal ID#>
> Specifies the journal number (0-255).

-ldev_id <ldev#> …
> Specifies the LDEV number (0-65279).
>
> Up to 64 LDEVs can be specified at a time. For example:
>
> - -ldev_id 200
> - -ldev_id 100-110
> - -ldev_id 100 -cnt 10

[-cnt <count>]
> Specifies the count (2-64).
>
> If this option is omitted, "1" is used as the count.

-grp_opt <group option>
> Specifies the device information extracted from the LDEV belonging to the device group. Specify "ldev" (fixed). The information of LDEV belonging to the device group is used.

-device_grp_name <device group name> [<device name>]
> Specifies the name of device group (maximum 32 characters) to be operated.
>
> To operate an LDEV in the device group, specify the device name of LDEV (maximum 32 characters) within the device group.
>
> If the device name is omitted, all of the LDEVs belonging in the device group are operated.

[-mp_blade_id <mp#>]
> Specifies the MP blade ID number (0-15).
>
> If this option is omitted, the device automatically allocates an MP blade ID number. For example:
>
> - -mp_blade_id 2

[-timer_type <timer type>]
> Specifies the timer type: system, local, or None. For example:
>
> - -timer_type system
> - -timer_type local
>
> For details, see the Universal Replicator manual for the storage system.

Examples

Examples for open systems:

Creating a journal #1 of LDEVs: 265, 266.

```
# raidcom add journal -journal_id 1 -ldev_id 265 266
```

Chapter 5: Configuration setting commands

Creating a journal #1 with an LDEV belonging to the device group: grp1.

```
# raidcom add journal -journal_id 1 -grp_opt ldev -device_grp_name grp1
```

Example for mainframe:

Creating a journal #1 with LDEVs: 265, 266. As a timer type, the system clock of the mainframe host is used. (When LDEV is added to the journal of M/F, it is required to specify the system clock of the mainframe host to the timer type.)

```
# raidcom add journal -journal_id 1 -ldev_id 265 266
-timer_type system
```

raidcom delete journal

Deletes a journal from the specified journal.

A device group can also be specified instead of an LDEV. If the LDEV and device group are not specified, a journal is deleted.

This command is executed asynchronously with the command input. Check the completion of this process using the `raidcom get command_status` command.

> 📄 **Note:** If a failure occurs in a processor on an MP blade or a controller, this command cannot be run even on an MP blade or a controller on which no failure occurs.

Syntax

```
raidcom delete journal -journal_id <journal ID#> [-ldev_id
   <ldev#> | -grp_opt <group option> -device_grp_name
   <device group name> [<device name>]]
```

Options and parameters

-journal <journal ID#>
Specifies the journal number (0-255).

[-ldev_id <ldev#>]
Specifies the LDEV number (0-65279). If the both LDEV and device group are not specified, journal is deleted. For example:

- -ldev_id 200

-grp_opt <group option>
Specifies the device information extracted from the LDEV belonging to the device group. Specify "ldev" (fixed). The information of LDEV belonging to the device group is used.

-device_grp_name <device group name> [<device name>]

> Specifies the name of device group (maximum 32 characters) to be operated.

> To operate an LDEV in the device group, specify the device name of LDEV (maximum 32 characters) within the device group.

> If the device name is omitted, all of the LDEVs belonging in the device group are operated. If the both LDEV and device group are not specified, journal is deleted.

Examples

Deleting the specified journal.

```
# raidcom delete journal -journal_id 6
```

Deleting the specified LDEV from the journal.

```
# raidcom delete journal -journal_id 6 -ldev_id 265
```

Deleting the LDEV belonging to the device group: grp1 from the journal.

```
# raidcom delete journal -journal_id 6 -grp_opt ldev
-device_grp_name grp1
```

raidcom get journal

Displays the information of registered journal.

Syntax

Displaying journal information

```
raidcom get journal [-key <keyword>]
```

Displaying timer-related information

```
raidcom get journalt
```

Options and parameters

[-key <keyword>]

> Specifies the display keyword. Specify opt as <keyword>.

Examples

Displaying journal information.

```
# raidcom get journal
```

```
JID MU CTG  JNLS AP U(%)  Q-Marker Q-CNT D-SZ(BLK)  Seq#  Num LDEV#
001 0  1    PJNN 4  21    43216fde 30    512345     62500 4   265
002 1  2    PJNF 4  95    3459fd43 52000 512345     62500 3   270
002 2  2    SJNS 4  95    3459fd43 52000 512345     62500 3   270
003 0  3    PJSN 4  0     -        -     512345     62500 1   275
004 0  4    PJSF 4  45    1234f432 78    512345     62500 1   276
005 0  5    PJSE 0  0     -        -     512345     62500 1   277
```

Displaying Timer related information of the journal.

```
# raidcom get journalt
```

```
JID MU CTG  JNLS AP U(%)  Q-Marker Q-CNT D-SZ(BLK)  Seq#  DOW PBW APW
001 0  1    PJNN 4  21    43216fde 30    512345     63528 20  300 40
002 1  2    PJNF 4  95    3459fd43 52000 512345     63528 20  300 40
003 0  3    PJSN 4  0     -        -     512345     63528 20  300 40
```

Displaying option information of the journal.

```
# raidcom get journal -key opt
```

```
JID MU CTG JNLS TYPE TTYPE MODE IF DOW(S) PBW(M) CR CS(bps) DM
MP# Seq# T RCMD#
000 0 0 SMPL OPEN - CACHE E 60 5 L 256 Y
0 302614 N NA
000 1 0 SMPL OPEN - CACHE E 60 5 L 256 Y
0 302614 N NA
000 2 0 SMPL OPEN - CACHE E 60 5 L 256 Y
0 302614 N NA
000 3 0 SMPL OPEN - CACHE E 60 5 L 256 Y
0 302614 N NA
001 0 1 PJSN OPEN - CACHE E 60 5 L 256 Y
4 302614 Y 100
001 1 0 SMPL OPEN - CACHE E 60 5 L 256 Y
4 302614 N NA
001 2 0 SMPL OPEN - CACHE E 60 5 L 256 Y
4 302614 N NA
001 3 0 SMPL OPEN - CACHE E 60 5 L 256 Y
4 302614 N NA
```

Chapter 5: Configuration setting commands

```
002 0 1 SJSN M/F system CACHE E 60 5 L 256 Y
0 302614 N NA
002 1 0 SMPL M/F system CACHE E 60 5 L 256 Y
0 302614 N NA
002 2 0 SMPL M/F system CACHE E 60 5 L 256 Y
0 302614 N NA
002 3 0 SMPL M/F system CACHE E 60 5 L 256 Y
0 302614 N NA
```

Description of each column in output example:

JID

> Journal number.

MU

> Mirror ID on Universal Replicator.

CTG

> Consistency group ID.

JNLS

> Status in the journal:
>
> - SMPL: a journal volume that does not have a pair, or is being deleted.
> - P(S)JNN: "P(S)vol Journal Normal Normal".
> - P(S)JNS: "P(S)vol Journal Normal Suspend" created with `-nocsus` option.
> - P(S)JSN: "P(S)vol Journal Suspend Normal".
> - P(S)JNF: "P(S)vol Journal Normal Full".
> - P(S)JSF: "P(S)vol Journal Suspend Full".
> - P(S)JSE: "P(S)vol Journal Suspend Error" including link failure.
> - P(S)JES: "P(S)vol Journal Error Suspend" created with `-nocsus` option.

AP

> Number of active link paths of Universal Replicator.

U(%)

> Usage rate of journal volumes assuming the entire relevant volume is 100%.

Q-Marker

> The P-VOL journal volume shows the newest sequential number (Q-marker) at the time of receiving WRITE data. The S-VOL journal volume shows the newest sequential number (Q-marker) written to the cache.

Q-CNT

> Number of Q-markers remaining in the P-VOL journal volume.

D-SZ(BLK)

> Capacity of the data block size of the journal volume in units of 512 bytes.
>
> For details about the displayed capacity, see *Hitachi Universal Replicator User Guide*.

Chapter 5: Configuration setting commands

Seq#

Serial number (Seq#).

Num

Number of LDEVs composing the journal volume.

LDEV#

LDEV number of a journal volume for storage systems that support reserve journal volumes:

- LDEV number of the master journal volume for the master journal

- LDEV number of the restore journal volume for the restore journal. Whether to support reserve journal volumes depends on the model of your storage system. For the support status, see the *Hitachi Universal Replicator User Guide* for your model.

DOW

Data Overflow Watch timer setting (in seconds) per the Journal.

PBW

Path Blockade Watch timer setting (in seconds) per the Journal. If the setting is more than 3600 seconds, it displays 6000 seconds.

APW

Active Path Watch timer setting (in seconds) to detect any link failures.

TYPE

Open or M/F system.

TTYPE

Timer type for mainframe: System or Local.

MODE

Status of the journal:

- HDD: Store the journal data to the journal volume (cache mode disabled).

- CACHE: Store the journal data to the cache (cache mode enabled).

IF: Inflow control for journal:

- E: Enable

- D: Disable

DOW(S)

Data Overflow Watch timer setting (in seconds) per the Journal.

PBW(M)

Path Blockade Watch timer setting (in minutes) per the Journal. If the setting is more than 60 minutes, it displays 100 minutes.

CR

Copy rate: L (Low), M (Medium), or H (High).

CS(bps)

Copy speed in Mbps: 3 (3 Mbps), 10 (10 Mbps), 100 (100 Mbps), or 256 (256 Mbps).

DM

Copy mode "Y" or "N" under failure of the delta resync:

- Y: copying ALL data.
- N: No copying.

MP#

MP blade ID.

T

Displays whether to transfer the path watch time for the master journal to the secondary side (RCU side) of the mirror.

- Y: Transfer.
- N: Do not transfer.
- - (hyphen): This information is not available for this journal. In addition, for VSP and HUS VM, a hyphen (-) is always displayed.

RCMD#

LDEV number of a remote command device. Displays "NA" if a remote command device is not allocated, or "NU" if a remote command device is allocated with no LDEV number specified. Displays a hyphen (-) if display of remote command device information is not supported. In addition, for VSP and HUS VM, a hyphen (-) is always displayed. If the -fx option is specified, an LDEV number is displayed in hexadecimal.

raidcom modify journal

Changes the option of Universal Replicator to be used at journal, and sets the specified control parameter to the journal.

If you specify the `-mp_blade_id` option, you cannot specify the other options (`-data_overflow_watch`, `-path_blocked_watch`, `-cache_mode`, or `-timer_type`).

The `-timer_type` option must be specified to the journal on mainframe systems. This means that this option cannot be used for changing from "Open System" to "M/F System".

> 📄 **Note:** If a failure occurs in a processor on an MP blade or a controller, this command cannot be run even on an MP blade or a controller on which no failure occurs.

Syntax

When changing the data overflow watching time of journal data area, cache mode or the timer type.

```
# raidcom modify journal -journal_id <journal ID#>
     {[-data_overflow_watch <time>][-cache_mode {y | n}]
     [-timer_type <timer type>]}
```

When changing the path blocked watch.

```
# raidcom modify journal -journal_id <journal ID#>
     -path_blocked_watch <time> [-mirror_id <mu#>]
```

> 📄 **Note:** This command is executed asynchronously with the command input. Check the completion of this process on the **raidcom get command_status** command.

When setting the path blocked watch, path watch time transfer, copy pace, transfer speed, and operation when delta resync fails by specifying the mirror ID.

```
raidcom modify journal -journal_id <journal ID#> [-mirror_id <mu#>] {[-
path_blocked_watch <time>] [-path_blocked_watch_transfer {y|n}] [-
copy_size <size>] [-transfer_speed <speed>] [-entire_copy {y|n}]}
```

When setting or releasing the remote command device.

```
raidcom modify journal -journal_id <journal ID#> [-mirror_id <mu#>]
{-command_device y [-ldev_id <ldev#>] | -command_device n}
```

When changing the MP blade ID.

```
# raidcom modify journal -journal_id <journal ID#>
     -mp_blade_id <mp#>
```

Options and parameters

-journal_id <journal ID#>
Specifies the journal number (0-255).

-data_overflow_watch <time(sec)>
The data overflow watch timer (0-600) (second).

Specifies the watch time for the journal data area being full.

If 0 is specified, the setting of the data overflow watch timer is changed and the inflow control for journals is set to disabled.

-path_blocked_watch <time(min)>
The watch for the path blockage (1-60) (minute).

Use Device Manager - Storage Navigator to set a value of more than 60 minutes.

If 0 is specified, the time of the watch for the path blockage does not change, and setting for transferring the path blockage time from the primary side of the mirror (MCU side) to the secondary side of the mirror (RCU side) is invalid.

If the path blockage time (1 to 60) (minutes) is specified, setting for transferring the path watch time from the primary side of the mirror (MCU side) to the secondary side of the mirror (RCU side) is enabled.

-cache_mode (y | n}

Specifies whether to use the cache mode.

- y: Cache mode enabled (E)

- n: Cache mode disabled (D)

[-timer_type <timer type>]

Specifies the timer type: system, local, or None. For example:

- -timer_type system

- -timer_type local

For details, see the Universal Replicator manual for the storage system.

-mp_blade_id <mp#>

Specifies the MP blade ID (0-15).

Do not change the MP blade ID during initial copying or I/O processing of Universal Replicator, ShadowImage, TrueCopy, or global-active device.

If you change MP blade IDs for multiple journals, wait at least 10 minutes before changing another MP blade ID.

To change the MP blade ID again for the same journal, wait for more than 30 minutes after changing the MP blade ID. For example:

- -mp_blade_id 2

[-mirror_id <mu#>]

Specifies the Mirror ID.

If the setting is omitted, "0" is used.

[-copy_size <size>]

Specify the <size> from 1 to 15 as follows. If you specify the large value as this option, the copy time shortens, but the I/O performance might deteriorate. If you do not use this option, the copy is performed at medium-speed.

- 1 or 2: low-speed

- 3: medium-speed

- More than 4: high-speed

-transfer_speed <speed>

Supported storage systems:

- VSP 5000 series

- VSP G1x00 and VSP F1500

- VSP E series

- VSP G130, G/F350, G/F370, G/F700, G/F900

Specify the data transfer speed by Mbps. You can specify 256, 100, 10 or 3.

-path_blocked_watch_transfer {y|n}

Supported storage systems:

- VSP 5000 series

- VSP G1x00 and VSP F1500

- VSP E series

- VSP G130, G/F350, G/F370, G/F700, G/F900

Specify whether to transfer the path watch time from the primary side of the mirror (MCU side) to the secondary side of the mirror (RCU side).

- y: Transfer

- n: Not transfer

-entire_copy {y|n}

Supported storage systems:

- VSP 5000 series

- VSP G1x00 and VSP F1500

- VSP E series

- VSP G130, G/F350, G/F370, G/F700, G/F900

Specify the operation when delta resync cannot be performed.

- y: All the data of the P-VOL is copied to the S-VOL.

- n: No operation. The S-VOL is not updated.

-command_device y [-ldev_id <ldev#>]

Supported storage systems:

- VSP 5000 series

- VSP G1x00 and VSP F1500

- VSP E series

- VSP G130, G/F350, G/F370, G/F700, G/F900

Sets a remote command device. A remote command device specified by -ldev_id <ldev#> is set for the specified mirror ID.

-command_device n

Supported storage systems:

- VSP 5000 series

- VSP G1x00 and VSP F1500

- VSP E series

- VSP G130, G/F350, G/F370, G/F700, G/F900

Releases a remote command device. The remote command device of the specified mirror ID is released.

Examples

Changing the data overflow watch time for journal "6" to 15 seconds.

```
# raidcom modify journal -journal_id 6 -data_overflow_watch 15
```

Changing the settings for journal 6: data overflow watch time to 15 seconds, and the timer type to the system clock of the mainframe host.

```
# raidcom modify journal -journal_id 6 -data_overflow_watch 15
-timer_type system
```

Changing the the MP blade ID for journal "6" to 2.

```
# raidcom modify journal -journal_id 6 -mp_blade_id 2
```

Changing the setting for journal "6" to store journal data in the secondary journal to the cache.

```
# raidcom modify journal -journal_id 6 -cache_mode y
```

Changing the path block monitoring time of mirror ID 1 for journal "6" to 59 minutes.

```
# raidcom modify journal -journal_id 6 -path_blocked_watch 59
-mirror_id 1
```

Changing the setting for journal "6" to set the remote command device of LDEV number 1 for mirror ID 1.

```
# raidcom modify journal -journal_id 6 -mirror_id 1 -command_device y -ldev_id 1
```

Changing the setting for journal "6" to release the remote command device of mirror ID 1.

```
# raidcom modify journal -journal_id 6 -mirror_id 1 -command_device n
```

Chapter 5: Configuration setting commands

raidcom add ldev

Adds an LDEV to the specified parity group or the external volume group. Or this adds V-VOL for Dynamic Provisioning, Dynamic Provisioning for Mainframe, Dynamic Tiering, active flash, Dynamic Tiering for Mainframe, active flash for mainframe, Thin Image, or Copy-on-Write Snapshot to the specified pool.

Alternatively, this command creates V-VOL for Dynamic Provisioning associated with a pool volume having the data direct mapping attribute.

This command is executed asynchronously with the command input. Check the completion of this process on the **raidcom get command_status** command.

> **Note:** If a failure occurs in a processor on an MP blade or a controller, this command cannot be run even on an MP blade or a controller on which no failure occurs.
>
> If cache memory is blocked, a DP-VOL cannot be created. If CHB, DKB, or cache memory is blocked, a V-VOL in Thin Image cannot be created.

Syntax

When creating an LDEV or V-VOL in the specified parity group, external volume group, or pool.

```
raidcom add ldev {-parity_grp_id <gno-sgno>|
   -external_grp_id <gno-sgno> | -pool {<pool ID#> |
    <pool naming> | snap}} {-ldev_id <ldev#> |
   -tse_ldev_id <ldev#> | -ese_ldev_id <ldev#> | -ldev_id auto -request_id auto
   [-resource_id <resource group id>] [-ldev_range <range>]}
   {-capacity <size> | -offset_capacity <size> | -cylinder <size>}
   [-emulation <emulation type>][-location <lba>]
   [-mp_blade_id <mp#>][-clpr <clpr#>]
   [-t10pi_enable]
   [-status {enable_fullallocation | disable_fullallocation} | -capacity_saving
{compression | deduplication_compression}
   [-capacity_saving_mode <saving_mode>] [-compression_acceleration {enable |
disable} -request_id auto] | [-drs -request_id auto]
   -capacity_saving disable]
```

When creating a V-VOL for Dynamic Provisioning associated with a pool volume having the data direct mapping attribute.

```
raidcom add ldev -ldev_id <ldev#> -mapping_ldev_id <ldev id>
```

Options and parameters

-parity_grp_id <gno-sgno>

Specifies the parity group number (gno:1-52, sgno:1-32) (for example, 3-1).

-external_grp_id <gno-sgno>

Specifies the external volume group number (gno:1-16384, sgno:1-4096) (for example, 52-11 ("E" is not required)).

-pool {<pool ID#> | <pool naming> | snap}

Specifies a Pool ID or Pool name for Dynamic Provisioning, Dynamic Provisioning for Mainframe, Dynamic Tiering, active flash, Dynamic Tiering for Mainframe, or active flash for mainframe when V-VOL is created to Dynamic Provisioning, Dynamic Provisioning for Mainframe, Dynamic Tiering, active flash, Dynamic Tiering for Mainframe, or active flash for mainframe pool.

When only a number is specified, it is recognized as a pool ID. Therefore, when you specify a pool that the pool name is only a number, specify the pool ID instead of the pool name.

When creating V-VOL for a Thin Image or Copy-on-Write Snapshot pool, specify "snap".

-ldev_id <ldev#> | -tse_ldev_id <ldev#> | -ese_ldev_id <ldev#>

Specifies the LDEV number (0-65279).

For VSP E590, the range is 0-32,767. For VSP E790, the range is 0-49,151.

If you specify the `-tse_ldev_id` option, the TSE-VOL that is used by FlashCopy SE is created. When you specify the `-ese_ldev_id` option, ESE-VOL is created. You can specify the `-ese_ldev_id` option to VSP 5000 series, and VSP G1x00 and VSP F1500 only. TSE-VOL and ESE-VOL can be created in the pool of Dynamic Provisioning for Mainframe only. If you want to create TSE-VOL and ESE-VOL, specify the pool of Dynamic Provisioning for Mainframe by using the `-pool` option. For example:

- -ldev_id 200

- -tse_ldev_id 400

- -ese_ldev_id 600

If you use the `-tse_ldev_id` or `-ese_ldev_id` option in a command other than the **raidcom add ldev** command, that option functions as the `-ldev_id` option.

-ldev_id auto -request_id auto [-resource_id <resource group id>] [-ldev_range <range>]

This option can be specified only when a V-VOL of which emulation type is OPEN-V is created in a Dynamic Provisioning or Dynamic Tiering pool.

Specify this option to automatically assign an LDEV number to the LDEV you created. From the unused numbers of the LDEVs for which the user has operation permission, the smallest LDEV number is assigned. To limit the range of the LDEV numbers to be assigned, use any of the following options:

- -request_id auto: Effective only when "auto" is specified as the -request_id option. If other option than "auto" is specified, EX_INVARG or EX_REQARG is replied.

- -resource_id <resource grp id>: Assigns the smallest LDEV number from the numbers of the unused LDEVs with the specified resource group ID.

- -ldev_range <range>: Assigns the smallest LDEV number from the numbers of the unused LDEVs with the LDEV number specified by <range>. If the -resource_id <resource grp id> is specified with this option, the smallest LDEV number from the numbers of the unused LDEVs that have the LDEV number in <range> and the specified resource group number is assigned.

Specify <range> in the format of *<starting-LDEV-number-ending-LDEV-number>*. LDEV numbers must be specified in decimal or hexadecimal. To specify in hexadecimal, add 0x at the beginning of the LDEV number.

For example, to specify LDEV numbers from 300 to 305:

- -ldev_range 300-305

- -ldev_range 0x12c-0x131

- -ldev_range 0x12c-305

  (You can combine decimal and hexadecimal numbers.)

You can check the selected LDEV numbers using the `raidcom get command_status` command.

The request ID is output in the following format after completion of the `raidcom add ldev` command:

- REQID : <request#>

  Where <request#> is a request ID assigned each time the command is executed. For details, see Request ID function (on page 246).

Unused LDEV numbers in the resource group where you create an LDEV and in the LDEV number range must satisfy the following requirements:

- The unused LDEV numbers are not in the same range (grouped by every 32 LDEV numbers) as the installed LDEV in an emulation group other than OPEN-V. If this requirement is not satisfied, the command might fail with EX_CMDRJE (SSB1=2E30, SSB2=0026).

- The unused LDEV numbers are not used as an alias of PAV. If this requirement is not satisfied, the command might fail with EX_CMDRJE (SSB1=2E30, SSB2=0025).

- The unused LDEV numbers can be used in the shared memory installed in the storage system. If this requirement is not satisfied, the command might fail with EX_CMDRJE (SSB1=2E23, SSB2=0001).

- When using a mainframe volume in the same storage system, SSID that is not used by the mainframe volume is assigned to the unused LDEV numbers.

The storage system might execute multiple raidcom add ldev commands at the same time. In such a batch operation, if one command fails, the others also fail. You can display command error information by executing raidcom get command_status -request_id <request#>, where you must specify the Request ID output when the command is executed. For errors caused due to batch operation, 'O' is displayed in column R, while for errors caused for other reason, 'T' is displayed in the same column. For example, if a user issues two raidcom add ldev -ldev_id auto commands at the same time to a storage system having only one free LDEV number, the system attempts to execute the commands in a batch, causing both of them to fail. Check command status by executing raidcom get command_status -request_id <request#> with the Request ID. If 'O' is displayed under column R, retry the command.

-capacity <size>

Specifies the capacity. The size can be specified in bytes or blocks. When specifying in bytes, the unit is t/T (terabyte), g/G (gigabyte), m/M (megabyte), or k/K (kilobyte). If this unit is omitted, block (512 bytes) is used.

When specifying in bytes, note the following:

- There is no size correction.

- If the capacity of LDEVs that are created by GUI and by CLI are the same, a copy pair might not be created. To create a pair with an LDEV that was created by GUI, create the LDEV by specifying blocks.

Example of specification:

1GB (gigabyte) is:

-capacity 1G, -capacity 1g, -capacity 1024M, -capacity 1024m, -capacity 1048576K, -capacity 1048576k, -capacity 2097152

In the case of volumes other than OPEN-V, more free space than the capacity that was actually specified is required. For details, see the *Provisioning Guide* for the storage system.

When you specify "all" instead of digits to create an OPEN-V LDEV and no LDEV has been created in the specified parity group* or the external volume group, the system allocates all empty space for the LDEV. If the empty space is larger than the maximum capacity of an LDEV, the system creates the LDEV with the maximum capacity and leaves the rest of the space as an empty space.

*: "all" cannot be specified depending on the combination of the drive type and drive level because of the LDEV control area allocation, and so on. In this case specify bytes or blocks.

-offset_capacity <size>

Specifies capacity. Corrects the size as well as GUI for the specified capacity. The capacity can be specified in bytes or blocks. When specifying byte, the unit is t/T (terabyte), g/G (gigabyte), m/M (megabyte), or k/K (kilobyte).

-cylinder <size>

Specifies the size in cylinder unit. When specifying, the unit is t/T (teracylinder), g/G (gigacylinder), m/M (megacylinder), or k/K (kilocylinder). Cylinder is applied if you do not specify the unit.

When an open-systems emulation type is specified, this option cannot be specified.

[-location <lba>]

Specifies the Location (the starting point of the LDEV to be created in the parity group/external volume group). If this specification is omitted, create a LDEV and close up in the free space.

[-emulation <emulation type>]

Specifies the emulation type (for example, OPEN-V).

If this specification is omitted, OPEN-V is specified.

If this specification is omitted when the virtual volume is crated, the following emulation type is specified.

- OPEN-V: Dynamic Provisioning, Dynamic Tiering, or active flash

- 3390-A: Dynamic Provisioning for Mainframe, Dynamic Tiering for Mainframe, or active flash for mainframe

The values that can be specified for <emulation type> are:

- OPEN-3, OPEN-8, OPEN-9, OPEN-E, OPEN-L, OPEN-V

- 3390-1, 3390-2, 3390-3, 3390-A, 3390-3A, 3390-3B, 3390-3C, 3390-3R, 3390-9, 3390-9A, 3390-9B, 3390-9C, 3390-L, 3390-LA, 3390-LB, 3390-LC, 3390-M, 3390-MA, 3390-MB, 3390-MC, 3390-V

- 3380-3, 3380-3A, 3380-3B, 3380-3C

Some emulation types cannot be specified depending on the emulation type.

> ⚠️ **Caution:**
>
> You can specify 3390-3 or 3390-3R as the emulation type, but these are unable to be mixed. You can specify the 3380 series or the 3390 series as the emulation type for each parity group, but they are unable to be mixed for each 32 address boundary because of the OS restriction.

[-mp_blade_id <mp#>]

Specifies the MP blade ID (0-15). If this specification is omitted, it is allocated automatically.

[-clpr <clpr#>]

When you create virtual volumes by specifying the `-pool <pool#>` option, specify the CLPR ID. If you omit this option, the ID number of CLPR to which the pool is allocated is used.

[-status {enable_fullallocation | disable_fullallocation}]

Specifies an availability of Full Allocation when the virtual volume is used for Dynamic Provisioning, Dynamic Tiering, or active flash. If this specification is omitted, Full Allocation is disabled.

- enable_fullallocation: Enables Full Allocation. If all areas equivalent to the sum of the pool capacities the specified volume requires can be reserved, it is guaranteed that all areas of DP-VOL are writable.

- disable_fullallocation: Disables Full Allocation.

[-t10pi_enable]

Enables the T10 PI attribute.

-mapping_ldev_id <ldev id>

Supported storage systems:

- VSP 5000 series

- VSP G1x00 and VSP F1500

- VSP E series

- VSP G130, G/F350, G/F370, G/F700, G/F900

- VSP G200, G400, G600, G800 and VSP F400, F600, F800

Creates the V-VOL for Dynamic Provisioning associated with the pool volume having the data direct mapping attribute. When you specify this option, the data direct mapping attribute is automatically set to the V-VOL to be created.

[-compression_acceleration {enable | disable} -request_id auto]

Supported storage systems:

- VSP 5000 series

- VSP E1090

Specifies whether compression accelerator is enabled or disabled. If you omit this option, it will be enabled when COMPRESSION_ACCELERATION of raidcom get system -key dedupe_compression is AVAILABLE. It will be disabled when COMPRESSION_ACCELERATION of raidcom get system -key dedupe_compression is UNAVAILABLE.

- enable: Compression accelerator is enabled.

- disable: Compression accelerator is disabled.

<request #> is the Request ID assigned for each command execution. For details, see .

[-capacity_saving <capacity saving>]

Specifies the capacity saving setting. If you omit this option, a volume is created with the capacity saving function disabled.

- compression: Enables compression.

- deduplication_compression: Enables deduplication and compression.

- disable: Disables capacity saving.

For VSP 5000 series and VSP E series, compressed volumes for accelerated compression, or volumes for deduplication and compression can be created only for pools whose subscription limit is set to 65535 (unlimited). An error occurs if compression or deduplication_compression is specified for a pool whose subscription limit is other than 65535 (unlimited).

[-capacity_saving_mode <saving mode>]

When you enable capacity saving, specifies the post-process mode or inline mode as the capacity saving processing mode for DP-VOLs. For details about the capacity saving processing modes, see the *Provisioning Guide* for the storage system.

- post_process: Sets the post-process mode for capacity saving processing.

- inline: Sets the inline mode for capacity saving processing.

For VSP G1x00, VSP F1500, VSP G200, G400, G600, G800, and VSP F400, F600, F800 post_process is the default. For VSP 5000 series, VSP E series, and VSP G/F350, G/F370, G/F700, G/F900 inline is the default.

[-drs -request_id auto]

Supported storage systems:

- VSP 5000 series

Specify this option to create a data reduction shared volume (DRS-VOL).

<request #> is the Request ID assigned for each command execution. For details, see .

Examples

Creating an LDEV: 100 of size 10 GB in a parity group: 5-2.

Location in the parity group: automatic allocation, LDEV Emulation type: OPEN-V

```
# raidcom add ldev -parity_grp_id 5-2 -ldev_id 100 -capacity 10G
```

Chapter 5: Configuration setting commands

Creating an LDEV with the following conditions: Parity group: 5-3, location of the parity group: automatic allocation, Emulation type: 3390-3, Size: 10 M cylinder (Cylinder specification), LDEV: 120.

```
# raidcom add ldev -parity_grp_id 5-3 -ldev_id 120 -cylinder 10m -emulation 3390-3
```

Creating an LDEV of position in the external volume: allocated automatically, emulation type: OPEN-V, External volume: 01-02, Capacity: 200 MB, and LDEV number: 200.

```
# raidcom add ldev -external_grp_id 01-02 -ldev_id 200 -capacity 200m
```

Creating an LDEV of external volume group: 01-03, position in the external volume: allocated automatically, emulation type: OPEN-V, External volume size: takeover, and LDEV number: 220.

```
# raidcom add ldev -external_grp_id 01-03 -ldev_id 220 -capacity all
```

Creating a V-VOL of Capacity: 300 MB and number: 300, to a pool for Dynamic Provisioning, Dynamic Provisioning for Mainframe, Dynamic Tiering, active flash, Dynamic Tiering for Mainframe, or active flash for mainframe.

```
# raidcom add ldev -pool 4 -ldev_id 300 -capacity 300m
```

Creating a V-VOL of Capacity: 300 MB and LDEV number: 400, to a pool for Thin Image or Copy-on-Write Snapshot.

```
# raidcom add ldev -pool snap -ldev_id 400 -capacity 300m
```

Creating a V-VOL of Capacity: 300 MB, LDEV number: 300, and Full Allocation: enable, to the pool ID: 4 for Dynamic Provisioning.

```
# raidcom add ldev -pool 4 -ldev_id 300 -capacity 300m -status enable_fullallocation
```

Creating a V-VOL (LDEV ID: 44:44) for Dynamic Provisioning associated with the pool volume (LDEV ID: 22:22) that has the data direct mapping attribute in the Dynamic Provisioning pool.

```
#raidcom add ldev -ldev_id 44:44 -mapping_ldev_id 22:22
```

Creating a V-VOL of Capacity: 300 MB, LDEV number: 300, and the T10 PI attribute: valid, in a pool (pool ID: 4) for Dynamic Provisioning, Dynamic Provisioning for Mainframe, Dynamic Tiering, active flash, Dynamic Tiering for Mainframe, or active flash for mainframe.

```
# raidcom add ldev -pool 4 -ldev_id 300 -capacity 300m -t10pi_enable
```

Creating a V-VOL of Capacity: 300 MB, LDEV number: 300, and capacity saving setting: compression, in the pool (ID: 4) for Dynamic Provisioning.

```
# raidcom add ldev -pool 4 -ldev_id 300 -capacity 300m -capacity_saving compression
```

Creating a V-VOL of Capacity: 300 MB, LDEV number: 300, capacity saving setting: compression, and compression accelerator: enable, in the pool (ID: 4) for Dynamic Provisioning.

```
# raidcom add ldev -pool 4 -ldev_id 300 -capacity 300m -capacity_saving compression -
compression_acceleration enable -request_id auto
```

Creating a V-VOL of Capacity: 300 MB, LDEV number: 300, capacity saving setting: deduplication and compression, and capacity saving mode: inline, in the pool (ID: 4) for Dynamic Provisioning.

```
#raidcom add ldev -pool 4 -ldev_id 300 -capacity 300m -capacity_saving
deduplication_compression -capacity_saving_mode inline
```

Creating a V-VOL of capacity: 10 GB for a Dynamic Provisioning or Dynamic Tiering pool (pool ID: 4), and assigning an unused LDEV number automatically.

```
# raidcom add ldev -pool 4 -ldev_id auto -request_id auto -capacity 10G
REQID : 1
```

Creating a DRS-VOL with capacity: 300 MB, LDEV number: 300, and capacity saving setting: deduplication and compression in a Dynamic Provisioning pool (pool ID: 4).

```
#raidcom add ldev -pool 4 -ldev_id 300 -capacity 300m -capacity_saving
deduplication_compression -drs -request_id auto
```

⚠️ **Caution:** The following message is shown whenever an LDEV larger than 4 TB is specified. It does not necessarily indicate that there is a problem:

```
raidcom : Please check if LDEV over 4 TB will be used for replication
products, and then check the supported capacity.
```

raidcom delete ldev

Deletes the specified LDEVs or V-VOLs. A device group can also be specified instead of an LDEV.

This command is executed asynchronously with the command input. Check the completion of this process on the raidcom get command_status command.

📄 **Note:** If a failure occurs in a processor on an MP blade or a controller, this command cannot be run even on an MP blade or a controller on which no failure occurs.

Syntax

```
raidcom delete ldev {-ldev_id <ldev#> | -grp_opt <group option>
    -device_grp_name <device group name>  [<device name>] }
    [-operation initialize_capacity_saving]
```

Options and parameters

-ldev_id <ldev#> (0-65279)

Specifies the LDEV number (0-65279). For example:

- -ldev_id 200

-grp_opt <group option>

Specifies the device information extracted from the LDEV belonging to the device group. Specify "ldev" (fixed). The information of LDEV belonging to the device group is used.

-device_grp_name <device group name> [<device name>]

Specifies the name of device group (maximum 32 characters) to be operated.

To operate an LDEV in the device group, specify the device name of LDEV (maximum 32 characters) within the device group.

If the device name is omitted, all the LDEVs belonging in the device group are operated.

[-operation initialize_capacity_saving]

Deletes an LDEV for which the capacity saving setting is enabled. When you specify this option, use the `raidcom get command_status` command to check if the LDEV deletion started, and then use the `raidcom get ldev` command to check if VOL_TYPE of the LDEV is changed from REMOVING to NOT DEFINED. If you perform another operation during or immediately after deletion of the LDEV, the operation might fail. In this case, wait a while and then retry the operation.

(VSP 5000 series, VSP E series, VSP G130, G/F350, G/F370, G/F700, G/F900) When there are no pool-associated LDEVs for which the capacity saving setting is deduplication_compression, the deduplication system data volumes are automatically deleted asynchronously with the deletion of the LDEVs. To verify that a deduplication system data volume is deleted, use the `raidcom get pool -key saving` command to check the LDEV number, and use the `raidcom get ldev command` to verify that VOL_TYPE of the LDEV is changed from REMOVING to NOT DEFINED.

> 📄 **Note:**
> - Deleting a deduplication volume using this command might take a while. Also, more pool/physical capacity might be used. You cannot stop the deletion operation while it is in process.
> - If you want to delete all of the pool-associated volumes for which capacity saving is enabled, first execute the **raidcom modify ldev** command to block both the pool-associated volumes and the deduplication system data volumes, and then execute the **raidcom initialize pool** command. By executing the **raidcom initialize pool** command before deleting the volumes, you can shorten the time for the volume deletion process and prevent the increase of pool capacity to be used.

Examples

Deleting an LDEV: 200.

```
raidcom delete ldev -ldev_id 200
```

Deleting an LDEV belonging to the device group: grp1.

```
raidcom delete ldev -grp_opt ldev -device_grp_name grp1
```

Deleting an LDEV: 200, for which capacity saving setting is enabled.

```
# raidcom delete ldev -ldev_id 200 -operation initialize_capacity_saving
```

raidcom extend ldev

Extends the capacity of a V-VOL for Dynamic Provisioning, Dynamic Provisioning for Mainframe, Dynamic Tiering, Dynamic Tiering for Mainframe, active flash, or active flash for mainframe.

If the specified LDEV is not a V-VOL of Dynamic Provisioning, Dynamic Provisioning for Mainframe, Dynamic Tiering, Dynamic Tiering for Mainframe, active flash, or active flash for mainframe, the command is rejected with EX_ENOOBJ.

A device group can also be specified instead of an LDEV.

When the -request_id auto option is specified, if a failure occurs in a processor on an MP blade or a controller, this command cannot be run even on an MP blade or a controller on which no failure occurs.

If you specify the -request_id option for the asynchronous command, the processing is run asynchronously with the command input. Verify that the processing is complete by using the **raidcom get command_status** command. There is a limit to the number of asynchronous commands that can be accepted simultaneously. For details, see <u>Maximum number of acceptable asynchronous commands (on page 228)</u>.

Syntax

```
raidcom extend ldev {-ldev_id <ldev#> | -grp_opt <group option>
    -device_grp_name <device group name> [<device name>]}
    -capacity <size> | -offset_capacity <size> | -cylinder <size> [-request_id auto]
```

Options and parameters

-ldev_id <ldev#> (0-65279)

Specifies the LDEV number (0-65279). Specify LDEVs for Dynamic Provisioning, Dynamic Provisioning for Mainframe, Dynamic Tiering, Dynamic Tiering for Mainframe, active flash, or active flash for mainframe. For example:

- -ldev_id 200

-grp_opt <group option>

Specifies the device information extracted from the LDEV belonging to the device group. Specify "ldev" (fixed). The information of LDEV belonging to the device group is used.

-device_grp_name <device group name> [<device name>]

Specifies the name of device group (maximum 32 characters) to be operated.

To operate an LDEV in the device group, specify the device name of LDEV (maximum 32 characters) within the device group.

If the device name is omitted, all the LDEVs belonging in the device group are operated.

-capacity <size>

The increment size of capacity can be specified in bytes or blocks. When specifying in bytes, the unit is t/T (terabyte), g/G (gigabyte), m/M (megabyte), or k/K (kilobyte). If this unit is omitted, block (512 bytes) is used.

Example of specification:

1GB (gigabyte) is:

-capacity 1G, -capacity 1g, -capacity 1024M, -capacity 1024m, -capacity 1048576K, -capacity 1048576k, -capacity 2097152

-offset_capacity <size>

Specifies capacity. Corrects the size as well as GUI for the specified capacity. The capacity can be specified in bytes or blocks.

-cylinder <size>

Specifies the size in cylinder unit. When specifying, the unit is g/G (gigabyte), m/M (megabyte), or k/K (kilobyte). Cylinder is applied if you do not specify the unit.

When the emulation type for the Open System is specified, this option cannot be specified.

-request_id auto

(VSP 5000 series, VSP G1x00, VSP F1500, VSP E series, and VSP G130, G/F350, G/F370, G/F700, G/F900 only) Specifying this option is recommended when running the command. Specify this option to expand the capacity of a DRS-VOL. <request#> is a Request ID assigned each time the command is run. For more information, see .

If you refresh the Device Manager - Storage Navigator window during the capacity expansion of a DRS-VOL, refreshing the window might not be performed. For more information, see the Troubleshooting section in the System Administrator Guide.

Examples

Extending the capacity of a Dynamic Provisioning or Dynamic Provisioning for Mainframe V-VOL: 200 by 10 GB.

```
# raidcom extend ldev -ldev_id 200 -capacity 10G
```

Extending the capacity of a Dynamic Provisioning or Dynamic Provisioning for Mainframe V-VOL: 201 by 200 MB.

```
# raidcom extend ldev -ldev_id 201 -capacity 200M
```

Extending the capacity of Dynamic Provisioning or Dynamic Provisioning for Mainframe V-VOL belonging to the device group: grp1 by 200 MB.

```
# raidcom extend ldev -grp_opt ldev -device_grp_name grp1 -capacity 200M
```

Extending the capacity of a Dynamic Provisioning or Dynamic Provisioning for Mainframe V-VOL: 201 by 200 MB.

```
# raidcom extend ldev -ldev_id 201 -capacity 200M -request_id auto
REQID : 1
```

raidcom get ldev

Displays the information of the specified LDEV or the device file. A device group can also be specified instead of an LDEV.

Syntax

```
raidcom get ldev {-ldev_id <ldev#> … [-cnt <count>] |
   -grp_opt <group option> -device_grp_name <device group
   name> [<device name>] | -ldev_list <ldev list option>}
   [-key <keyword>][{-check_status | -check_status_not}
   <string>... [-time <time>]] [-time_zone <time zone>]
```

Options and parameters

-ldev_id <ldev#> …

Specifies the LDEV number (0-65279). For example:

- -ldev_id 200

- -ldev_id 100-110

- -ldev_id 100 -cnt 10

[-cnt <count>]

Specifies the count (2-65280).

If this option is omitted, the count is set to one.

-grp_opt <group option>

Specifies the device information extracted from the LDEV belonging to the device group. Specify "ldev" (fixed). The information of LDEV belonging to the device group is used.

-device_grp_name <device group name> [<device name>]

Specifies the name of device group (maximum 32 characters) to be operated.

To operate an LDEV in the device group, specify the device name of LDEV (maximum 32 characters) within the device group.

If the device name is omitted, all the LDEVs belonging in the device group are operated.

-ldev_list <ldev list option>

Specifies the type of LDEVs to display. Specify one of the following LDEV list options. If the LDEV that meets the specified condition does not exist, `[EX_ENODEV] No such device` is displayed.

- defined: Displays all implemented LDEVs.

- dp_volume: Displays LDEVs that have Dynamic Provisioning attributes.

  Combining this with -pool_id <pool id> specification displays LDEVs that have Dynamic Provisioning attributes related to the specified POOL.

- external_volume: Displays external volumes.

- undefined: Displays all LDEV numbers that are not implemented.

- mapped: Displays all LDEVs to which LU paths are defined.

  Combining this with -pool_id <pool id> specification displays LDEVs (defined the LU path) relate to the specified POOL.

- mapped_nvme: Displays all LDEVs that are assigned to namespaces.

- unmapped: Displays all LDEVs to which LU paths are defined. However, LDEVs, mainframe volumes, and namespaces that are not implemented are not displayed because they cannot specify the LU path.

  Combining this with -pool_id <pool id> specification displays LDEVs (not defined the LU path) relate to the specified POOL.

Chapter 5: Configuration setting commands

- journal -journal_id <journal id>: Displays LDEVs that belong to the specified journal.

- pool -pool_id <pool id>: Displays LDEVs that belong to the specified pool.

  If a -pool_id option is non-numeric, the specified value is recognized as a pool name to identify the pool ID.

- parity_grp -parity_grp_id <parity group id>: Displays LDEVs that belong to the specified parity group.

- mp_blade -mp_blade_id <mp#>: Displays LDEVs that are set to the specified MP blades.

- quorum: Displays LDEVs set for the quorum disks.

- clpr -clpr_id <clpr#> : Displays LDEVs belonging to the specified CLPR.

  Supported storage systems:

  - VSP 5000 series

  - VSP G1x00 and VSP F1500

  - VSP E series

  - VSP G/F350, G/F370, G/F700, G/F900

[-key <keyword>]

Specifies a display keyword.

If this option is omitted, basic LDEV information is displayed. If this option is specified, the following information is displayed. The following display keywords can be specified:

- front_end: Front-end information

- front_end_nvme: Front-end NVMe (NVMe-oF) information

- parity_grp: Parity group information

- external: External volume information

- tier: Tier information for the Dynamic Tiering or active flash V-VOL.

- software_saving: Information on the capacity saved by the capacity saving feature for volumes with the capacity saving feature enabled.

- naa: The NAA identifier of the LU (LUN WWN)

- nguid: Namespace identifier (NGUID: Namespace Global Unique Identifier)

- qos: Specifies the qos when displaying the QoS setting information. This display keyword can be specified only for VSP 5000 series, VSP E series, VSP G130, G/F350, G/F370, G/F700, G/F900.

- qos_monitor: Specifies the qos_monitor when displaying the performance monitor information related to the QoS of the target LDEV. This display keyword can be specified only for VSP 5000 series, VSP E series, VSP G130, G/F350, G/F370, G/F700, G/F900.

Chapter 5: Configuration setting commands

[-check_status <string>... [-time <time>]]

Check if the LDEV is in the same state as the specified in <string>. If the option contains multiple states, the OR condition check is performed and verifies that the LDEV is in one of the states contained in the option.

The following strings are specified in the <string>.

- STS
  - NML: Normal
  - BLK: Blocked
  - BSY: Status is changing
- OPE_TYPE
  - FMT: Formatting
  - QFMT: Quick formatting
  - CCOPY: Data copying (Correction copy / Copy back / Drive copy / Dynamic sparing)
  - CACCS: Accessing to collections
  - NONE: Not in operation
  - SHRD: Shredding
  - ZPD: Page discarding
  - SHRPL: Deleting from the pool
  - RLC: Pool relocating
  - RBL: Pool rebalancing
- VOL_TYPE
  - NOT DEFINED: An LDEV is not installed
  - DEFINING: An LDEV is being created
  - REMOVING: An LDEV is being deleted

If "-time" is specified, the status of the LDEV is checked every three seconds until the end of the specified <time> (seconds).

When this option is specified, the returned values are as follows.

- The LDEV is in one of the specified states: 0
- The LDEV is in none of the specified states (without -time option): 1
- The LDEV is in none of the specified states (when the specified <time> passed): EX_EWSTOT

[-check_status_not <string>... [-time <time>]]

Check that the LDEV is not in the same state as the specified in <string>. If the option contains multiple states, the NOR condition check is performed and verifies that the LDEV is not in any of the states contained in the option.

The following strings are specified in the <string>.

- STS
  - NML: Normal
  - BLK: Blocked
  - BSY: Status is changing
- OPE_TYPE
  - FMT: Formatting
  - QFMT: Quick formatting
  - CCOPY: Data copying (Correction copy / Copy back / Drive copy / Dynamic sparing)
  - CACCS: Accessing to collections
  - NONE: Not in operation
  - SHRD: Shredding
  - ZPD: Page discarding
  - SHRPL: Deleting from the pool
  - RLC: Pool relocating
  - RBL: Pool rebalancing
- VOL_TYPE
  - NOT DEFINED: An LDEV is not installed
  - DEFINING: An LDEV is being created
  - REMOVING: An LDEV is being deleted

If "-time" is specified, the status of the LDEV is checked every three seconds until the end of the specified <time> (seconds).

When this option is specified, the returned values are as follows.

- The LDEV is not in any of the specified states: 0
- The LDEV is in one of the specified states (without -time option): 1
- The LDEV is in one of the specified states (when the specified <time> passed): EX_EWSTOT

[-time_zone <time zone>]
Supported storage systems:

- VSP 5000 series
- VSP E series
- VSP G130, G/F350, G/F370, G/F700, G/F900

Specifies the time zone displayed in storage systems. If this option is omitted, the time in the time zone set for the storage system is displayed. This option is valid only when it is specified with the `-key qos` option or `-key qos_monitor` option.

- utc: Displays the time in Coordinated Universal Time (UTC).

A hyphen (-) is displayed if the DKCMAIN microcode version is not supported.

Some keywords might not be displayed depending on the LDEV attribute as shown below.

LDEV attribute		Front end	Parity group	External	Tier
Normal volume	Internal volume	Y	Y	N	Y
	External volume	Y	N	Y	Y
POOL Volume for Thin Image, Copy-on-Write Snapshot, Dynamic Provisioning, Dynamic Provisioning for Mainframe, Dynamic Tiering, active flash, Dynamic Tiering for Mainframe, or active flash for mainframe	Internal volume	N	Y	N	N
	External volume	N	N	Y	N
Journal volume	Internal volume	N	Y	N	N
	External volume	N	N	Y	N
Legend Y: Displayed, N: Not displayed					

Examples

- Internal volume examples
- External volume examples
- Dynamic Provisioning V-VOL examples
- Deduplication system data volume example
- Volume deletion examples
- Dynamic Tiering V-VOL examples
- Thin Image primary volume example
- Pool volume example
- rmawk command examples
- Capacity saving example

Internal volume examples

Example 1

Displaying the information of the LDEV number 577 (internal VOL).

```
# raidcom get ldev -ldev_id 577
```

```
Serial# : 63502 PHY_Serial#  : 302594
LDEV : 577 PHY_LDEV : 600
SL : 0
CL : 0
VOL_TYPE : OPEN-V-CVS
VOL_Capacity(BLK) : 2181120
VOL_Capacity (cyl): 22720 (Displayed only for the mainframe volume)
NUM_LDEV : 1
LDEVs : 577
NUM_PORT : 2
PORTs : CL2-E-0 1 Linux_X86 : CL2-E-1 1 Solaris
F_POOLID : NONE
VOL_ATTR : CVS
RAID_LEVEL : RAID1
RAID_TYPE : 2D+2D
NUM_GROUP : 1
RAID_GROUPs : 02-01
CMP : Enable
EXP_SPACE : V

DRIVE_TYPE : DKS2C-K072FC
DRIVE_Capa : 141822798
LDEV_NAMING : Oracle_data_1
STS : NML
OPE_TYPE : QFMT
OPE_RATE : 100
MP# : 2
SSID:001F
ALUA : Enable
RSGID : 0
PWSV_S : -
CL_MIG : Y
```

Example 2

Displaying the NAA identification information of the LDEV number 577 (internal VOL).

```
# raidcom get ldev -ldev_id 577 -key naa
```

```
Serial# : 63502 PHY_Serial# : 302594
```

Chapter 5: Configuration setting commands

```
LDEV : 577 PHY_LDEV : 600
SL : 0
CL : 0
VOL_TYPE : OPEN-V-CVS
VOL_Capacity(BLK) : 2181120
VOL_Capacity(cyl) : 22720  (Displayed for the mainframe volume only)
NUM_LDEV : 1
LDEVs : 577
NUM_PORT : 2
PORTs : CL2-E-0 1 Linux_X86 : CL2-E-1 1 Solaris
LUN-WWN : 60060e8016016400000010164000000893
F_POOLID : NONE
VOL_ATTR : CVS
RAID_LEVEL : RAID1
RAID_TYPE : 2D+2D
NUM_GROUP : 1
RAID_GROUPs : 02-01
CMP : Enable
EXP_SPACE : V
DRIVE_TYPE : DKS2C-K072FC
DRIVE_Capa : 141822798
LDEV_NAMING : Oracle_data_1
STS : NML
OPE_TYPE : QFMT
OPE_RATE : 100
MP# : 2
SSID : 001F
ALUA : Enable
RSGID : 0
PWSV_S : -
CL_MIG : Y
```

Description of each column in output example 1:

Serial#

Product serial number.

The virtual serial number is displayed when you specify a volume that is virtualized by the global storage virtualization function using the virtual LDEV number.

PHY_Serial#

If you specify a volume that is virtualized by the global storage virtualization function using the virtual LDEV number, the serial number of the storage system is displayed. In other cases, this item is not displayed.

LDEV

LDEV number.

PHY_LDEV

- LDEV number of VSP G1x00, VSP F1500, VSP Gx00 models and VSP Fx00 models or VSP 5000 series. This item is displayed only when you specify a volume that was virtualized by the global storage virtualization function using the virtual LDEV number.

- LDEV number of VSP or HUS VM. This item is displayed only when you specify a volume that was virtualized by the Resource ID takeover function using the virtual LDEV number when you input the **raidcom get ldev** command. This item and VIR_LDEV are not displayed at the same time.

VIR_LDEV

- (VSP G1x00, VSP F1500, VSP Gx00 models and VSP Fx00 models or VSP 5000 series): Virtual LDEV number. This item is displayed only when you specify a volume that was virtualized by the global storage virtualization function using the LDEV number and it is different from the virtual LDEV number. When the virtual LDEV number is not given to the volume, "FF:FE(65534)" is displayed. When you set the reserve attribute of global-active device to the volume, "FF:FF(65535)" is displayed. This item is displayed only when the volume is virtualized by global storage virtualization, and this item is displayed on the place where PHY_LDEV is displayed in the example (Displaying the information of the LDEV number 577 (internal VOL)). This item and PHY_LDEV are not displayed at the same time.

- (VSP or HUS VM) Virtual LDEV number. This item is displayed only when you specify a volume that was virtualized by the Resource ID takeover function using the LDEV number when you input the raidcom get ldev command. If the virtual LDEV number is not given, "FF:FE(65534)" is displayed. This item is displayed in the same position as PHY_LDEV. This item and PHY_LDEV are not displayed at the same time.

SL

SLPR information ("0" is displayed for all storage systems except USP V/VM).

CL

CLPR information.

VOL_TYPE

Emulation type of the relevant LDEV. Displays the same name as the product ID of the Inquiry command. For open-systems LDEVs, a string is appended to the emulation type to indicate the LDEV attribute:

- *n: LUN Expansion (LUSE)

- -CVS: Virtual LVI/LUN

- -A: ALU

- -S: SLU

- -CM: command device

> 📄 **Note:** You can check the emulation type of a volume that was virtualized by the global storage virtualization function by running the command specifying the following instances:
>
> - The instance for which the command device belonging to the same virtual storage machine as the virtualized volume is specified for HORCM_CMD in the configuration definition file
>
> - The instance for which the serial number of the virtual storage machine that the virtualized volume belongs to is defined for HORCM_VCMD in the configuration definition file

If the LDEV is not installed or is in the process of being created or deleted, the status is one of the following:

- NOT DEFINED: An LDEV is not installed.

- DEFINING: An LDEV is being created.

- REMOVING: An LDEV is being deleted.

VOL_Capacity (BLK)
   Capacity of LDEV in block size.

VOL_Capacity(cyl)
   Capacity of LDEV in cylinder size. Displayed only when the attribute is MF-VOL.

NUM_LDEV
   Number of LDEVs that compose the LU where the specified LDEV belongs.

LDEVs
   Number of LDEVs in the LU.

NUM_PORT
   Number of ports defined to the paths for relevant LDEV. When the relevant LDEV is a mainframe volume or a multiplatform volume without the LU path definition, the number of ports for the dummy LU is displayed. For details about the dummy LU, see the *Command Control Interface User and Reference Guide*.

PORTs
   Ports defined to the paths for the relevant LDEV. It lists up the ports defined to the paths for the relevant LDEV. When the relevant LDEV is a mainframe volume or a multiplatform volume without the LU path definition, the port number of the dummy LU is displayed.

   If the number of characters for the host group name is 17 or more, a maximum of 16 characters can be displayed. Execute the `raidcom get host_grp` command to display the host group name with 17 characters or more.

LUN-WWN

The NAA identifier (LUN WWN) of the first LU defined for the LDEV. A hyphen (-) is displayed for either of the cases below. This information is displayed only when the -key naa option is specified.

- An LU is not defined for the LDEV.

- When the microcode or firmware version that supports the virtual LDEV ID is used, and the virtual LDEV ID is not assigned to the LDEV.

F_POOLID

Pool ID if the LDEV is a component of the pool. In other cases, "NONE" is displayed.

VOL_ATTR

Attributes of the LDEV.

- CMD: Command device (open-systems only)

- CLUN: Cache LUN (DCR)

- CVS: CVS volume

- LUSE: LUSE volume

- ALUN: Volume Migration volume

- ELUN: External volume

- OLG: Open LDEV Guard volume

- VVOL: V-VOL

- HORC: The remote copy pair volumes:

    - TrueCopy/TrueCopy for Mainframe, Universal Replicator/Universal Replicator for Mainframe, global-active device: P-VOL or S-VOL

- MRCF: ShadowImage volume (P-VOL or S-VOL)

- QS: Thin Image or Copy-on-Write Snapshot volume (P-VOL or S-VOL)

- JNL: JNL volume

- HDP: volume for Dynamic Provisioning or Dynamic Provisioning for Mainframe

- HDT: volume for Dynamic Tiering (HDT), Dynamic Tiering for Mainframe, active flash, or active flash for mainframe

- POOL: POOL volume

- QRD: Quorum disk

- ENCD: Encryption disk

- SYSD: System disk

- TSE: Dynamic Provisioning for Mainframe volumes that are used in FlashCopy SE

- ESE: a virtual volume capable of page release by the User Directed Space Release function

- GAD: Global-active device volume

- HNASS: a volume used as the system LU of Hitachi NAS

- HNASU: a volume used as a user LU of Hitachi NAS

- MG: a volume for the data migration

- T10PI: a volume of which T10 PI is enabled

- DSD: Deduplication system data volume

- DS: Deduplication system data volume that stores duplicated data (data store)

- RCMD: Remote command device

RAID_LEVEL

RAID level (RAID1, RAID5, or RAID6).

RAID_TYPE

Configuration of the drives.

NUM_GROUP

Number of parity groups to which the relevant LDEV belongs.

RAID_GROUPs

Parity groups to which the relevant LDEVs belongs.

CMP

Accelerated compression setting for the parity group to which the LDEV belongs.

This item is displayed only for the VSP Fx00 models and VSP Gx00 models.

- Enable: Accelerated compression is enabled.

- Disable: Accelerated compression is disabled.

- NA: Accelerated compression is not supported.

- - (hyphen): This information is invalid for this LDEV.

For VSP F1500, VSP G1x00, VSP, and HUS VM, a hyphen (-) is always displayed.

EXP_SPACE

Displays whether the LDEV uses the accelerated compression area of the parity group.

This item is displayed only for the VSP Fx00 models and VSP Gx00 models.

- V: The LDEV uses the accelerated compression.

- R: The LDEV does not use accelerated compression.

- - (hyphen): This information is invalid for this LDEV.

For VSP F1500, VSP G1x00, VSP, and HUS VM, a hyphen (-) is always displayed.

DRIVE_TYPE

Drive type code which is set when the parity group to which the LDEV belongs is set.

To view the drive type code of the drive in the parity group, execute the raidcom get drive command.

Chapter 5: Configuration setting commands

DRIVE_Capa

Capacity of relevant HDD in the number of block (512 bytes) (decimal number).

LDEV_NAMING

Nickname of the LDEV.

STS

Status of the LDEV:

- NML: Normal

- BLK: Blocked

- BSY: Status is changing

- NONE: unknown state (not supported)

OPE_TYPE

Current operation.

- FMT: Formatting

- QFMT: Quick formatting

- CCOPY: Data copying (Correction copy / Copy back / Drive copy / Dynamic sparing)

- CACCS: Accessing to collections

- NONE: Not in operation

- SHRD: Shredding

- ZPD: Page Discarding

- SHRPL: Deleting from the pool

- RLC: Pool relocating

- RBL: Pool rebalancing

OPE_RATE

Progress of the format or shred operation. When the status is other than formatting or shredding, 100 is displayed. If the process ended abnormally, "BSY" is displayed on the "STS".

MP#

MP blade ID.

SSID

Storage subsystem ID number (hexadecimal). VSP E series, VSP Gx00 models, VSP Fx00 models, and Unified Storage VM do not support SSID but display the specified value.

> 📄 **Note:** You can check the SSID of a volume that was virtualized by the global storage virtualization function by running the command specifying the following instances:
>
> - The instance for which the command device belonging to the same virtual storage machine as the virtualized volume is specified for HORCM_CMD in the configuration definition file
>
> - The instance for which the serial number of the virtual storage machine that the virtualized volume belongs to is defined for HORCM_VCMD in the configuration definition file

ALUA

ALUA mode

- Enable: The ALUA mode is enabled.

- Disable: The ALUA mode is disabled.

RSGID

Displays the resource group ID of the resource group to which the LDEV belongs.

CL_MIG

Displays whether the allocation of LDEV to CLPR is changed.

- Y: The allocation of LDEV is being changed.

- N: The allocation of LDEV is not changed.

Example 3

Displaying front-end information of the LDEV number 577

```
# raidcom get ldev -ldev_id 577 -key front_end
```

```
 Serial# LDEV# SL CL VOL_TYPE  VOL_Cap(BLK) PID ATTRIBUTE Ports
  PORT_No:LU#:GRPNAME …
63502    577   0  0 OPEN-V-CVS 2181120     -        CVS 2
  CL2-E-0:1:Linux_X86 CL2-E-1:1:Solaris
```

Description of each column in output example 3

PID

Pool ID for a virtual volume. If the volume is not a virtual volume, a hyphen (-) is displayed.

ATTRIBUTE

LDEV attribute. Same as VOL_ATTR.

Chapter 5: Configuration setting commands

Example 4

Displaying front-end nvme (NVMe-oF) information of the LDEV number 577

```
# raidcom get ldev -ldev_id 577 -key front_end_nvme
Serial# LDEV# SL CL VOL_TYPE     VOL_Cap(BLK)  PID ATTRIBUTE NVMSS_ID
NS_ID Ports  PORT_No
63502    577 0  0 OPEN-V-CVS   2181120 -   CVS       1
1 2    CL2-E CL2-F
```

Description of each column in output example 4

NVMSS_ID

NVM subsystem ID.

NSID

Namespace ID.

PORT_No

NVM subsystem port number.

If the NVM subsystem port number is not assigned, a hyphen (-) is displayed.

Example 5

Displaying back-end (parity group) information of the LDEV number 577.

```
# raidcom get ldev -ldev_id 577 -key parity_grp
```

```
 Serial# LDEV# SL CL PID ATTRIBUTE R_LVL RAID_TYPE DRV_TYPE DRV_Cap
GRPs RAID_GRP …
   63502 577 0 0 - CVS RAID1 2D+2D DKS2C-K072FC 141822798 1 02-01
```

Example 6

Check if the LDEV number 577 is in normal status.

```
# raidcom get ldev -ldev_id 577 -check_status NML
```

Example 7

Check if the LDEV number 577 is in blocked status.

```
# raidcom get ldev -ldev_id 577 -check_status BLK
```

Example 8

Set in wait status for 30 seconds until the formatting process of LDEV number 577 is complete.

```
# raidcom get ldev -ldev_id 577 -check_status_not FMT -time 30
raidcom:[EX_EWSTOT]Timeout waiting for specified status
```

Refer to the command log(/HORCM/log0/horcc_rmhost.log) for details.

(The command fails if the FMT does not change (is not complete) in 30 seconds.)

External volume examples

Example 1

Displaying the information of the LDEV number 160 (external VOL).

```
# raidcom get ldev -ldev_id 160
Serial# : 63502
LDEV : 160
SL : 0
CL : 0
VOL_TYPE : OPEN-V
VOL_Capacity(BLK) : 4385280
NUM_PORT : 0
PORTs :
F_POOLID : NONE
VOL_ATTR : ELUN : RCMD

E_VendorID : HITACHI
E_ProductID : OPEN-V
E_VOLID :
484954414348492052353030463830453030364100000000000000000000000000000000
E_VOLID_C : HITACHI R500F80E006A................
NUM_E_PORT : 1
E_PORTs : CL2-G-0 0 50060e8004f80e34
LDEV_NAMING : Oracle_data_1
STS : NML
OPE_TYPE : QFMT
OPE_RATE : 70
MP# : 2
SSID:001F
ALUA : Enable
RSGID : 0
CL_MIG : Y
```

Description of columns in output example 1:

E_VendorID
> Vendor name that the external storage system's volume notifies to the host.

E_ProductID

System name that the external storage system's volume notifies to the host.

E_VOLID

Number to identify the external volume (hexadecimal).

E_VOLID_C

Number to identify the external volume (ASCII display).

NUM_E_PORT

Number of alternate paths.

E_PORTs:

List of defined alternate paths.

Example 2

If the LDEV number 160 is an external volume, its back-end (RAID group) information is displayed.

```
# raidcom get ldev -ldev_id 160 -key external
```

```
Serial#  LDEV#  SL CL PID ATTRIBUTE  E_VendorID  E_ProductID
E_VOLID
"E_VOLID_C"                          E_PORTs PORT_NO:LU#:WWN ...
  63502    160   0  0   - ELUN        HITACHI    OPEN-V
48495441434849205235303046383045303036410000000000000000000000000000000 "HITACHI
R500F80E006A..............."      1 CL2-G-0:0:50060e8004f80e34
```

Example 3

Displaying the information of the LDEV number 39320 (quorum disk).

```
# raidcom get ldev -ldev_id 39320
```

```
Serial#  : 302656
LDEV : 39320
SL : 0
CL : 0
VOL_TYPE : OPEN-V-CVS
VOL_Capacity(BLK) : 31457280
NUM_PORT : 0
PORTs :
F_POOLID : NONE
VOL_ATTR : CVS : ELUN : QRD
E_VendorID : HITACHI
E_ProductID : OPEN-V
E_VOLID : 48495441434849202353033330304135363132303000000000000000000000000000000000
E_VOLID_C : HITACHI 50300A561200...............
```

Chapter 5: Configuration setting commands

```
NUM_E_PORT : 1
E_PORTs : CL1-B-0 0 50060e80070a5630
LDEV_NAMING :
STS : BLK
OPE_TYPE : NONE
OPE_RATE : 100
MP# : 0
SSID : 004D
QRDID : 31
QRP_Serial# : 302646
QRP_ID : R8
ALUA : Disable
RSGID : 0
CL_MIG : Y
```

Description of a column in output example 2 and 3

QRDID

> ID of the quorum disk for HAM or GAD.

QRP_Serial#

> Shows the serial number of the storage system when the external volume is the quorum disk for HAM or GAD.

QRP_ID

> Shows the ID for identifying storage systems when the external volume is the quorum disk for HAM or GAD.
>
> - R9: VSP 5000 series
>
> - R8: VSP G1x00 and VSP F1500
>
> - R7: VSP
>
> - M8: VSP E series, VSP G/F350, G/F370, G/F700, G/F900, VSP G200, G400, G600, G800, and VSP F400, F600, F800
>
> - M7: HUS VM

Dynamic Provisioning V-VOL examples

Example

Displaying the information of the LDEV number 4368 (V-VOL of Dynamic Provisioning).

```
# raidcom get ldev -ldev_id 4368
```

```
Serial#  : 302614
LDEV : 4368
SL : 0
CL : 0
VOL_TYPE : OPEN-V-CVS
```

Chapter 5: Configuration setting commands

```
VOL_Capacity(BLK) : 16777216
NUM_PORT : 0
PORTs :
F_POOLID : NONE
VOL_ATTR : CVS : HDP
B_POOLID : 0
S_POOLID : 6
LDEV_NAMING :
STS : BLK
OPE_TYPE : NONE
OPE_RATE : 100
MP# : 0
SSID : 0006
Used_Block(BLK) : 0
FLA(MB) : Disable
RSV(MB) : 0
CSV_Status : ENABLING
CSV_PROGRESS(%): 60
CSV_Mode : DEDUP+COMPRESS
COMPRESSION_ACCELERATION : ENABLED
COMPRESSION_ACCELERATION_STATUS : ENABLED
CSV_PROCESS_MODE : POST_PROCESS
DEDUPLICATION_DATA : ENABLED
ALUA : Disable
RSGID : 0
DM_LDEV : 4096
PWSV_S:PS
CL_MIG : Y
```

Description of columns in output example:

B_POOL ID

Pool ID to which the LDEV is associated.

S_POOLID

Snapshot pool ID to which the LDEV is associated. This item is displayed only for the S-VOL whose LDEV attribute is HDP and QS. When S_POOLID is displayed, B_POOL ID indicates the ID of the pool for Dynamic Provisioning, Dynamic Tiering, or active flash.

Used_Block(BLK)

Number of blocks used in the pool. This number of blocks includes the reserved blocks by Full Allocation. When Full Allocation is changed to enabled from disabled, the value of the Used_Block(BLK) will be increased by an amount equal to the number of blocks which are reserved by Full Allocation.

FLA(MB)

Shows the capacity which is reserved by Full Allocation or Proprietary Anchor. When Full Allocation is disabled, "Disable" is displayed.

RSV(MB)

Shows the capacity that is reserved by Full Allocation or Proprietary Anchor.

CSV_Status

Shows the status of capacity saving:

- DISABLED: Capacity saving is disabled. For VSP and HUS VM, DISABLED is always displayed.

- ENABLED: Capacity saving is enabled.

- ENABLING: Capacity saving is being enabled.

- REHYDRATING: Capacity saving is being disabled.

- DELETING: A volume which has capacity saving enabled is being deleted.

- FAILED: Consistency of the data in the deduplication system data volume cannot be guaranteed.

- CONVERTING: Compression method is being changed.

CSV_PROGRESS(%)

Shows the progress rate when the capacity saving status is ENABLING, REHYDRATING, or DELETING. For other statuses, a hyphen (-) is displayed. In addition, for VSP and HUS VM, a hyphen (-) is always displayed.

> **Note:** When the status is REHYDRATING, post-processing takes time, and a hyphen (-) might be displayed for CSV_PROGRESS(%) for a while.

CSV_Mode

Shows the capacity saving setting.

- DISABLED: Capacity saving is disabled. For VSP and HUS VM, DISABLED is always displayed.

- COMPRESS: Compression is enabled.

- DEDUP+COMPRESS: Deduplication and compression are enabled.

CSV_PROCESS_MODE

Shows the capacity saving processing mode.

- POST_PROCESS: Post-process mode.

- INLINE: Inline mode.

- - (hyphen): Capacity saving is disabled. For VSP and HUS VM, a hyphen (-) is always displayed.

COMPRESSION_ACCELERATION

Shows whether the compression accelerator setting is enabled or disabled. Microcode versions that do not support compression accelerator display - (hyphen).

- ENABLED: Compression accelerator is enabled.

- DISABLED: Compression accelerator is disabled.

- - (hyphen): This information is invalid for this LDEV.

Chapter 5: Configuration setting commands

COMPRESSION_ACCELERATION_STATUS

Shows whether compression accelerator is applied or not to the data in the volume. Microcode versions that do not support compression accelerator display - (hyphen).

- ENABLED: Compression accelerator is applied to all data on the volume.

- DISABLED: Compression accelerator is not applied to any data on the volume.

- HYBRID: The volume contains data that is compressed with compression accelerator and data that is compressed without using compression accelerator.

- - (hyphen): This information is invalid for this LDEV.

DEDUPLICATION_DATA

Shows the applied status for deduplication.

- DISABLED: Deduplication function is not applied to a virtual volume.

- ENABLED: Deduplication function is applied to a virtual volume.

  If DEDUP+COMPRESS is selected as the capacity saving setting for a virtual volume, ENABLED is shown even when the amount of used capacity is 0. Also, when the capacity saving setting is being disabled and the capacity saving status for a virtual volume is DISABLED, ENABLED is shown when deduplicated data is remaining.

- - (hyphen): The volume is not a virtual volume, or the microcode or firmware version does not support the deduplication function. For VSP and HUS VM, a hyphen (-) is always displayed.

DM_LDEV

Supported storage systems:

- VSP 5000 series

- VSP G1x00 and VSP F1500

- VSP E series

- VSP G/F350, G/F370, G/F700, G/F900

- VSP G200, G400, G600, G800 and VSP F400, F600, F800

Shows the LDEV number of the pool volume associated with a V-VOL for Dynamic Provisioning that has the data direct mapping attribute. All LBAs of the V-VOL for Dynamic Provisioning are mapped to LBAs of the pool volume one-on-one. DM_LDEV is displayed only when the LDEV has the data direct mapping attribute.

Dynamic Provisioning V-VOL namespace examples

Examples

Displaying the information about the LDEV number 4368 (namespace of Dynamic Provisioning V-VOL) (VSP 5000 series and VSP E1090).

```
# raidcom get ldev -ldev_id 4368
```

```
Serial# : 302614
LDEV : 4368
SL : 0
CL : 0
VOL_TYPE : OPEN-V-CVS
VOL_Capacity(BLK) : 16777216
NUM_PORT : 0
PORTs : -
NVM_NUM_PORT : 2
NVM_PORTs : CL1-A CL2-B
NSID : 2
NVMSS_ID : 1
F_POOLID : NONE
VOL_ATTR : CVS : HDP
B_POOLID : 0
S_POOLID : 6
LDEV_NAMING :
STS : BLK
OPE_TYPE : NONE
OPE_RATE : 100
MP# : 0
CURRENT_MP# : 0
SSID : 0006
Used_Block(BLK) : 0
FLA(MB) : Disable
RSV(MB) : 0
CSV_Status : ENABLING
CSV_PROGRESS(%): 60
CSV_Mode : DEDUP+COMPRESS
COMPRESSION_ACCELERATION : ENABLED
COMPRESSION_ACCELERATION_STATUS : ENABLED
CSV_ PROCESS_MODE : POST_PROCESS
DEDUPLICATION_DATA : ENABLED
ALUA : Disable
RSGID : 0
DM_LDEV:4096
PWSV_S : PS
CL_MIG : Y
```

Chapter 5: Configuration setting commands

Displaying the NGUID of the LDEV number 4368 (namespace of Dynamic Provisioning V-VOL) (VSP 5000 series and VSP E1090).

```
# raidcom get ldev -ldev_id 4368 -key nguid
```

```
Serial# : 302614
LDEV : 4368
SL : 0
CL : 0
VOL_TYPE : OPEN-V-CVS
VOL_Capacity(BLK) : 16777216
NUM_PORT : 0
PORTs : -
NVM_NUM_PORT : 2
NVM_PORTs : CL1-A CL2-B
NSID : 2
NVMSS_ID : 1
LDEV_NGUID : 12345…
F_POOLID : NONE
VOL_ATTR : CVS : HDP
B_POOLID : 0
S_POOLID : 6
LDEV_NAMING :
STS : BLK
OPE_TYPE : NONE
OPE_RATE : 100
MP# : 0
CURRENT_MP# : 0
SSID : 0006
Used_Block(BLK) : 0
FLA(MB) : Disable
RSV(MB) : 0
CSV_Status : ENABLING
CSV_PROGRESS(%): 60
CSV_Mode : DEDUP+COMPRESS
COMPRESSION_ACCELERATION : ENABLED
COMPRESSION_ACCELERATION_STATUS : ENABLED
CSV_ PROCESS_MODE : POST_PROCESS
DEDUPLICATION_DATA : ENABLED
ALUA : Disable
RSGID : 0
DM_LDEV:4096
PWSV_S : PS
CL_MIG : Y
```

Description of columns in output example:

NVM_NUM_PORT

Number of NVM subsystem ports to which the LDEV is defined.

NVM_PORTs

> NVM subsystem port number to which the LDEV is defined.

LDEV_NGUID

> Namespace identifier (NGU ID). "Unknown" is displayed if the NGUID is not assigned.

NSID

> Namespace ID.

NVMSS_ID

> NVM subsystem ID.

Deduplication system data volume example

Example

Displaying the information of the LDEV number 4368 (deduplication system data volume).

```
# raidcom get ldev -ldev_id 4368
```

```
Serial#  : 302614
LDEV : 4368
SL : 0
CL : 0
VOL_TYPE : OPEN-V-CVS
VOL_Capacity(BLK) : 16777216
NUM_PORT : 0
PORTs :
F_POOLID : NONE
VOL_ATTR : CVS : HDP : DSD
B_POOLID : 0
LDEV_NAMING :
STS : BLK
OPE_TYPE : NONE
OPE_RATE : 100
MP# : 0
SSID : 0006
Used_Block(BLK) : 0
FLA(MB) : Disable
RSV(MB) : 0
CSV_Status : DISABLED
CSV_PROGRESS(%): -
CSV_Mode : DISABLED
COMPRESSION_ACCELERATION : -
COMPRESSION_ACCELERATION_STATUS : -
CSV_PROCESS_MODE : -
DEDUPLICATION_DATA : DISABLED
ALUA : Disable
```

```
RSGID : 0
CL_MIG : Y
```

Volume deletion examples

Example 1

Displaying the information of the LDEV being deleted (VOL_ATTR is other than HDP).

```
# raidcom get ldev -ldev_id 4096
```

```
Serial# : 64568
LDEV : 4096
SL : -
CL : -
VOL_TYPE : REMOVING
SSID : 0005
RSGID : 0
```

Example 2

Displaying the information of the LDEV being deleted by the raidcom delete ldev command without the `-operation initialize_capacity_saving` option (VOL_ATTR is HDP).

```
# raidcom get ldev -ldev_id 4096
```

```
Serial# : 64568
LDEV : 4096
SL : -
CL : -
VOL_TYPE : REMOVING
CSV_Status : DISABLED
CSV_PROGRESS(%) : -
SSID : 0005
RSGID : 0
```

Example 3

Displaying the information of the LDEV being deleted by the raidcom delete ldev command with the `-operation initialize_capacity_saving` option (VOL_ATTR is HDP).

```
# raidcom get ldev -ldev_id 4096
```

```
Serial# : 64568
LDEV : 4096
SL : -
```

```
CL : -
VOL_TYPE : REMOVING
CSV_Status : DELETING
CSV_PROGRESS(%) : 30
SSID : 0005
RSGID : 0
```

Dynamic Tiering V-VOL examples

Example 1

Displaying the information of the LDEV number 640 (V-VOL of Dynamic Tiering).

```
# raidcom get ldev -ldev_id 640
```

```
Serial# : 63502
LDEV : 640
SL : 0
CL : 0
VOL_TYPE : OPEN-V-CVS
VOL_Capacity(BLK) : 2181120
NUM_LDEV : 1
LDEVs : 640
NUM_PORT : 1
PORTs : CL2-E-0 14 Linux_X86
F_POOLID : NONE
VOL_ATTR : CVS : HDP : HDT
B_POOLID : 5
LDEV_NAMING : Oracle_data_1
STS : NML
OPE_TYPE : NONE
OPE_RATE : 100
MP# : 2
SSID : 001F
Used_Block(BLK) : 218112
TIER_Relocation : Enable
TIER_LEVEL: 6
TIER#1(MB) : 1120
TIER#2(MB) : 3000
    :
    :
TIER_Alloc_level : H
TIER#1_Alloc_rate : MAX : 50 : MIN : 30
TIER#3_Alloc_rate : MAX : 50 : MIN : 30
FLA(MB) : 980
RSV(MB) : 980
CSV_Status : DISABLED
CSV_PROGRESS(%): -
CSV_Mode : DISABLED
```

Chapter 5: Configuration setting commands

```
COMPRESSION_ACCELERATION : -
COMPRESSION_ACCELERATION_STATUS : -
CSV_PROCESS_MODE : -
DEDUPLICATION_DATA : DISABLED
ALUA : Enable
RSGID : 0
CL_MIG : Y
```

Description of columns in output example 1:

TIER_Relocation

Status of the relocation setting:

- Enable: The status where the relocation is enabled

- Disable: The status where the relocation is disabled

TIER_LEVEL

Shows the level of tiering policy that is used for reallocation.

- all: Shows that all tiers of the pool to which the relevant LDEV is allocated are being used.

- 1-5: Shows the tiering policy level that is configured to the relevant LDEV.

- 6-31: Shows the tiering policy (customized policy (1-26)) that is configured to the relevant LDEV. For details, see the *Provisioning Guide* or *LUN Manager User Guide* for the storage system.

TIER#n(MB)

Shows the capacity (in MB) allocated to each tier.

TIER_Alloc_level

Shows the tier level of the new mapped page.

- H: High

- M: Middle

- L: Low

TIER#1_Alloc_rate

Shows the Tier1 Max or Min value that is set in the tiering policy.

TIER#3_Alloc_rate

Shows the Tier3 Max or Min value that is set in the tiering policy.

FLA(MB)

Shows the capacity that is reserved by Full Allocation. When Full Allocation is disabled, "Disable" is displayed.

Example 2

Displays the tier information about the LDEV number 640 (V-VOL of Dynamic Tiering).

```
# raidcom get ldev -ldev_id 640 -key tier
```

```
Serial# LDEV#  SL  CL VOL_TYPE VOL_Cap(BLK) PID ATTRIBUTE VOL_Used(BLK) TR TL T#1(MB)
T#2(MB) ...
63502 640 0 0 OPEN-V-CVS 204800 13 CVS|2181120 - CVS|VVOL|HDT 218112 E 6 1120 3000
```

Description of each column in output example 2:

TR

> Displays the enabled or disabled of the tier relocation.

> - E: Enabled

> - D: Disabled

TL

> Displays the tier level.

> - 0: ALL

> - 1-5: Level (1-5)

> - 6-31: Customized policy (1-26)

T#x(MB)

> Displays the allocated volume of the relevant LDEV for Tier x. Displays up to 5 tiers. If there is no tier, "0" is displayed.

Thin Image primary volume example

Example

Displays the information about the LDEV number 1000 (primary volume of Thin Image).

```
# raidcom get ldev -ldev_id 1000
```

```
Serial#  : 64568
LDEV : 1000
SL : 0
CL : 0
VOL_TYPE : OPEN-V-CVS
VOL_Capacity(BLK) : 204800
NUM_PORT : 2
PORTs : CL1-A-0 0 1A-G00 : CL5-B-0 1 5B-G00
F_POOLID : NONE
VOL_ATTR : CVS : QS
```

Chapter 5: Configuration setting commands

```
RAID_LEVEL  : RAID1
RAID_TYPE   : 2D+2D
NUM_GROUP : 1
RAID_GROUPs : 01-02
DRIVE_TYPE  : DKR2G-K146SS
DRIVE_Capa : 285177528
LDEV_NAMING :
STS : NML
OPE_TYPE : NONE
OPE_RATE : 100
MP# : 2
SSID : 0009
ALUA : Enable
RSGID : 0
Snap_Used_Pool(MB) : 100
CL_MIG : Y
SNAP_USED(MB) : 126
SNAP_GARBAGE(MB) : 15
DELETING_SNAP_GARBAGE : NONE
DELETING_SNAP_GARBAGE(%) : -
```

Description of column in output example:

Snap_Used_Pool(MB)

For the Thin Image root volume, displays the capacity (in MB) in the root volume used for the snapshot data out of the snapshot capacity assigned from the pool. If the used capacity is less than 1 MB, the displayed value is rounded up.

SNAP_USED(MB)

For the Thin Image root volume, displays the capacity (in MB) in the root volume for the snapshot data area. If the used capacity is less than 1 MB, the displayed value is rounded up. This item is displayed only for VSP 5000 series and VSP G1x00, VSP F1500 . For a Thin Image Advanced pair, a hyphen (-) is always displayed.

SNAP_GARBAGE(MB)

Displays the garbage data volume (in MB) in the snapshot data area of the specified Thin Image root volume. If the garbage data volume is less than 1 MB, the displayed value is rounded up. While the snapshot garbage data is being deleted (defragmentation is in progress), a hyphen (-) is displayed. This item is displayed only for VSP 5000 series. For other models, this item is not displayed, or a hyphen (-) is always displayed. For a Thin Image Advanced pair, a hyphen (-) is always displayed.

DELETING_SNAP_GARBAGE

Displays the execution status of the snapshot garbage data deletion process (defragmentation process).

This item is displayed only for VSP 5000 series. For other models, this item is not displayed, or a hyphen (-) is always displayed.

- NONE: Not in operation

- PROCESSING: Processing is in progress.

  The execution status is PROCESSING even when processing is suspended.

  See the Hitachi Thin Image User Guide for the conditions when the processing is suspended.

- STOPPING: Defragmentation is being stopped (transition from PROCESSING to NONE).

For a Thin Image Advanced pair, a hyphen (-) is always displayed.

DELETING_SNAP_GARBAGE(%)

Displays the progress rate of the snapshot garbage data deletion processing (defragmentation processing). This item is displayed only for VSP 5000 series. For other models, this item is not displayed, or a hyphen (-) is always displayed. For a Thin Image Advanced pair, a hyphen (-) is always displayed.

Pool volume example

Example

Displays the information about the LDEV number 4096 (pool volume).

```
# raidcom get ldev -ldev_id 4096
```

```
Serial#  : 64568
LDEV : 4096
SL : 0
CL : 0
VOL_TYPE : OPEN-V-CVS
VOL_Capacity(BLK) : 16777216
NUM_LDEV : 1
LDEVs : 4096
NUM_PORT : 0
PORTs :
F_POOLID : 127
VOL_ATTR : CVS : POOL
RAID_LEVEL  : RAID5
RAID_TYPE   : 3D+1P
NUM_GROUP : 1
RAID_GROUPs : 01-01
DRIVE_TYPE  : DKR2G-K146SS
DRIVE_Capa : 285177528
LDEV_NAMING :
STS : BLK
OPE_TYPE : NONE
```

```
OPE_RATE : 100
MP# : 0
SSID : 0005
ALUA : Disable
RSGID : 0
DM_LDEV : 640
CL_MIG : Y
```

Description of column in output example:

DM_LDEV

> Shows the LDEV number of the V-VOL for Dynamic Provisioning associated with a pool volume that has the data direct mapping attribute. When any V-VOL for Dynamic Provisioning is not associated with the pool volume, "NONE" is displayed. DM_LDEV is displayed only when the LDEV has the data direct mapping attribute.

rmawk command examples

Example 1

Displays the information of the used Tier 1 capacity for the pool ID 73 of the HDT volume (using the **rmawk** command).

```
# raidcom get ldev -ldev_list dp_volume -pool_id 73 -key tier |rmawk -EC @L-ne:0 @12?
=ad:@12 -n  exe="print Total = @12?"


Total = 8064
```

Example 2

Displays the information of each used Tier capacity for the pool ID 73 of the HDT volume (using the **rmawk** command).

```
# raidcom get ldev -ldev_list dp_volume -pool_id 73 -key tier |rmawk -EC @L-ne:0
@12=ad:@12? @13=ad:@13? @14=ad:@14? exe="print @0" -n  exe="print " exe="print Total
= T#1(MB): @12? T#2(MB): @13? T#3(MB): @14?"
```

```
Serial# LDEV# SL CL VOL_TYPE    VOL_Cap(BLK) PID ATTRIBUTE      VOL_ Used(BLK) TR TL
T#1(MB)  T#2(MB)...
64558   29440 0  0  OPEN-V-CVS 4042752       73  CVS|HDP|HDT     404 2752      E  5
0        0         1974    0       0
64558   29441 0  0  OPEN-V-CVS 10924032      73  CVS|HDP|HDT    1092 4032      E  6
5334     0         0       0       0
64558   29442 0  0  OPEN-V-CVS 10924032      73  CVS|HDP|HDT    1092 4032      E  13
1596     1638      2100    0       0
64558   29443 0  0  OPEN-V-CVS 10924032      73  CVS|HDP|HDT    1092 4032      E  18
1134     3654      546     0       0
```

Chapter 5: Configuration setting commands

```
Total = T#1(MB):8064 T#2(MB):5292 T#3(MB):4620
```

Example 3

Displays information of the LDEV (when VOL_ATTR is other than HDP) which is being deleted.

```
# raidcom get ldev -ldev_id 4096
Serial# : 64568
LDEV : 4096
SL : -
CL : -
VOL_TYPE : REMOVING
SSID : 0005
RSGID : 0
```

Example 4

Displays information of the LDEV which is being deleted by the **raidcom delete ldev** command without -operation initialize_capacity_saving (when VOL_ATTR is other than HDP).

```
# raidcom get ldev -ldev_id 4096
Serial# : 64568
LDEV : 4096
SL : -
CL : -
VOL_TYPE : REMOVING
CSV_Status : DISABLED
CSV_PROGRESS(%) : -
SSID : 0005
RSGID : 0
```

Example 5

Displays information of the LDEV which is being deleted by the **raidcom delete ldev** command without -operation initialize_capacity_saving (when VOL_ATTR is other than HDP).

```
# raidcom get ldev -ldev_id 4096
Serial# : 64568
LDEV : 4096
SL : -
CL : -
VOL_TYPE : REMOVING
CSV_Status : DELETING
CSV_PROGRESS(%) : 30
SSID : 0005
RSGID : 0
```

Capacity saving examples

Example

Displays the information on the capacity saved by the capacity saving feature for the LDEV number 18 (volume with the capacity saving feature enabled).

For details about the capacity saved by the capacity saving feature, see the *Provisioning Guide* for your storage system.

```
# raidcom get ldev -ldev_id 18 -key software_saving
LDEV#    TLS_R    TOTAL_SAVING(BLK)    CMP(BLK)    DDP(BLK)
RECLAIM(BLK)    SYSTEM(BLK)    PRE_USED(BLK)    POOL_USED(BLK)
18    4.00    15728640    6291456    8388608
2097152    1048576    20971520    5242880
```

Description of columns in output example:

TLS_R

> Displays the ratio of the amount of data reduced by the capacity saving feature against the capacity before reduction.
>
> For example, if the capacity before and after reduction is 4.00 to 1, 4.00 is displayed.

TOTAL_SAVING(BLK)

> Displays the capacity reduced by the capacity saving feature in blocks. The capacity includes the amount of zero data reduction, metadata, and garbage data.

CMP(BLK)

> Displays the capacity reduced by compression of the capacity saving feature in blocks. The reduced capacity does not include the amount of metadata and garbage data.

DDP(BLK)

> Displays the capacity reduced by deduplication of the capacity saving feature in blocks. The reduced capacity does not include the amount of metadata and garbage data.

RECLAIM(BLK)

> Displays the capacity reduced by reclaiming the specified data pattern using the capacity saving feature in blocks. The reduced capacity does not include the amount of metadata and garbage data.

SYSTEM(BLK)

> Displays the amount of consumed system data (the amount of metadata and garbage data) of the capacity saving feature in blocks. The amount of metadata capacity and garbage data for deduplication system data volumes is not included .

PRE_USED(BLK)

> Displays the capacity before capacity saving of data to be deleted by the capacity saving feature in blocks.

POOL_USED(BLK)

> Displays the used capacity of the pool volumes used by volumes.
>
> For VSP G1x00, VSP F1500, a hyphen (-) is always displayed.

QoS examples

Displays the QoS setting information of LDEV number 200 (VSP 5000 series, VSP E series, VSP G130, G/F350, G/F370, G/F700, G/F900)

Nothing is displayed if QoS is not set. If the acquisition of QoS setting information fails, "NA" is displayed for items other than LDEV number.

```
# raidcom get ldev -ldev_id 200 -key qos
LDEV# UPPER_THROUGHPUT(IOps) UPPER_DATA_TRANS(MBps) UPPER_ALERT(s)
UPPER_ALERT_NOTICE LOWER_THROUGHPUT(IOps) LOWER_DATA_TRANS(MBps)
LOWER_ALERT(s) LOWER_ALERT_NOTICE PRIORITY RESPONSE_TARGET(ms)
RESPONSE_ALERT(s) RESPONSE_ALERT_NOTICE
```

Chapter 5: Configuration setting commands

```
200                          80                      100              30
2019-07-31T10:15:20                     20
-               0 2019-07-31T10:15:20         1
80              100  2019-07-31T10:15:20
```

Description of each column in output example:

UPPER_THROUGHPUT(IOps)

Displays the upper limit of the throughput.

For more information on the upper limit of the throughput per second, see the *Performance Guide*.

UPPER_DATA_TRANS(MBps)

Displays the upper limit of the data transfer volume.

For more information on the upper limit of the data transfer volume, see the *Performance Guide*.

UPPER_ALERT(s)

Displays the following alert notification time (second):

- The alert notification time (second) when the number of received commands continuously exceeds the upper limit of the throughput per second.

- The alert notification time (second) when the transfer volume of received commands continuously exceeds the upper limit of the data transfer volume per second.

UPPER_ALERT_NOTICE

Displays the last alert time of the following alert.

- The alert when the number of received commands continuously exceeds the upper limit of the throughput per second.

- The alert when the transfer volume of received commands continuously exceeds the upper limit of the data transfer volume per second.

  If the last alert time does not exist, a hyphen (-) is displayed.

LOWER_THROUGHPUT(IOps)

Displays the lower limit of the throughput.

For more information on the lower limit of the throughput per second, see the *Performance Guide* .

LOWER_DATA_TRANS(MBps)

Displays the lower limit of the data transfer volume.

For more information on the lower limit of the data transfer volume per second, see the *Performance Guide* .

LOWER_ALERT(s)

Displays the following alert notification times (second):

- The alert notification time (second) when the number of received commands has not continuously reached the lower limit of the throughput per second.

- The alert notification time (second) when the transfer volume of received commands has not continuously reached the lower limit of the data transfer volume per second.

LOWER_ALERT_NOTICE

Displays the last alert time of the following alerts:

- The alert when the number of received commands has not continuously reached the lower limit of the throughput per second.

- The alert when the transfer volume of received commands has not continuously reached the lower limit of the data transfer volume per second.

A hyphen (-) is displayed if the last alert time does not exist.

PRIORITY

Displays the I/O processing priority. 1 represents the lowest priority and a greater number represents a higher priority.

RESPONSE_TARGET(ms)

Displays the target response time.

RESPONSE_ALERT(s)

Displays the alert notification time (second) when the target response time is not reached.

If the target response time has not been continuously reached for longer than the alert notification time, an alert is issued.

RESPONSE_ALERT_NOTICE

Displays the last alert time when the target response time has not been continuously reached.

A hyphen (-) is displayed if the last alert time does not exist.

Displays I/Os monitoring information of the target LDEV (VSP 5000 series, VSP E series, VSP G130, G/F350, G/F370, G/F700, G/F900)

Nothing is displayed if all the monitoring information is invalid. If the acquisition of QoS monitoring information fails, "NA" is displayed for items other than LDEV number.

```
# raidcom get ldev -ldev_id 200 -key qos_monitor
LDEV# RCV_CMD(IOps) RCV_CMD_TRANS(KBps) THROUGHPUT(IOps)
DATA_TRANS(KBps) RESPONSE(us) MONITOR_TIME
200           3840                 15              2650
10             5   2019-07-31T10:15:20
```

Description of each column in output example:

RCV_CMD(IOps)
>Displays the number of received commands.

RCV_CMD_TRANS(KBps)
>Displays the data transfer volume of received commands.

THROUGHPUT(IOps)
>Displays the throughput.

DATA_TRANS(KBps)
>Displays the data transfer volume.

RESPONSE(us)
>Displays the response time.

MONITOR_TIME
>Displays the monitoring collection time.

raidcom initialize ldev

Formats LDEVs. You can specify Quick Format, Normal Format, or Shredding.

A device group can also be specified instead of an LDEV.

This command is executed asynchronously with the command input. Check the completion of this process on the `raidcom get command_status` command.

If a failure occurs in a processor on an MP blade or a controller, this command cannot be run even on an MP blade or a controller on which no failure occurs.

> ⚠ **Caution:**
>
> - Formatting a deduplication volume might take a while. Also, more pool/ physical capacity might be used. You cannot stop the formatting operation while it is in process.
>
> - If you want to format all of the pool-associated volumes for which capacity saving is enabled, first execute the `raidcom modify ldev` command to block the deduplication system data volumes, and then execute the `raidcom initialize pool` command. By executing the `raidcom initialize pool` command before formatting the volumes, you can shorten the time for the formatting operation and prevent the increase of pool capacity to be used.

> ⚠️ **Caution:**
>
> - Make sure to run the **`raidcom initialize ldev`** command after the completion of the asynchronous commands including the **`raidcom initialize ldev`** command. The status of the asynchronous commands can be verified by using the **`raidcom get command_status`** command. For more information about whether the command is run either synchronously or asynchronously with the command input, see each command section.
>
>   If the **`raidcom initialize ldev`** command is run before the completion of the asynchronous commands, the formatting or shredding operation might not be complete.
>
> - If the Normal Format operation is performed by using CCI, the other asynchronous commands might be in the standby status until the formatting is complete.

> 📄 **Note:** If a failure occurs in a processor on an MP blade or a controller, this command cannot be run even on an MP blade or a controller on which no failure occurs.

Syntax

```
raidcom initialize ldev {-ldev_id <ldev#>
     | -grp_opt <group option> -device_grp_name <device group
     name> [<device name>]} -operation <type>
```

Options and parameters

-ldev_id <ldev#>

Specifies the LDEV number (0-65279). For example:

- -ldev_id 200

-grp_opt <group option>

Specifies the device information extracted from the LDEV belonging to the device group. Specify "ldev" (fixed). The information of LDEV belonging to the device group is used.

-device_grp_name <device group name> [<device name>]

Specifies the name of device group (maximum 32 characters) to be operated.

To operate an LDEV in the device group, specify the device name of LDEV (maximum 32 characters) within the device group.

If the device name is omitted, all the LDEVs belonging in the device group are operated.

-operation <type>

Instructs the operation.

The following operations can be specified.

- fmt: Normal Format

- qfmt: Quick Format

- shrd [<pattern>]: Shredding

  Formats the LDEV three times according to the shredding pattern which is specified to "pattern".

  If <pattern> is specified, format the LDEV three times according to the following order.

  - 0x00000000

  - The specified shredding pattern

  - 0x00000000

  If <pattern> is omitted, format the LDEV three times according to the following order.

  - 0x00000000

  - 0xFFFFFFFF

  - 0x00000000

- stop: Stops shredding. The processing for all LDEVs stops. However, normal and quick format processing cannot be stopped.

Examples

Performing Quick Format for an LDEV: 200.

```
# raidcom initialize ldev -operation qfmt -ldev_id 200
```

Performing Quick Format for an LDEV belonging to the device group: grp1.

```
# raidcom initialize ldev -operation qfmt -grp_opt ldev
-device_grp_name grp1
```

Performing Normal Format for an LDEV: 200.

```
# raidcom initialize ldev -operation fmt -ldev_id 200
```

Performing Shredding (Pattern: 0x55aa55aa) for an LDEV: 200.

```
# raidcom initialize ldev -operation shrd 0x55aa55aa -ldev_id 200
```

Stopping Shredding.

```
# raidcom initialize ldev -operation stop -ldev_id 200
```

raidcom modify ldev

Changes the following LDEV attributes:

- Blocking LDEV and restoring LDEV.

- Setting LDEV nickname.

- Setting MP blade ID of LDEV.

- Setting the Tiering policy, the new page assignment tier, or enabling or disabling of the tier relocation for the Dynamic Tiering, active flash, Dynamic Tiering for Mainframe, or active flash for mainframe V-VOL.

- Instructing the page discarding of V-VOL for Dynamic Provisioning, Dynamic Provisioning for Mainframe, Dynamic Tiering, active flash, Dynamic Tiering for Mainframe, or active flash for mainframe.

- Setting the SSID.

- Setting the command device attribute.

- Setting the quorum disk and releasing setting of the quorum disk.

- Setting capacity saving.

- Setting the capacity saving processing mode (post process or inline).

- Setting compression accelerator (enable or disable) (only for VSP 5000 series and VSP E1090).

- Setting QoS (only for VSP 5000 series, VSP E series, VSP G130, G/F350, G/F370, G/F700, G/F900).

- ESE-VOL setting

When blocking LDEV, restoring LDEV, setting the quorum disk, releasing setting of the quorum disk, setting capacity saving, setting compression accelerator, or setting QoS, this command is executed asynchronously with the command input. Check the completion of this process using the `raidcom get command_status` command.

> 📄 **Note:** If a failure occurs in a processor on an MP blade or a controller, this command cannot be run even on an MP blade or a controller on which no failure occurs.

There is a limit to the number of asynchronous commands that can be accepted simultaneously. For details, see .

Syntax for specifying LDEV ID

```
raidcom modify ldev -ldev_id <ldev#> {-status <status> [<level>]
   [-forcible -password <One Time Password>]
   | -ldev_name <ldev naming> | -mp_blade_id <mp#> | -ssid <value> |
   -command_device {y | n} [Security value] | -quorum_enable <serial#>
   <id> -quorum_id <quorum id> | -quorum_disable |
   -alua {enable|disable} | -capacity_saving <capacity saving>
```

```
    | -capacity_saving_mode <saving mode>| -compression_acceleration {enable |
disable} -request_id auto | -upper_throughput_io<upper throughput io> -request_id
auto | -upper_data_trans_mb <upperdata trans mb> -request_id auto | -upper_alert_time
<upper alert time> -request_id auto | -lower_throughput_io <lower throughput io> -
request_id
auto | -lower_data_trans_mb <lower data trans mb> -request_id auto | -
lower_alert_time <lower alert time> -request_id auto | -
response_priority <priority> -request_id auto | -response_alert_time
<response alert time> -request_id auto  | -ese {enable | disable} -request_id auto}
```

Syntax for specifying pool ID

```
raidcom modify ldev -pool {<pool ID#> | <pool naming>} -compression_acceleration
{enable | disable} -request_id auto
```

Options and parameters

-ldev_id <ldev#>

Specifies the LDEV number (0-65279). For example:

- -ldev_id 200

-status <status> [<level>]

Specifies the LDEV status. The following LDEV statuses can be specified.

- nml: Changes the LDEV status to Normal

- blk: Changes the LDEV status to Blockade

- {enable_reallocation [<level>]|enable_relocation [<level>]}: Relocation of LDEV (V-VOL for Dynamic Tiering, active flash, Dynamic Tiering for Mainframe, or active flash for mainframe) enabled. Defines the tier that is used for reallocation depending on the value of the level.

  - all: Uses all tiers in the pool.

  - 1-5: Specifies the level of the tier for use. For details, see the *Provisioning Guide* for your storage system.

- disable_reallocation | disable_relocation: Relocation of LDEV (V-VOL for Dynamic Tiering, active flash, Dynamic Tiering for Mainframe, or active flash for mainframe) disabled

- discard_zero_page: Discarding 0 page of LDEV(V-VOL) for Dynamic Provisioning, Dynamic Provisioning for Mainframe, Dynamic Tiering, active flash, Dynamic Tiering for Mainframe, or active flash for mainframe.

- stop_discard_zero_page (only for VSP 5000 series, VSP E series, VSP G130, G/F350, G/F370, G/F700, G/F900): Stops the discard of LDEV (virtual volume) pages for Dynamic Provisioning, Dynamic Provisioning for Mainframe, Dynamic Tiering, active flash, Dynamic Tiering for Mainframe, or active flash for mainframe.

Chapter 5: Configuration setting commands

- enable_relocation_policy <policy_id>: Enables the LDEV relocation for Dynamic Tiering or active flash, and sets the level or customized policy of the Tier to be used for the relocation by Policy ID.

  <policy_id>:

  - all: Uses all tiers in the pool.

  - 1-5: Specifies the level of the tier for use.

  - 6-31: Specifies the customized policy of the tier for use.

- new_page_allocation: Sets the Tier when new page is allocated to the LDEV for Dynamic Tiering or active flash (V-VOL). The value is one of high/middle/low.

- enable_fullallocation: Enables Full Allocation. If all areas equivalent to the sum of the pool capacities the specified volume requires can be reserved, it is guaranteed that all areas of DP-VOL are writable.

- disable_fullallocation: Disables Full Allocation.

For example:

- -status nml

[-forcible -password <One Time Password>]

Supported storage systems:

- VSP 5000 series

- VSP E series

- VSP G130, G/F350, G/F370, G/F700, G/F900

Specify this option with a one-time password to restore an LDEV status forcibly without considering data consistency. For details about the password, contact customer support. This option is valid only if the `-status nml` option is specified.

-ldev_name <ldev naming>

Specifies the LDEV nickname (maximum 32 characters).

-mp_blade_id <mp#>

Specifies the MP blade ID (0-15).

Do not change the MP blade ID during initial copying or I/O processing of TrueCopy, ShadowImage, Universal Replicator, or global-active device.

To change the MP blade ID again for the same LDEV, wait for more than 30 minutes after changing the MP blade ID. For example:

-mp_blade_id 2

When you change the MP blade ID where the LDEV is allocated, both before and after the changing should be executed during as off-peak hours for the rate of write pending data of the MP blade as possible. It is recommended to execute when the rate of write pending data of all CLPRs is less than 50%.

Do not change the MP blade ID in regard to a lot of LDEVs at the same time. The number of LDEVs and I/O workload for which you can change the MP blade ID at the same time is lower than 10% of the total number of LDEVs where the same MP blade ID is allocated as a guideline.

-ssid <value>

Specifies SSID (hexadecimal number).

When specifying SSID, specify not only unallocated SSID but also LDEV ID. In this case, LDEV ID must be ID for the undefined LDEV in the area where SSID is not allocated. For example:

-ssid 0x1234 -ldev_id 200

> 📄 **Note:** This option is for enterprise storage systems. You do not need to specify SSID for VSP E series, VSP G130, G/F350, G/F370, G/F700, G/F900, VSP G200, G400, G600, G800 and VSP F400, F600, F800, or HUS VM, but this option is enabled.

-command_device {y | n} [Security value]

Configures command device attribute.

y: Command device attribute enabled.

n: Command device attribute disabled.

Specifies the value of command device security (0-7).

You can specify 0-7 to Security value as the command device security setting:

- 0: Security: OFF, User authentication: OFF, Group information acquisition: OFF
- 1: Security: OFF, User authentication: OFF, Group information acquisition: ON
- 2: Security: OFF, User authentication: ON, Group information acquisition: OFF
- 3: Security: OFF, User authentication: ON, Group information acquisition: ON
- 4: Security: ON, User authentication: OFF, Group information acquisition: OFF
- 5: Security: ON, User authentication: OFF, Group information acquisition: ON
- 6: Security: ON, User authentication: ON, Group information acquisition: OFF
- 7: Security: ON, User authentication: ON, Group information acquisition: ON

-quorum_enable <serial#> <id>

Sets quorum disk for global-active device configuration. You must also set the -quorum_id parameter.

serial#: Serial number (of the migration source storage system)

- When specifying the serial number for VSP 5000 series, add a "5" at the beginning of the serial number. For example, for serial number 12345, enter 512345.
- When specifying the serial number for VSP G1x00 and VSP F1500, add a "3" at the beginning of the serial number. For example, for serial number 12345, enter 312345.

Chapter 5: Configuration setting commands

id: The identifier of the source storage system

- R900: VSP 5000 series

- R800: VSP G1x00, VSP F1500

- R700: VSP

- M800: VSP E series, VSP G/F350, G/F370, G/F700, G/F900, VSP G200, G400, G600, G800, VSP F400, F600, F800

- M700: HUS VM

-quorum_id <quorum id>
> Specifies the quorum ID for setting the quorum disk.

-quorum_disable
> Releases setting of the quorum disk.

-alua {enable|disable}
> Specifies the ALUA mode. You must enable ALUA mode only when you use ALUA by global-active device.

- enable: The ALUA mode is enabled.

- disable: The ALUA mode is disabled.

[-capacity_saving <capacity saving>]
> Specifies the capacity saving setting:

- disable: Capacity saving is disabled.

- compression: Compression is enabled.

- deduplication_compression: Deduplication and compression are enabled.

(VSP 5000 series, VSP E series, VSP G130, G/F350, G/F370, G/F700, G/F900) When there are no pool-associated LDEVs for which the capacity saving setting is deduplication_compression, the deduplication system data volumes are automatically deleted asynchronously with the deletion of the LDEVs. To verify that a deduplication system data volume is deleted, use the `raidcom get pool -key saving` command to check the LDEV number, and use the `raidcom get ldev command` to verify that VOL_TYPE of the LDEV is changed from REMOVING to NOT DEFINED.

> 📄 **Note:** Disabling the capacity saving setting for a deduplication volume might take a while to complete due to its data extension operation. Also, more pool/physical capacity might be used. You cannot stop the operation of disabling the capacity saving setting while it is in process.

For VSP 5000 series and VSP E series, compressed volumes for accelerated compression, or volumes for deduplication and compression can be created only for pools whose subscription limit is set to 65535 (unlimited). An error occurs if compression or deduplication_compression is specified for a pool whose subscription limit is other than 65535 (unlimited).

[-capacity_saving_mode <saving mode>]

Specifies post-process mode or inline mode as the capacity saving processing mode for DP-VOLs. For details about the capacity saving processing modes, see the *Provisioning Guide* for the storage system.

- post_process: Sets the post-process mode for capacity saving processing.

- inline: Sets the inline mode for capacity saving processing.

For VSP G1x00, VSP F1500, VSP G200, G400, G600, G800, and VSP F400, F600, F800 post_process is the default. For VSP 5000 series, VSP E series, and VSP G/F350, G/F370, G/F700, G/F900, inline is the default.

-upper_throughput_io <upperthroughput io> -request_id auto

Specifies the upper limit of the throughput per second. To disable it, specify 0. This option is available only for VSP 5000 series, VSP E series, VSP G130, G/F350, G/F370, G/F700, G/F900. For more information on the upper limit of the throughput per second, see the *Performance Guide*.

The request ID is output in the following format when the `raidcom modify ldev` command ends.

- REQID : <request#>

Where <request#> is a request ID assigned each time the command is executed. For details, see Request ID function (on page 246).

-upper_data_trans_mb <upper datatrans mb> -request_id auto

Specifies the upper limit of the data transfer volume (MB) per second. To disable it, specify 0. This option is available only for VSP 5000 series, VSP E series, VSP G130, G/F350, G/F370, G/F700, G/F900. For more information on the upper limit of the data transfer volume (MB) per second, see the *Performance Guide*.

The request ID is output in the following format when the `raidcom modify ldev` command ends.

- REQID : <request#>

Where <request#> is a request ID assigned each time the command is executed. For details, see Request ID function (on page 246).

-upper_alert_time <upper alerttime> -request_id auto

Specifies the following alert notification time (second). To disable it, specify 0. This option is available only for VSP 5000 series, VSP E series, VSP G130, G/F350, G/F370, G/F700, G/F900. If either of the following conditions has continued for the alert notification time or longer, an alert is issued.

- The alert notification time (second) when the number of received commands continuously exceeds the upper limit of the throughput per second.

- The alert notification time (second) when the transfer volume of received commands continuously exceeds the upper limit of the data transfer volume per second.

The request ID is output in the following format when the execution of the **raidcom modify ldev** command ends:

- REQID : <request#>

Where <request#> is a request ID assigned each time the command is executed. For details, see Request ID function (on page 246).

-lower_throughput_io <lowerthroughput io> -request_id auto
(For VSP 5000 series, VSP E series, VSP G130, G/F350, G/F370, G/F700, G/F900) Specifies the lower limit of the throughput per second. To disable it, specify 0. For more information on the lower limit of the throughput per second, see the *Performance Guide*.

The request ID is output in the following format when the **raidcom modify ldev** command ends.

- REQID : <request#>

Where <request#> is a request ID assigned each time the command is executed.

For details, see Request ID function (on page 246).

- lower_data_trans_mb <lower data trans mb> -request_id auto
(For VSP 5000 series, VSP E series, VSP G130, G/F350, G/F370, G/F700, G/F900) Specifies the lower limit of the data transfer volume (MB) per second. To disable it, specify 0. For more information on the lower limit of the data transfer volume (MB) per second, see the *Performance Guide*.

The request ID is output in the following format when the **raidcom modify ldev** command ends.

- REQID : <request#>

Where <request#> is a request ID assigned each time the command is executed.

For details, see Request ID function (on page 246).

- lower_alert_time <lower alert time> -request_id auto
(For VSP 5000 series, VSP E series, VSP G130, G/F350, G/F370, G/F700, G/F900) Specifies the following alert notification time (second). To disable it, specify 0.

- The alert notification time (second) when the number of received commands has not continuously reached the lower limit of the throughput per second.

- The alert notification time (second) when the transfer volume of received commands has not continuously reached the lower limit of the data transfer volume per second.

For more information on the lower limit of the alert notification time per second, see the *Performance Guide*.

The request ID is output in the following format when the execution of the **raidcom modify ldev** command ends.

- REQID : <request#>

Where <request#> is a request ID assigned each time the command is executed.

For details, see Request ID function (on page 246).

-response_priority <priority> -request_id auto

(For VSP 5000 series, VSP E series, VSP G130, G/F350, G/F370, G/F700, G/F900) Specifies the I/O processing priority. The valid ranges are 1-3. To disable it, specify 0. For more information on the I/O processing priority, see the *Performance Guide*.

The request ID is output in the following format when the execution of the `raidcom modify ldev` command ends.

- REQID : <request#>

Where <request#> is a request ID assigned each time the command is executed.

For details, see Request ID function (on page 246).

-response_alert_time <responsealert time> -request_id auto

(For VSP 5000 series, VSP E series, VSP G130, G/F350, G/F370, G/F700, G/F900) Specifies the alert notification time (second) when the target response time is not reached . To disable it, specify 0. If the target response time has not been continuously reached for more than the alert notification time, an alert is issued. For more information on the alert notification time (second) when the target response time is not reached , see the *Performance Guide*.

The request ID is output in the following format when the execution of the `raidcom modify ldev` command ends.

- REQID : <request#>

Where <request#> is a request ID assigned each time the command is executed.

For details, see Request ID function (on page 246).

-compression_acceleration {enable | disable} -request_id auto

Supported storage systems:

- VSP 5000 series

- VSP E1090

Specifies whether compression accelerator is enabled or disabled.

If you change -capacity_saving from disable to compression or deduplication_compression, and omit this this option, it will be enabled when COMPRESSION_ACCELERATION of raidcom get system -key dedupe_compression is AVAILABLE. It will be disabled when COMPRESSION_ACCELERATION of raidcom get system -key dedupe_compression is UNAVAILABLE.

- enable: Compression accelerator is enabled.

- disable: Compression accelerator is disabled.

<request #> is the Request ID assigned for each command execution. For details, see Request ID function (on page 246).

-ese {enable | disable} -request_id auto

Supported storage systems:

- VSP 5000 series

- VSP E1090

Specifies whether to set or cancel the ESE attribute for the LDEV.

- enable: Set the ESE attribute

- disable: Cancel the ESE attribute

<request#> is the request ID assigned each time the command is run. For details, see Request ID function (on page 246)

Examples

Restoring the LDEV: 200.

```
# raidcom modify ldev -status nml -ldev_id 200
```

Blocking the LDEV: 200.

```
# raidcom modify ldev -status blk -ldev_id 200
```

Assigning an LDEV nickname: my_volume to LDEV: 200

```
# raidcom modify ldev -ldev_id 200 -ldev_name my_volume
```

Setting the LDEV owner MP blade ID of LDEV: 200 to 2

```
# raidcom modify ldev -ldev_id 200 -mp_blade_id 2
```

Enabling relocation of LDEV (Dynamic Tiering, active flash, Dynamic Tiering for Mainframe, or active flash for mainframe V-VOL): 200

```
# raidcom modify ldev -ldev_id 200 -status enable_reallocation
```

Disabling relocation of LDEV (Dynamic Tiering, active flash, Dynamic Tiering for Mainframe, or active flash for mainframe V-VOL): 200

```
# raidcom modify ldev -ldev_id 200 -status disable_reallocation
```

Discarding the zero page of LDEV (Dynamic Provisioning, Dynamic Provisioning for Mainframe, Dynamic Tiering, active flash, Dynamic Tiering for Mainframe, or active flash for mainframe V-VOL): 200

```
# raidcom modify ldev -ldev_id 200 -status discard_zero_page
```

Specifying 0x1234 to LDEV SSID of LDEV: 200.

Chapter 5: Configuration setting commands

📄 **Note:**

Specify undefined LDEV number: 200 in the area that is not allocated SSID, and assign new SSID: 0x1234. If you specify an LDEV number that is already defined or an SSID that is already registered, an error occurs.

```
# raidcom modify ldev -ssid 0x1234 -ldev_id 200
```

Enabling LDEV command device attribute of LDEV: 200. Specifies 2 to the value of command device security.

```
# raidcom modify ldev -command_device y 2 -ldev_id 200
```

Disabling LDEV command device attribute of LDEV: 200.

```
# raidcom modify ldev -command_device n -ldev_id 200
```

Setting the relocation of LDEV for LDEV:200 by the customized policy 6.

```
# raidcom modify ldev -ldev_id 200 -status enable_relocation_policy 6
```

Setting the tier from which the new mapped page of LDEV: 200 is allocated to High.

```
# raidcom modify ldev -ldev_id 200 -status new_page_allocation high
```

Setting LDEV: 200 as the quorum disk whose quorum ID is 10.

```
# raidcom modify ldev -ldev_id 200 -quorum_enable 65384 R700 -quorum_id 10
```

Enabling the ALUA mode for LDEV: 200.

```
# raidcom modify ldev -ldev_id 200 -alua enable
```

Setting the capacity saving setting to compression for LDEV: 200.

```
# raidcom modify ldev -ldev_id 200 -capacity_saving compression
```

Enabling the compression accelerator setting for LDEV: 200.

```
# raidcom modify ldev -ldev_id 200 -compression_acceleration enable -request_id auto
```

Enabling all compression accelerator settings associated with pool ID: 1.

```
# raidcom modify ldev -pool 1 -compression_acceleration enable -request_id auto
```

Chapter 5: Configuration setting commands

Enabling the compression accelerator settings for all LDEVs associated with from pool ID: 0 to maximum 128 pools.

```
# raidcom get dp_pool | rmawk @L-gt:1 exe="raidcom modify ldev -pool @1
-compression_acceleration enable -request_id auto"
```

Changing the capacity saving processing mode for LDEV: 0 to the inline mode.

```
# raidcom modify ldev -ldev_id 0 -capacity_saving_mode inline
```

Changing the capacity saving processing mode for LDEV: 0 to the post-process mode.

```
# raidcom modify ldev -ldev_id 0 -capacity_saving_mode post_process
```

Setting the upper limit of the data transfer volume to 100 MB/s for LDEV ID: 200.

```
# raidcom modify ldev -ldev_id 200 -upper_data_trans_mb 100 -request_id
auto
```

(For VSP 5000 series, VSP E series, VSP G130, G/F350, G/F370, G/F700, G/F900) Setting the lower limit of the throughput per second to 1500 IOPS for LDEV ID: 200.

```
# raidcom modify ldev -ldev_id 200 -lower_throughput_io 1500 -request_id auto
```

(For VSP 5000 series, VSP E series, VSP G130, G/F350, G/F370, G/F700, G/F900) Setting the I/O processing priority to 1 for LDEV ID: 200.

```
# raidcom modify ldev -ldev_id 200 -response_priority 1 -request_id auto
```

(For VSP 5000 series, VSP E series, VSP G130, G/F350, G/F370, G/F700, G/F900) Setting the alert notification time when the lower limit is not reached to 10 seconds for LDEV ID: 200.

```
# raidcom modify ldev -ldev_id 200 -lower_alert_time 10 -request_id auto
```

(For VSP 5000 series) Enabling ESE for LDEV: 200.

```
# raidcom modify ldev -ldev_id 200 -ese enable -request_id auto
```

raidcom add license

Supported storage systems:

- VSP 5000 series
- VSP E series
- VSP G130, G/F350, G/F370, G/F700, G/F900
- VSP G200, G400, G600, G800 and VSP F400, F600, F800

This command installs the license.

> 📄 **Note:** If a failure occurs in a processor on an MP blade or a controller, this command cannot be run even on an MP blade or a controller on which no failure occurs.

Syntax

```
raidcom add license -keycode <key code>
```

Options and parameters

-keycode <key code>
    Specify the key code of the license.

Examples

Specify PXPQRS275WMYZ as the key code of the license, and then install the license key.

```
#raidcom add license -keycode PXPQRS275WMYZ
```

raidcom delete license

Supported storage systems:

- VSP 5000 series
- VSP E series
- VSP G130, G/F350, G/F370, G/F700, G/F900
- VSP G200, G400, G600, G800 and VSP F400, F600, F800

This command removes the license.

> 📄 **Note:** If a failure occurs in a processor on an MP blade or a controller, this command cannot be run even on an MP blade or a controller on which no failure occurs.

Syntax

```
raidcom delete license -product_id <product ID>
```

Options and parameters

-product_id <product ID>
    Specify the software ID.

Chapter 5: Configuration setting commands

Examples

Remove the license of ID: 4102.

```
#raidcom delete license -product_id 4102
```

raidcom modify license

Supported storage systems:

- VSP 5000 series
- VSP E series
- VSP G130, G/F350, G/F370, G/F700, G/F900
- VSP G200, G400, G600, G800 and VSP F400, F600, F800

This command changes the term license status.

> 📄 **Note:** If a failure occurs in a processor on an MP blade or a controller, this command cannot be run even on an MP blade or a controller on which no failure occurs.

Syntax

```
raidcom modify license -product_id <product ID> -license_status <status>
```

Options and parameters

-product_id <product ID>
Specify the software ID.

-license_status <status>
Specify the term license status.

- enable: Enables the term license.
- disable: Disables the term license.

Examples

Enable the term license of ID: 4102.

```
#raidcom modify license -product_id 4102 -license_status enable
```

raidcom get license

This command acquires the license information.

Chapter 5: Configuration setting commands

Syntax

```
raidcom get license [-key opt]
```

Options and parameters

[-key opt]

Displays the capacity of the internal and external volumes created in the storage system in GB, and the serial number.

Examples

Display the license information.

```
#raidcom get license
```

```
PRO_ID  STS Type L Cap_Perm(TB) Cap_Used(GB) - Term  Name
34049   INS PER  L          50           10 -    - "Cache Residency Manager"
34055   INS TEM  U           -            - -    - "Dynamic Provisioning"

#raidcom get license -key opt
Serial# : 302656
Cap_Mounted(GB) : 1229000
```

Description of each column in output example:

PRO_ID

Displays the software ID.

STS

Displays the installation status of the software. For details about installation statuses, see the *System Administrator Guide*.

- INS: Installed

- DIS: Installed but the license is disabled

- NIN: Not installed

- NEL: Installed but the license capacity is insufficient

- GRP: The license capacity is insufficient because an LDEV was added, a copy pair was created, or a pool volume was added. The license will expire in 30 days.

- EXT: The temporary key is expired

Type

Displays the license key type. For details about license key types, see the *System Administrator Guide*.

- PER: Permanent

- TER: Term

- TEM: Temporary

- EME: Emergency

- - (hyphen): License is not installed. But if the license period is still available and you removed temporary key, TEM is displayed.

L

Displays if there is an upper limit for the installed permitted capacity.

- U: There is no upper limit

- L: There is an upper limit

- T: There is no upper limit. Displayed only when the license key type is temporary or emergency and the license key does not overwrite the installed license key.

Cap_Perm(TB)

Displays the installed permitted capacity in TB. If the software does not have the capacity limit, or if the license is not installed, a hyphen is displayed.

Cap_Used(GB):

Displays the size of the volume used by the software in GB. If the licensed capacity type is other than Used capacity, or if the license is not installed, a hyphen is displayed. For details about the license capacity, see the *System Administrator Guide*.

Term

Displays the remaining days until the term key, temporary key, or emergency key expires. If the license period is still available and you removed the temporary key, Term displays the remaining days of the license period. After the temporary key expires, Term displays the remaining days until the license can be installed again. If the term is unlimited and the license is not installed, a hyphen (-) is displayed.

Name

Displays the product name. Double quotation marks (") are added at the beginning and end of the product name.

Serial#

Displays the serial number of the unit.

Cap_Mounted(GB)

Displays the capacity of the internal and external volumes created in the storage system in GB.

raidcom add quorum

Supported storage systems:

- VSP 5000 series

- VSP G1x00 and VSP F1500

- VSP E series

- VSP G/F350, G/F370, G/F700, G/F900

Sets a quorum disk. If an LDEV number is not allocated to a quorum disk, no volume is set for the quorum disk.

This command is executed asynchronously with the command input. Check the completion of this process using the `raidcom get command_status` command.

📄 **Note:** If a failure occurs in a processor on an MP blade or a controller, this command cannot be run even on an MP blade or a controller on which no failure occurs.

Syntax

```
raidcom add quorum -quorum_id <quorum id> -request_id auto - remote_storage <serial#>
<id> [-ldev_id <ldev#>]
```

Options and parameters

-quorum_id <quorum id> - request_id auto

Specifies the quorum ID.

The request ID is output in the following format when the execution of the `raidcom add quorum` command ends:

- REQID : <request#>

<request#> is a request ID assigned each time the command is executed. For details, see Request ID function (on page 246).

-remote_storage <serial#> <id>

Specifies the remote storage system for GAD pairs monitored by the quorum disk.

- serial#: Specifies the product serial number of the remote storage system.

  - When specifying the serial number for VSP 5000 series, add a "5" at the beginning of the serial number. For example, for serial number 12345, enter 512345.

  - When specifying the serial number for VSP G1x00 and VSP F1500, add a "3" at the beginning of the serial number. For example, for serial number 12345, enter 312345.

- id: Specifies the type identifier of the remote storage system.

  - R800: VSP G1x00, VSP F1500

  - R900: VSP 5000 series

  - M800: VSP E series, VSP G/F350, G/F370, G/F700, G/F900, VSP G200, G400, G600, G800, VSP F400, F600, F800

Chapter 5: Configuration setting commands

[-ldev_id <ldev#>]

Specifies an LDEV number of a quorum disk. If an LDEV number is not allocated to a quorum disk, no volume is set for the quorum disk.

Examples

Setting a quorum disk whose quorum ID is 0, remote storage serial number is 310001, type identifier is R800, and LDEV number is 0x0010.

```
#raidcom add quorum -quorum_id 0 -request_id auto -remote_storage 310001
R800 -ldev_id 0x0010
REQID : 0
```

Setting a quorum disk whose quorum ID is 0, remote storage serial number is 310001, type identifier is R800 without setting an LDEV.

```
#raidcom add quorum -quorum_id 0 -request_id auto -remote_storage 310001
R800
REQID : 0
```

raidcom delete quorum

Supported storage systems:

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

Deletes a quorum disk.

This command is executed asynchronously with the command input. Check the completion of this process using the **raidcom get command_status** command.

> 📄 **Note:** If a failure occurs in a processor on an MP blade or a controller, this command cannot be run even on an MP blade or a controller on which no failure occurs.

Syntax

```
raidcom delete quorum -quorum_id <quorum id> -request_id auto
```

Options and parameters

-quorum_id <quorum id> -request_id auto

Specifies the quorum ID.

The request ID is output in the following format when the execution of the **raidcom delete quorum** command ends:

- REQID : <request#>

<request#> is a request ID assigned each time the command is executed. For details, see .

Example

Deleting a quorum disk whose quorum ID is 0.

```
# raidcom delete quorum -quorum_id 0 -request_id auto
```

raidcom modify quorum

For the quorum disk, this command sets the period of Read Response Guaranteed Time When Quorum Monitoring Stopped.

> 📄 **Note:** If a failure occurs in a processor on an MP blade or a controller, this command cannot be run even on an MP blade or a controller on which no failure occurs.

Syntax

```
raidcom modify quorum -quorum_id <quorum id> -timeout <timeout>
```

Options and parameters

-quorum_id <quorum id>

Specifies a quorum ID (0 to 31) in decimal or hexadecimal. For hexadecimal IDs, prefix 0x.

Example:

- -quorum_id 10

- -quorum_id 0x0a

-timeout <timeout>

Specifies the period of time in seconds (5 to 100) for Read Response Guaranteed Time When Quorum Monitoring Stopped.

Examples

For the quorum disk (quorum ID: 1), specify 40 seconds as the period of Read Response Guaranteed Time When Quorum Monitoring Stopped.

```
#raidcom modify quorum -quorum_id 1 -timeout 40
```

raidcom get quorum

This command displays information about the quorum disk.

Syntax

```
raidcom get quorum [-quorum_id <quorum id>] [-fx]
```

Options and parameters

[-quorum_id <quorum id>]

Specifies a quorum ID (0 to 31) in decimal or hexadecimal. For hexadecimal IDs, prefix 0x. If this option is omitted, information of all quorum disks is displayed.

Example:

- -quorum_id 10

- -quorum_id 0x0a

[-fx]

Specifies this option to display an LDEV number in hexadecimal.

Examples

Displays the information of the quorum disk (quorum ID: 1).

```
#raidcom get quorum -quorum_id 1
```

```
QRDID : 1
LDEV : 2045
QRP_Serial# : 302646
QRP_ID : R8
Timeout(s) : 30
STS: REPLACING
```

Description of each column in output example:

QRDID

Displays the quorum disk ID for GAD.

LDEV

Displays the LDEV number of a quorum disk for GAD. If no volume is set for the quorum disk, no LDEV number is allocated to the quorum disk. Therefore, 65535 (0xffff) is displayed.

QRP_Serial#

Displays the serial number of a remote storage system assigned to serve as a quorum disk of GAD devices.

When the storage system is VSP G1x00 or VSP F1500, 3 is added to the beginning of the serial number.

QRP_ID

Displays the identifier of a remote storage system assigned to serve as a quorum disk of GAD devices.

- R8: VSP G1x00 or VSP F1500

- R9: VSP 5000 series

- M8: VSP E series, VSP G/F350, G/F370, G/F700, G/F900, VSP G200, G400, G600, G800, VSP F400, F600, F800

Timeout(s)

Displays the time in seconds for Read Response Guaranteed Time When Quorum Monitoring Stopped set for the quorum disk of GAD.

STS

Displays the status of a quorum disk for GAD. A hyphen (-) is displayed if no volume is set for the quorum disk.

- NORMAL: The quorum disk is in normal status.

- TRANSITIONING: The status of the quorum disk is being changed.

- BLOCKED: The quorum disk is blocked.

- REPLACING: The quorum disk is being replaced.

- FAILED: The quorum disk is in abnormal state.

- - (hyphen): This information is not available for this quorum disk.

raidcom replace quorum

This command replaces a quorum disk.

📄 **Note:** If a failure occurs in a processor on an MP blade or a controller, this command cannot be run even on an MP blade or a controller on which no failure occurs.

Syntax

```
raidcom replace quorum -quorum_id <quorum id> -ldev_id <ldev#>
```

Options and parameters

-quorum_id <quorum id>

Specifies the quorum ID.

-ldev_id <ldev#>

Specifies the LDEV number (0 to 65279). For example:

- -ldev_id 200

Examples

Change the LDEV number of the quorum disk (LDEV) associated with the quorum ID 1 to the LDEV number 200.

```
#raidcom replace quorum -quorum_id 1 -ldev_id 200
```

raidcom modify local_replica_opt

This command sets the local replica option. For details about the local replica option, see the user manuals for ShadowImage, ShadowImage for Mainframe, Thin Image, Compatible FlashCopy® , Volume Migration, and nondisruptive migration.

📄 **Note:** If a failure occurs in a processor on an MP blade or a controller, this command cannot be run even on an MP blade or a controller on which no failure occurs.

Syntax

```
raidcom modify local_replica_opt -opt_type <option type>
 {-set_system_opt <system option> | -reset_system_opt <system option>}
```

Options and parameters

-opt_type <option type>

Specify the local replication type to be set.

- open: Local replication option for ShadowImage, Thin Image, Volume Migration, nondisruptive migration.

- mainframe: Local replication option for ShadowImage for Mainframe, Compatible FlashCopy®, Volume Migration.

-set_system_opt <system option>

Specifies the local replica option ID that you are enabling. For details about the local replica option ID, see the applicable user manual.

-reset_system_opt <system option>

Specifies the local replica option ID that you are enabling. For details about the local replica option ID, see the applicable user manual.

Examples

Enable local replica option: 1

```
raidcom modify local_replica_opt -opt_type open -set_system_opt 1
```

Disable local replica option: 1

```
raidcom modify local_replica_opt -opt_type open -reset_system_opt 1
```

raidcom get local_replica_opt

This command lets you view the local replica option.

Syntax

```
raidcom get local_replica_opt -opt_type <option type>
```

Options and parameters

-opt_type <option type>

Specifies the local replica option type that you are going to view:

- open: Local replication option for ShadowImage, Thin Image, Volume Migration, and nondisruptive migration.

- mainframe: Local replication option for ShadowImage for Mainframe, Compatible FlashCopy®, and Volume Migration.

Examples

View the local replica option of ShadowImage, Thin Image, Volume Migration, and nondisruptive migration.

```
#raidcom get local_replica_opt -opt_type open
```

```
Serial# : 563528
Type : open
Option : 1 14
```

View the local replica option of ShadowImage for Mainframe, Compatible FlashCopy®, and Volume Migration.

```
#raidcom get local_replica_opt -opt_type mainframe
```

```
Serial# : 563528
```

Chapter 5: Configuration setting commands

```
Type : mainframe
Option : 1 14
```

Description of each column in output example:

Serial#
　　　　Displays the serial number.

Type
　　　　Displays the local replica option type.

Option
　　　　Displays the local replica option that is set to enable.

raidcom get remote_replica_opt

Supported storage systems:

- VSP 5000 series

- VSP E series

- VSP G/F350, G/F370, G/F700, G/F900

Displays remote replication options.

Syntax

```
raidcom get remote_replica_opt [-opt_type <option type>
[-key cu [-cu <cu#>]]]
```

Options and parameters

[-opt_type <option type>]
　　　　Specify the remote replication type to be displayed.

- tc: TrueCopy

- ur: Universal Replicator

- gad: global-active device

[-key cu]
　　　　Specify this option to view the remote replication option of each CU. This option is
　　　　valid when the -opt_type tc option is specified.

[-cu <cu#>]
　　　　Specify this option to view the remote replication option of a specific CU number.

Examples

Displaying information about path blocked watch and path blocked watch SIM.

```
#raidcom get remote_replica_opt
Serial# : 563528
PBW(s) : 40
PBW_SIM(s) : 70
```

Viewing the remote replication options of TrueCopy.

```
# raidcom get remote_replica_opt -opt_type  tc
Serial# : 563528
Type :  TC
CAS : CU
MAX_CA : 2
```

Viewing the remote replication options of Universal Replicator.

```
# raidcom get remote_replica_opt -opt_type  ur
Serial# : 563528
Type :  UR
MAX_CA : 2
```

Viewing the remote replication options of global-active device.

```
# raidcom get remote_replica_opt -opt_type  gad
Serial# : 563528
Type :  GAD
MAX_CA : 2
```

Description of each column in the output examples:

- Serial#: Product serial number.
- Type: Remote replication option type.
- CAS: Unit of setting the maximum number of initial copy activities
  - System: Set per system.
  - CU: Set per CU.
- MAX_CA: Maximum number of initial copy activities.
- PBW(s): Path blocked watch time (shown in seconds)
- PBW_SIM(s): Path blocked watch sim time (shown in seconds)

Viewing the remote replication options of TrueCopy for each CU.

```
# raidcom get remote_replica_opt -opt_type  tc -key cu
CU# MAX_CA
 00 1
 01 1
```

Chapter 5: Configuration setting commands

```
…
 fe 1
```

Viewing the remote replication options of TrueCopy for CU number: 1.

```
# raidcom get remote_replica_opt -opt_type  tc -key cu -cu 1
CU# MAX_CA
 01 1
```

Description of each column in output example :

- CU#: CU number. The number is displayed in hexadecimal notation, and not displayed if it cannot be used.

- MAX_CA: The maximum number of initial copy activities. Displayed only when the unit for the maximum number of initial copy activities is CU.

raidcom modify remote_replica_opt

Supported storage systems:

- VSP 5000 series

- VSP E series

- VSP G/F350, G/F370, G/F700, G/F900

Sets remote replication options.

> 📄 **Note:** If a failure occurs in a processor on an MP blade or a controller, this command cannot be run even on an MP blade or a controller on which no failure occurs.

Syntax

Setting the `path_blocked_watch` or `path_blocked_watch_sim` option.

```
raidcom modify remote_replica_opt { -path_blocked_watch <time(sec)>
 | -path_blocked_watch_sim <time(sec)>}
```

Setting other remote replication options.

```
raidcom modify remote_replica_opt -opt_type <option type> {-copy_activity_setting
{system|cu}
 | -copy_activity <number of activities> [-cu <cu#>]}
```

Options and parameters

-opt_type <option type>

Specifies the remote replication type to be set.

- tc: TrueCopy

- ur: Universal Replicator

- gad: global-active device

-copy_activity_setting {system|cu}

Specifies the unit (per system or CU) of managing the maximum number of initial copy activities.

- system: Managed per system.

- cu: Managed per CU.

-copy_activity <number of activities>

Specifies the maximum number of initial copy activities.

[-cu <cu#>]

Specifies the CU number for which the `-copy_activity <number of activities>` option is set, in decimal or hexadecimal. For hexadecimal numbers, prefix 0x. For example:

- `-cu 128`

- `-cu 0x80`

-path_blocked_watch <time(sec)>

Specifies the path blocked watch time in seconds.

-path_blocked_watch_sim <time(sec)>

Specifies the path blocked watch SIM time in seconds.

Examples

Specifying 30 seconds as the path blocked watch time .

```
# raidcom modify remote_replica_opt -path_blocked_watch 30
```

Managing the maximum number of initial copy activities of TrueCopy per system.

```
# raidcom modify remote_replica_opt -opt_type tc -copy_activity_setting system
```

Specifying 2 as the maximum number of initial copy activities of TrueCopy for CU number: 1.

```
# raidcom modify remote_replica_opt -opt_type tc -copy_activity 2 -cu 1
```

Specifying 2 as the maximum number of initial copy activities of Universal Replicator.

```
# raidcom modify remote_replica_opt -opt_type ur -copy_activity 2
```

Chapter 5: Configuration setting commands

raidcom add lun

To set the LU path, this maps the specified LDEV to a LUN on a host group on the specified port and creates an LU path or alternate path. A device group can also be specified instead of an LDEV.

If the specified port or host group does not exist, this command is rejected with EX_ENOOBJ(EX_CMDRJE).

If the specified LUN or LDEV already exists, this command is ignored.

If a LUN is not specified, an empty LUN is assigned automatically.

Not allowed:

- Mapping the same LDEV to another LUN in the same host group.
- Overwriting the same LUN to another LDEV.

When the `-request_id auto` option is specified, if a failure occurs in a processor on an MP blade or a controller, this command cannot be run even on an MP blade or a controller on which no failure occurs.

If you specify the `-request_id` option for the asynchronous command, the processing is run asynchronously with the command input. Verify that the processing is complete by using the **raidcom get command_status** command. There is a limit to the number of asynchronous commands that can be accepted simultaneously. For details, see <u>Maximum number of acceptable asynchronous commands (on page 228)</u>.

> ⚠️ **Caution:**
>
> - LDEVs in LUSE volumes cannot be configured with this command.
> - If you execute this command on an LDEV with the command device attribute already set and the LDEV already has a defined/configured path, the command device attribute will be released.

Syntax

To set LU path by specifying the LDEV:

```
raidcom add lun -port <port#> [<host group name>]
   -ldev_id <ldev#> [-lun_id <lun#> | -lun_id auto
   -request_id auto]
```

To set LU path by specifying the device group:

```
raidcom add lun -port <port#> [<host group name>]
   -grp_opt ldev -device_grp_name <device group name> [<device name>]
```

To set LU path by specifying multiple port numbers simultaneously:

```
raidcom add lun -port <port#> [<host group name>]
   {-ldev_id <ldev#> [-lun_id <lun#> | -lun_id auto
 -request_id auto ] -additional_port <additional port>...}
```

Options and parameters

-port <port#> [<host group name>]

Specifies the Port number and the host group. You can specify the host group ID or the host group name for the host group. If you specify neither the host group ID nor the host group name, the host group 0 is used. For example:

- CL1-A-g (g is from 0 to 254)

- CL1-A Linux_X86

-ldev_id <ldev#>

Specifies the LDEV number (0-65279). For example:

- -ldev_id 200

-grp_opt ldev

The information of LDEV belonging to the device group is used. Specify 'ldev' whenever.

-device_grp_name <device group name> [<device name>]

Specifies the name of device group (maximum 32 characters) to be operated.

To operate an LDEV in the device group, specify the device name of LDEV (maximum 32 characters) within the device group.

If the device name is omitted, all the LDEVs belonging in the device group are operated.

[-lun_id <lun#> | -lun_id auto -request_id auto]

Supported storage systems:

- VSP 5000 series

- VSP G1x00 and VSP F1500

- VSP E series

- VSP G130, G/F350, G/F370, G/F700, G/F900

- VSP G200, G400, G600, G800 and VSP F400, F600, F800

Chapter 5: Configuration setting commands

Specifies the LU number. If this option is omitted, a free LU number is assigned automatically.

- -lun_id <lun#>: Specifies the LU number (0 to 2047). For VSP 5000 series, the LUN that can be specified is (0 to 4095).

- -lun_id auto: The storage system assigns an available LU number.

- -request_id auto: Effective only when "auto" is specified as the `-request_id` option. If an option other than "auto" is specified, EX_INVARG or EX_REQARG is replied.

The request ID is output in the following format when the execution of the **raidcom add lun** command ends:

- REQID : <request#>

<request#> is a request ID assigned each time the command is executed. For details, see Request ID function (on page 246).

If multiple commands without LU numbers are executed simultaneously for the same host group, the same LU number might be assigned to multiple newly created LU paths. To avoid this behavior, perform one of the following:

- Specify the `-lun_id auto` option.

- Use the **raidcom lock resource** command to lock the resource group to which the host group belongs. The host group sets the LU number.

If you lock the resource group, commands executed by other users cannot take out unused LU numbers from the resource group. If you set multiple LU paths for the same host group, execute the next **raidcom add lun** command after the execution of the current **raidcom add lun** command is complete.

After the execution of all the **raidcom add lun** commands are complete, use the **raidcom unlock resource** command to unlock the resource group.

> 📄 **Note:**
> - You cannot specify this option if you want to specify a device group as the operation target.
> - You cannot omit this option if you specify multiple port numbers at the same time.

additional_port <additional port>...
Supported storage systems:

- VSP 5000 series

- VSP E series

- VSP G130, G/F350, G/F370, G/F700, G/F900

- VSP G200, G400, G600, G800 and VSP F400, F600, F800

Up to 5 port numbers can be specified to set the LU path additionally.

Chapter 5: Configuration setting commands

The LU path is set to the host group that you specified in the `-port <port#>` option. When you add ports, you cannot specify the host group name by using the [<host group name>] option.

If the `-lun_id auto` option is specified, a common LU number that is not used by the host groups to which LU paths are added is selected. If there are multiple unused LU numbers, the smallest LU number is selected.

Examples

Map the LDEV: 200 to the LU numbers: 1 for the port: CL1-A, the host group #0.

```
# raidcom add lun -port CL1-A-0 -lun_id 1 -ldev_id 200
```

Map the LDEV: 200 for the port: CL1-A, the host group #0. The LU numbers are automatically assigned.

```
# raidcom add lun -port CL1-A-0 -ldev_id 200
```

Map the LDEV for the port: CL1-A, the host group #0, and the one belonging to the device group: grp1. The LU numbers are automatically assigned.

```
# raidcom add lun -port CL1-A-0 -grp_opt ldev -device_grp_name grp1
```

Map the LDEV: 200 for the port: CL1-A, CL2-A, CL3-A, the hosts group #0, and the LU number 1.

```
# raidcom add lun -port CL1-A-0 -ldev_id 200 -lun_id 1 -additional_port CL2-A CL3-A
```

raidcom delete lun

Deletes the LU path on the host group on the specified port. An LDEV or a device group can also be specified instead of a LUN. If an LDEV does not exist on the specified port /host group/LUN, this command is rejected with EX_ENLDEV or EX_ENOOBJ.

When deleting the LU path, stop the I/O for the LU path to be deleted. In the last one path, LDEV must be specified as SMPL volume.

> 📋 **Note:**
>
> LDEVs in LUSE volumes cannot be configured with this command. Do not execute this command to an LDEV whose command device attribute is set. If the command is executed, the command device attribute is released.

Syntax

To delete LU path with specifying the LUN:

```
raidcom delete lun -port <port#> [<host group name>]
   -lun_id <lun#>
```

To delete LU path with specifying the LDEV:

```
raidcom delete lun -port <port#> [<host group name>]
   -ldev_id <ldev#>
```

To delete LU path with specifying the device group:

```
raidcom delete lun -port <port#> [<host group name>]
   -grp_opt <group option> -device_grp_name <device group
   name> [<device name>]
```

To delete LU path with specifying the multiple port numbers simultaneously :

Supported storage systems:

▪ VSP 5000 series

▪ VSP E series

▪ VSP G130, G/F350, G/F370, G/F700, G/F900

▪ VSP G200, G400, G600, G800 and VSP F400, F600, F800

```
raidcom delete lun -port <port#> [<host group name>]
   {-lun_id <lun#> | -ldev_id <ldev#>}
   -additional_port <additional port>...
```

Options and parameters

-port <port#>[<host group name>]
> Specifies the Port number and the host group. You can specify the host group ID or the host group name for the host group. If you specify neither the host group ID nor the host group name, the host group 0 is used. For example:
>
> ▪ CL1-A-g (g is from 0 to 254)
>
> ▪ CL1-A Linux_X86

-lun_id <lun#>
> Specifies the LUN number (0-2047). For VSP 5000 series, the LUN that can be specified is (0 to 4095).

-ldev_id <ldev#>
> Specifies the LDEV number (0-65279). For example:
>
> ▪ -ldev_id 200

-grp_opt ldev
> The information of LDEV belonging to the device group is used. Specify 'ldev' whenever.

-device_grp_name <device group name> [<device name>]
> Specifies the name of device group (maximum 32 characters) to be operated.

To operate an LDEV in the device group, specify the device name of LDEV (maximum 32 characters) within the device group.

If the device name is omitted, all of the LDEVs belonging in the device group are operated.

-additional_port <additional port>...

Supported storage systems:

- VSP 5000 series

- VSP E series

- VSP G130, G/F350, G/F370, G/F700, G/F900

- VSP G200, G400, G600, G800 and VSP F400, F600, F800

Up to 5 port numbers can be specified to delete the LU path additionally.

The LU path is deleted from the host group that you specified into `-port <port#>` `[<host group name>]` option. If you use this option, you cannot specify the host group name by using the `-port <port#>` `[<host group name>]` option.

If you use this option with the `-ldev_id <ldev#>` option, all LUNs of the LU path which you delete must be the same.

If you use this option with the `-lun_id <lun#>` option, all LU path which you delete must be set to the same LDEV.

Examples

Deleting LUN: 1(LDEV number 200) on the port: CL1-A-0

```
# raidcom delete lun -port CL1-A-0 -lun_id 1
# raidcom delete lun -port CL1-A-0 -ldev_id 200
```

Deleting LDEV belonging to the device group: grp1 on the port: CL1-A-0

```
# raidcom delete lun -port CL1-A-0 -grp_opt ldev -device_grp_name grp1
```

Deleting LUN: 1(LDEV number 200) on the port CL1-A, CL2-A, CL3-A of the host group number #0

```
# raidcom delete lun -port CL1-A-0 -lun_id 1 -additional_port CL2-A CL3-A
# raidcom delete lun -port CL1-A-0 -ldev_id 200 -additional_port CL2-A CL3-A
```

raidcom discover lun

Searches external volumes. Displays a list of LUs that can be referred to from the External port of a specific external storage system.

This command is rejected by EX_ENOOBJ in the following cases:

▪ The specified iSCSI port cannot be found.

▪ The iSCSI virtual port mode is enabled, but the specified iSCSI virtual port ID is not correct.

To fix this error, specify the correct port and iSCSI virtual port ID.

If you execute the **raidcom discover lun** command while the **raidcom discover external_storage** command or another **raidcom discover lun** command is being executed, the LU of the external storage system might not be displayed. If this happens, confirm that the storage system in which the command is being executed and the external storage system are connected correctly, and the LU of the external storage system is configured correctly.

If both the storage system and the external storage system have no problem, wait a while, and then execute one command at a time for a storage system.

> **Note:**
>
> If the local storage system port is a Fiber Channel Bidirectional port, it is recommended that you always execute the command by specifying the `-safety_check enable` option unless instructed by the user guide.
>
> If you do not specify the `-safety_check enable` option, the following problems might occur. For the details, see the Hitachi Universal Volume Manager User Guide.
>
> ▪ If you use a non-Hitachi external storage system:
>
>   The I/O path from the external storage system to the local storage system that uses the route between the specified local storage system port and the searched external storage system port might be disconnected.
>
> ▪ If you use the Hitachi external storage system:
>
>   The external or remote path connection that uses the route between the specified local storage system port and the searched external storage system port might be temporarily disconnected (the external or remote path will be reconnected immediately and no blockage occurs).

Syntax

```
raidcom discover lun -port <port#> {-external_wwn <wwn strings>
   | -external_iscsi_name <external iscsi name>
   -external_address <IP address>
   [-iscsi_virtual_port_id <iSCSI virtual port ID>]} [-safety_check enable]
```

Options and parameters

-port <port#>
> Specifies the Port number. It specifies a port of which attribute is External. For example:
>
> - CL1-A
>
> > 📄 **Note:** When you specify a non-bidirectional port with an attribute other than External for the `-port` option, the **raidcom discover lun** command works as the **raidcom get lun** command.

-external_wwn <wwn strings>
> Specifies the WWN (16-digit hexadecimal value) of the external storage system. The 17th digit and later are ignored. The value can be split in units of 4 bytes by "," (comma). For example:
>
> - 210000e08b0256f8
> - 210000e0,8b0256f8

-external_iscsi_name <external iscsi name>
> Specifies the iSCSI name of the iSCSI target on the external storage system in an iqn or eui format.
>
> - iqn format: `iqn.` and the subsequent maximum 219 characters.
> - eui format: `eui.` and the subsequent 16 characters in hexadecimal notation.
>
> > 📄 **Note:** You can specify the pseudo WWN of the iSCSI target to the `-external_wwn` option instead of specifying the iSCSI target by using the `-external_iscsi_name` option or the `-external_address` option. The pseudo WWN corresponds to the iSCSI name and IP address of the iSCSI target on the external storage system. By executing the **raidcom get external_iscsi_name** command, you can confirm the suitability between the pseudo WWN and the iSCSI name and IP address of the iSCSI target on the external storage system.

-external_address <IP address>
> Specifies the IP address of the iSCSI target on the external storage system. An IPv4 address or an IPv6 address can be specified.
>
> The following addresses cannot be specified when you specify the IPv4 address:
>
> - Network address (for example: 192.168.10.0, 0.120.10.1)
> - Broadcast address (for example: 255.255.255.255, 10.1.255.255)
> - Loop back address (for example: 127.0.0.1)
>
> The following addresses cannot be specified when you specify the IPv6 address:
>
> - Unspecified address (for example: ::)
> - Multicast address (for example: ff:1024:1215::01)
> - Loop back address (for example: ::1)

📝 **Note:** You can specify the pseudo WWN of the iSCSI target to the -`external_wwn` option instead of specifying the iSCSI target by using the -`external_iscsi_name` option or the -`external_address` option. The pseudo WWN corresponds to the iSCSI name and IP address of the iSCSI target on the external storage system. By executing the **raidcom get external_iscsi_name** command, you can confirm the suitability between the pseudo WWN and the iSCSI name and IP address of the iSCSI target on the external storage system.

[-iscsi_virtual_port_id <iSCSI virtual port ID>]

Specifies the iSCSI virtual port ID when the specified port is iSCSI and the iSCSI virtual port mode is enabled. If this option is omitted, 0 is assumed.

An error occurs if you specify this option in any of the following cases:

- The virtual port of the specified iSCSI virtual port ID is not found.

- The iSCSI virtual port mode is disabled.

- The device does not support the iSCSI virtual port mode.

[-safety_check enable]

Specify the Fiber Channel bidirectional port for the -`port` option. If you specify any other port, this option is ignored. When this option is specified, processing that might cause the disconnection between the Bidirectional port specified by the -`port` option and the Fiber Channel port of the external storage system specified by the -`external_wwn` option is suppressed. When processing is suppressed, the LU information is not displayed.

Example 1

Displaying LUs defined to the external storage system port: 50060e80, 06fc3850 connected to port: CL5-A (External port) when the external storage system is connected to a FIBRE port.

```
# raidcom discover lun -port CL5-A -external_wwn 50060e8006fc3850 -safety_check enable
```

```
PORT  WWN              LUN VOL_Cap(BLK) PRODUCT_ID E_VOL_ID_C
CL5-A 50060e8006fc3850  0 102400        OPEN-V    HITACHI R500FC381000
CL5-A 50060e8006fc3850  1 102400        OPEN-V    HITACHI R500FC381001
CL5-A 50060e8006fc3850  2 102400        OPEN-V    HITACHI R500FC381002
CL5-A 50060e8006fc3850  3 102400        OPEN-V    HITACHI R500FC381003
CL5-A 50060e8006fc3850  4 102400        OPEN-V    HITACHI R500FC381004
CL5-A 50060e8006fc3850  5 102400        OPEN-V    HITACHI R500FC381005
CL5-A 50060e8006fc3850  6 102400        OPEN-V    HITACHI R500FC381006
CL5-A 50060e8006fc3850  7 102400        OPEN-V    HITACHI R500FC381007
CL5-A 50060e8006fc3850  8 102400        OPEN-V    HITACHI R500FC381008
CL5-A 50060e8006fc3850  9 102400        OPEN-V    HITACHI R500FC381009
CL5-A 50060e8006fc3850 10 102400        OPEN-V    HITACHI R500FC38100A
```

Description of each column in output example 1:

PORT

External port number.

WWN

WWN on the external storage system. When the external storage system is connected to an iSCSI port, this item displays the pseudo WWN of the external storage system.

LUN

LUN of the port on external storage system.

VOL_Cap (BLK)

Capacity of the external volume in units of block (1 block = 512 bytes).

PRODUCT_ID

Product_ID included in the SCSI Inquiry command response of the external volume. The displayed contents and format depend on the connected external volume.

For LUs that cannot be used as an external volume, "OTHER" is displayed.

E_VOL_ID_C

Volume identifier included in the SCSI inquiry command of the external volume. The displayed contents and format depend on the connected external volume.

Example 2

Displaying LUs defined to the iSCSI target of the external storage system (iSCSI name: iqn.z2, IP address: 158.214.135.100) connected to the port: CL1-B, iSCSI virtual port ID: 2.

```
# raidcom discover lun -port CL1-B -iscsi_virtual_port_id 2 -external_iscsi_name
iqn.z2 -external_address 158.214.135.100
```

```
PORT                WWN LUN VOL_Cap(BLK)  PRODUCT_ID          E_VOL_ID_C
CL1-B 50060e8006fc3d60  16      545280      OPEN-V HITACHI R500FC3D0210
CL1-B 50060e8006fc3d60  17      545280      OPEN-V HITACHI R500FC3D0211
CL1-B 50060e8006fc3d60  18      545280      OPEN-V HITACHI R500FC3D0212
CL1-B 50060e8006fc3d60  19      545280      OPEN-V HITACHI R500FC3D0213
```

Example 3

Displaying LUs defined to the iSCSI target of the external storage system (iSCSI name: iqn.z2, IP address: 158.214.135.100) connected to port: CL5-A when the external storage system is connected to an iSCSI port.

```
# raidcom discover lun -port CL5-A -external_iscsi_name iqn.z2
```

```
-external_address 158.214.135.100
PORT      WWN               LUN  VOL_Cap(BLK)  PRODUCT_ID  E_VOL_ID_C
CL1-B     50060e8006fc3d60  16        545280   OPEN-V       HITACHI R500FC3D0210
```

Chapter 5: Configuration setting commands

```
CL1-B     50060e8006fc3d60  17        545280  OPEN-V        HITACHI R500FC3D0211
CL1-B     50060e8006fc3d60  18        545280  OPEN-V        HITACHI R500FC3D0212
CL1-B     50060e8006fc3d60  19        545280  OPEN-V        HITACHI R500FC3D0213
```

Getting the external storage system's iSCSI target information corresponding to the pseudo WWN

Supported storage systems:

- VSP 5000 series
- VSP G1x00 and VSP F1500
- VSP E series
- VSP G130, G/F350, G/F370, G/F700, G/F900
- VSP G200, G400, G600, G800 and VSP F400, F600, F800

To get the iSCSI name and IP address of the iSCSI target of the external storage system which corresponds to the pseudo WWN, execute the **`raidcom get external_iscsi_name`** command.

```
# raidcom get external_iscsi_name
```

```
PORT   Serial# IP_ADDR        IQN      WWN(pseudo)      AMD  D CHAP_user     Sec
CL4-E    63528 158.214.135.100 iqn.z1   50060e80070a3640 CHAP D Win_SQL_EX    *
CL2-E    63528 158.214.135.102 iqn.z3   50060e80070a3642 CHAP S -             -
CL1-B    63528 158.214.135.100 iqn.z2   50060e8006fc3d60 CHAP S -             -
```

The pseudo WWN is managed by each storage system. Therefore, when two storage systems share the iSCSI target of an external storage system as shown in the following figure, the pseudo WWN corresponding to iSCSI target 1 of the storage system A is different from the pseudo WWN corresponding to iSCSI target 1 of the storage system B.

Chapter 5: Configuration setting commands
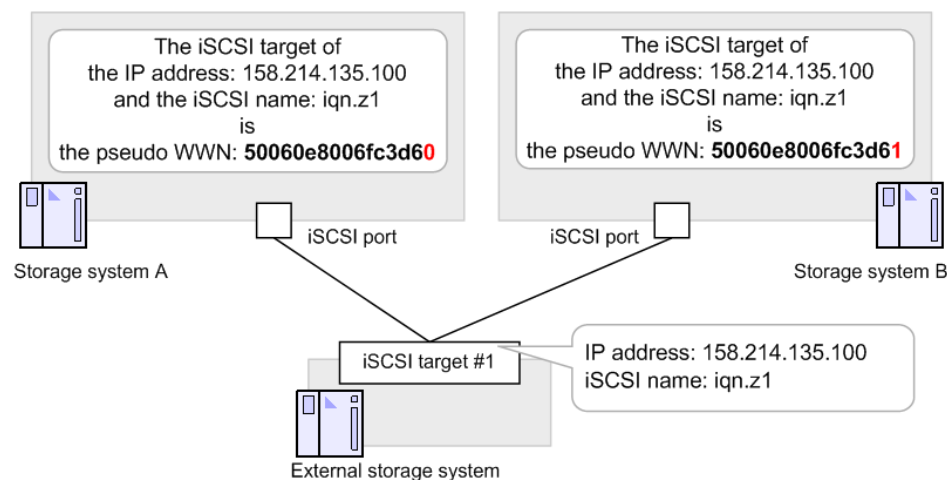
Examples

The following example shows, in the configuration shown above, how to get the iSCSI name and the IP address of the external storage system, and how to display the list of LUNs, by specifying the pseudo WWN which is managed by the storage system A.

```
# raidcom get external_iscsi_name -s 64562 | rmawk @5-eq:50060e8006fc3d60
exe="raidcom discover lun -s 34562 -port CL1-b -external_address@3 -iscsi_name @4"
```

The following examples show how to display the list of LUNs by the storage system B by getting the pseudo WWN of the storage system B which corresponds to the pseudo WWN managed by the storage system A.

Windows example

```
C:\horcm\etc>raidcom get external_iscsi_name -s 64562 | rmawk @5-eq: 50060e8006fc3d60
exe="raidcom get external_iscsi_name -s 34562 | rmawk @@3-eq:@3 -a @@4-eq:@4 exe=
\"raidcom discovery lun -s 34562 -port CL1-b -external_wwn @@5\""
```

UNIX example

```
# raidcom get external_iscsi_name -s 64562 | rmawk @5-eq: 50060e8006fc3d60
exe="raidcom get external_iscsi_name -s 34562 | rmawk @@3-eq:@3 -a @@4-eq:@4
exe="raidcom discovery lun -s 34562 -port CL1-b -external_wwn @@5""
```

raidcom get lun

Displays the LU path information defined in the specified port and host group.

If the specified port does not exist, this command is rejected with EX_ENOOBJ. If an external port is specified, it is rejected with EX_REQARG.

Syntax

```
raidcom get lun -port <port#> <host group name>
   [-key <keyword>]
```

Options and parameters

-port <port#>[<host group name>]
> Specifies the port number, host group ID, and host group name. If the host group name is more than 64 characters, use the host group ID. If the host group ID and host group name are omitted, the LU path information for host group ID 0 is displayed.
>
> For example,
>
> - CLI-A
> - CL1-A-g (g: 0-254)
> - CL1-A Linux_X86

-key <keyword>]

Specifies the information about the LU to be displayed.

- ▪ opt: Displays the status of the host reservation.

- ▪ opt_page1: Displays the information about the specified ALUA mode.

Examples

Displaying the LU information defined to the port CL4-E, the host group #0.

```
# raidcom get lun -port CL4-E-0
```

```
PORT GID      HMD LUN NUM LDEV CM Serial# HMO_BITs
CL4-E 0 LINUX/IRIX  0   1    0 CM   63528    2 13
CL4-E 0 LINUX/IRIX  2   1    2  -   63528    2 13
CL4-E 0 LINUX/IRIX  3   1    3  -   63528    2 13
CL4-E 0 LINUX/IRIX  4   1    4  -   63528    2 13
CL4-E 0 LINUX/IRIX  5   1  992  -   63528    2 13
CL4-E 0 LINUX/IRIX  6   1  993  -   63528    2 13
```

```
#raidcom get lun -port CL4-E-0 -key opt
```

```
PORT GID HMD       LUN NUM LDEV CM Serial# OPKMA HMO_BITs
CL4-E  0 LINUX/IRIX 0   1    0 CM  63528  -Y---   2 13
CL4-E  0 LINUX/IRIX 2   1    2  -  63528  -Y---   2 13
CL4-E  0 LINUX/IRIX 3   1    3  -  63528  -Y---   2 13
CL4-E  0 LINUX/IRIX 6   1  993  -  63528  -Y---   2 13
```

```
# raidcom get lun -port CL4-E-0 -key opt_page1
```

```
PORT   GID  HMD         LUN  NUM   LDEV  CM   Serial#  AL AAS
CL4-E   0   LINUX/IRIX    0   1      0   CM    63528   E  AO
CL4-E   0   LINUX/IRIX    2   1      2   -     63528   D  AO
CL4-E   0   LINUX/IRIX    3   1      3   -     63528   E  AO
```

Description of each column in output example:

PORT

Displays the port number.

GID

Displays the host group ID on the port.

HMD

Displays the following HOST MODE for the host adapter setting on host group.

- HP-UX, SOLARIS, AIX, WIN, LINUX/IRIX, TRU64, DYNIX, OVMS, NETWARE, HI-UX

- VMWARE, HP-XP, VMWARE_EX, WIN_EX, UVM

LUN

Displays LUN number on host group mapping LDEV.

NUM

Displays the number of LDEVs composing an LUSE.

LDEV

Displays the LDEV number.

CM

Displays the command device.

Serial#

Product serial number (Serial#).

OPKMA

Displays the host-reserved LU.

- O: The LU is reserved by an open system[1].

- P: The LU is reserved by a persistent group[1].

- K: The LU is registered by a PGR key[2].

- M: The LU is reserved by a mainframe[1].

- A: The LU is reserved by ACA[1].

  1. If Y is displayed under each character, the LU is reserved. If a hyphen (-) is displayed, the LU is not reserved.
  2. If Y is displayed under each character, the PGR key is registered. If a hyphen (-) is displayed, the PGR key is not registered.

HMO_BITs

Displays the host mode options of the host groups.

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

AL

Displays the information about the ALUA mode.

- E: The ALUA mode is enabled.

- D: The ALUA mode is disabled.

AAS

Displays the setting value of the asymmetric access state for ALUA. The displayed setting value shows whether or not the LU is accessed from the host preferentially. If the ALUA mode is enabled, the setting value displayed under AAS is reported to the host as the value of the asymmetric access state.

- AO: Active and optimized LU. The host accesses the LU preferentially.

- AN: Active and non-optimized LU. When an LU whose setting value of the asymmetric access state is AO cannot be used, the host accesses the LU.

- - (hyphen) : The setting for the asymmetric access state is not supported.

raidcom modify lun

Modifies the LU attribute.

When releasing LU host reserves, this command is executed asynchronously with command input. Check the completion of the process using the **`raidcom get command_status`** command.

> ⚠ **Caution:** When releasing the LU host reserve, do not operate from Device Manager - Storage Navigator and CCI simultaneously. If you do so, the LU host reserve might not be released. If you fail to release the host reserve, specify the LU again and retry the operation from CCI.

> ⚠ **Caution:**
>
> If you set the asymmetric access state for LU with the **`raidcom modify lun`** command, an EX_CMDIOE error might occur. If the EX_CMDIOE error occurs, wait for approximately 40 seconds, and then execute the **`raidcom get lun - key opt_page1`** command to verify that the specified value is set in the asymmetric access state. If the specified value is in the asymmetric access state, the **`raidcom modify lun`** command is ended correctly.
>
> If the specified value is not in the asymmetric access state, set the timeout value in the configuration definition file to 80 seconds or more, and then retry the **`raidcom modify lun`** command. If the error occurs again, see the command error messages in the Command Control Interface User and Reference Guide to take appropriate actions. After the asymmetric access state migration is completed correctly, return the timeout value in the configuration definition file to the original value.

Syntax

```
raidcom modify lun {-port <port#> [<host group name>]
   -lun_id {all | <lun#>} | -ldev_id <ldev#>} {-asymmetric_access_state {optimized |
non_optimized}
   | -reservation release}
```

Chapter 5: Configuration setting commands

Options and parameters

-port <port#> [<host group name>]

Specifies the port number, and host group ID or host group name (iSCSI target alias if iSCSI is used). If the number of characters for the host group name is more than 64, specify the host group ID or the iSCSI target alias.

For example,

- CL1-A-g (g is from 0 to 254)

- CL1-A Linux_X86

- CL1-A Target00

-lun_id {all | <lun#>}

Specifies the LU.

- all: Specifies all LUs in the specified host group. Users who execute the command must have authority to the specified host group, and to all the LDEVs mapped to the LUs of the host group.

- <lun#>:

  Supported storage systems:

  - VSP 5000 series

  - VSP G1x00 and VSP F1500

  - VSP E series

  - VSP G130, G/F350, G/F370, G/F700, G/F900

  - VSP G200, G400, G600, G800 and VSP F400, F600, F800

  Specifies the LUN number (0 to 2047). For VSP 5000 series, the LUN that can be specified is (0 to 4095). Users who execute the command must have authority to the specified host group, and to all the LDEVs mapped to the specified LUs. <lun#> cannot be specified together with the `-asymmetric_access_state` option.

-ldev_id <ldev#>

Specifies the LDEV ID. When releasing the host reserve by specifying this option, specify an LDEV with an emulation type of OPEN-X (X is optional).

Supported storage systems:

- VSP 5000 series

- VSP G1x00 and VSP F1500

- VSP E series

- VSP G130, G/F350, G/F370, G/F700, G/F900

- VSP G200, G400, G600, G800 and VSP F400, F600, F800

For example:

- -ldev_id 200

- -ldev_id 0xc8

- -ldev_id 00:c8

Users who execute the command must have the following authority:

- The authority of LDEVs.

- The authority of the host group connected to the LU path.

- The authority of the iSCSI target.

If the LU path of the target LDEV is changed during the command execution, releasing the LU path host reserve of the non-target LDEV might be executed. To avoid this behavior, use the **raidcom lock resource** command in advance to lock the host group or iSCSI target connected to the LU path. As a result, this prevents LU path change operations of other users. After the execution of all the **raidcom modify lun** commands are complete, use the **raidcom unlock resource** command to unlock the host group or iSCSI target.

When releasing the host reserve is executed by specifying the `-ldev_id` option, the **host reserve release** command is executed by the `-port -lun_id` option for all the LUs of the LDEV specified internally.

This internally executed command is processed asynchronously with the completion of the **raidcom modify lun** command specifying the `-ldev_id` option. When releasing the host reserve of the LDEV where the LU exceeding the following number is defined, release the host reserve by specifying the `-lun_id` option. The number of asynchronous commands that can be accepted by the storage system is as follows.

- VSP 5000 series: 2,560

- VSP G1x00 and VSP F1500: 1,280

- VSP G130, G/F350, G/F370, G/F700, G/F900, VSP E series: 2,560

- VSP G200, G400, G600, G800 and VSP F400, F600, F800: 2,560

When releasing the host reserve by specifying the `-ldev_id` option, perform the following procedure:

1. Every time the **raidcom modify lun** command is executed, execute the **raidcom get command_status** command to confirm the completion of the host reserve release.

2. If asynchronous commands that exceed the number of acceptable asynchronous commands are executed for the storage system, an EX_CMDRJE (SSB1 = 2E11, SSB2 = 9400) error is displayed. If this error is displayed, wait a while, and then retry the command.

-asymmetric_access_state {optimized | non_optimized}

Specifies the asymmetric access state for the LU. For the LU mapped to an LDEV whose ALUA mode is enabled, the value specified by the `-asymmetric_access_state` option is reported to the host as the value of the asymmetric access state.

Chapter 5: Configuration setting commands

Specify the `-asymmetric_access_state` option together with the `-lun_id all` option.

- optimized: Active and optimized LU. The host accesses the LU preferentially.

- non_optimized: Active and non-optimized LU. When an LU whose asymmetric access state is "optimized" cannot be used, the host accesses the LU.

-reservation release

Supported storage systems:

- VSP 5000 series

- VSP G1x00 and VSP F1500

- VSP E series

- VSP G130, G/F350, G/F370, G/F700, G/F900

- VSP G200, G400, G600, G800 and VSP F400, F600, F800

Releases the host reserve of the LU specified by the `-lun_id` option or `-ldev_id` <ldev#>.

Examples

Set the asymmetric access state for all LU on the host group (host group ID: 2) on the port (port ID: CL4-E) to optimized.

```
# raidcom modify lun -port CL4-E-2 -lun_id all -asymmetric_access_state optimized
```

Release the host reserve of all LUs on the host group (host group ID: 2) on the port (port ID: CL4-E).

```
# raidcom modify lun -port CL4-E-2 -lun_id all -reservation release
```

Release the host reserve of all LUs on the host group (host group ID: 0x01)

```
# raidcom modify lun -ldev_id 1 -reservation release
```

raidcom add path

Adds and changes an external path to an external volume. Only one path is operated in one operation. The order of priority for the path is allocated in accordance with the order of adding paths.

This command is executed asynchronously with the command input. Check the completion of this process on the **raidcom get command_status** command.

This command is rejected by EX_ENOOBJ in the following cases:

- The specified iSCSI port cannot be found.

- The iSCSI virtual port mode is enabled, but the specified iSCSI virtual port ID is not correct.

Chapter 5: Configuration setting commands

To fix this error, specify the correct port and iSCSI virtual port ID.

> 📄 **Note:** If a failure occurs in a processor on an MP blade or a controller, this command cannot be run even on an MP blade or a controller on which no failure occurs.

> 📄 **Note:** If the local storage system port is a Fiber Channel Bidirectional port, it is recommended that you always execute the command by specifying the `-safety_check enable` option unless instructed by the Command Control Interface User and Reference Guide .
>
> If you do not specify the `-safety_check enable` option, the following problems might occur. For the details, see the Hitachi Universal Volume Manager User Guide.
>
> - If you use a non-Hitachi external storage system:
>
>   The I/O path from the external storage system to the local storage system that uses the route between the specified local storage system port and the external storage system port might be disconnected.
>
> - If you use a Hitachi external storage system:
>
>   The external or remote path connection that uses the route between the specified local storage system port and the external storage system port might be temporarily disconnected (the external or remote path will be reconnected immediately and no blockage occurs).

Syntax

```
raidcom add path -path_grp <path group#> -port <port#>
   {-external_wwn <wwn strings> | -external_iscsi_name
   <external iscsi name> -external_address <IP address>
   [-iscsi_virtual_port_id <iSCSI virtual port ID>]}[-safety_check enable]
```

Options and parameters

-path_grp <path group#>
> Specifies the external VOL path group number (0-63231).

-port <port#>
> Specifies the Port number. Specifies the number of the port whose attribute is External. For example:
>
> - CL1-A

-external_wwn <wwn strings>
> Specifies the WWN (16-digit hexadecimal value) of the external storage system. The 17th digit and later are ignored. The value can be split in units of 4 bytes by "," (comma). For example:
>
> - 210000e08b0256f8
>
> - 210000e0,8b0256f8

-external_iscsi_name <external iscsi name>

Specifies the iSCSI name of the iSCSI target on the external storage system in an iqn format or an eui format.

- iqn format: `iqn.` and the subsequent maximum 219 characters.

- eui format: `eui.` and the subsequent 16 characters in hexadecimal notation.

> 📄 **Note:**
>
> You can specify the pseudo WWN of the iSCSI target to the `-external_wwn` option instead of specifying the iSCSI target by using the `-external_iscsi_name` option and the -external_address option. The pseudo WWN corresponds to the iSCSI name and IP address of the iSCSI target on the external storage system. By executing the **raidcom get external_iscsi_name** command, you can confirm the suitability between the pseudo WWN and the iSCSI name and IP address of the iSCSI target on the external storage system.

-external_address <IP address>

Specifies the IP address of the iSCSI target on the external storage system. An IPv4 address or an IPv6 address can be specified.

Note that the following addresses cannot be specified when you specify the IPv4 address:

- Network address (for example: 192.168.10.0, 0.120.10.1)

- Broadcast address (for example: 255.255.255.255, 10.1.255.255)

- Loop back address (for example: 127.0.0.1)

Note that the following addresses cannot be specified when you specify the IPv6 address:

- Unspecified address (for example: ::)

- Multicast address (for example: ff:1024:1215::01)

- Loop back address (for example: ::1)

> 📄 **Note:**
>
> You can specify the pseudo WWN of the iSCSI target to the `-external_wwn` option instead of specifying the iSCSI target by using the `-external_iscsi_name` option and the -external_address option. The pseudo WWN corresponds to the iSCSI name and IP address of the iSCSI target on the external storage system. By executing the **raidcom get external_iscsi_name** command, you can confirm the suitability between the pseudo WWN and the iSCSI name and IP address of the iSCSI target on the external storage system.

[-iscsi_virtual_port_id <iSCSI virtual port ID>]

Specifies the iSCSI virtual port ID when the specified port is iSCSI and the iSCSI virtual port mode is enabled. If this option is omitted, 0 is assumed.

An error occurs if you specify this option in any of the following cases:

- The virtual port of the specified iSCSI virtual port ID is not found.

- The iSCSI virtual port mode is disabled.

- The device does not support the iSCSI virtual port mode.

[-safety_check enable]

Specify the Fiber Channel bidirectional port for the `-port` option. If you specify any other port, this option is ignored. When this option is specified, processing that might cause the disconnection between the Bidirectional port specified by the `-port` option and the Fiber Channel port of the external storage system specified by the `-external_wwn` option is suppressed. When processing is suppressed, the external path is not added.

Examples

Adding a path of External port CL1-A, external storage system port 50060e80,05fa0f36 to an external volume path group number: 1.

```
# raidcom add path -path_grp 1 -port CL1-A -external_wwn 50060e80,05fa0f36 -
safety_check enable
```

Adding the path between the external storage system iSCSI target (iSCSI name: iqn.z2, IP address: 158.214.135.100) and the iSCSI port: CL1-A of the local storage system to the external volume path group number: 1.

```
# raidcom add path -path_grp 1 -port CL1-A -external_iscsi_name iqn.z2 -
external_address 158.214.135.100
```

Adding the path between the external storage system iSCSI target (iSCSI name: iqn.z2, IP address: 158.214.135.100) and the iSCSI port: CL1-A, iSCSI virtual port ID: 2 of the local storage system to the external volume path group number: 1.

```
# raidcom add path -path_grp 1 -port CL1-A -iscsi_virtual_port_id 2 -
external_iscsi_name iqn.z2 -external_address 158.214.135.100
```

raidcom get path

Displays the external path information to an external volume.

Syntax

```
raidcom get path [-path_grp <path group#> |
   -external_grp_id <gno-sgno> | -ldev_id <ldev#>]
   [{-check_status | -check_status_not}  <string>...
   [-time <time>]]
```

Options and parameters

[-path_grp <path group#>]

Specifies the external VOL path group number (0-63231).

If it is omitted, all groups are displayed.

[-external_grp_id <gno-sgno>]

Specifies the external volume group number (gno:1-16384, sgno: 1-4096). For example:

- 52-11

When you specify this option, only the external path information of the specified external volume group is displayed.

[-ldev_id<ldev#>]

Specifies the LDEV number (0-65279). For example:

- -ldev_id 200

When you specify this option, only the external path information to the external volume of the specified LDEV is displayed.

[-check_status <string>... [-time <time>]]

Check if the external volume is in the same state as specified in <string>. If the option contains multiple states, the OR condition check is performed and verifies that the external volume is in one of the states contained in the option. You can check the state of the external volume displayed at the beginning when executing the **raidcom get path** command without specifying the -check_status option.

The following strings are specified in the <string>:

- NML: "Normal"
- CHK: "Checking"
- SYN: "Cache Destage"
- DSC: "Disconnect"
- BLK: "Blockading" in the external path for the external volume
- UNK: "Unknown"
- WAR: "Warning"

If you specify the -time option, this command checks the status of the external volume every three seconds until the end of the specified <time> (seconds).

When this option is specified, the returned values are as follows:

- The external volume is not in any of the specified states: 0
- The external volume is in one of the specified states (without -time option): 1
- The external volume is in one of the specified states (when the specified <time> passed): EX_EWSTOT

[-check_status_not <string> [-time <time>]]

Check if the external volume is not in the same state as specified in <string>. If the option contains multiple states, the NOR condition check is performed and verifies that the external volume is not in any of the states contained in the option. You can check the state of the external volume displayed at the beginning when executing the **raidcom get path** command without specifying the -check_status_not option.

The following strings are specified in the <string>.

- NML: "Normal"

- CHK: "Checking"

- SYN: "Cache Destage"

- DSC: "Disconnect"

- BLK: "Blockading" in the external path for the external volume

- UNK: "Unknown"

- WAR: "Warning"

If you specify the -time option, this command checks the status of the external volume every three seconds until the end of the specified <time> (seconds).

When this option is specified, the returned values are as follows:

- The external volume is in one of the specified states: 0

- The external volume is in none of the specified states (without -time option): 1

- The external volume is in none of the specified states (when the specified <time> passed): EX_EWSTOT

Example 1

Displaying the external path (group) information to the external volume

```
# raidcom get path
```

```
PHG GROUP STS CM IF MP# PORT WWN PR LUN PHS Serial# PRODUCT_ID LB PM DM QD TO(s)
PBW(s)
1 1-1 NML E D 2 CL1-A 50060e8005fa0f36 1 3 NML 60010 VSP N M E 8 15 10
1 1-1 NML E D 2 CL2-A 50060e8005fa0f38 2 3 NML 60010 VSP N M D 8 15 10
5 2-1 NML E D 0 CL3-B 50060e8006fc3222 1 0 NML 64562 VSP N M E 8 15 10
5 2-1 NML E D 0 CL5-A 50060e8006fc4150 - - UNK 64562 VSP N M E 8 15 10
5 2-2 NML E D 2 CL3-B 50060e8006fc3222 - - UNK 64562 VSP N A D 8 15 10
5 2-2 NML E D 2 CL5-A 50060e8006fc4150 1 0 NML 64562 VSP N A D 8 15 10
```

Displaying the information of the external volume path group number: 1

```
# raidcom get path -path_grp 1
```

```
PHG GROUP STS CM IF MP# PORT WWN PR LUN PHS Serial# PRODUCT_ID LB PM DM QD TO(s)
PBW(s)
1 1-1 NML E D 2 CL1-A 50060e8005fa0f36 1 3 NML 60010 VSP N M E 8 15 10
1 1-1 NML E D 2 CL2-A 50060e8005fa0f38 2 3 NML 60010 VSP N M D 8 15 10
```

Displaying the information of the external volume path group number: 5

```
# raidcom get path -path_grp 5
```

```
PHG GROUP STS CM IF MP# PORT WWN PR LUN PHS Serial# PRODUCT_ID LB PM DM QD TO(s)
PBW(s)
5 1-1 NML E D 0 CL3-B 50060e8006fc3222 1 0 NML 64562 VSP N M E 8 15 10
5 1-1 NML E D 0 CL5-A 50060e8006fc4150 - - UNK 64562 VSP N M E 8 15 10
5 1-2 NML E D 2 CL3-B 50060e8006fc3222 - - UNK 64562 VSP N A D 8 15 10
5 1-2 NML E D 2 CL5-A 50060e8006fc4150 1 0 NML 64562 VSP N A D 8 15 10
```

Description of each column in output example 1:

PHG

Displays the path group number for the external volume.

GROUP

Displays the external volume group number.

STS

Displays the following status of the external volume.

- NML: "Normal"
- CHK: "Checking"
- SYN: "Cache Destage"
- DSC: "Disconnect"
- BLK: "Blockading" in the external path for the external volume.
- UNK: "Unknown"
- WAR: "Warning"

CM

Displays the cache mode for external volume.

- E: Write cache enabled.
- D: Write cache disabled.

- EM: NDM attribute and Write cache enabled.

- DM: NDM attribute and Write cache disabled.

- TM: NDM attribute and cache through mode.

- SM: NDM attribute and Write sync mode.

IF

Displays the cache inflow control for external volume.

E: Enable, D: Disable

MP#

Displays the MP blade ID for the external volume owner.

PORT

Displays the port number.

WWN

Displays the target wwn on the external storage system. For VSP E series, VSP G130, G/F350, G/F370, G/F700, G/F900, VSP G200, G400, G600, G800 and VSP F400, F600, F800, when the external storage system is connected by iSCSI, this item displays the pseudo WWN of the external storage system.

PR

Displays the priority number in the external volume path group.

LUN

Displays the LUN in the port on the external storage system side.

PHS

Displays the following status of the external path.

- NML: "Normal" status in external path.

- CHK: "temporary blockading " status in external path.

- BLK: "blockading" status in external path.

- DSC: "disconnected" status in external path.

- UNK: "Unknown" status in external path.

Serial#

Displays the serial number of external storage system.

PRODUCT_ID

Displays the PRODUCT_ID of the external storage system.

LB

Displays the following I/O load balance mode for the external storage system.

- N: "normal round robin" mode.

- E: "extended round robin" mode.

- D: Executes with one path without load balance mode.

Chapter 5: Configuration setting commands

If Single is used for the path mode or the load balance mode is not supported, a hyphen (-) is displayed.

PM

Displays the path mode for the external storage system.

- M: Multiple path mode

- S: Single path mode

- A: APLB mode

- AL: ALUA mode

- MA: Multiple path mode (changeable to ALUA mode)

- SA: Single path mode (changeable to ALUA mode)

DM

Displays whether the data direct mapping attribute is set to the external volume group.

- E: The data direct mapping attribute is set.

- D: The data direct mapping attribute is not set.

QD

Displays the number of **read-write** commands that can be issued (queued) at a time to an external volume. (For VSP 5000 series, VSP E series, VSP G/F350, G/F370, G/F700, G/F900)

- - (hyphen): This information is invalid for this external path.

For VSP G1x00 and VSP F1500, a hyphen (-) is always displayed.

TO (s)

Displays the setting value of the I/O time over to the external volume in seconds. (For VSP 5000 series, VSP E series, VSP G/F350, G/F370, G/F700, G/F900)

- - (hyphen): This information is invalid for this external path.

For VSP G1x00 and VSP F1500, a hyphen (-) is always displayed.

PBW (s)

Displays the time in seconds after all the paths connected to the external volume are disconnected until the external volume is blocked. For VSP 5000 series, VSP E series, VSP G/F350, G/F370, G/F700, G/F900)

- - (hyphen): This information is invalid for this external path.

For VSP G1x00 and VSP F1500, a hyphen (-) is always displayed.

Example 2

Waiting until the status of the external volume #1-1 changes to DSC.

When the status has changed to DSC, the command ends with the return value 0. If the status does not change to DSC within 30 minutes, the command times out, and ends with the return value EX_EWSTOT.

```
# raidcom get path -external_grp_id 1-1 -check_status DSC -time 1800
```

Example 3

Waiting until the status of the LDEV #0x10 of the external volume changes to DSC.

When the status has changed to DSC, the command ends with the return value 0. If the status does not change to DSC within 30 minutes, the command times out, and ends with the return value EX_EWSTOT.

```
# raidcom get path -ldev_id 0x10 -check_status DSC -time 1800
```

Example 4

Checking if the external volume#1-1 is in DSC status.

When the status is in DSC, the command ends with the return value 0. If the status is not in DSC, the command ends with the return value 1.

```
# raidcom get path -external_grp_id 1-1 -check_status DSC
```

Example 5

Executing the `raidcom disconnect external_grp` command to the external volume #1-1, and then waiting until the status of the external volume #1-1 changes to DSC.

```
# raidcom disconnect external_grp -external_grp_id 1-1
# raidcom get path -external_grp_id 1-1 -check_status DSC -time 1800
```

raidcom check_ext_storage path

Restores an external path to external VOLs. Only one path is operated in one operation.

This command is executed asynchronously with the command input. Check the completion of this process on the `raidcom get command_status` command.

This command is rejected by EX_ENOOBJ in the following cases:

- The specified iSCSI port cannot be found.

- The iSCSI virtual port mode is enabled, but the specified iSCSI virtual port ID is not correct.

To fix this error, specify the correct port and iSCSI virtual port ID.

> 📄 **Note:** If a failure occurs in a processor on an MP blade or a controller, this command cannot be run even on an MP blade or a controller on which no failure occurs.

Syntax

```
raidcom check_ext_storage path -path_grp <path group#>
    -port <port#> {-external_wwn <wwn strings> |
    -external_iscsi_name <external iscsi name>
```

```
-external_address <IP address>
[-iscsi_virtual_port_id <iSCSI virtual port ID>]}
```

Options and parameters

-path_grp <path group#>
> Specifies the external VOL (0-63231) path group number.

-port <port#>
> Specifies the port number. Specifies the number of the port whose attribute is ELUN (External). For example:

> ▪ CL1-A

-external_wwn <wwn strings>
> Specifies the WWN (16-digit hexadecimal value) of the external storage system. The 17th digit and later are ignored. The value can be split in units of 4 bytes by "," (comma). For example:

> ▪ 210000e08b0256f8

> ▪ 210000e0,8b0256f8

-external_iscsi_name <external iscsi name>
> Specifies the iSCSI name of the iSCSI target on the external storage system in an iqn format or an eui format.

> ▪ iqn format: `iqn.` and the subsequent maximum 219 characters.

> ▪ eui format: `eui.` and the subsequent 16 characters in hexadecimal notation.

> 📄 **Note:**
>
> You can specify the pseudo WWN of the iSCSI target to the `-external_wwn` option instead of specifying the iSCSI target by using the `-external_iscsi_name` option and the `-external_address` option. The pseudo WWN corresponds to the iSCSI name and IP address of the iSCSI target on the external storage system. By executing the **raidcom get external_iscsi_name** command, you can confirm the suitability between the pseudo WWN and the iSCSI name and IP address of the iSCSI target on the external storage system.

-external_address <IP address>
> Specifies the IP address of the iSCSI target on the external storage system. An IPv4 address or an IPv6 address can be specified.

> Note that the following addresses cannot be specified when you specify the IPv4 address:

> ▪ Network address (for example: 192.168.10.0, 0.120.10.1)

> ▪ Broadcast address (for example: 255.255.255.255, 10.1.255.255)

> ▪ Loop back address (for example: 127.0.0.1)

Note that the following addresses cannot be specified when you specify the IPv6 address:

- Unspecified address (for example: ::)
- Multicast address (for example: ff:1024:1215::01)
- Loop back address (for example: ::1)

> 📄 **Note:**
>
> You can specify the pseudo WWN of the iSCSI target to the `-external_wwn` option instead of specifying the iSCSI target by using the `-external_iscsi_name` option and the `-external_address` option. The pseudo WWN corresponds to the iSCSI name and IP address of the iSCSI target on the external storage system. By executing the **raidcom get external_iscsi_name** command, you can confirm the suitability between the pseudo WWN and the iSCSI name and IP address of the iSCSI target on the external storage system.

[-iscsi_virtual_port_id <iSCSI virtual port ID>]

Specifies the iSCSI virtual port ID when the specified port is iSCSI and the iSCSI virtual port mode is enabled. If this option is omitted, 0 is assumed.

An error occurs if you specify this option in any of the following cases:

- The virtual port of the specified iSCSI virtual port ID is not found.
- The iSCSI virtual port mode is disabled.
- The device does not support the iSCSI virtual port mode.

Examples

Restoring a path of external volume path group number: 1, External port CL1-A, and external storage system port 50060e80,05fa0f36.

```
# raidcom check_ext_storage path -path_grp 1 -port CL1-A
     -external_wwn 50060e80,05fa0f36
```

Restoring the path of the external volume path group number: 1, iSCSI port: CL1-A of the local storage system, and the external storage system iSCSI target (iSCSI name: iqn.z2, IP address: 158.214.135.100).

```
# raidcom check_ext_storage path -path_grp 1 -port CL1-A -external_iscsi_name iqn.z2 -
external_address 158.214.135.100
```

Restoring a path of the external volume path group number: 1, the iSCSI port: CL1-A, iSCSI virtual port ID: 2 of the local storage system, and the external storage system iSCSI target (iSCSI name: iqn.z2, IP address: 158.214.135.100).

```
# raidcom check_ext_storage path -path_grp 1 -port CL1-A -iscsi_virtual_port_id 2 -
external_iscsi_name iqn.z2 -external_address 158.214.135.100
```

Chapter 5: Configuration setting commands

raidcom delete path

Deletes the external path or alternative path to an external volume. Only one path is operated in one operation.

This command is executed asynchronously with the command input. Check the completion of this process on the `raidcom get command_status` command.

This command is rejected by EX_ENOOBJ in the following cases:

- The specified iSCSI port cannot be found.

- The iSCSI virtual port mode is enabled, but the specified iSCSI virtual port ID is not correct.

To fix this error, specify the correct port and iSCSI virtual port ID.

> 📄 **Note:** If a failure occurs in a processor on an MP blade or a controller, this command cannot be run even on an MP blade or a controller on which no failure occurs.

Syntax

```
raidcom delete path -path_grp <path group#> -port <port#>
    {-external_wwn <wwn strings> | -external_iscsi_name
    <external iscsi name> -external_address <IP address>
    [-iscsi_virtual_port_id <iSCSI virtual port ID>]}
```

Options and parameters

-path_grp <path group#>
Specifies the external VOL path group number (0-63231).

-port <port#>
Specifies the Port number. Specifies the number of the port whose attribute is ELUN (External). For example:

- CL1-A

-external_wwn <wwn strings>
Specifies the WWN (16-digit hexadecimal value) of the external storage system. The 17th digit and later are ignored. The value can be split in units of 4 bytes by "," (comma). For example:

- 210000e08b0256f8

- 210000e0,8b0256f8

-external_iscsi_name <external iscsi name>
Specifies the iSCSI name of the iSCSI target on the external storage system in an iqn format or an eui format.

- iqn format: `iqn.` and the subsequent maximum 219 characters.

- eui format: `eui.` and the subsequent 16 characters in hexadecimal notation.

> 📄 **Note:**
>
> You can specify the pseudo WWN of the iSCSI target to the `-external_wwn` option instead of specifying the iSCSI target by using the `-external_iscsi_name` option and the `-external_address` option. The pseudo WWN corresponds to the iSCSI name and IP address of the iSCSI target on the external storage system. By executing the **raidcom get external_iscsi_name** command, you can confirm the suitability between the pseudo WWN and the iSCSI name and IP address of the iSCSI target on the external storage system.

-external_address <IP address>

Specifies the IP address of the iSCSI target on the external storage system. An IPv4 address or an IPv6 address can be specified.

Note that the following addresses cannot be specified when you specify the IPv4 address:

- Network address (for example: 192.168.10.0, 0.120.10.1)

- Broadcast address (for example: 255.255.255.255, 10.1.255.255)

- Loop back address (for example: 127.0.0.1)

Note that the following addresses cannot be specified when you specify the IPv6 address:

- Unspecified address (for example: ::)

- Multicast address (for example: ff:1024:1215::01)

- Loop back address (for example: ::1)

> 📄 **Note:**
>
> You can specify the pseudo WWN of the iSCSI target to the `-external_wwn` option instead of specifying the iSCSI target by using the `-external_iscsi_name` option and the `-external_address` option. The pseudo WWN corresponds to the iSCSI name and IP address of the iSCSI target on the external storage system. By executing the **raidcom get external_iscsi_name** command, you can confirm the suitability between the pseudo WWN and the iSCSI name and IP address of the iSCSI target on the external storage system.

[-iscsi_virtual_port_id <iSCSI virtual port ID>]

Specifies the iSCSI virtual port ID when the specified port is iSCSI and the iSCSI virtual port mode is enabled. If this option is omitted, 0 is assumed.

An error occurs if you specify this option in any of the following cases:

- The virtual port of the specified iSCSI virtual port ID is not found.

- The iSCSI virtual port mode is disabled.

- The device does not support the iSCSI virtual port mode.

Chapter 5: Configuration setting commands

Examples

Deleting a path of External port CL1-A, and external storage system port 50060e80,05fa0f36 from the external volume path group number: 1.

```
# raidcom delete path -path_grp 1 -port CL1-A -external_wwn 50060e80,05fa0f36
```

Deleting the path of the external storage system iSCSI target (iSCSI name: iqn.z2, IP address: 158.214.135.100) and the iSCSI port: CL1-A of the local storage system from the external volume path group number: 1.

```
# raidcom delete path -path_grp 1 -port CL1-A -external_iscsi_name iqn.z2 -
external_address 158.214.135.100
```

Deleting the path of the external storage system iSCSI target (iSCSI name: iqn.z2, IP address: 158.214.135.100), the iSCSI port: CL1-A, iSCSI virtual port ID: 2 of the local storage system from the external volume path group number: 1.

```
# raidcom delete path -path_grp 1 -port CL1-A -iscsi_virtual_port_id 2 -
external_iscsi_name iqn.z2 -external_address 158.214.135.100
```

raidcom disconnect path

Blocks the usage of external paths to external VOLs. Only one path is operated in one operation.

This command is executed asynchronously with the command input. Check the completion of this process on the **raidcom get command_status** command.

This command is rejected by EX_ENOOBJ in the following cases:

- The specified iSCSI port cannot be found.

- The iSCSI virtual port mode is enabled, but the specified iSCSI virtual port ID is not correct.

To fix this error, specify the correct port and iSCSI virtual port ID.

> 📄 **Note:** If a failure occurs in a processor on an MP blade or a controller, this command cannot be run even on an MP blade or a controller on which no failure occurs.

Syntax

```
raidcom disconnect path -path_grp <path group#>
   -port <port#> {-external_wwn <wwn strings> |
   -external_iscsi_name <external iscsi name>
   -external_address <IP address>
    [-iscsi_virtual_port_id <iSCSI virtual port ID>]}
```

Chapter 5: Configuration setting commands

Options and parameters

-path_grp <path group#>
>   Specifies the external VOL path group number (0-63231).

-port <port#>
>   Specifies the port number. Specifies the number of the port whose attribute is ELUN (External). For example:

- CL1-A

-external_wwn <wwn strings>
>   Specifies the WWN (16-digit hexadecimal value) of the external storage system. The 17th digit and later are ignored. The value can be split in units of 4 bytes by "," (comma). For example:

- 210000e08b0256f8

- 210000e0,8b0256f8

-external_iscsi_name <external iscsi name>
>   Specifies the iSCSI name of the iSCSI target on the external storage system in an iqn format or an eui format.

- iqn format: `iqn.` and the subsequent maximum 219 characters.

- eui format: `eui.` and the subsequent 16 characters in hexadecimal notation.

> 📄 **Note:**
>
>   You can specify the pseudo WWN of the iSCSI target to the `-external_wwn` option instead of specifying the iSCSI target by using the `-external_iscsi_name` option and the `-external_address` option. The pseudo WWN corresponds to the iSCSI name and IP address of the iSCSI target on the external storage system. By executing the **raidcom get external_iscsi_name** command, you can confirm the suitability between the pseudo WWN and the iSCSI name and IP address of the iSCSI target on the external storage system.

-external_address <IP address>
>   Specifies the IP address of the iSCSI target on the external storage system. An IPv4 address or an IPv6 address can be specified.

>   Note that the following addresses cannot be specified when you specify the IPv4 address:

- Network address (for example: 192.168.10.0, 0.120.10.1)

- Broadcast address (for example: 255.255.255.255, 10.1.255.255)

- Loop back address (for example: 127.0.0.1)

Note that the following addresses cannot be specified when you specify the IPv6 address:

- Unspecified address (for example: ::)

- Multicast address (for example: ff:1024:1215::01)

- Loop back address (for example: ::1)

> 📄 **Note:**
>
> You can specify the pseudo WWN of the iSCSI target to the `-external_wwn` option instead of specifying the iSCSI target by using the `-external_iscsi_name` option and the `-external_address` option. The pseudo WWN corresponds to the iSCSI name and IP address of the iSCSI target on the external storage system. By executing the **`raidcom get external_iscsi_name`** command, you can confirm the suitability between the pseudo WWN and the iSCSI name and IP address of the iSCSI target on the external storage system.

[-iscsi_virtual_port_id <iSCSI virtual port ID>]

Specifies the iSCSI virtual port ID when the specified port is iSCSI and the iSCSI virtual port mode is enabled. If this option is omitted, 0 is assumed.

An error occurs if you specify this option in any of the following cases:

- The virtual port of the specified iSCSI virtual port ID is not found.

- The iSCSI virtual port mode is disabled.

- The device does not support the iSCSI virtual port mode.

Examples

Blocks the usage of a path of the external volume path group number: 1, the External port CL1-A, and the external storage port 50060e80,05fa0f36.

```
# raidcom disconnect path -path_grp 1 -port CL1-A -external_wwn 50060e80,05fa0f36
```

Blocks the usage of the path between the external storage system iSCSI target (iSCSI name: iqn.z2, IP address: 158.214.135.100), and the iSCSI port: CL1-A of the local storage system, and the external volume path group number: 1.

```
# raidcom disconnect path -path_grp 1 -port CL1-A -external_iscsi_name iqn.z2 -
external_address 158.214.135.100
```

Blocks the usage of the path for the external storage system iSCSI target (iSCSI name: iqn.z2, IP address: 158.214.135.100), the iSCSI port: CL1-A, iSCSI virtual port ID: 2 of the local storage system, and the external volume path group number: 1.

```
# raidcom disconnect path -path_grp 1 -port CL1-A -iscsi_virtual_port_id 2 -
external_iscsi_name iqn.z2 -external_address 158.214.135.100
```

raidcom modify path

Supported storage systems:

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

Changes the values set for the external path to the external volume.

📄 **Note:** If a failure occurs in a processor on an MP blade or a controller, this command cannot be run even on an MP blade or a controller on which no failure occurs.

If HIE or a logical path inside the storage is blocked, the operation cannot be performed.

Syntax

```
raidcom modify path -external_wwn <WWN strings> {-qdepth <Q depth>
 | -timeout <timeout> | -path_blocked_watch <time(sec)>}
```

Options and parameters

-external_wwn <WWN strings>*

Specifies the WWN (16-digit hexadecimal value) of the external storage system. The value can be split in units of 4 bytes by a "," (comma).

For example:

- 210000e08b0256f8
- 210000e0, 8b0256f8

-qdepth <Q depth>

Specifies the number of Read/Write commands that can be issued (queued) at a time for external volumes. The same setting value is applied to all the external paths that use the WWN or iSCSI name of the specified external storage system.

-timeout <timeout>

Specifies timeout values in seconds for I/O to external volumes.

-path_blocked_watch <time(sec)>

Specifies a time (in seconds) from when all the paths to the external volumes are disconnected until when all the external volumes are blocked. The same setting value is applied to all the external paths that use the WWN or iSCSI name of the specified external storage system.

* Instead of specifying an iSCSI target by using the `-external_iscsi_name` option and the `-external_address` option, you can specify the pseudo WWN of the iSCSI target to the `-external_wwn` option. The pseudo WWN corresponds to the iSCSI name and IP address of the iSCSI target on the external storage system. By executing the **raidcom get external_iscsi_name** command, you can confirm the suitability between the pseudo WWN and the iSCSI name and IP address of the iSCSI target on the external storage system.

Example

Specifying 15 as the timeout value for I/O to external volumes, for all the external paths that use the port 50060e80 or 05fa0f36 of the external storage subsystem.

```
# raidcom modify path -external_wwn 50060e80,05fa0f36
-timeout 15
```

raidcom delete pool

Deletes the specified Pool for Thin Image, Copy-on-Write Snapshot, Dynamic Provisioning, Dynamic Provisioning for Mainframe, Dynamic Tiering, Dynamic Tiering for Mainframe, active flash, or active flash for mainframe.

When LDEVs or device groups are specified, LDEVs are deleted from the specified pools.

This command is executed asynchronously with the command input. Check the completion of this process on the raidcom get command_status command.

> **Note:** If a failure occurs in a processor on an MP blade or a controller, this command cannot be run even on an MP blade or a controller on which no failure occurs.
>
> If HIE or a logical path inside the storage is blocked, the operation cannot be performed.

Syntax

```
raidcom delete pool -pool {<pool ID#> | <pool naming>}
 [-ldev_id <ldev#> |-grp_opt <group option> -device_grp_name <device group name>
[<device name>] | -delete_volume {yes|no}]
```

Options and parameters

-pool {<pool ID#> | <pool naming>}

Specifies the Pool ID (0-127) or pool name for Thin Image, Copy-on-Write Snapshot, Dynamic Provisioning, Dynamic Provisioning for Mainframe, Dynamic Tiering, Dynamic Tiering for Mainframe, active flash, or active flash for mainframe.

When you specify only a number, it is recognized as a pool ID. Therefore, to specify a pool whose name is a number, use the pool ID instead of the pool name.

-ldev_id <ldev#>

Specifies the LDEV number (0 to 65279), for example:

- -ldev_id 200

-grp_opt <group option>

Specifies the device information extracted from the LDEV belonging to the device group. Specify "ldev" (fixed). The information of LDEV belonging to the device group is used.

-device_grp_name <device group name> [<device name>]

Specifies the name of device group (maximum 32 characters) to be operated.

To operate an LDEV in the device group, specify the device name of the LDEV (maximum 32 characters).

If the device name is omitted, all the LDEVs belonging in the device group are operated.

-delete_volume {yes|no}

Supported storage systems:

- VSP E series

- VSP G130, G/F350, G/F370, G/F700, G/F900

Specifies whether to delete the pool volume when deleting a pool.

- yes (recommended): The pool volume is deleted when a pool is deleted.

- no (default): The pool is deleted, but the pool volume is not deleted.

The external volume is not deleted even if "yes" is specified. To delete it, first delete the pool using the `raidcom delete pool` command, and then execute the `raidcom delete ldev` command.

Examples

Deleting a pool of pool ID 5:

```
# raidcom delete pool -pool 5
```

Deleting a pool of pool name "my_aou_pool":

```
# raidcom delete pool -pool my_aou_pool
```

Deleting a pool volume when deleting a pool of pool name "my_aou_pool":

```
# raidcom delete pool -pool my_aou_pool -delete_volume yes
```

📄 **Note:** "Aou" (allocation on use) refers to Dynamic Provisioning.

raidcom get pool

Displays pool information for Thin Image, Copy-on-Write Snapshot, Dynamic Provisioning, Dynamic Provisioning for Mainframe, Dynamic Tiering, active flash, Dynamic Tiering for Mainframe, and active flash for mainframe.

Syntax

```
raidcom get pool [[-pool {<pool ID#> | <pool naming>}]-key <keyword>] [-fx] [-pcap][-
time_zone <time zone>]
```

Options and parameters

[-key <keyword>]

Specify this option to display the following information that can be specified for *<keyword>*:

- opt: Pool name.

- basic: Basic information about the pool.

- fmc: Information about parity groups that support accelerated compression.

    FMC stands for flash module compression (a drive in a parity group that supports accelerated compression).

- software_saving: Information about the amount of capacity saved by deduplication and compression.

    Supported storage systems:

    - VSP 5000 series

    - VSP G1x00 and VSP F1500

    - VSP E series

    - VSP G130, G/F350, G/F370, G/F700, G/F900

    If the microcode version does not support this option, nothing is displayed when this option is specified.

    If you use the capacity saving function, the saving ratio is calculated against the amount of metadata and garbage data generated internally in addition to user data.

- saving: Information about the amount of saved capacity.

- total_saving: Information about the amount of capacity saved by deduplication and compression, and accelerated compression of the parity group.

    Supported storage systems:

    - VSP 5000 series

    - VSP G1x00 and VSP F1500

    - VSP E series

    - VSP G130, G/F350, G/F370, G/F700, G/F900

- If the microcode version does not support this option, nothing is displayed when this option is specified.

    - If you use the capacity saving function, the saving ratio is calculated against the amount of metadata and garbage data generated internally in addition to user data.

    - This option is ignored when the microcode version does not support the option.

- resource: Information about the parity group configuring the pool or external volume group. When specifying this display keyword, specify the `-pool` option. Otherwise, the `-pool` option cannot be specified.

    Supported storage systems:

    - VSP 5000 series

    - VSP G1x00 and VSP F1500

    - VSP E series

    - VSP G130, G/F350, G/F370, G/F700, G/F900

    If the microcode version does not support this keyword, nothing is displayed when this keyword is specified.

- efficiency: Information about the data saving effect, snapshot effect, provisioning effect, and their total for the pool.

    Supported storage systems:

    - VSP E series

    - VSP G130, G/F350, G/F370, G/F700, G/F900

[-fx]

Displays the LDEV number in hexadecimal.

[-pcap]

Supported storage systems:

- VSP 5000 series

- VSP G1x00 and VSP F1500

- VSP E series

- VSP G130, G/F350, G/F370, G/F700, G/F900

Displays the information about the actual capacity. This option is valid only when the `-key basic` option is specified.

For pools for which accelerated compression is enabled, and for pools to which pool volumes can be added automatically, specify this option to check the capacity information.

The following GUI displays the actual capacity information as the capacity of the accelerated compression enabled pool. Specify this option to confirm the capacity equivalent to the capacity displayed in the GUI.

- Device Manager - Storage Navigator (VSP 5000 series, VSP E series, VSP G130, G/F350, G/F370, G/F700, G/F900)

- Storage Advisor Embedded

This option is ignored when the microcode version does not support the option.

[-time_zone <time zone>]

Supported storage systems:

- VSP 5000 series

- VSP E series

- VSP G130, G/F350, G/F370, G/F700, G/F900

Specifies the time zone displayed in storage systems. If this option is omitted, the time in the time zone set for the storage system is displayed.

This option is valid only when it is specified with the -key efficiency option.

utc: Displays the time in Coordinated Universal Time (UTC).

Example1

Displaying pool information.

```
# raidcom get pool
```

```
PID POLS U(%) SSCNT Available(MB) Capacity(MB) Seq# Num LDEV# H(%) FMT_CAP(MB)
001 POLN 10   330   10000000      1000000000   62500 2   365   80   100
002 POLF 95   9900  100000        1000000000   62500 3   370   70   100
003 POLS 100  10000 100           1000000000   62500 1   375   70   100
004 POLE 0    0     0             0            62500 0   0     80   100
005 POLN 10   330   10000000      1000000000   62500 2   365   80   100
```

Example 2

Displaying pool name.

```
# raidcom get pool -key opt
```

```
PID POLS U(%) POOL_NAME Seq# Num LDEV# H(%) VCAP(%) TYPE PM
PT AUTO_ADD_PLV
001 POLN 10 my_aou_pool 62500 2 265 80 65500 OPEN S
HDP E
002 POLF 95 New_Pool_2 62500 3 270 70 65534 OPEN S
HDP -
003 POLS 100 my_ss_pool 62500 1 275 70 - OPEN N
TI D
004 POLN 0 New_Pool_4 62500 2 280 80 0 M/F N
CW -
005 POLE 0 New_Pool_5 62500 4 0 80 100 M/F S
DM -
```

Description of each column in output examples:

PID
> Displays the pool ID.

POLS
> Displays the status of the pool.
> - POLN: "Pool Normal" (The pool is in the normal status.)
> - POLF: "Pool Full" (The pool is in the overflow status exceeding the threshold.)
> - POLS: "Pool Suspend" (The pool is in the overflow status exceeding the threshold and is blocked.)
> - POLE: "Pool failure" (The pool is blocked in the failure status.) In this status, the pool information cannot be displayed.

U(%)
> Displays the usage rate of the pool.

SSCNT
> Displays the number of volumes in the pool.

POOL_NAME
> Displays the pool name.

Available (MB)
> Displays the capacity available to the volume data in the pool.

Capacity (MB)
> Displays the total capacity of the pool.

Seq#
> Displays the product serial number (Seq#).

Num

Displays the number of LDEVs composing the pool.

LDEV#

Displays the first number of LDEV composing the pool. 65535(ffff) is displayed if the pool is being created.

H(%)

Displays threshold for the pool.

VCAP(%)

Displays the percentage of the subscription limit of V-VOL and Thin Image pair to the pool capacity. For VSP G130, G/F350, G/F370, G/F700, G/F900, a hyphen (-) indicating invalid is displayed.

TYPE

Displays the platform type of pools.

- Open: Shows that it is a Dynamic Provisioning pool.

- M/F: Shows that it is a Dynamic Provisioning for Mainframe pool.

PM

Displays the pool status.

- N: Normal status.

- S: Shrinking or rebalancing.

- NT: The pool for Thin Image is in the normal status.

- ST: (VSP G1x00 and VSP F1500, VSP Gx00 models, and VSP Fx00 models only) The pool for Thin Image is shrinking or rebalancing.

PT

Displays the pool type. Any one of the following types is displayed.

- HDP: Pool for Dynamic Provisioning

- HDT: Pool for Dynamic Tiering

- RT: Pool for active flash

- TI: Pool for Thin Image

- CW: Pool for Hitachi Copy-on-Write Snapshot

- DM: Pool for Dynamic Provisioning that has the data direct mapping attribute

FMT_CAP(MB)

Displays the formatted pool capacity. If there is no valid information for the pool, a hyphen (-) is displayed. In addition, for VSP and HUS VM, a hyphen (-) is always displayed.

AUTO_ADD_PLV

Displays whether to automatically add pool volumes according to the compression ratio of a parity group for which accelerated compression is enabled.

- E: Pool volumes are added automatically according to the compression ratio of the accelerated compression enabled parity group.

- D: Pool volumes in the parity group for which accelerated compression is enabled are not added automatically.

- - (hyphen): This information is invalid for this pool. For VSP and HUS VM, a hyphen (-) is always displayed.

Example 3

Displaying basic information about a pool

```
# raidcom get pool -key basic
```

```
PID POLS U(%)  LCNT  SSCNT Available(MB) Capacity(MB) Snap_Used(MB)
TL_CAP(MB) BM TR_CAP(MB) RCNT   Seq# Num LDEV# W(%) H(%) STIP
VCAP(%) TYPE PM  PT  POOL_NAME
000 POLN   0 11001  11001        46998       46998            0
2432398 NB       0    0 300050   1    0   70   80  YES
UNLIMITED OPEN  N HDP dp_ti_pool
001 POLN   0    - 11001          46998       46998            -
- -          -    - 300050   1    1    -   80    -
- OPEN  N  TI tipool
```

Displaying basic information about a pool. The actual capacity is displayed.

```
# raidcom get pool -key basic -pcap
```

```
PID POLS U(%) LCNT SSCNT Available(MB) Capacity(MB) Snap_Used(MB)
TL_CAP(MB) BM TR_CAP(MB) RCNT Seq# Num LDEV# W(%) H(%) STIP VCAP(%)
TYPE PM PT POOL_NAME
000 POLN 0 11001 11001 46998 46998 -
2432398 NB 0 0 300050 1 0 70 80 YES -
OPEN N HDP dp_ti_pool
001 POLN 0 - 11001 46998 46998
- - - - - 300050 1 1 - 80
- - OPEN N TI tipool
```

Description of each column in output example 3:

PID

Displays the pool ID. Three digits padded with leading zeros.

POLS

    Displays the status of the pool.

- POLN: "Pool Normal" (The pool is in the normal status.)

- POLF: "Pool Full" (The pool is in the overflow status exceeding the threshold.)

- POLS: "Pool Suspend" (The pool is in the overflow status exceeding the threshold and is blocked.)

- POLE: "Pool failure" (The pool is blocked in the failure status.) In this status, the pool information cannot be displayed.

U(%)

    Displays the usage rate of the pool.

    Displays the actual usage rate of the pool when the -pcap option is specified.

LCNT

    Displays the total number of Dynamic Provisioning virtual volumes mapped to the pool. A hyphen (-) is displayed if the information is not available for this pool.

SSCNT

    Displays the total number of snapshot data items mapped to the pool. A hyphen (-) is displayed if the information is not available for this pool.

Available(MB)

    Displays the available capacity for the volume data in the pool.

    Displays the actual capacity that can be used for volume data in the pool when the -pcap option is specified.

Capacity(MB)

    Displays the total capacity of the pool. When the -pcap option is specified, displays the total actual capacity of the pool.

Snap_Used(MB)

    Displays the capacity used for Thin Image data in megabytes. If the used capacity is less than 1 MB, the value is rounded up. A hyphen (-) is displayed if the information is not available for this pool.

    Displays a hyphen (-) when the -pcap option is specified.

TL_CAP(MB)

    Displays the total capacity of all Dynamic Provisioning virtual volumes and Thin Image pairs mapped to the pool. A hyphen (-) is displayed if the information is not available for this pool.

BM

Displays the I/O Blocking Mode of the pool.

- PF (Pool Full): If the pool is full, you cannot read from or write to the target DP-VOL. If the pool VOL is blocked, you can read from or write to the target DP-VOL.

- PB (Pool vol Blockade): If the pool VOL is blocked, you cannot read from or write to the target DP-VOL. If the pool is full, you can read from or write to the target DP-VOL.

- FB (Full or Blockade): If the pool is full or pool VOL is blocked, you cannot read from or write to the target DP-VOL.

- NB (No Blocking): If the pool is full or pool VOL is blocked, you can read from or write to the target DP-VOL.

- - (Not supported): The configuration does not support the I/O Blocking Mode.

TR_CAP(MB)

Displays the sum of the pool capacities reserved for the volumes for which Full Allocation or Proprietary Anchor is enabled. For a configuration that does not support Full Allocation, a hyphen (-) is displayed.

RCNT

Displays the number of volumes for which Full Allocation is enabled that are mapped to a pool. For a configuration that does not support Full Allocation, a hyphen (-) is displayed.

Seq#

Displays the serial number (Seq#).

Num

Displays the number of LDEVs belonging to the pool.

LDEV#

Displays the number of the first LDEV in the pool. "65535 (ffff)" is shown while the pool is being created.

W(%)

Displays the threshold value for WARNING set for the pool. A hyphen (-) is displayed if the information is not available for this pool.

H(%)

Displays the threshold value set for the pool as the high water mark.

STIP

Displays the setting for Thin Image pairs when the high water mark threshold is exceeded.

- YES: Thin Image pairs are suspended.

- NO: Thin Image pairs are not suspended.

- - (hyphen): The information is not available for this pool.

Chapter 5: Configuration setting commands

VCAP(%)

Displays the percentage of the subscription limit of V-VOL and Thin Image to the pool capacity. For VSP G130, G/F350, G/F370, G/F700, G/F900, a hyphen (-) indicating invalid is displayed.

- UNLIMITED: Unlimited.

- - (hyphen): The information is not available for this pool.

TYPE

Displays the platform type of the pool.

- OPEN: Pool for open systems

- M/F: Pool for mainframe systems

PM

Displays the pool status.

- N: Normal status.

- S: Shrinking or rebalancing.

- NT: The pool for Thin Image is in the normal status.

- ST: The pool for Thin Image is shrinking or rebalancing.

PT

Displays the pool type. One of the following types is displayed.

- HDP: Pool for Dynamic Provisioning

- HDT: Pool for Dynamic Tiering

- RT: Pool for active flash

- TI: Pool for Thin Image

- CW: Pool for Hitachi Copy-on-Write Snapshot

- DM: Pool for Dynamic Provisioning that has the data direct mapping attribute

POOL_NAME

Displays the pool name.

Example 4

Displaying information about an accelerated compression-enabled parity group that configures a pool. This information is not displayed if the functionality to show this information is not supported.

For details about the capacity saved by the capacity saving feature, see the *Provisioning Guide* for your storage system.

```
# raidcom get pool -key fmc
```

```
PID U(%) ACT_AV(MB) ACT_TP(MB) FR(%) FMC_LOG_USED(BLK)
FMC_ACT_USED(BLK) FMC_ACT_TP(BLK) FMC_PLV_USED(BLK) FMC_LOG_TP(BLK)
001 90 10000000 100000000 0 0
```

Chapter 5: Configuration setting commands

```
0 0 0 0
002 99 10000 100000000 200 100000000
100000000 51200000000 409608601600 409600000000
004 90 10000000 100000000 150 100000000
100000000 38400000000 307201720320 307200000000
```

Description of each column in output example 4:

PID

> Displays the pool ID.

U(%)

> Displays the actual usage rate of the pool.

ACT_AV(MB)

> Displays the available actual capacity of the volume mapped to this pool for Dynamic Provisioning, Dynamic Provisioning for Mainframe, and Dynamic Tiering.

ACT_TP(MB)

> Displays the total actual capacity of the pool.

FR(%)

> Displays the percentage of the logical capacity against the actual capacity of an accelerated compression-enabled parity group in the pool. 0 is displayed if the pool does not contain an accelerated compression-enabled parity group.

FMC_LOG_USED(BLK)

> Displays the logically used capacity of an accelerated compression-enabled parity group in the pool, in the unit of 512 bytes. 0 is displayed if the pool does not contain an accelerated compression-enabled parity group.

FMC_ACT_USED(BLK)

> Displays the actually used capacity of an accelerated compression-enabled parity group, in the unit of 512 bytes. 0 is displayed if the pool does not contain an accelerated compression-enabled parity group.

FMC_ACT_TP(BLK)

> Displays the total physical capacity of the pool volume that supports accelerated compression used in the pool. 0 is displayed if the pool does not contain a parity group that supports accelerated compression.

FMC_PLV_USED(BLK)

> Displays the logically used page capacity of the pool volume that supports accelerated compression in the pool in the unit of 512 bytes. 0 is displayed if the pool does not contain a parity group that supports accelerated compression.

FMC_LOG_TP(BLK)

> Displays the total logical capacity of the pool volume that supports accelerated compression used in the pool. 0 is displayed if the pool does not contain a parity group that supports accelerated compression.

Example 5

Displaying information about the reduced capacity of data in a pool.

```
# raidcom get pool -key saving
```

```
PID    SE_SAVING(BLK)    PL_SAVING(BLK)    PL_PRE_USED(BLK)    SES(%)    DDP(%)
CMP(%)    PLS(%)    Num    LDEV#
000    12582912    9437184    28311552    20    15    5    33    1    0
001            0    0    -    -    -    -    0    -
```

Description of each column in output example 5:

PID

 Displays the pool ID.

SE_SAVING(BLK)

 Displays the capacity reduced by both capacity saving and accelerated compression of the parity group in number of blocks of the parity group.

 The saved capacity does not include the amount of zero data reduction, metadata, garbage data, metadata reduction, and garbage data reduction.

 For VSP 5000 series, VSP F1500, VSP G1x00, VSP E series, and VSP G/F350, G/F370, G/F700, G/F900, this information is provided for compatibility. See the description of SE_SAVING(BLK) displayed by **# raidcom get pool -key total_saving**.

PL_SAVING(BLK)

 Displays the capacity reduced by capacity saving in number of blocks.

 The saved capacity does not include the amount of zero data reduction, metadata, garbage data, metadata reduction, and garbage data reduction.

 For VSP 5000 series, VSP F1500, VSP G1x00, VSP E series, and VSP G/F350, G/F370, G/F700, G/F900, this information is provided for compatibility. See the description of PL_SAVING(BLK) displayed by **# raidcom get pool -key software_saving**.

PL_PRE_USED(BLK)

 Displays the capacity of data (prior to reduction) to be reduced by capacity saving in number of blocks.

 The capacity prior to reduction does not contain the amount of zero data.

 For VSP 5000 series, VSP F1500, VSP G1x00, VSP E series, and VSP G/F350, G/F370, G/F700, G/F900, this information is provided for compatibility. See the description of PL_PRE_USED(BLK) displayed by **# raidcom get pool -key software_saving**.

SES(%)

Displays the percentage of capacity reduced by both capacity saving and accelerated compression against the capacity used by the virtual volume mapped to this pool before the reduction.

The saved capacity does not include the amount of zero data reduction, metadata, garbage data, metadata reduction, and garbage data reduction.

For VSP 5000 series, VSP F1500, VSP G1x00, VSP E series, and VSP G/F350, G/F370, G/F700, G/F900, this information is provided for compatibility. See the description of `SES(%)` displayed by **# raidcom get pool -key total_saving**.

DDP(%)

Displays the percentage of capacity reduced by deduplication of capacity saving against the capacity used by the virtual volume mapped to this pool before the reduction.

The saved capacity does not include the amount of metadata and garbage data.

For VSP 5000 series, VSP F1500, VSP G1x00, VSP E series, and VSP G/F350, G/F370, G/F700, G/F900, this information is provided for compatibility.

CMP(%)

Displays the percentage of capacity reduced by compression of capacity saving or the percentage of capacity of the parity group reduced by accelerated compression against the capacity used by the virtual volume mapped to this pool before the reduction.

The saved capacity does not include the amount of metadata, garbage data, metadata reduction, and garbage data reduction.

For VSP 5000 series, VSP F1500, VSP G1x00, VSP E series, and VSP G/F350, G/F370, G/F700, G/F900, this information is provided for compatibility.

> **Note:** The capacity used by the virtual volume before reduction, which is the denominator for calculating SES(%), DDP(%), CMP(%), and PLS(%), does not include the capacity used by or reserved for the deduplication system data volume.

PLS(%)

Displays the percentage of capacity reduced by capacity saving against the amount before the reduction, which is calculated as follows:

`PL_SAVING(BLK) / PL_PRE_USED(BLK) × 100[%]`

The saved capacity does not include the amount of zero data reduction, metadata, and garbage data.

For VSP 5000 series, VSP F1500, VSP G1x00, VSP E series, and VSP G/F350, G/F370, G/F700, G/F900, this information is provided for compatibility. See the description of `PLS(%)` displayed by **# raidcom get pool -key software_saving**.

Num

Displays the number of deduplication system data volumes in the pool.

LDEV#

Displays the LDEV number of the deduplication system data volume in the pool. A hyphen (-) is displayed if the number of deduplication system data volumes is zero.

If there is more than one deduplication system data volume, LDEV numbers are separated by a space.

Example 6

Displaying information about the capacity saving effect achieved by the capacity saving function and accelerated compression of a parity group.

For details about the capacity saving effect achieved by the capacity saving feature and accelerated compression of a parity group, see the *Provisioning Guide* for your storage system.

```
# raidcom get pool -key total_saving
PID SE_SAVING(BLK) SES(%) DATAVOL_USED(BLK)
000 12582912 20 56623104
001 0 0 42467328
```

Description of each column in output example:

- SE_SAVING(BLK): Displays the capacity saved by both capacity saving and accelerated compression of the parity group in blocks. The saved capacity includes the amount of zero data reduction, metadata, garbage data, metadata reduction, and garbage data reduction.

  A hyphen (-) indicating an invalid value might be displayed if the amount of used data volume before the reduction is smaller than the used pool capacity.

- SES(%): Displays the percentage of capacity saved by both capacity saving and accelerated compression of the parity group against the capacity used by the virtual volume mapped to this pool before the reduction, which is calculated as follows:

  ```
  SE_SAVING(BLK) / DATAVOL_USED(BLK) × 100[%]
  ```

  When the capacity saving function is used, the saving ratio is calculated against the amount of metadata and garbage data generated internally in addition to user data.

  A hyphen (-) indicating an invalid value might be displayed if the amount of used data volume before the reduction is smaller than the used pool capacity.

- DATAVOL_USED(BLK): Displays the capacity used by a virtual volume in blocks. The capacity used by the deduplication system data volume is not included.

Example 7

Displaying information about the capacity of data stored in a pool saved by the capacity saving function.

For details about the capacity saved by the capacity saving feature, see the *Provisioning Guide* for your storage system.

```
# raidcom get pool -key software_saving
PID PLS(%) PL_SAVING(BLK) CMP(BLK) DDP(BLK) RECLAIM(BLK)
SYSTEM(BLK) PL_PRE_USED(BLK) PRE_CMP_USED(BLK) PRE_DDP_USED(BLK)
000 52 100931731456 0 64424509440 42949672960
6442450944 193273528320 0 193273528320
```

Description of each column in output example:

- PLS(%): Displays the percentage of capacity compressed by capacity saving against the amount of data before the compression, which is calculated as follows:

  ```
  PL_SAVING(BLK) / PL_PRE_USED(BLK) × 100[%]
  ```

  When the capacity saving function is used, the saving ratio is calculated against the amount of metadata and garbage data generated internally in addition to user data.

  A hyphen (-) indicating an invalid value might be displayed if the amount of used data volume before the reduction is smaller than the used pool capacity.

- PL_SAVING(BLK): Displays the capacity saved by capacity saving in blocks. The saved capacity includes the amount of zero data reduction, metadata, and garbage data.

  A hyphen (-) indicating an invalid value might be displayed if the amount of used data volume before the reduction is smaller than the used pool capacity.

- CMP(BLK): Displays the capacity compressed by capacity saving in blocks. The compressed amount does not include the amount of metadata and garbage data.

- DDP(BLK): Displays the capacity saved by deduplication of capacity saving in blocks. The saved amount does not include the amount of metadata and garbage data.

- RECLAIM(BLK): Displays the capacity saved by reclaiming the specified data pattern using capacity saving in blocks. The saved amount does not include the amount of metadata and garbage data.

- SYSTEM(BLK): Displays the amount of consumed system data (metadata and garbage data) for capacity saving in blocks.

- PL_PRE_USED(BLK): Displays the capacity of data (before reduction) to be reduced by capacity saving in blocks.

- PRE_CMP_USED(BLK): Displays the capacity of data (before compression) to be compressed by capacity saving in blocks.

- PRE_DDP_USED(BLK): Displays the capacity of data (before deduplication) to be deduplicated by capacity saving in blocks.

Displaying information about the data saving effect , snapshot effect, provisioning effect, and their total for the pool.

For details about each effect, see the *Provisioning Guide* for your storage system.

```
#raidcom get pool -key efficiency
PID TOTAL_EFF_R TLS_R PLS_R PLS_CMP_R
PLS_DDP_R PLS_RECLAIM_R FMD_SAVING_R FMD_CMP_R
```

Chapter 5: Configuration setting commands

```
FMD_RECLAIM_R SNAPSHOT_EFF_R PROVISIONING_EFF(%)
CALCULATION_START CALCULATION_END
000 15.18 3.79 2.25 1.34
1.64 1.09 2.22 2.18
1.06 7.25 80
2015-05-15T10:05 2015-05-15T10:38
001 9.78 1.67 1.44 1.21
1.20 1.01 1.26 1.26
1.00 5.85 40
2015-05-15T10:05 2015-05-15T10:38
```

> **Note:** The following values are the maximum of each saving effect ratio available to display in HDvM - SN or CCI:
>
> - HDvM - SN: 99999999999999.99
>
> - CCI: 92233720368547758.07 (a hundredth value of hexadecimal 0x7FFFFFFFFFFFFFFF.)
>
> If no data is stored in the pool as shown in the following example, the saving effect ratio displays the maximum available value. If data is stored, the saving effect ratio calculated from the amount of actual data reduction is displayed. The following examples show the conditions where the maximum values are displayed:
>
> Example 1. Total Efficiency Ratio after the pool and DP-VOL creations, but before data writing.
>
> - HDvM - SN: 99999999999999.99
>
> - CCI: 92233720368547758.07
>
> Example 2. Snapshot Efficiency Ratio applied when the pool includes no Thin Image pairs, but only Thin Image Advanced pairs in the PAIR status or ones with no differential data.
>
> - HDvM - SN: 99999999999999.99
>
> - CCI: 92233720368547758.07

- TOTAL_EFF_R

  Displays the ratio for the total (excluding system data) of the data saving effect, snapshot effect, and provisioning effect.

  - NC: Calculated information is not available.

  - - (hyphen): The information is not available for this pool.

- TLS_R

  Displays the ratio of the data saving effect (excluding system data) achieved by the capacity saving feature and accelerated compression of the parity group.

  - NC: Calculated information is not available.

  - - (hyphen): The information is not available for this pool.

- **PLS_R**

  Displays the ratio of the data saving effect (excluding system data) of the capacity saving feature.

  - NC: Calculated information is not available.

  - - (hyphen): The information is not available for this pool.

- **PLS_CMP_R**

  Displays the ratio of the data saving effect (excluding system data) achieved by compression of the capacity saving feature.

  - NC: Calculated information is not available.

  - - (hyphen): The information is not available for this pool.

- **PLS_DDP_R**

  Displays the ratio of the data saving effect (excluding system data) achieved by deduplication of the capacity saving feature.

  - NC: Calculated information is not available.

  - - (hyphen): The information is not available for this pool.

- **PLS_RECLAIM_R**

  Displays the ratio of the saving effect (excluding system data) achieved by reclaiming the specified data pattern using the capacity saving feature.

  - NC: Calculated information is not available.

  - - (hyphen): The information is not available for this pool.

- **FMD_SAVING_R**

  Displays the ratio of the data saving effect (excluding system data) achieved by accelerated compression of a parity group.

  - NC: Calculated information is not available.

  - - (hyphen): The information is not available for this pool.

- **FMD_CMP_R**

  Displays the ratio of the data saving effect (excluding system data) achieved by compression of accelerated compression for a parity group.

  - NC: Calculated information is not available.

  - - (hyphen): The information is not available for this pool.

- **FMD_RECLAIM_R**

  Displays the ratio of the data saving effect (excluding system data) achieved by reclaiming the specified data pattern using accelerated compression of a parity group.

  - NC: Calculate information is not available.

  - - (hyphen): The information is not available for this pool.

Chapter 5: Configuration setting commands

- **SNAPSHOT_EFF_R**

  Displays the ratio of the snapshot effect (excluding system data).

  - NC: Calculated information is not available.
  - - (hyphen): The information is not available for this pool.

- **PROVISIONING_EFF (%)**

  Displays the percentage of the provisioning effect (excluding system data).

  - NC: Calculated information is not available.
  - - (hyphen): The information is not available for this pool.

- **CALCULATION_START**

  Displays the start date and time for calculating the data saving effect, snapshot effect, provisioning effect, and their total (excluding system data).

  - NC: Calculated information is not available.
  - - (hyphen): The information is not available for this pool.

- **CALCULATION_END**

  Displays the end date and time for calculating the data saving effect, snapshot effect, provisioning effect, and their total (excluding system data).

  - NC: Calculated information is not available.
  - - (hyphen): The information is not available for this pool.

Displaying information about the parity group configuring the pool and the external volume group.

```
# raidcom get pool -pool 0 -key resource
PID T GROUP AV_CAP(GB) TOTAL_CAP(GB) R_LVL R_TYPE TNO V E AUTO_ADD_PLV
000 R 5-2 300 300 RAID1 2D+2D 1 D D -
000 E 1-1 300 300 - - 2 D - -
```

Description of each column in output example:

- T: Displays the type of the volume group.

  - R: Parity group

  - E: External volume group

- Group: Displays the parity group number or the external volume group number.

- AV_CAP(GB): Displays the remaining available capacity (physically available area) of the parity group or the external volume group. A value less than 1 GB is rounded down and 0 is displayed.

- TOTAL_CAP(GB): Displays the total capacity that can be used by the parity group or external volume group.

- R_LVL: Displays the RAID level of the parity group or external volume group. A hyphen is displayed for external volume groups.

- R_TYPE: Displays the RAID type of the parity group or external volume group. A hyphen is displayed for external volume groups.

- TNO: Displays the number of the tier to which the parity group or external volume group belongs. The tier number is displayed only for Dynamic Tiering pools. A hyphen (-) is displayed for other pools.

- V: Displays whether capacity virtualization of the parity group is enabled or disabled.

  - E: Capacity virtualization is enabled.

  - D: Capacity virtualization is disabled.

- E: Displays whether encryption is enabled or disabled.

  - E: Encryption is enabled.

  - D: Encryption is disabled.

  - - (hyphen): Displaying encryption status is not supported.

- AUTO_ADD_PLV: Displays whether to automatically add pool volumes for which accelerated compression is enabled according to the compression ratio of a parity group for which accelerated compression is enabled.

  - E: Automatically adds pool volumes for which accelerated compression is enabled according to the compression ratio of a parity group for which accelerated compression is enabled.

  - D: Pool volumes for which accelerated compression is enabled are not added automatically.

  - - (hyphen): This information is invalid for this parity group. If accelerated compression is disabled for the parity group, a - (hyphen) is displayed.

raidcom modify pool

Sets the options of a pool for Thin Image, Copy-on-Write Snapshot, Dynamic Provisioning, Dynamic Provisioning for Mainframe, Dynamic Tiering, active flash, Dynamic Tiering for Mainframe, or active flash for mainframe. This option also changes the setting of the data direct mapping attribute of the pool for Dynamic Provisioning.

When the `-status` or `-deduplication` option is specified, the operation is executed asynchronously with the command input. Check the completion of this process by using the **raidcom get command_status** command.

📄 **Note:** If a failure occurs in a processor on an MP blade or a controller, this command cannot be run even on an MP blade or a controller on which no failure occurs.

Syntax

```
raidcom modify pool -pool {<pool ID#> | <pool naming>}
   {-status {nml|stop_shrinking} | -user_threshold <threshold_1>
   [<threshold_2>] | -tier <Tier number> [<ratio>]
   [-tier_buffer_rate <%>] | -subscription <%> |
   -pool_attribute <pool_attribute> | -monitor_mode <Monitor mode> |
   -blocking_mode <IO blocking mode> | -data_direct_mapping {enable|disable}
   | -deduplication yes -ldev_id <ldev#>… [-ssids <ssid> …] |
   -deduplication no | -suspend_tipair {yes | no} |
   -delete dsd_volumes} | -auto_add_poolvol enable
   [-password <One Time Password>] | -auto_add_poolvol disable [-password
   <One Time Password>]
```

Options and parameters

-pool {<pool ID#> | <pool naming>}

Specifies the Pool ID (0-127) or pool name for Thin Image, Copy-on-Write Snapshot, Dynamic Provisioning, Dynamic Provisioning for Mainframe, Dynamic Tiering, active flash, Dynamic Tiering for Mainframe, or active flash for mainframe.

When only a number is specified, it is recognized as a pool ID. Therefore, to specify a pool whose name is a number, use the pool ID instead of the pool name.

-status <status>

Supported storage systems:

▪ VSP G200, G400, G600, G800 and VSP F400, F600, F800

▪ VSP

▪ HUS VM

Specifies the status of the pool. To restore the pool status, specify "nml".

Chapter 5: Configuration setting commands

-status {nml|stop_shrinking}

> Supported storage systems:
>
> - VSP 5000 series
> - VSP E series
> - VSP G130, G/F350, G/F370, G/F700, G/F900
>
> Specifies the status of the pool:
>
> - nml: Pool status is restored.
> - stop_shrinking: Stops shrinking the pool.

-user_threshold <threshold_1> [<threshold_2>]

> Sets a user-defined threshold.
>
> - For Dynamic Provisioning, Dynamic Provisioning for Mainframe, Dynamic Tiering, active flash, Dynamic Tiering for Mainframe, and active flash for mainframe, you may specify two user-defined thresholds, and their valid ranges are 1-100%, respectively. If you specify <threshold_1> and <threshold_2>:
>
>   the value of <threshold_1> is set as the threshold for WARNING specified to a pool.
>
>   the value of <threshold_2> is set as the threshold for High water mark specified to a pool.
>
>   If you specify only <threshold_1>, your specified value and the system default value (80%) are applied.
>
>   The setting in which you specify only <threshold_1> is supported to maintain backward compatibility with microcode before 70-02-0x-xx/xx. Once you specify two user-defined thresholds, you must continue to specify two user-defined thresholds thereafter.
>
> - The valid range for Thin Image or Copy-on-Write Snapshot is 20-95%. You may specify only <threshold_1>. Even if you specify the value for <threshold_2>, the value is ignored.

-tier <Tier number> [<ratio>]

> Specifies the pool information for Dynamic Tiering, active flash, Dynamic Tiering for Mainframe, or active flash for mainframe. This is specified when changing the free space percentage for a new allocation by each tier. When this option is specified, pool attribute changes into manual relocation.
>
> <Tier number>:Tier number (1-3)
>
> <ratio>: Free space percentage for new allocation(0-50) [%]

[-tier_buffer_rate <%>]

> Specifies the pool information for Dynamic Tiering, active flash, Dynamic Tiering for Mainframe, or active flash for mainframe. Specifies the amount of reallocation buffer (2-40) for each tier in percent (%).

-subscription <%>

> For VSP G130, G/F350, G/F370, G/F700, G/F900, an error occurs if you specify this option.

You can specify this option for VSP 5000 series and VSP E series. However, an error occurs if a value other than 65535 (unlimited) is specified for the following pools:

- The pool that contains virtual volumes for which capacity saving feature, deduplication and compression are enabled.

- The pool that contains pool volumes that belongs to a parity group for which accelerated compression is enabled.

Sets the percentage (0-65535) of the subscription limit of V-VOL and Thin Image pair to the pool capacity.

- 0 to 65534: Specified percentage

- 65535: Unlimited

-pool_attribute <pool_attribute>

Specifies when changing the pool from Dynamic Provisioning/Dynamic Provisioning for Mainframe to Dynamic Tiering/Dynamic Tiering for Mainframe, or from Dynamic Tiering/Dynamic Tiering for Mainframe to Dynamic Provisioning/Dynamic Provisioning for Mainframe.

Specifies the following value depend on the type of pool desired to be changed.

- dt_manual: Changes the pool from Dynamic Provisioning/Dynamic Provisioning for Mainframe to Dynamic Tiering/Dynamic Tiering for Mainframe. It is changed to manual relocation.

- dp: Changes the pool from Dynamic Tiering/Dynamic Tiering for Mainframe to Dynamic Provisioning/Dynamic Provisioning for Mainframe.

-monitor_mode <monitor mode>

Specifies the monitoring mode for a Dynamic Tiering/Dynamic Tiering for Mainframe pool. Or specifies the availability of active flash.

- period: Performs monitoring periodically.

- continuous: Performs monitoring continuously.

- realtime_tiering: Enables active flash.

- non_realtime_tiering: Disables active flash.

-blocking_mode <IO blocking mode>

Sets the I/O activity (availability for read/write access) when the pool for Dynamic Provisioning, Dynamic Tiering, or active flash is full and when the pool for Dynamic Provisioning, Dynamic Tiering, or active flash is blocked. This option is ignored when the microcode version does not support the option.

> 📄 **Note:** If the microcode version does not support this option, "-" is displayed in the BM column when you execute the `raidcom get dp_pool` command.

- pool_full: If the pool is full, read/write access for the target DP-VOL is disabled. If the pool is blocked, read/write access for the target DP-VOL is enabled.

- pool_vol_blockade: If the pool-VOL is blocked, read/write access for the target DP-VOL is disabled. If the pool-VOL is full, read/write access for the target DP-VOL is enabled.

- full_or_blockade: If the pool is full or blocked, read/write access for the target DP-VOL is disabled.

- no_blocking: If the pool is full and/or blocked, read/write access for the target DP-VOL is enabled.

-data_direct_mapping {enable| disable}

For VSP 5000 series, VSP G1x00, VSP F1500, VSP E series, VSP G130, G/F350, G/F370, G/F700, G/F900, VSP G200, G400, G600, G800 and VSP F400, F600, F800: Changes the setting of the data direct mapping attribute of the pool for Dynamic Provisioning.

- enable: Changes a pool for Dynamic Provisioning to a pool for Dynamic Provisioning that has the data direct mapping attribute.

- disable: Changes a pool for Dynamic Provisioning that has the data direct mapping attribute to a pool for Dynamic Provisioning.

-deduplication yes -ldev_id <ldev#> … [-ssids <ssid> …]

Enables deduplication of the pool. When this option is specified, a deduplication system data volume that has an LDEV number (0 to 65279) specified for "-ldev_id <ldev#> …" is created.

For VSP 5000 series, VSP E series, VSP G/F350, G/F370, G/F700, G/F900, an error occurs if you specify this option.

If you want to assign an SSID to the deduplication system data volume to be created, specify "-ssids <ssid> …" in hexadecimal. SSIDs specified by "-ssids <ssid> …" are assigned to deduplication system data volumes specified for "-ldev_id <ldev#> …" in the order they were specified. If you omit this option, or specify auto for <ssid>, SSIDs are assigned automatically to the corresponding deduplication system data volumes.

If you specify multiple LDEVs, an error occurs because the current microcode does not support this function. If this happens, deduplication of the pool is not enabled.

-deduplication no

Disables deduplication of the pool. When this option is specified, a deduplication system data volume in this pool is deleted.

For VSP 5000 series, VSP E series, VSP G/F350, G/F370, G/F700, G/F900, an error occurs if you specify this option.

[-suspend_tipair {yes | no}]

> Specifies or cancels the setting for suspending Thin Image pairs when the high water mark threshold is exceeded.

> ▪ yes: Suspends Thin Image pairs.

> ▪ no: Cancels the setting of suspending Thin Image pairs.

-delete dsd_volumes

> Supported storage systems:

> ▪ VSP 5000 series

> ▪ VSP E series

> ▪ VSP G130, G/F350, G/F370, G/F700, G/F900

> Deletes the deduplication system data volume.

-auto_add_poolvol enable [-password <One Time Password>]

> Supported storage systems:

> ▪ VSP 5000 series

> ▪ VSP G1x00 and VSP F1500

> ▪ VSP G/F350, G/F370, G/F700, G/F900

> Automatically adds pool volumes for which accelerated compression of the parity group is enabled according to the compression ratio of the parity group for which accelerated compression is enabled.

-auto_add_poolvol disable [-password <One Time Password>]

> Supported storage systems:

> ▪ VSP 5000 series

> ▪ VSP G1x00 and VSP F1500

> ▪ VSP G/F350, G/F370, G/F700, G/F900

> Pool volumes for which accelerated compression of the parity group is enabled are not added automatically.

Examples

Restoring the status of a pool ID: 6.

```
# raidcom modify pool -pool 6 -status nml
```

Restoring the status of a pool name: my_ss_pool.

```
# raidcom modify pool -pool my_ss_pool -status nml
```

Changing the user-defined thresholds of the pool ID:6 of the pool for Dynamic Provisioning, Dynamic Provisioning for Mainframe, Dynamic Tiering, active flash, Dynamic Tiering for Mainframe, or active flash for mainframe, WARNING to 70% and High water mark to 80%.

```
# raidcom modify pool -pool 6 -user_threshold 70 80
```

Changing the user-defined threshold of the pool ID: 6 of the pool for Thin Image or Copy-on-Write Snapshot to 80%.

```
# raidcom modify pool -pool 6 -user_threshold 80
```

Changing the free space percentage for a new allocation to the tier number 1 of the pool ID:6 for Dynamic Tiering, active flash, Dynamic Tiering for Mainframe, or active flash for mainframe to 30%.

```
# raidcom modify pool -pool 6 -tier 1 30
```

Changing the amount of reallocation buffer to the tier number 1 of the pool ID:6 for Dynamic Tiering, active flash, Dynamic Tiering for Mainframe, or active flash for mainframe to 20%.

```
# raidcom modify pool -pool 6 -tier 1 -tier_buffer_rate 20
```

Changing a Dynamic Provisioning/Dynamic Provisioning for Mainframe pool of the pool ID:6 to a Dynamic Tiering/Dynamic Tiering for Mainframe pool.

```
# raidcom modify pool -pool 6 -pool_attribute dt_manual
```

Changing a Dynamic Tiering/Dynamic Tiering for Mainframe pool of the pool ID:6 to a Dynamic Provisioning/Dynamic Provisioning for Mainframe pool.

```
# raidcom modify pool -pool 6 -pool_attribute dp
```

Changing the automatic relocation of the pool (Pool name: my_pool) for Dynamic Provisioning to manual relocation.

```
# raidcom modify pool -pool my_pool -pool_attribute dt_manual
```

Changing the monitoring mode of Dynamic Tiering/Dynamic Tiering for Mainframe pool of the pool ID: 6 to continuous.

```
# raidcom modify pool -pool 6 -monitor_mode continuous
```

When the pool is full, changing read/write of Dynamic Provisioning pool ID: 6 to rejected.

```
# raidcom modify pool -pool 6 -blocking_mode pool_full
```

Enabling active flash of Dynamic Tiering pool ID: 6.

```
# raidcom modify pool -pool 6 -monitor_mode realtime_tiering
```

Changing a pool (pool ID: 6) for Dynamic Provisioning to a pool for Dynamic Provisioning that has the data direct mapping attribute.

```
# raidcom modify pool -pool 6 -data_direct_mapping enable
```

Enabling deduplication of the pool (pool ID: 6) for Dynamic Provisioning to create a deduplication system data volume of LDEV: 400. For VSP 5000 series, VSP E series, VSP G/F350, G/F370, G/F700, G/F900, an error occurs.

```
# raidcom modify pool -pool 6 -deduplication yes -ldev_id 400
```

Enabling deduplication of the pool (pool ID: 6) for Dynamic Provisioning to create a deduplication system data volume of LDEV: 400, and assigning SSID: 0x0004 to the deduplication system data volume. For VSP 5000 series, VSP E series, VSP G/F350, G/F370, G/F700, G/F900, an error occurs.

```
# raidcom modify pool -pool 6 -deduplication yes -ldev_id 400 -ssids 0x0004
```

Disabling deduplication of the pool (pool ID: 6) for Dynamic Provisioning to delete a deduplication system data volume. For VSP 5000 series, VSP E series, VSP G/F350, G/F370, G/F700, G/F900, an error occurs.

```
# raidcom modify pool -pool 6 -deduplication no
```

Deleting the deduplication system data volume of the pool (pool ID: 6) for Dynamic Provisioning.

```
# raidcom modify pool -pool 6 -delete dsd_volumes
```

Suspending Thin Image pairs when the high water mark threshold of the pool (pool ID: 6) for Dynamic Provisioning is exceeded.

```
# raidcom modify pool -pool 6 -suspend_tipair yes
```

Canceling the setting for suspending Thin Image pairs when the high water mark threshold of the pool (pool ID: 6) for Dynamic Provisioning is exceeded.

```
# raidcom modify pool -pool 6 -suspend_tipair no
```

Disabling the setting for automatically adding pool volumes for which accelerated compression is enabled in the pool (pool ID: 6).

```
# raidcom modify pool -pool 6 -auto_add_poolvol disable -password
<One Time Password>
```

Disabling the setting for automatically adding pool volumes for which accelerated compression is enabled in the pool (pool ID: 6).

```
# raidcom modify pool -pool 6 -auto_add_poolvol disable
```

Enabling the setting for automatically adding pool volumes for which accelerated compression is enabled in the pool (pool ID: 6).

```
# raidcom modify pool -pool 6 -auto_add_poolvol enable -password <One Time Password>
```

Enabling the setting for automatically adding pool volumes for which accelerated compression is enabled in the pool (pool ID: 6).

```
# raidcom modify pool -pool 6 -auto_add_poolvol enable
```

raidcom monitor pool

Sets the start or stop of performance monitoring of a pool for Dynamic Tiering, Dynamic Tiering for Mainframe, active flash, or active flash for mainframe.

> **Note:** If a failure occurs in a processor on an MP blade or a controller, this command cannot be run even on an MP blade or a controller on which no failure occurs.

Syntax

```
raidcom monitor pool -pool {<pool ID#> | <pool naming>}
    -operation <type>
```

Options and parameters

-pool {<pool ID#> | <pool naming>}
Specifies the pool ID (0-127) or pool name of a pool for Dynamic Tiering, Dynamic Tiering for Mainframe, active flash, or active flash for mainframe.

When specifying just a number, it is recognized as a pool ID. Therefore, when you specify a pool that the pool name is only a number, operate by specifying pool ID instead of specifying pool name.

-operation <type>
Instructs the operation of performance monitoring.

The operational types that can be specified are shown below.

- start: Start the performance monitoring.

- stop: Stop the performance monitoring.

Examples

Starting the performance monitoring of a pool for Dynamic Tiering, Dynamic Tiering for Mainframe, active flash, or active flash for mainframe with Pool ID: 20.

```
# raidcom monitor pool -pool 20 -operation start
```

Starting the performance monitoring of a pool for Dynamic Tiering, Dynamic Tiering for Mainframe, active flash, or active flash for mainframe with Pool Name: my_tier_pool.

```
# raidcom monitor pool -pool my_tier_pool -operation start
```

Stopping the performance monitoring of a pool for Dynamic Tiering, Dynamic Tiering for Mainframe, active flash, or active flash for mainframe with Pool ID: 20.

```
# raidcom monitor pool -pool 20 -operation stop
```

Stopping the performance monitoring of a pool for Dynamic Tiering, Dynamic Tiering for Mainframe, active flash, or active flash for mainframe with Pool Name: my_tier_pool.

```
# raidcom monitor pool -pool my_tier_pool -operation stop
```

raidcom reallocate pool

Sets the start or stop of tier relocation of a pool for Dynamic Tiering, Dynamic Tiering for Mainframe, active flash, or active flash for mainframe.

📄 **Note:** If a failure occurs in a processor on an MP blade or a controller, this command cannot be run even on an MP blade or a controller on which no failure occurs.

Syntax

```
raidcom reallocate pool -pool {<pool ID#> | <pool naming>}
    -operation <type>
```

Options and parameters

-pool {<pool ID#> | <pool naming>}

Specifies the pool ID (0-127) or pool name of a Dynamic Tiering, Dynamic Tiering for Mainframe, active flash, or active flash for mainframe pool.

When only a number is specified, it is recognized as a pool ID. Therefore, when you specify a pool whose name is only a number, use the pool ID instead of the pool name.

-operation <type>

Relocation operation instruction:

- start: Start the tier relocation.

- stop: Stop the tier relocation.

Examples

Starting the tier relocation of a pool for Dynamic Tiering, Dynamic Tiering for Mainframe, active flash, or active flash for mainframe with Pool: 20.

```
# raidcom reallocate pool -pool 20 -operation start
```

Starting the tier relocation of a pool for Dynamic Tiering, Dynamic Tiering for Mainframe, active flash, or active flash for mainframe with Pool Name: my_tier_pool.

```
# raidcom reallocate pool -pool my_tier_pool -operation start
```

Stopping the relocation of a pool for Dynamic Tiering, Dynamic Tiering for Mainframe, active flash, or active flash for mainframe with Pool: 20.

```
# raidcom reallocate pool -pool 20 -operation stop
```

Stopping the relocation of a pool for Dynamic Tiering, Dynamic Tiering for Mainframe, active flash, or active flash for mainframe with Pool Name: my_tier_pool.

```
# raidcom reallocate pool -pool my_tier_pool -operation stop
```

raidcom rename pool

Changes the pool name of a pool for Thin Image, Copy-on-Write Snapshot, Dynamic Provisioning, Dynamic Provisioning for Mainframe, Dynamic Tiering, Dynamic Tiering for Mainframe, active flash, or active flash for mainframe.

> **Note:** If a failure occurs in a processor on an MP blade or a controller, this command cannot be run even on an MP blade or a controller on which no failure occurs.

Syntax

```
raidcom rename pool -pool_id <pool ID#> -pool_name <pool naming>
```

Options and parameters

-pool_id <pool ID#>

Specifies a pool ID (0-127).

Chapter 5: Configuration setting commands

If a -pool_id option is non-numeric, the specified value is recognized as a pool name to identify the pool ID.

-pool_name <pool naming>

Specifies a new pool name. You can specify up to 32 characters.

The -pool_name option cannot be specified by configuring only in numeric because numeric values are preferentially identified as a pool ID. Specifies a pool ID with the '-pool_id<pool ID#>' option.

Examples

Changing the pool name of the pool (ID: 1) to my_pool.

```
# raidcom rename pool -pool_id 1 -pool_name my_pool
```

raidcom initialize pool

This command initializes deduplication system data volumes and volumes that contain deduplicated data. This command is processed asynchronously with the command input. Use the `raidcom get command_status` command to check if the `raidcom initialize pool` command completes.

> 📄 **Note:** If a failure occurs in a processor on an MP blade or a controller, this command cannot be run even on an MP blade or a controller on which no failure occurs.

Syntax

```
raidcom initialize pool -pool {<pool ID#> | <pool naming>} -operation <type>
```

Options and parameters

-pool {<pool ID#> | <pool naming>}

Specifies the pool ID (0 to 127) or pool name.

-operation <type>

Supported storage systems:

- VSP G1x00 and VSP F1500

- VSP G/F350, G/F370, G/F700, G/F900

Specifies the operation to perform. Specify the following:

initialize_deduplication: Initializes the deduplication system data volume specified by the -pool option and the volumes that contain deduplicated data.

-operation initialize_deduplication

Supported storage systems:

- VSP G200, G400, G600, G800, and VSP F400, F600, F800

Initializes the deduplication system data volume specified by the `-pool` option and the volumes that contain deduplicated data.

Examples

Initializing the deduplication system data volume and the volumes that contain deduplicated data in the pool with pool ID: 1.

```
#raidcom initialize pool -pool 1 -operation initialize_deduplication
```

raidcom modify system

Supported storage systems:

- VSP 5000 series

- VSP G1x00 and VSP F1500

- VSP E series

- VSP G130, G/F350, G/F370, G/F700, G/F900

Sets the description of the storage system.

> 📄 **Note:** If a failure occurs in a processor on an MP blade or a controller, this command cannot be run even on an MP blade or a controller on which no failure occurs.

Syntax

```
raidcom modify system -system_operation {set_description -description <description> |
reset_description}
```

Options and parameters

-system_operation set_description -description <description>

Sets the description of the storage system using 1 to 255 characters. You can specify all characters that can be used for raidcom commands. For details about the characters, see Supported characters (on page 14).

-system_operation reset_description

Deletes the description of the storage system.

Example 1

Setting the character string, "This is a storage system." as the description of the storage system:

```
raidcom modify system -system_operation set_description -description "This is a
storage system."
```

Example 2

Deleting the description of the storage system.

```
raidcom modify system -system_operation reset_description
```

raidcom get system

Displays the data saving effect, snapshot effect, provisioning effect, and availability of compression accelerator.

Syntax

```
# raidcom get system [-key <keyword>][-time_zone <time zone>]
```

Options and parameters

[-key <keyword>]

Specify the following keyword for *<keyword>*.

- ▪ [efficiency] (VSP 5000 series, VSP E series, VSP G130, G/F350, G/F370, G/F700, G/F900)

  Displays the data saving, snapshot, and provisioning effects of the system.

- ▪ [information] (VSP 5000 series, VSP G1x00, VSP F1500, VSP E series, VSP G130, G/F350, G/F370, G/F700, G/F900)

  Displays the information set for the storage system.

- ▪ [dedupe_compression] (VSP 5000 series, VSP E1090)

  Displays the information on the availability of compression accelerator.

[-time_zone <time zone>]

Specifies the time zone displayed in storage systems. If this option is omitted, the time in the time zone set for the storage system is displayed.

(VSP 5000 series, VSP E series, VSP G130, G/F350, G/F370, G/F700, G/F900) This option is valid only when it is specified with the -key efficiency option.

utc: Displays the time in Coordinated Universal Time (UTC).

Chapter 5: Configuration setting commands

Example 1

Displays the status of the system (VSP 5000 series, VSP E series, VSP G/F350, G/F370, G/F700, G/F900, VSP G200, G400, G600, G800, VSP F400, F600, F800).

```
#raidcom get system
Serial# : 12345
SYSTEM_TIME : 2016-09-21T08:30:20
AVE(W) : 500
MODEL : RH10HG
```

Description of each column in output examples:

Serial#

Displays the storage system serial number:

- The serial number for VSP 5000 series is displayed with a "5" at the beginning ("512345" = serial number 12345).

- The serial number for VSP G1x00, VSP F1500 is displayed with a "3" at the beginning ("312345" = serial number 12345).

SYSTEM_TIME

Displays the controller system time when the acquisition instruction is issued. It is represented in yyyy-mm-ddThh:mm:ss format:

- yyyy (year): 4-digit decimal number in A.D.

- mm (month): 2-digit decimal number (01-12)

- dd (day): 2-digit decimal number (01-31)

- hh (hour): Time in 24 hours, 2-digit decimal number (00-23)

- mm (minute): 2-digit decimal number (00-59)

- ss (second): 2-digit decimal number (00-59)

AVE(W)

Displays the average power consumption per minute collected at the time of command execution.

If an invalid value is returned, - (hyphen) is displayed.

MODEL

Displays the model information of storage system.

If the storage system model information display is not supported, - (hyphen) is displayed.

The following model information is displayed:

- R9F: VSP 5100, VSP 5500

- R9G: VSP 5100H, VSP 5500H

- RH10HF: VSP 5200, VSP 5600

- RH10HG: VSP 5200H, VSP 5600H

Chapter 5: Configuration setting commands

Example 2

Displays information about the data saving effect, snapshot effect, provisioning effect, and their total.

For details about each effect, see the *Provisioning Guide* for your storage system.

```
#raidcom get system -key efficiency
```

```
Serial# TOTAL_EFF_R TLS_R PLS_R PLS_CMP_R
PLS_DDP_R PLS_RECLAIM_R FMD_SAVING_R FMD_CMP_R
FMD_RECLAIM_R SNAPSHOT_EFF_R PROVISIONING_EFF(%)
CALCULATION_START CALCULATION_END
12345 11.64 1.84 1.52 1.15
1.34 1.11 2.21 2.14
1.08 10.37 70
2015-05-15T10:05 2015-05-15T10:38
```

Description of each column in output examples:

> **Note:** The following values are the maximum of each saving effect ratio available to display in HDvM - SN or CCI:
>
> - HDvM - SN: 99999999999999.99
>
> - CCI: 92233720368547758.07 (a hundredth value of hexadecimal 0x7FFFFFFFFFFFFFFF.)
>
> If no data is stored in the pool as shown in the following example, the saving effect ratio displays the maximum available value. If data is stored, the saving effect ratio calculated from the amount of actual data reduction is displayed. The following examples show the conditions where the maximum values are displayed:
>
> Example 1. Total Efficiency Ratio after the pool and DP-VOL creations, but before data writing.
>
> - HDvM - SN: 99999999999999.99
>
> - CCI: 92233720368547758.07
>
> Example 2. Snapshot Efficiency Ratio applied when the pool includes no Thin Image pairs, but only Thin Image Advanced pairs in the PAIR status or ones with no differential data.
>
> - HDvM - SN: 99999999999999.99
>
> - CCI: 92233720368547758.07

TOTAL_EFF_R

Displays the ratio for the total (excluding system data) of data saving effect, snapshot effect, and provisioning effect.

- NC: Calculated information is not available.

- - (hyphen): The information is not available for this pool.

TLS_R

Displays the ratio of the data saving effect (excluding system data) achieved by the capacity saving feature and accelerated compression of a parity group.

- NC: Calculated information is not available.

- - (hyphen): The information is not available for this pool.

PLS_R

Displays the ratio of the data saving effect (excluding system data) achieved by the capacity saving feature.

- NC: Calculated information is not available.

- - (hyphen): The information is not available for this pool.

PLS_CMP_R

Displays the ratio of the data saving effect (excluding system data) achieved by compression of the capacity saving feature.

- NC: Calculated information is not available.

- - (hyphen): The information is not available for this pool.

PLS_DDP_R

Displays the ratio of the data saving effect (excluding system data) achieved by deduplication of the capacity saving feature.

- NC: Calculated information is not available.

- - (hyphen): The information is not available for this pool.

PLS_RECLAIM_R

Displays the ratio of the saving effect (excluding system data) achieved by reclaiming the specified data patterns using the capacity saving feature.

- NC: Calculated information is not available.

- - (hyphen): The information is not available for this pool.

FMD_SAVING_R

Displays the ratio of the data saving effect (excluding system data) achieved by accelerated compression of a parity group.

- NC: Calculated information is not available.

- - (hyphen): The information is not available for this pool.

FMD_CMP_R

Displays the ratio of the data saving effect (excluding system data) achieved by compression of accelerated compression for a parity group.

- NC: Calculated information is not available.

- - (hyphen): The information is not available for this pool.

FMD_RECLAIM_R

Displays the ratio of the data saving effect (excluding system data) achieved by reclaiming the specified data patterns using accelerated compression of a parity group.

- NC: Calculated information is not available.

- - (hyphen): The information is not available for this pool.

SNAPSHOT_EFF_R

Displays the ratio of the snapshot effect (excluding system data).

- NC: Calculated information is not available.

- - (hyphen): The information is not available for this pool.

PROVISIONING_EFF(%)

Displays the percentage of the provisioning effect (excluding system data).

- NC: Calculated information is not available.

- - (hyphen): The information is not available for this pool.

CALCULATION_START

Displays the start date and time for calculating the data saving effect, snapshot effect, provisioning effect, and their total (excluding system data).

- NC: Calculated information is not available.

- - (hyphen): The information is not available for this pool.

CALCULATION_END

Displays the end date and time for calculating the data saving effect, snapshot effect, provisioning effect, and their total (excluding system data).

- NC: Calculated information is not available.

- - (hyphen): The information is not available for this pool.

Example 3

Displays the information set for the storage system.

```
#raidcom get system -key information
```

```
DESCRIPTION: This is a storage system.
```

Description of each column in output examples:

DESCRIPTION

Displays the description of the storage system.

Example 4

Displays availability of compression accelerator.

```
#raidcom get system -key dedupe_compression
COMPRESSION_ACCELERATION: AVAILABLE
```

Description of each column in output examples:

COMPRESSION_ACCELERATION

Displays availability of compression accelerator.

- AVAILABLE: Compression accelerator is available.

- UNAVAILABLE: Compression accelerator is unavailable.

raidcom get port

Displays port information.

This queries the setting information on all ports.

Syntax

```
raidcom get port  [-key detail]
```

```
raidcom get port [-port <port#> [-key <keyword> [-iscsi_virtual_port_id
<iSCSI virtual port ID>]]]
```

Options and parameters

[-port <port#>]

Specifies the port number (for example, CL1-A). The port type you specify must be FIBRE, FCoE, or iSCSI. If you specify a port for which LUN security is enabled, the following items are displayed:

- If the target port type is FIBRE or FCoE: LOGIN_WWN

- If the target port type is iSCSI: LOGIN_ISCSI_NAME

[-key <keyword>]

Specifies one of the following values:

- opt: Displays the detailed information of FIBRE, FCoE, or iSCSI. To specify this option, you must specify the `-port <port>` option at the same time.

- login_host_nqn: Specifying this option for a port whose operation mode is NVMe mode displays the login status of the host NQN that is logged in to the target port. To specify this option, you must specify the `-port <port>` option at the same time.

- detail: Displays a column showing the additional information. To specify this option, you must not specify the `-port <port>` option at the same time.

[-iscsi_virtual_port_id <iSCSI virtual port ID>]

Displays detailed information when the specified port is iSCSI and the iSCSI virtual port mode is enabled. If this option is omitted, information about the virtual port ID: 0 is displayed.

An error occurs if you specify this option in any of the following cases:

- The virtual port of the specified iSCSI virtual port ID is not found.

- The iSCSI virtual port mode is disabled.

- The device does not support the iSCSI virtual port mode.

Example 1

Displaying port information.

For a mainframe port (ESCON, FICON®), "-" is displayed for the values from SPD to SSW. If the DKCMAIN microcode supports display of transfer speed, SPD displays the setting value of the transfer speed for the FICON® port.

For a port used for Hitachi NAS (HNASS or HNASU), "-" is displayed for the values from ATTR to SSW and for WWN.

```
# raidcom get port
```

```
PORT   TYPE ATTR SPD LPID FAB CONN SSW SL Serial# WWN       PHY_PORT
CL1-A FIBRE TAR AUT  EF   N   FCAL  N   0 64568  50060e8006fc3800 -
CL1-B FIBRE TAR AUT  EF   N   FCAL  N   0 64568  50060e8006fc3801 -
CL1-C FICON TAR  -   -    -   -     -   0 64568  50060e8006fc3802 -
CL1-D FICON TAR  -   -    -   -     -   0 64568  50060e8006fc3803 -
CL2-A FIBRE TAR AUT  D9   N   FCAL  N   0 64568  50060e8006fc3810 -
CL2-B FIBRE TAR AUT  D3   N   FCAL  Y   0 64568  50060e8006fc3811 -
CL2-C FICON TAR  -   -    -   -     -   0 64568  50060e8006fc3812 -
CL2-D FICON TAR  -   -    -   -     -   0 64568  50060e8006fc3813 -
CL3-A FIBRE MCU AUT  E8   N   FCAL  N   0 64568  50060e8006fc3820 -
CL3-B FIBRE TAR AUT  E0   N   FCAL  Y   0 64568  50060e8006fc3821 -
CL3-J ISCSI TAR AUT  00   N   UNKN  N   0 64015  -                -
CL4-A HNASS -    -   -    -   -     -   0 64568  -                -
CL4-B HNASU -    -   -    -   -     -   0 64568  -                -
```

Chapter 5: Configuration setting commands

Description of each column in output example:

PORT

Displays the port numbers.

TYPE

Displays the following port type: FIBRE, SCSI, ISCSI, ENAS, ESCON, FICON®, FCoE, HNASS, HNASU.

- HNASS means the port is for the system LU of Hitachi NAS.

- HNASU means the port is for a user LU of Hitachi NAS.

ATTR

Displays the following attribute setting on a port. One of the following items is displayed. If the port is a bidirectional port, the following four attributes are all displayed for each port. If you specify raidcom modify port-port_attribute TAR for the bidirectional port in VSP 5000 series, only the attribute TAR is displayed. When the raidcom modify port-port attribute ALL is specified, the following attributes are displayed.

- TAR: Fibre target port (target port)

- MCU: MCU initiator port (initiator port)

- RCU: RCU target port (RCU target port)

- ELUN: External initiator port (External port)

SPD

Displays the transfer rate setting on a port. One of the following values is displayed: AUT(AUTO), 1G, 2G, 4G, 8G, 10G…

LPID

Displays the ALPA (Arbitrated Loop Physical Address) setting value on a port.

FAB

Displays the fabric mode setting on a port as Y(YES) or N(NO).

CONN

Displays the following topology setting on a port: FCAL/PtoP/UNKN/-. If the port does not support the topology setting, UNKN or a hyphen is displayed.

SSW

Displays the LUN security setting on a port as Y (enabled) or N (disabled).

SL

Displays the SLPR number to which the port belongs.

Serial#

Product serial number (Serial#).

WWN

Displays the external WWN.

PHY_PORT

> Displays the port number of the physical port in a resource ID takeover configuration.

Example 2

Displaying the additional port information

```
# raidcom get port -key detail
PORT TYPE ATTR SPD LPID FAB CONN SSW SL Serial#
WWN PHY_PORT PORT_MODE
CL1-A FIBRE TAR AUT EF N FCAL N 0 64568
50060e8006fc3800 - SCSI
CL1-B FIBRE TAR AUT EF N FCAL N 0 64568
50060e8006fc3801 - NVME
CL1-C FICON TAR - - - - - 0 64568
50060e8006fc3802 - -
CL1-D FICON TAR - - - - - 0 64568
50060e8006fc3803 - -
CL2-A FIBRE TAR AUT D9 N FCAL N 0 64568
50060e8006fc3810 - SCSI
CL2-B FIBRE TAR AUT D3 N FCAL Y 0 64568
50060e8006fc3811 - SCSI
CL2-C FICON TAR - - - - - 0 64568
50060e8006fc3812 - -
CL2-D FICON TAR - - - - - 0 64568
50060e8006fc3813 - -
CL3-A FIBRE MCU AUT E8 N FCAL N 0 64568
50060e8006fc3820 - SCSI
CL3-B FIBRE TAR AUT E0 N FCAL Y 0 64568
50060e8006fc3821 - SCSI
CL3-J ISCSI TAR 10G 00 N UNKN N 0 64015
- - -
CL4-A HNASS - - - - - - 0 64568
- - -
CL4-B HNASU - - - - - - 0 64568
- - -
```

Description of each column in output example:

PORT_MODE

> Operation mode of the port. Displays one of the following:
>
> - SCSI: SCSI mode
>
> - NVME: NVMe mode
>
> - – (hyphen): The information is not available for this port.

Example 3

Displaying the information of the port CL4-E (in case the TYPE is FIBRE or FCoE)

```
# raidcom get port -port CL4-E
```

```
PORT  LOGIN_WWN       Serial# -
CL4-E 210000e08b0256f8 63528   OLA_NODE0_CTL_0
CL4-E 210000e08b039c15 63528   OLA_NODE1_CTL_0
```

Description of each column in output example:

PORT

Displays the port numbers.

LOGIN_WWN

Displays the WWN of the host adapter login to this port.

The WWN nickname specified in Performance Manager, or the SPM name specified in Server Priority Manager is displayed.

The WWN nickname specified with the `raidcom set hba_wwn` command is not displayed.

> 📄 **Note:** Only the currently connected WWN is displayed unlike the display on Device Manager - Storage Navigator.

Serial#

Product serial number (Serial#).

Example 4

Displaying the information of the port CL4-E (in case of TYPE is ISCSI)

```
# raidcom get port -port CL4-E
```

```
PORT   LOGIN_IQN   Serial#   -
CL4-E  iqn.z1…     63528     OLA_NODE0_CTL_0
CL4-E  iqn.z2…     63528     OLA_NODE1_CTL_0
```

Description of each column in output example:

PORT

Displays the port numbers.

LOGIN_IQN

Displays the iSCSI name for the host adapter log-in to this port that is currently being connected.

CCI displays only the iSCSI name of the port that is currently being connected, differently from Storage Navigator.

Serial#
Product serial number.

Example 5

Example for getting information on FCoE option:

```
# raidcom get port -port CL4-E -key opt
```

```
PORT  ENMA               VLAN_ID FPMA               VPS  VP_I FCF_I
CL4-E e3:00:00:e0:8b:02 0x03fe  e2:00:00:e0:8b:02 DWN  0x00 0x0000
```

Description of each column in output example:

PORT
Displays the port number.

ENMA
Displays Enode MAC address setting to this port.

VLAN_ID
Displays the VLAN identifier.

FPMA
Displays FP MAC address setting to this port.

VPS
Displays the virtual port status.

- DWN: the Link status is in Link Down.

- LOT: the Link status is in Link Up and Log-Out.

- LIN: the Link status is in Link Up and Log-In.

VP_I
Displays the virtual port index (zero is currently displayed).

FCF_I
Displays the FCoE index (zero is currently displayed).

Example 6

Displaying the example of iSCSI.

```
# raidcom get port -port CL4-E -key opt
```

```
PORT : CL4-E
```

Chapter 5: Configuration setting commands

```
TCP_OPT : IPV6_E : SACK_E : DACK_E : INS_E : VTAG_E
TCP_MTU : 1500
WSZ : 64KB
KA_TIMER : 30
TCP_PORT : 3260
IPV4_ADDR : 158.214.135.100
IPV4_SMSK : 255.255.255.255
IPV4_GWAD : 158.214.135.101
IPV6_ADDR_INF : STS : AM : fe80::209:6bff:febe:3c17
IPV6_GADR_INF : STS : AM : fe80::209:6bff:febe:3c17
IPV6_GWAD_INF : STS : fe80::209:6bff:febe:3c17 : fe80::209:6bff:febe:3c17
ISNS_PORT : 3260
ISNS_ADDR : 158.214.135.101
VLAN_ID : 0001
ISCSI_VP_MODE : E
ISCSI_VP_ID : 0 1 15
IPV6_GADR2_INF : STS : AM : fe80::209:6bff:febe:3c18
MAC_ADDR : e3:00:00:e0:8b:02
```

Description of each column in output example:

PORT

Displays the port numbers.

TCP_OPT

Displays whether each option for iSCSI communication (IPv6 mode, Selective Ack mode, delayed ACK mode, iSNS mode, and Tag VLAN) is enabled or disabled. The meanings of the displayed value are as follows.

- IPV6_E: IPv6 mode is enabled.
- IPV6_D: IPv6 mode is disabled.
- SACK_E: Selective ACK mode is enabled.
- SACK_D: Selective ACK mode is disabled.
- DACK_E: Delayed ACK mode is enabled.
- DACK_D: Delayed ACK mode is disabled.
- INS_E: iSNS service is enabled.
- INS_D: iSNS service is disabled.
- VTAG_E: Tag VLAN is enabled.
- VTAG_D: Tag VLAN is disabled.

TCP_MTU

Displays the MTU value for iSCSI communication.

WSZ

Displays the window size for iSCSI communication.

KA_TIMER

Displays the Keep Alive Timer value for iSCSI communication.

TCP_PORT

Displays the TCP port number for iSCSI communication.

IPV4_ADDR

Displays IPv4 address.

IPV4_SMSK

Displays IPv4 subnet mask.

IPV4_GWAD

Displays IPv4 address of the gateway to use for iSCSI communication.

IPV6_ADDR_INF

Displays the status of IPv6 link local address, and the value of address acquiring mode and address. The details of STS and AM in the displayed example are described in the table below.

Item in the displayed example	Meanings of the item	Displayed value	Meanings of the value
STS	Status	INV	Invalid
		VAL	Valid
		ACQ	acquiring
		DUP	duplicated
AM	address acquiring mode	AM	Automatic acquiring
		MM	Manual acquiring

IPV6_GADR_INF

Displays the status of IPv6 Global address, and the value of address acquiring mode and address. The details of STS and AM in the displayed example are described in the table below.

Item in the displayed example	Meanings of the item	Displayed value	Meanings of the value
STS	Status	INV	Invalid
		VAL	Valid
		ACQ	acquiring

Item in the displayed example	Meanings of the item	Displayed value	Meanings of the value
		DUP	duplicated
AM	address acquiring mode	AM	Automatic acquiring
		MM	Manual acquiring

IPV6_GWAD_INF
Displays the IPv6 Global address of the gateway to use for the iSCSI communication. The values of address are displayed in order of address and current address. The details of STS in the displayed example are described in the table below.

Item in the displayed example	Meanings of the item	Displayed value	Meanings of the value
STS	Status	INV	Invalid
		VAL	Valid
		ACQ	acquiring
		DUP	duplicated

ISNS_PORT
Displays the TCP port number of iSNS server. If iSNS is invalid, this item is not displayed.

ISNS_ADDR
Displays the address of iSNS server. If iSNS is invalid, this item is not displayed.

VLAN_ID
Displays VLAN ID in decimal. If VLAN ID is not set, a hyphen (-) is displayed.

ISCSI_VP_MODE
Displays the iSCSI virtual port mode status.

- E: Enabled
- D: Disabled

When the iSCSI virtual port mode is not supported, a hyphen (-) is displayed.

ISCSI_VP_ID
Lists the iSCSI virtual port IDs defined for the specified ports in decimal. When the iSCSI virtual port mode is disabled, this item is not displayed.

IPV6_GADR2_INF

Displays the IPv6 Global address2 status, and the value of address acquiring mode and address. Enabled for the virtual port of which physical port or iSCSI virtual port ID is 0. When the iSCSI virtual port ID is from 1 to 15, "- : - : -" is displayed. The details of STS and AM in the displayed example are described in the following table.

Item in the displayed example	Meanings of the item	Displayed value	Meanings of the value
STS	Status	INV	Invalid
		VAL	Valid
		ACQ	Acquiring
		DUP	Duplicated
AM	Address acquiring mode	AM	Automatic acquiring
		MM	Manual acquiring

MAC_ADDR

Displays the MAC address of the port. When the MAC address output function is not supported, a hyphen (-) is displayed.

Example 7

Example of detailed information when the port is iSCSI and the iSCSI virtual port mode is enabled.

```
# raidcom get port -port CL4-E -key opt -iscsi_virtual_port_id 15
```

```
PORT : CL4-E
ISCSI_VP_ID : 15
TCP_OPT : IPV6_E : SACK_E : DACK_E : INS_E : VTAG_E
TCP_MTU : 1500
WSZ : 64KB
KA_TIMER : 30
TCP_PORT : 3260
IPV4_ADDR : 158.214.135.100
IPV4_SMSK : 255.255.255.255
IPV4_GWAD : 158.214.135.101
IPV6_ADDR_INF : STS : AM : fe80::209:6bff:febe:3c17
IPV6_GADR_INF : STS : AM : fe80::209:6bff:febe:3c17
IPV6_GWAD_INF : STS : fe80::209:6bff:febe:3c17 : fe80::209:6bff:febe:3c17
ISNS_PORT : 3260
ISNS_ADDR : 158.214.135.101
```

```
VLAN_ID : 0001
IPV6_GADR2_INF : - : - : -
```

Description of each column in output example:

PORT

Displays the port numbers.

ISCSI_VP_ID

Displays the iSCSI virtual port ID.

TCP_OPT

Displays whether each option for iSCSI communication (IPv6 mode, Selective ACK mode, delayed ACK mode, iSNS mode, and Tag VLAN) is enabled or disabled. The meanings of the displayed values are as follows:

- IPV6_E: IPv6 mode is enabled.

- IPV6_D: IPv6 mode is disabled.

- SACK_E: Selective ACK mode is enabled.

- SACK_D: Selective ACK mode is disabled.

- DACK_E: Delayed ACK mode is enabled.

- DACK_D: Delayed ACK mode is disabled.

- INS_E: iSNS service is enabled.

- INS_D: iSNS service is disabled.

- VTAG_E: Tag VLAN is enabled.

- VTAG_D: Tag VLAN is disabled.

TCP_MTU

Displays the MTU value for iSCSI communication.

WSZ

Displays the window size for iSCSI communication.

KA_TIMER

Displays the Keep Alive Timer value for iSCSI communication.

TCP_PORT

Displays the TCP port number for iSCSI communication.

IPV4_ADDR

Displays IPv4 address.

IPV4_SMSK

Displays IPv4 subnet mask.

IPV4_GWAD:

Displays IPv4 address of the gateway to use for iSCSI communication.

IPV6_ADDR_INF

Displays the status of IPv6 link local address, and the value of address acquiring mode and address. The details of STS and AM in the displayed example are described in the table below.

Item in the displayed example	Meanings of the item	Displayed value	Meanings of the value
STS	Status	INV	Invalid
		VAL	Valid
		ACQ	Acquiring
		DUP	Duplicated
AM	Address acquiring mode	AM	Automatic acquiring
		MM	Manual acquiring

IPV6_GADR_INF

Displays the status of IPv6 Global address, and the value of address acquiring mode and address. The details of STS and AM in the displayed example are described in the following table.

Item in the displayed example	Meanings of the item	Displayed value	Meanings of the value
STS	Status	INV	Invalid
		VAL	Valid
		ACQ	Acquiring
		DUP	Duplicated
AM	Address acquiring mode	AM	Automatic acquiring
		MM	Manual acquiring

IPV6_GWAD_INF

Displays the IPv6 Global address of the gateway to use for the iSCSI communication. The values of address are displayed in order of address and current address. The details of STS in the displayed example are described in the following table.

Item in the displayed example	Meanings of the item	Displayed value	Meanings of the value
STS	Status	INV	Invalid
		VAL	Valid
		ACQ	Acquiring
		DUP	Duplicated

ISNS_PORT
　　Displays the TCP port number of the iSNS server. If iSNS is disabled, this item is not displayed.

ISNS_ADDR
　　Displays the address of the iSNS server. If iSNS is invalid, this item is not displayed.

VLAN_ID
　　Displays the VLAN ID in decimal. If the VLAN ID is not set, a hyphen (-) is displayed.

IPV6_GADR2_INF
　　Displays the IPv6 Global address2 status, and the value of address acquiring mode and address. Enabled for the virtual port of which physical port or iSCSI virtual port ID is 0. When the iSCSI virtual port ID is from 1 to 15, "- : - : -" is displayed. The details of STS and AM in the displayed example are described in the following table.

Item in the displayed example	Meanings of the item	Displayed value	Meanings of the value
STS	Status	INV	Invalid
		VAL	Valid
		ACQ	Acquiring
		DUP	Duplicated
AM	Address acquiring mode	AM	Automatic acquiring
		MM	Manual acquiring

Example 8

Displaying the example of FIBRE.

```
# raidcom get port -port CL4-E -key opt
```

```
PORT   S  LKNSPD  CURADR  T
CL4-E  U      16  821A00  D
```

Description of each column in output example:

PORT

Displays the port numbers.

S

Displays the link status of the port.

- U: The port is link up status.

- D: The port is not link up status.

- - (hyphen): The port does not support to display the link status.

The item is displayed for VSP 5000 series , VSP E series, VSP G130, G/F350, G/F370, G/F700, G/F900, VSP G200, G400, G600, G800 and VSP F400, F600, F800.

LNKSPD

Displays the present transfer speed of the port by Gbps. If the port does not support to display the present transfer speed or the port is not link up status, a hyphen is displayed.

CURADR

Displays the present port address of the port by hexadecimal number. If the port does not support to display the present port address or the port is not link up status, a hyphen is displayed.

T

Displays the setting of the T10 PI mode of the port.

- E: T10 PI mode is enabled.

- D: T10 PI mode is disabled.

- - (hyphen): T10 PI is not supported.

Example 9

Displaying the login status of the host NQN that is logged in to the port CL4-E.

```
# raidcom get port -port CL4-E -key login_host_nqn
PORT LOGIN_STATUS HOST_NQN
CL4-E LOGIN
nqn.2014-08.org.nvmexpress:uuid:24ffc04c-74ed-413d-9261-0a40bfb4585f
```

Chapter 5: Configuration setting commands

```
CL4-E LOGOUT
nqn.2014-08.org.nvmexpress:uuid:24ffc04c-74ed-413d-9261-0a40bfb4584f
```

Description of each column in output example:

HOST_NQN
> Displays Host NQN.

raidcom modify port

Sets the attribute of the specified port.

When you set a port attribute by options other than a `-port_attribute` option, the port type of the specified port must be one of the following:

- FIBRE

- FCoE

- iSCSI

- FICON®

If these conditions are not satisfied, this command is rejected with EX_ENOOBJ.

When using an FCoE package, specify Port speed: 10G, Port topology: f_port. In this case, the port attribute cannot be changed.

When you set the T10 PI mode or specify the `-request_id` option, this command is executed asynchronously with the command input. Check the completion of this process using the **raidcom get command_status** command.

There is a limit to the number of asynchronous commands that can be accepted simultaneously. For details, see Maximum number of acceptable asynchronous commands (on page 228).

If the port type is FICON®, you can change only the port speed of 32 Gbps FICON® port (shown as "4Mx32(Mfibre)" in the applicable windows). To check the CHB type, see the **Port Condition** window of Device Manager - Storage Navigator. For details about how to display the **Port Condition** window and the contents displayed in the window, see the *System Administrator Guide*.

The syntax is separated into 6 groups as follows. Syntaxes 2 through 6 can be used only for iSCSI ports. If you execute a command that includes options in syntax 3 with options in syntax 2, 4, 5, or 6, an error occurs.

When the `{-t10pi {enable|disable} | -iscsi_virtual_port_mode {enable | disable} | -port_mode {scsi | nvme} -request_id auto | -delete_login_host_nqn}` option is specified, if a failure occurs in a processor on an MP blade or a controller, this command cannot be run even on an MP blade or a controller on which no failure occurs.

Syntax

Syntax 1

```
raidcom modify port -port <port#>{[-port_speed <value>]
   [-loop_id <value>][-topology <topology>][-security_switch
   {y | n}] | -port_attribute <port attribute> | -t10pi {enable|disable}
    | -iscsi_virtual_port_mode {enable|disable}| -port_mode {scsi | nvme} -request_id
auto |
-delete_login_host_nqn}
```

Syntax 2

```
raidcom modify port -port <port#> [-mtu <value>]
   [-vlan_tagging_mode {enable|disable}] [-add_vlan_id <value>]
   [-delete_vlan_id <value>] [-modify_vlan_id <new value>]
   [-ipv4_address <address>] [-ipv4_subnetmask <subnet mask>]
   [-ipv4_gateway_address <address>] [-ipv6_mode {enable|disable}]
   [-ipv6_local_address {auto|<address>}]
   [-ipv6_global_address {auto|<address>}]
   [-ipv6_gateway_address <address>] [-tcp_port <value>]
   [-selective_ack_mode {enable|disable}]
   [-delayed_ack_mode {enable|disable}]
   [-window_size <size>] [-keep_alive_timer <value>]
   [-ipv6_global_address2 <address>]
```

Syntax 3

```
raidcom modify port -port <port#> [-isns_mode {enable|disable}]
   [-isns_server_address <IPv4 address|IPv6 address>]
   [-isns_port <iSNS TCP Port number>]
```

Syntax 4

```
raidcom modify port -port <port> -add_iscsi_virtual_port <value>
   -ipv6_mode {enable|disable} [-mtu <value>]
   [-vlan_tagging_mode {enable|disable}] [-add_vlan_id <value>]
   [-ipv4_address <address>] [-ipv4_subnetmask <subnet mask>]
   [-ipv4_gateway_address <address>]
   [-ipv6_local_address {auto|<address>}]
   [-ipv6_global_address {auto|<address>}]
   [-ipv6_gateway_address <address>] [-tcp_port <value>]
   [-selective_ack_mode {enable|disable}]
   [-delayed_ack_mode {enable|disable}]
   [-window_size <size>] [-keep_alive_timer <value>]
```

Syntax 5

```
raidcom modify port -port <port> -modify_iscsi_virtual_port <value>
   [-mtu <value>] [-vlan_tagging_mode {enable|disable}]
```

Chapter 5: Configuration setting commands

```
[-add_vlan_id <value>] [-delete_vlan_id <value>]
[-modify_vlan_id <new value>]
[-ipv4_address <address>] [-ipv4_subnetmask <subnet mask>]
[-ipv4_gateway_address <address>] [-ipv6_mode {enable|disable}]
[-ipv6_local_address {auto|<address>}] [-ipv6_global_address {auto|<address>}]
[-ipv6_gateway_address <address>] [-tcp_port <value>]
[-selective_ack_mode {enable|disable}] [-delayed_ack_mode {enable|disable}]
[-window_size <size>] [-keep_alive_timer <value>]
[-ipv6_global_address2 <address>]
```

Syntax 6

```
raidcom modify port -port <port> -delete_iscsi_virtual_port <value>
```

Options and parameters

-port <port#>

Specifies the port number. For example:

- CL1-A

[-port_speed <value>]

Specifies the port speed (0, 1, 2, 4, 8, 10, 16, 32) as follows.

- 0: AUTO

- 1: 1G

- 2: 2G

- 4: 4G

- 8: 8G

- 10: 10G

- 16: 16G

- 32: 32G

The specifiable values differ depending on the type of the channel board or channel adapter.

[-loop_id <alpa value>]

Specifies the Loop ID (0x01-0xEF) of the Port.

[-topology <topology>]

Specifies the topology of the Port as follows:

- fl_port: fabric on and fcal

- f_port: fabric on and PtoP

- nl_port: fabric off and fcal

- n_port: fabric off and PtoP

When you change the topology setting of a port, delete the remote path and external path defined for that port beforehand. If an LU path is defined for the port, make sure in advance that the host is not performing I/O processing and the state is not the host reserve (mounted).

[-security _switch {y | n}]

Specifies whether to use the security switch or not.

-port_attribute <port attribute>

Specifies the Port attribute as follows:

- TAR: Fibre target port (target port)

- MCU: MCU initiator port (initiator port)

- RCU: RCU target port (RCU target port)

- ELUN: External initiator port (External port)

- ALL: Resets (removes the TAR attribute from) a bidirectional port.

If the port is a bidirectional port of VSP 5000 series, VSP E series, VSP G130, G/F350, G/F370, G/F700, G/F900, VSP G200, G400, G600, G800 and VSP F400, F600, F800, the storage system detects the appropriate attribute of the port, and then operates the port with the detected attribute. If TAR is specified as the attribute value for the bidirectional port, the port attribute is changed to TAR (Target), and the storage system treats the port as a target port. By specifying ALL as the attribute value for the port specified in the TAR attribute, the user can restore the operation with the appropriate port attribute. If the package type of VSP G130, G/F350, G/F370, G/F700, G/F900, VSP G200, G400, G600, G800 and VSP F400, F600, F800 is iSCSI, the user cannot set the port attribute.

If this port attribute is changed from Target or RCU Target to Initiator or External, the host group belonging to this port belongs to meta_resource.

-t10pi {enable|disable}

Specifies the T10 PI mode as follows:

- enable: T10 PI mode is enable

- disable: T10 PI mode is disable

When you change the T10 PI mode, the T10 PI mode of the ports which share the settings with the port are also changed. The user who executes the command must have authority of the specified port and the ports which share the settings with the port. Details about the ports which share the settings with the specified port, see the *Provisioning Guide* for the storage system.

-iscsi_virtual_port_mode {enable|disable}

Specifies the iSCSI virtual port mode as follows:

- enable: iSCSI virtual port mode is enabled

- disable: iSCSI virtual port mode is disabled

The setting of the iSCSI virtual port mode is shared among all ports on the channel adapter (CHA) or channel board (CHB) to which ports specified for -port <port> belong.

Chapter 5: Configuration setting commands

When you change the iSCSI virtual port mode, the iSCSI virtual port mode of the ports that share the settings with the port is also changed. The user who executes the command must have authority of the specified port and the ports that share the settings with the port. For details about the ports that share the settings with the specified port, see the *Provisioning Guide* for the storage system.

If you specify this option when the device does not support the iSCSI virtual port, an error occurs.

-port_mode {scsi | nvme} -request_id auto
Changes the operation mode of the port.

▪ scsi: SCSI mode

▪ nvme: NVMe mode

<request#> is a request ID assigned each time the command is run. For more information, see Request ID function (on page 246).

-delete_login_host_nqn
Deletes the host NQN login information.

[-mtu <value>]
Specifies the MTU value (1500/4500/9000) used during iSCSI communication. When you omit the specification, the MTU value is not changed.

> 📝 **Note:**
>
> If this option is specified, the communications through the port are interrupted.
>
> When you set the attribute value related to iSCSI, the communications through the specified port are interrupted, and an I/O error is reported to the SCSI command which is being executed. Therefore, if the attribute value related to iSCSI is set more than once continuously, troubles on I/O, for example the decreasing I/O throughput, and the failure in reading and writing, might occur. We recommend that you set the attribute value for iSCSI only when the I/O from the hosts does not exist.

[-vlan_tagging_mode {enable|disable}]
Specifies the availability of Tag VLAN. When you omit the specification, the settings for the availability of Tag VLAN is not changed.

▪ enable: enable Tag VLAN

▪ disable: disable Tag VLAN

📄 **Note:**

If this option is specified, the communications through the port are interrupted.

When you set the attribute value related to iSCSI, the communications through the specified port are interrupted, and an I/O error is reported to the SCSI command which is being executed. Therefore, if the attribute value related to iSCSI is set more than once continuously, troubles on I/O, for example the decreasing I/O throughput, and the failure in reading and writing, might occur. We recommend that you set the attribute value for iSCSI only when the I/O from the hosts does not exist.

[-add_vlan_id <value>]

Specifies the adding VLAN ID (1-4094). When you omit the specification, the VLAN ID is not added.

📄 **Note:**

If this option is specified, the communications through the port are interrupted.

When you set the attribute value related to iSCSI, the communications through the specified port are interrupted, and an I/O error is reported to the SCSI command which is being executed. Therefore, if the attribute value related to iSCSI is set more than once continuously, troubles on I/O, for example the decreasing I/O throughput, and the failure in reading and writing, might occur. We recommend that you set the attribute value for iSCSI only when the I/O from the hosts does not exist.

If you specify a virtual port, Tag VLAN is enabled.

[-delete_vlan_id <value>]

Specifies the deleting VLAN ID (1-4094). Also, Tag VLAN is disabled. When you omit the specification, the VLAN ID is not deleted.

📄 **Note:**

If this option is specified, the communications through the port are interrupted.

When you set the attribute value related to iSCSI, the communications through the specified port are interrupted, and an I/O error is reported to the SCSI command which is being executed. Therefore, if the attribute value related to iSCSI is set more than once continuously, troubles on I/O, for example the decreasing I/O throughput, and the failure in reading and writing, might occur. We recommend that you set the attribute value for iSCSI only when the I/O from the hosts does not exist.

[-modify_vlan_id <new value>]

Specifies the VLAN ID (1 to 4094). The VLAN ID is changed to the specified ID. When you omit this specification, the VLAN ID is not changed.

Chapter 5: Configuration setting commands

> 📄 **Note:**
>
> If this option is specified, the communications through the port are interrupted.
>
> When you set the attribute value related to iSCSI, the communications through the specified port are interrupted, and an I/O error is reported to the SCSI command which is being executed. Therefore, if the attribute value related to iSCSI is set more than once continuously, troubles on I/O, for example the decreasing I/O throughput, and the failure in reading and writing, might occur. We recommend that you set the attribute value for iSCSI only when the I/O from the hosts does not exist.

If you specify this option for a port or virtual port for which VLAN ID is not set, an error occurs.

[-ipv4_address <address>]

Specifies the IPv4 address. When you omit the specification, the IPv4 address is not changed.

This option is required if you specify both -add_iscsi_virtual_port and -ipv6_mode disable.

You cannot specify following IPv4 addresses:

- Network address (For example 192.168.10.0 or 0.120.10.1)

- Broadcast address (For example 255.255.255.255 or 10.1.255.255)

- Loopback address (For example 127.0.0.1)

> 📄 **Note:**
>
> If this option is specified, the communications through the port are interrupted.
>
> When you set the attribute value related to iSCSI, the communications through the specified port are interrupted, and an I/O error is reported to the SCSI command which is being executed. Therefore, if the attribute value related to iSCSI is set more than once continuously, troubles on I/O, for example the decreasing I/O throughput, and the failure in reading and writing, might occur. We recommend that you set the attribute value for iSCSI only when the I/O from the hosts does not exist.

[-ipv4_subnetmask <subnet mask>]

Specifies the IPv4 subnet mask. When you omit the specification, the IPv4 subnet mask is not changed.

This option is required if you specify both -add_iscsi_virtual_port and -ipv6_mode disable.

> 📄 **Note:**
>
> If this option is specified, the communications through the port are interrupted.
>
> When you set the attribute value related to iSCSI, the communications through the specified port are interrupted, and an I/O error is reported to the SCSI command which is being executed. Therefore, if the attribute value related to iSCSI is set more than once continuously, troubles on I/O, for example the decreasing I/O throughput, and the failure in reading and writing, might occur. We recommend that you set the attribute value for iSCSI only when the I/O from the hosts does not exist.

[-ipv4_gateway_address <address>]

Specifies the IPv4 default gateway address. When you omit the specification, the IPv4 default gateway address is not changed.

> 📄 **Note:**
>
> If this option is specified, the communications through the port are interrupted.
>
> When you set the attribute value related to iSCSI, the communications through the specified port are interrupted, and an I/O error is reported to the SCSI command which is being executed. Therefore, if the attribute value related to iSCSI is set more than once continuously, troubles on I/O, for example the decreasing I/O throughput, and the failure in reading and writing, might occur. We recommend that you set the attribute value for iSCSI only when the I/O from the hosts does not exist.

[-ipv6_mode {enable|disable}]

Specifies the availability of IPv6 mode. When you omit the specification, the settings for the availability of IPv6 mode is not changed.

- enable: enable IPv6 mode

- disable: disable IPv6 mode

> 📄 **Note:**
>
> If this option is specified, the communications through the port are interrupted.
>
> When you set the attribute value related to iSCSI, the communications through the specified port are interrupted, and an I/O error is reported to the SCSI command which is being executed. Therefore, if the attribute value related to iSCSI is set more than once continuously, troubles on I/O, for example the decreasing I/O throughput, and the failure in reading and writing, might occur. We recommend that you set the attribute value for iSCSI only when the I/O from the hosts does not exist.

This option is required if you specify `-add_iscsi_virtual_port`. The operation modes for iSCSI virtual ports are as follows:

▪ enable: IPv6 address is enabled.

▪ disable: IPv4 address is enabled. In this case, specify both -ipv4_address <address> and -ipv4_subnetmask <subnet mask>.

[-ipv6_local_address {auto|<address>}]

Specifies the IPv6 link local address. When you omit the specification, the IPv6 link local address is not changed.

When you specify "auto", the address is set automatically. When you set the address manually, enter the IPv6 address for <address>.

You cannot specify following IPv6 addresses:

▪ Not set (For example ::)

▪ Multicast address (For example ff00:1024:1215::01)

▪ Loopback address (For example ::1)

> **📄 Note:**
>
> If this option is specified, the communications through the port are interrupted.
>
> When you set the attribute value related to iSCSI, the communications through the specified port are interrupted, and an I/O error is reported to the SCSI command which is being executed. Therefore, if the attribute value related to iSCSI is set more than once continuously, troubles on I/O, for example the decreasing I/O throughput, and the failure in reading and writing, might occur. We recommend that you set the attribute value for iSCSI only when the I/O from the hosts does not exist.

[-ipv6_global_address {auto|<address>}]

Specifies the IPv6 global address. When you omit the specification, the IPv6 global address is not changed.

When you specify "auto", the address is set automatically. When you set the address manually, enter the IPv6 address for <address>.

You cannot specify following IPv6 addresses:

▪ Multicast address (For example ff00:1024:1215::01)

▪ Loopback address (For example ::1)

> 📄 **Note:**
>
> If this option is specified, the communications through the port are interrupted.
>
> When you set the attribute value related to iSCSI, the communications through the specified port are interrupted, and an I/O error is reported to the SCSI command which is being executed. Therefore, if the attribute value related to iSCSI is set more than once continuously, troubles on I/O, for example the decreasing I/O throughput, and the failure in reading and writing, might occur. We recommend that you set the attribute value for iSCSI only when the I/O from the hosts does not exist.

[-ipv6_gateway_address <address>]

Specifies the IPv6 gateway address. When you omit the specification, the IPv6 gateway address is not changed.

You cannot specify following IPv6 addresses:

- Multicast address (For example ff00:1024:1215::01)

- Loopback address (For example ::1)

> 📄 **Note:**
>
> If this option is specified, the communications through the port are interrupted.
>
> When you set the attribute value related to iSCSI, the communications through the specified port are interrupted, and an I/O error is reported to the SCSI command which is being executed. Therefore, if the attribute value related to iSCSI is set more than once continuously, troubles on I/O, for example the decreasing I/O throughput, and the failure in reading and writing, might occur. We recommend that you set the attribute value for iSCSI only when the I/O from the hosts does not exist.

[-tcp_port <value>]

Specifies the TCP port number (1 - 65535) during iSCSI communication. When you omit the specification, the TCP port number is not changed.

> 📄 **Note:**
>
> If this option is specified, the communications through the port are interrupted.
>
> When you set the attribute value related to iSCSI, the communications through the specified port are interrupted, and an I/O error is reported to the SCSI command which is being executed. Therefore, if the attribute value related to iSCSI is set more than once continuously, troubles on I/O, for example the decreasing I/O throughput, and the failure in reading and writing, might occur. We recommend that you set the attribute value for iSCSI only when the I/O from the hosts does not exist.

Chapter 5: Configuration setting commands

[-selective_ack_mode {enable|disable}]

Specifies the availability of selective ACK. When you omit the specification, the settings of the selective ACK is not changed.

- enable: enable selective ACK

- disable: disable selective ACK

> 📄 **Note:**
>
> If this option is specified, the communications through the port are interrupted.
>
> When you set the attribute value related to iSCSI, the communications through the specified port are interrupted, and an I/O error is reported to the SCSI command which is being executed. Therefore, if the attribute value related to iSCSI is set more than once continuously, troubles on I/O, for example the decreasing I/O throughput, and the failure in reading and writing, might occur. We recommend that you set the attribute value for iSCSI only when the I/O from the hosts does not exist.

[-delayed_ack_mode {enable|disable}]

Specifies the availability of delayed ACK. When you omit the specification, the settings of the delayed ACK is not changed.

- enable: enable delayed ACK

- disable: disable delayed ACK

> 📄 **Note:**
>
> If this option is specified, the communications through the port are interrupted.
>
> When you set the attribute value related to iSCSI, the communications through the specified port are interrupted, and an I/O error is reported to the SCSI command which is being executed. Therefore, if the attribute value related to iSCSI is set more than once continuously, troubles on I/O, for example the decreasing I/O throughput, and the failure in reading and writing, might occur. We recommend that you set the attribute value for iSCSI only when the I/O from the hosts does not exist.

[-window_size <size>]

Specifies the size of window. You can specify 64KB, 128KB, 256KB, 512KB, or 1024KB. The specifiable unit is m or M for mega byte, and k or K for kilo byte. When you omit the unit, Block (512 byte) is used. For example:

- When you specify 1,024KB:

  ```
  -window_size 1M, -window_size 1m, -window_size 1024K, -window_size 1024k, or -window_size 2048
  ```

- When you specify 256KB:

  ```
  -window_size 256K, -window_size 256k, -window_size 512
  ```

Chapter 5: Configuration setting commands

> 📄 **Note:**
>
> If this option is specified, the communications through the port are interrupted.
>
> When you set the attribute value related to iSCSI, the communications through the specified port are interrupted, and an I/O error is reported to the SCSI command which is being executed. Therefore, if the attribute value related to iSCSI is set more than once continuously, troubles on I/O, for example the decreasing I/O throughput, and the failure in reading and writing, might occur. We recommend that you set the attribute value for iSCSI only when the I/O from the hosts does not exist.

[-keep_alive_timer <value>]

Specifies the value of the Keep Alive Timer (30 - 64800 seconds) during iSCSI communication.

When you omit the specification, the value of the Keep Alive Timer is not changed.

[-ipv6_global_address2 <address>]

Specifies the IPv6 global address2. You can specify this option for virtual ports for which the physical port ID or iSCSI virtual port ID is 0. If a value from 1 to 15 is specified as the iSCSI virtual port ID, an error occurs. When you omit this option, the IPv6 global address2 is not changed.

If you specify the IPv6 address automatically, specify auto for -ipv6_global_address. An error occurs if you specify auto again when it is already specified.

If you want to specify IPv6 global address2 manually, you must also specify the IPv6 global address manually. When the IPv6 global address is already set manually, you do not need to change it. Type the IPv6 address for <address>.

The following IPv6 addresses cannot be set:

- Multicast Address (Example: ff00:1024:1215::01)

- Loopback Address (Example: ::1)

[-isns_mode {enable|disable}]

Specifies the availability of iSNS service.

- enable: enable iSNS service

- disable: disable iSNS service

When you omit the specification, the settings of the iSNS service is not changed.

[-isns_server_address <IPv4 address|IPv6 address>]

Specifies the IP address of the iSNS server. You can specify both IPv4 and IPv6 address.

You cannot specify following IPv4 addresses:

- Broadcast address (For example 255.255.255.255 or 10.1.255.255)

- Loopback address (For example 127.0.0.1)

Chapter 5: Configuration setting commands

You cannot specify following IPv6 addresses:

- Not set (For example ::)

- Multicast address (For example ff00:1024:1215::01)

- Loopback address (For example ::1)

When you omit the specification, the IP address of the iSNS server is not changed.

[-isns_port <iSNS TCP Port number>]

Specifies the value of the TCP port in the iSNS server (1 - 65535).

When you omit the specification, the value of the TCP port in the iSNS server is not changed.

-add_iscsi_virtual_port <value>

Adds the virtual port to the port for which iSCSI virtual port mode is enabled. Specify the iSCSI virtual port ID (1 to 15) for <value>. Either IPv4 or IPv6 must be set. When you specify IPv4, select disable for -ipv6_mode, and specify both -ipv4_address <address> and -ipv4_subnetmask <subnet mask>. For IPv6, select enable for -ipv6_mode.

An error occurs if you specify this option in any of the following cases:

- The specified iSCSI virtual port is already registered.

- The iSCSI virtual port mode is disabled.

- The device does not support the iSCSI virtual port mode.

-modify_iscsi_virtual_port <value>

Edits virtual port information of the port for which iSCSI virtual port mode is enabled.

Specify the iSCSI virtual port ID (0 to 15) for <value>.

When you change settings from IPv6 to IPv4 for iSCSI virtual ports whose IDs are from 1 to 15, select disable for -ipv6_mode, then set both -ipv4_address <address> and -ipv4_subnetmask <subnet mask>.

An error occurs if you specify this option in any of the following cases:

- The virtual port of the specified iSCSI virtual port ID is not found.

- The iSCSI virtual port mode is disabled.

- The device does not support the iSCSI virtual port mode.

-delete_iscsi_virtual_port <value>

Deletes the virtual port from the port for which iSCSI virtual port mode is enabled.

Specify the iSCSI virtual port ID (1 to 15) for <value>.

An error occurs if you specify this option in any of the following cases:

- The virtual port of the specified iSCSI virtual port ID is not found.

- The iSCSI virtual port mode is disabled.

- The device does not support the iSCSI virtual port mode.

Returned values

One of the values shown below is returned to exit() so that you can check the execution results using a user program or a script.

- **0:** Normal termination

- **EX_ENOSUP:** The host does not support IPv6. For details, see the *Command Control Interface Installation and Configuration Guide*.

- **Other than 0 and EX_ENOSUP:** Abnormal termination. For details, see the *Command Control Interface User and Reference Guide*.

Examples

Changing the port attributes (the Loop ID and the topology of the port) of a port CL3-E.

```
# raidcom modify port -port CL3-E -loop_id 0xAB -topology fl_port
```

Changing the port attributes of the port CL3-E to the External initiator port (ELUN).

```
# raidcom modify port -port CL3-E -port_attribute ELUN
```

When you enable the IPv6 of the port CL3-E, and automate the settings of the Global address.

```
# raidcom modify port -port CL3-E -ipv6_mode enable
-ipv6_global_address auto
```

When you enable the iSNS service of the port CL3-E, and set the IP address to the iSNS server:

```
# raidcom modify port -port CL3-E -isns_mode enable
-isns_server_address fe80::209:6bff:febe:3c17
```

When you enable the T10 PI mode of the port CL3-E:

```
# raidcom modify port -port CL3-E -t10pi enable
```

When you enable the virtual port mode of the port CL3-E:

```
# raidcom modify port -port CL3-E -iscsi_virtual_port_mode enable
```

When you add virtual port 1 to the port CL3-E using IPv4 to set IP address:10.213.46.63, subnet mask: 255.255.254.0:

```
#raidcom modify port -port CL3-E -add_iscsi_virtual_port 1 -ipv6_mode disable -
ipv4_address 10.213.46.63 -ipv4_subnetmask 255.255.254.0
```

When you change the operation mode of the port CL3-E to NVMe mode:

```
# raidcom modify port -port CL3-E -port_mode nvme -request_id auto
```

Chapter 5: Configuration setting commands

When you delete the login information about the host NQN that is logged in to the port CL3-E:

```
# raidcom modify port -port CL3-E -delete_login_host_nqn
```

raidcom add parity_grp

Supported storage systems:

- VSP E series

- VSP G130, G/F350, G/F370, G/F700, G/F900

- VSP G200, G400, G600, G800 and VSP F400, F600, F800

Creates parity groups.

This command is executed asynchronously with the command input. Use the **raidcom get command_status** command to check if the command is completed.

> 📄 **Note:** If a failure occurs in a processor on an MP blade or a controller, this command cannot be run even on an MP blade or a controller on which no failure occurs.
>
> If a drive, cache memory, HIE, or a logical path inside the storage is blocked, the operation cannot be performed.

Syntax

```
raidcom add parity_grp {-parity_grp_id <gno-sgno> |
-concatenated_parity_grp_id <gno-sgno>... } -drive_location <drive location>...
-raid_type <raid type> [-encryption {enable|disable}]
[-copy_back {enable|disable}] [-accelerated_compression {enable|disable}]
[-clpr <clpr#>]
```

Options and parameters

-parity_grp_id <gno-sgno>

Specifies the parity group number (gno: 1-52, sgno: 1-32).

Example:

3-1

-concatenated_parity_grp_id <gno-sgno>…

Specifies the number of the parity group to be concatenated if a distributed parity group is created.

Drives specified by **-drive_location <drive location>**… are used in the order they were specified to create parity groups to be concatenated.

Parity group numbers specified by **-concatenated_parity_grp_id <gno-sgno>**… are assigned in the order they were specified to parity groups to be created.

-drive_location <drive location>…

Specifies the drive location.

Example:

When specifying drive location HDD00-01:

0-1

-raid_type <raid type>

Specifies the RAID type. The following types can be specified:

- For RAID type 2D+2D: 2D2D
- For RAID type 3D+1P: 3D1P
- For RAID type 4D+1P: 4D1P
- For RAID type 6D+1P: 6D1P
- For RAID type 7D+1P: 7D1P
- For RAID type 6D+2P: 6D2P
- For RAID type 12D+2P: 12D2P
- For RAID type 14D+2P: 14D2P

[-encryption {enable|disable}]

Enables or disables encryption. If you omit this option, disable is set.

- enable: Encryption is enabled.
- disable: Encryption is disabled.

[-copy_back {enable|disable}]

Enables or disables the copy back mode. If you omit this option, enable is set.

- enable: Copy back is enabled.
- disable: Copy back is disabled.

For VSP E series, VSP G130, G/F350, G/F370, G/F700, G/F900, specifying `copy_back disable` is recommended.

[-accelerated_compression {enable|disable}]

Enables or disables accelerated compression of parity groups. If you omit this option, disable is set.

- enable: accelerated compression is enabled.
- disable: accelerated compression is disabled.

[-clpr <clpr#>]

Specifies the CLPR number (0 to 31). If you omit this option, 0 is set.

Example:

-clpr 2

Chapter 5: Configuration setting commands

Examples

Creating parity group 1-1 with RAID type 3D+1P using drives of which drive locations are HDD00-00, HDD00-01, HDD00-02, HDD00-03.

```
# raidcom add parity_grp -parity_grp_id 1-1 -drive_location 0-0 0-1 0-2 0-3 -
raid_type 3D1P
```

Creating parity groups 1-1, 1-2, 1-3, 1-4 with RAID type 7D+1P using drives of which drive locations are from HDD00-00 to HDD00-31, and configuring distributed parity groups. This command uses HDD00-00 to HDD00-07 to create parity group 1-1, HDD00-08 to HDD00-15 to create parity group 1-2, HDD00-16 to HDD00-023 to create parity group 1-3, and HDD00-24 to HDD00-31 to create parity group 1-4.

```
# raidcom add parity_grp -concatenated_parity_grp_id 1-1 1-2 1-3 1-4
-drive_location 0-0 0-1 0-2 0-3 0-4 0-5 0-6 0-7 0-8 0-9 0-10 0-11 0-12 0-13 0-14
0-15 0-16 0-17 0-18 0-19 0-20 0-21 0-22 0-23 0-24 0-25 0-26 0-27 0-28 0-29 0-30
0-31 -raid_type 7D1P
```

raidcom delete parity_grp

Supported storage systems:

- VSP E series
- VSP G130, G/F350, G/F370, G/F700, G/F900
- VSP G200, G400, G600, G800 and VSP F400, F600, F800

Deletes parity groups.

This command is executed asynchronously with the command input. Use the `raidcom get command_status` command to check if the command is completed.

> **Note:** If a failure occurs in a processor on an MP blade or a controller, this command cannot be run even on an MP blade or a controller on which no failure occurs.
>
> If a drive, cache memory, HIE, or a logical path inside the storage is blocked, the operation cannot be performed.

📄 **Note:** Do not delete a parity group that is being moved to a CLPR. If you delete the parity group being moved to a CLPR, the deletion of the source and destination CLPRs might fail with the error of SSB=2E11, 7004. If the error, SSB=2E11, 7004, occurs, perform the following steps.

1. Use the **raidcom get parity_grp** command to make sure there are no parity groups being moved to CLPRs.

   The M column of the command output shows information about whether or not the parity group is being moved to CLPR.

2. Wait for a while (about 20 minutes) and then execute the **raidcom delete clpr** command.

   If you get the same error again, repeat step 2.

Syntax

```
raidcom delete parity_grp -parity_grp_id <gno-sgno>
```

Options and parameters

-parity_grp_id <gno-sgno>

Specifies the parity group number (gno: 1 to 52, sgno: 1 to 32). If the specified parity group is concatenated with other parity groups, all the concatenated parity groups are deleted.

Example:

3-1

Examples

Deleting parity group 1-1.

```
# raidcom delete parity_grp -parity_grp_id 1-1
```

raidcom get parity_grp

Displays parity group information.

Syntax

```
raidcom get parity_grp [-parity_grp_id <gno-sgno> | -key <keyword> ] [-pcap]
```

Options and parameters

[-parity_grp_id <gno-sgno>]

Specifies the parity group number (gno: 1 to 52, sgno: 1 to 32).

If this option is specified, the LDEV and free space information defined in the specified parity group is displayed. For example:

- 3-1

> 📄 **Note:** If this option is omitted, the list of parity groups defined in the storage system is displayed.

[-key <keyword>]

Specifies a display keyword. Specify opt for <keyword>. When this option is specified, the following information is displayed:

- Total capacity of the parity group

- Information whether accelerated compression setting for the parity group is enabled or disabled

- Information whether copy back mode is enabled or disabled

- Information whether encryption is enabled or disabled

- When a parity group is contained in a distributed parity group, the number of the parity group that has the smallest number and the operation status of the capacity saving feature.

> 📄 **Note:** If this option is omitted, the list of parity groups defined in the storage system is displayed.

[-pcap]

Supported storage systems:

- VSP 5000 series

- VSP E series

- VSP G130, G/F350, G/F370, G/F700, G/F900

Displays the actual capacity. This option is valid only when the `-key <keyword>` option is not specified, or when the `-key opt` option is specified. For VSP 5000 series, VSP E series, VSP G130, G/F350, G/F370, G/F700, G/F900, specify this option to check the capacity information.

If the microcode version does not support this option, nothing is displayed when this option is specified.

Example 1

Displaying parity group information.

```
# raidcom get parity_grp
```

```
T GROUP   Num_LDEV  U(%) AV_CAP(GB) R_LVL R_TYPE SL CL DRIVE_TYPE
M E_TYPE
R 5-2          4    45     140000  RAID1 2D+2D  0  0 DKS2C-K072FC
Y OPEN-V
```

```
R 5-3        4   45     140000  RAID1  2D+2D  0   0 DKS2C-K072FC
N OPEN-V
```

```
# raidcom get parity_grp -key opt
T    GROUP    TOTAL_CAP(GB)    V    C    E    C_GROUP
R    1-1    300         D    E    E    1-1
R    1-2    300         D    E    E    1-1
R    1-3    300         D    E    E    1-1
R    1-4    300         D    E    E    1-1
R    2-16     500         E    D    D    -
R 3-1    300 D E E - PS
```

Description of each column in output example:

T

>   Displays the type of the volume group.
>
>   Where R is the parity group

GROUP

>   Displays the parity group number.

Num_LDEV

>   Displays the number of LDEVs assigned to this parity group.

U(%)

>   Displays the usage rate of this parity group. Displays the actual usage rate of the parity group when the -pcap option is specified.

AV_CAP(GB)

>   Displays the available capacity (free space) for this parity group. A value less than 1 GB is rounded down and 0 is displayed. When the -pcap option is specified, displays the actual capacity that can be used for the parity group (free area).

R_LVL

>   Displays the RAID level of the parity group.

R_TYPE

>   Displays the RAID type of the parity group.

SL

>   Displays the SLPR number to which the parity group belongs.

CL

>   Displays the CLPR number to which the parity group belongs.

DRIVE_TYPE

>   Drive type code which is set when the parity group is set.
>
>   To view the drive type code of the drive in the parity group, execute the **raidcom get drive** command.

M

Displays whether the allocation of parity group to CLPR is changed.

- ▪ Y: The allocation of parity group is being changed.

- ▪ N: The allocation of parity group is not changed.

- ▪ - (hyphen): a hyphen (-) is always displayed for VSP, HUS VM, VSP G200, G400, G600, G800, and VSP F400, F600, F800.

E_TYPE

Displays the base emulation type of the parity group.

TOTAL_CAP(GB)

Displays the total capacity of the parity group.

For VSP 5000 series, VSP G1x00, and VSP F1500, VSP E series, VSP G130, G/F350, G/F370, G/F700, G/F900, displays the actual usage rate of the parity group when the -pcap option is specified.

V

Displays whether accelerated compression of the parity group is enabled or disabled.

- ▪ E: Accelerated compression is enabled.

- ▪ D: Accelerated compression is disabled.

C

Displays whether the copy back mode is enabled or disabled.

- ▪ E: Copy back mode is enabled.

- ▪ D: Copy back mode is disabled.

- ▪ - (hyphen): Displaying copy back mode is not supported.

E

Displays whether encryption is enabled or disabled.

- ▪ E: Encryption is enabled.

- ▪ D: Encryption is disabled.

- ▪ - (hyphen): Displaying the encryption status is not supported.

C_GROUP

When a parity group is contained in a distributed parity group, displays the smallest parity group number in the distributed parity group. If parity groups are not contained in distributed parity groups, or the function for displaying the parity group numbers is not supported, a hyphen (-) is displayed.

Example 2

Displaying information on LDEVs and the free space defined in the parity group.

```
# raidcom get parity_grp -parity_grp_id 5-2
```

```
T GROUP P_NO LDEV#  STS        LOC_LBA          SIZE_LBA Serial# SP
R 5-2      0     -  NML  0x000000000000  0x000000003f00    64034  -
R 5-2      1   100  NML  0x000000003f00  0x000000010000    64034  R
R 5-2      2   101  REG  0x000000013f00  0x000000010000    64034  V
R 5-2      3     -  DEL  0x000000023f00  0x0000f0000000    64034  -
```

Description of each column in output example:

T

> Displays the type of the volume group.

> Where R is the parity group

GROUP

> Displays the parity group number.

P_NO

> Displays the partition number partitioning this parity group.

LDEV#

> Displays LDEV number.

STS

> Displays the following status.

> - NML: LDEV is installed or free space is settled.

> - REG: LDEV is being created.

> - DEL: LDEV is being deleted.

LOC_LBA

> Displays the Start of LBA for this partition on this parity group, in blocks (512 bytes).

SIZE_LBA

> Displays the size for this partition on this parity group, in blocks (512 bytes).

Serial#

> Product serial number.

SP

> Displays whether the LDEV uses the expanded space of the parity group.

> - V: LDEV uses the expanded space.

> - R: LDEV does not use the expanded space.

> - -: LDEVs are not mounted.

Chapter 5: Configuration setting commands

raidcom initialize parity_grp

Supported storage systems:

- VSP 5000 series

- VSP E series

- VSP G130, G/F350, G/F370, G/F700, G/F900

- VSP G200, G400, G600, G800 and VSP F400, F600, F800

This command formats all areas in the drives of the specified parity group.

This command is executed asynchronously with the command input. Check the completion of this process using the `raidcom get command_status` command.

> 📄 **Note:** If a failure occurs in a processor on an MP blade or a controller, this command cannot be run even on an MP blade or a controller on which no failure occurs.

> ⚠️ **Caution:** This is a destructive operation. Verify the correct parity group ID before issuing this command. The user is responsible for backing up the data in the parity group, if necessary, before issuing this command.

Syntax

```
raidcom initialize parity_grp -parity_grp_id <gno-sgno> -operation <type>
```

Options and parameters

-parity_grp_id <gno-sgno>

Specifies the parity group ID of the drives to be formatted (gno:1-52, sgno:1-32). Example:

- 3-1

You can see the progress of the format by OPE_RATE of the `raidcom get ldev` command.

-operation <type>

Specifies "fmt" in <type> to format all areas in the drives of the specified parity group.

Examples

Formats all areas in the drives of parity group: 1-1:

```
# raidcom initialize parity_grp -parity_grp_id 1-1 -operation fmt
```

raidcom modify parity_grp

Enables or disables accelerated compression of a parity group.

Chapter 5: Configuration setting commands

This command is executed asynchronously with the command input. Check the completion of this process using the **raidcom get command_status** command.

> 📄 **Note:** If a failure occurs in a processor on an MP blade or a controller, this command cannot be run even on an MP blade or a controller on which no failure occurs.

Syntax

```
raidcom modify parity_grp -parity_grp_id <gno-sgno>
   -accelerated_compression <accelerated compression>
```

Options and parameters

-parity_grp_id <gno-sgno>
Specifies the parity group number (gno:1-52, sgno:1-32) (for example, 3-1).

-accelerated_compression <enable | disable>
Enables or disables accelerated compression of the specified parity group.

For VSP 5000 series, an error occurs if enable is specified for the parity group that belongs to the pool whose subscription limit is other than 65535 (unlimited).

- enable: Enables accelerated compression.
- disable: Disables accelerated compression.

Examples

Enables accelerated compression of parity group: 1-1:

```
# raidcom modify parity_grp -parity_grp_id 1-1 -accelerated_compression enable
```

raidcom add rcu

Registers RCUs.

This command is executed asynchronously with the command input. Check the completion of this process on the **raidcom get command_status** command.

> 📄 **Note:** If a failure occurs in a processor on an MP blade or a controller, this command cannot be run even on an MP blade or a controller on which no failure occurs.

Syntax

```
raidcom add rcu {-rcu <serial#> <mcu#> <rcu#> <id> -ssid <ssid>
    | -cu_free <serial#> <id> <pid> -mcu_port <port#>
    -rcu_port <port#>
```

Options and parameters

-rcu <serial#> <mcu#> <rcu#> <id>

Specifies the CU specified by serial number, <mcu#>, <rcu#>, and <id>. You can use hexadecimal numbers (add "0x" prefix) or decimal numbers for <mcu#> and <rcu#>.

<serial#> specifies the storage system serial number:

- When specifying the serial number for VSP 5000 series, add a "5" at the beginning of the serial number. For example, for serial number 12345, enter 512345.

- When specifying the serial number for VSP G1x00 and VSP F1500, add a "3" at the beginning of the serial number. For example, for serial number 12345, enter 312345.

<id> specifies the storage system type as follows:

- R900: VSP 5000 series

- R800: VSP G1x00 and VSP F1500

- R700: VSP

- M800: VSP E series, VSP G/F350, G/F370, G/F700, G/F900, VSP G200, G400, G600, G800 and VSP F400, F600, F800

- M700: HUS VM

-ssid <ssid>

Specifies storage subsystem IDs. Up to 4 SSIDs can be specified.

You can use hexadecimal numbers (add "0x" prefix) or decimal numbers for <ssid>.

Note: This option is for enterprise storage systems. You need not specify the -ssid option for the Unified Storage VM, but this option is enabled.

-cu_free <serial#><id><pid>

This parameter is used to specify CU free specified by serial number, ID, and PID.

<serial#> specifies the storage system serial number:

- When specifying the serial number for VSP 5000 series, add a "5" at the beginning of the serial number. For example, for serial number 12345, enter 512345.

- When specifying the serial number for VSP G1x00 and VSP F1500, add a "3" at the beginning of the serial number. For example, for serial number 12345, enter 312345.

<id> specifies the storage system:

- R900: VSP 5000 series

- R800: VSP G1x00 and VSP F1500

- R700: VSP

- M800: VSP E series, VSP G/F350, G/F370, G/F700, G/F900, VSP G200, G400, G600, G800 and VSP F400, F600, F800

- M700: HUS VM

<pid> specifies the path group ID (1-255).

> 📄 **Note:**
>
> The RCU is displayed by default on the GUI of TrueCopy, Universal Replicator, or Universal Replicator for Mainframe when "0" is specified for <pid>.

-mcu_port <port#>

Specifies the port number of the MCU.

-rcu_port <port#>

Specifies the port number on the RCU (storage system port on the remote side).

Specify the port that the attribute is MCU Initiator port (MCU) or RCU Target port (RCU).

Examples

Register VSP of serial number: 64034 with CU free. Sets the path group ID: 0, the port on MCU: CL1-A, and the port on RCU: CL1-B.

```
# raidcom add rcu -cu_free 64034 R700 0 -mcu_port CL1-A -rcu_port CL1-B
```

raidcom delete rcu

Deletes the RCU.

This command is executed asynchronously with the command input. Check the completion of this process on the **raidcom get command_status** command.

> 📄 **Note:** If a failure occurs in a processor on an MP blade or a controller, this command cannot be run even on an MP blade or a controller on which no failure occurs.

Syntax

```
raidcom delete rcu {-rcu <serial#> <mcu#> <rcu#> -ssid <ssid> |
   -cu_free <serial#> <id> <pid>
```

Options and parameters

-rcu <serial#> <mcu#> <rcu#>

Specifies the CU specified by serial number, <mcu#>, <rcu#>. You can use hexadecimal numbers (add "0x" prefix) or decimal numbers for <mcu#> and <rcu#>.

<serial#> specifies the storage system serial number:

- When specifying the serial number for VSP 5000 series, add a "5" at the beginning of the serial number. For example, for serial number 12345, enter 512345.

- When specifying the serial number for VSP G1x00 and VSP F1500, add a "3" at the beginning of the serial number. For example, for serial number 12345, enter 312345.

-ssid <ssid>

Specifies the storage subsystem ID.

You can use hexadecimal numbers (add "0x" prefix) or decimal numbers for <ssid>.

Note: This option is for enterprise storage systems. You need not specify the -ssid option for the Unified Storage VM, but this option is enabled.

-cu_free <serial#><id><pid>

This parameter is used to specify CU free specified by serial number, ID, and PID.

<serial#> specifies the storage system serial number:

- When specifying the serial number for VSP 5000 series, add a "5" at the beginning of the serial number. For example, for serial number 12345, enter 512345.

- When specifying the serial number for VSP G1x00 and VSP F1500, add a "3" at the beginning of the serial number. For example, for serial number 12345, enter 312345.

<id> specifies the storage system:

- R900: VSP 5000 series

- R800: VSP G1x00 and VSP F1500

- R700: VSP

- M800: VSP E series, VSP G/F350, G/F370, G/F700, G/F900, VSP G200, G400, G600, G800 and VSP F400, F600, F800

- M700: HUS VM

<pid> specifies the path group ID (1-255).

> 📄 **Note:**
>
> The RCU is displayed by default on the GUI of TrueCopy, Universal Replicator, or Universal Replicator for Mainframe when "0" is specified for <pid>.

Examples

Deleting the VSP for which the product number: 64034, the RAID type: R700 and the path group ID: 1 are set.

```
# raidcom delete rcu -cu_free 64034 R700 1
```

raidcom get rcu

Displays MCU/RCU information.

Syntax

```
raidcom get rcu [-rcu <serial#> <mcu#> <rcu#> [-ssid <ssid>] |
   -cu_free <serial#> <id> <pid>]
```

Chapter 5: Configuration setting commands

Options and parameters

-rcu <serial#> <mcu#> <rcu#>

Specifies the CU specified by serial number, <mcu#>, <rcu#>. You can use hexadecimal numbers (add "0x" prefix) or decimal numbers for <mcu#> and <rcu#>.

<serial#> specifies the storage system serial number:

- When specifying the serial number for VSP 5000 series, add a "5" at the beginning of the serial number. For example, for serial number 12345, enter 512345.

- When specifying the serial number for VSP G1x00 and VSP F1500, add a "3" at the beginning of the serial number. For example, for serial number 12345, enter 312345.

[-ssid <ssid>]

Specifies the storage subsystem ID.

You can use hexadecimal numbers (add "0x" prefix) or decimal numbers for <ssid>.

Note: This option is for enterprise storage systems. You need not specify the -ssid option when you use Unified Storage VM, but this option is enabled.

-cu_free <serial#><id><pid>

This parameter is used to specify CU free specified by serial number, ID, and PID.

<serial#> specifies the storage system serial number:

- When specifying the serial number for VSP 5000 series, add a "5" at the beginning of the serial number. For example, for serial number 12345, enter 512345.

- When specifying the serial number for VSP G1x00 and VSP F1500, add a "3" at the beginning of the serial number. For example, for serial number 12345, enter 312345.

<id> specifies the storage system:

- R900: VSP 5000 series

- R800: VSP G1x00 and VSP F1500

- R700: VSP

- M800: VSP E series, VSP G/F350, G/F370, G/F700, G/F900, VSP G200, G400, G600, G800 and VSP F400, F600, F800

- M700: HUS VM

<pid> specifies the path group ID (1-255).

> 📄 **Note:**
>
> The RCU is displayed by default on the GUI of TrueCopy, Universal Replicator, or Universal Replicator for Mainframe when "0" is specified for <pid>.

Example 1

Displaying RCU information.

```
# raidcom get rcu
```

```
Serial# ID PID MCU RCU M/R T STS MP NP IM FZ RTO(s) RTT(ms)
  64034 R7 - 1C 23 RCU F NML 4 8 MR D 15 20
  64034 R7 - 1C 23 RCU F NML 4 8 RO E 15 20
  64034 R7 1  -  - MCU E UNK 4 8  - - 15 20
```

Description of each column in output example:

Serial#

    Displays the product serial number.

ID

    Displays the ID for identifying RAID type:

- R9: VSP 5000 series

- R8: VSP G1x00 and VSP F1500

- R7: VSP

- M8: VSP E series, VSP G/F350, G/F370, G/F700, G/F900, VSP G200, G400, G600, G800, VSP F400, F600, F800

- M7: HUS VM

PID

    Displays a path group ID. If CU is specified for RCU, "-" is displayed.

MCU

    Displays the CU number (hexadecimal) as MCU. If CU free is specified for RCU, "-" is displayed.

RCU

    Displays the CU number (hexadecimal) as RCU. If CU free is specified for RCU, "-" is displayed.

M/R

    Displays the CU type as MCU/RCU.

    MCU is displayed in the storage system that has TrueCopy or TrueCopy for Mainframe S-VOL after the initial copy was performed.

T

    Displays the type of physical path:

- F: Fibre

- E: ESCON

- I: iSCSI

- M: path types are mixed

- - (hyphen): cannot identify the path type

For VSP 5000 series, VSP E series, VSP G130, G/F350, G/F370, G/F700, G/F900, VSP G200, G400, G600, G800 and VSP F400, F600, F800, a hyphen (-) is always displayed if M/R is MCU. For other storage systems, a hyphen (-) is displayed if M/R is MCU and MCU might be connected via a path other than the Fibre Channel. To determine if MCU is connected via a path other than Fibre Channel, check the existence of an RCU target port other than FIBRE in the storage system.

STS

Displays the following status of the CU.

- NML: Normal

- WAR: Warning

- ERR: Failing

- UNK: Unknown, displayed when the target of the pair is MCU.

MP

Displays the number of path as minimum.

NP

Displays the number of path setting between MCU and RCU.

IM

Displays the incident mode setting to RCU.

- MR: it sends incident to MCU host and RCU host

- RO: it sends incident only to RCU host

- If CU free is specified for RCU, "-" is displayed.

FZ

Displays the freeze option.

- D: the freeze option is disabled.

- E: the freeze option is enabled.

- -: "-" is displayed when CU free is specified for RCU.

RTO (s)

Displays the timeout value for RIO (Remote IO) setting between MCU and RCU.

RTT (ms)

Displays the round trip time value between MCU and RCU.

Example 2

```
# raidcom get rcu -cu_free 64034 R700 1
```

```
Serial# ID PID MCU RCU M/R T PNO MPORT RPORT STS_CD SSIDs …
64034   R7  1   -   - RCU F   1 CL1-A CL1-B NML_01    -
64034   R7  1   -   - RCU F   2 CL1-A CL1-B NML_01    -
64034   R7  1   -   - RCU F   3 CL1-A CL1-B NML_01    -
```

Description of each column in output example:

Serial#

Product serial number.

ID

Displays the ID for identifying RAID type:

- R9: VSP 5000 series

- R8: VSP G1x00 and VSP F1500

- R7: VSP

- M8: VSP E series, VSP G/F350, G/F370, G/F700, G/F900, VSP G200, G400, G600, G800, VSP F400, F600, F800

- M7: HUS VM

PID

Displays the path group ID. If CU is specified for RCU, "-" is displayed.

MCU

Displays the CU number (hexadecimal) as MCU. If CU free is specified for RCU, "-" is displayed.

RCU

Displays the CU number (hexadecimal) as RCU. If CU free is specified for RCU, "-" is displayed.

M/R

Displays the CU type as MCU/RCU.

T

Displays the type of physical path:

- F: Fibre

- E: ESCON

- I: iSCSI

- M: path types are mixed

- - (hyphen): cannot identify the path type

Chapter 5: Configuration setting commands

For VSP 5000 series, VSP E series, VSP G130, G/F350, G/F370, G/F700, G/F900, VSP G200, G400, G600, G800 and VSP F400, F600, F800, a hyphen (-) is always displayed if M/R is MCU. For other storage systems, a hyphen (-) is displayed if M/R is MCU and MCU might be connected via a path other than the Fibre Channel. To determine if MCU is connected via a path other than Fibre Channel, check the existence of an RCU target port other than FIBRE in the storage system.

PNO

Displays the path number.

MPORT

Displays the MCU port number.

RPORT

Displays the RCU port number.

STS_CD

Displays the following path status:

- NML_01: Normal

- ERR_02: Initialization failed

- ERR_03: Communication timeout

- ERR_04: Logical blockade

- ERR_05: Resource Shortage

- ERR_06: Serial Number Mismatch

- ERR_10: Invalid Port

- ERR_80: RCU Port Number Mismatch

- ERR_81: RCU Port Type Mismatch

- ERR_82: Communication Failed.

- If path creation or path deletion is in progress, "-" is displayed.

SSIDs

Displays the SSIDs (hexadecimal) setting to RCU. If CU free is specified for RCU, "-" is displayed.

raidcom modify rcu

This sets the control parameters to specified CU that is specified using two way.

📄 **Note:** If a failure occurs in a processor on an MP blade or a controller, this command cannot be run even on an MP blade or a controller on which no failure occurs.

Syntax

```
raidcom modify rcu {-rcu <serial#> <mcu#> <rcu#> -ssid <ssid> |
    -cu_free <serial#> <id> <pid>} -rcu_option <mpth> <rto>
    <rtt> [fzd | fze]
```

Options and parameters

-rcu <serial#> <mcu#> <rcu#>

Specifies the CU specified by serial number, <mcu#>, and <rcu#>. You can use hexadecimal numbers (add 0x prefix) or decimal numbers for the <mcu#> and <rcu#>.

<serial#> specifies the storage system serial number:

- When specifying the serial number for VSP 5000 series, add a "5" at the beginning of the serial number. For example, for serial number 12345, enter 512345.

- When specifying the serial number for VSP G1x00 and VSP F1500, add a "3" at the beginning of the serial number. For example, for serial number 12345, enter 312345.

-ssid <ssid>

Specifies the storage subsystem ID.

You can use hexadecimal numbers (add 0x prefix) or decimal numbers for the <ssid>.

Note: This option is for enterprise storage systems. You need not specify the -ssid option for the Unified Storage VM, but this option is enabled.

-cu_free <serial#><id><pid>

This parameter is used to specify CU free specified by serial number, ID, and PID.

<serial#> specifies the storage system serial number:

- When specifying the serial number for VSP 5000 series, add a "5" at the beginning of the serial number. For example, for serial number 12345, enter 512345.

- When specifying the serial number for VSP G1x00 and VSP F1500, add a "3" at the beginning of the serial number. For example, for serial number 12345, enter 312345.

<id> specifies the storage system:

- R900: VSP 5000 series

- R800: VSP G1x00 and VSP F1500

- R700: VSP

- M800: VSP E series, VSP G/F350, G/F370, G/F700, G/F900, VSP G200, G400, G600, G800 and VSP F400, F600, F800

- M700: HUS VM

<pid> specifies the path group ID (1-255).

📄 **Note:**

The RCU is displayed by default on the GUI of TrueCopy, Universal Replicator, or Universal Replicator for Mainframe when "0" is specified for <pid>.

-rcu_option <mpth> <rto> <rtt> [fzd | fze]

Specifies CU control parameters:

- <mpth>: minimum number of paths (1-8)

- <rto>: RIO timeout value (10-100) (second) for RIO (Remote IO) setting between MCU and RCU.

- <rtt> is used to set the round trip time value (1-500) (millisecond) between MCU and RCU.

- [fzd | fze]: Specify `fze` to enable the freeze option, or `fzd` to disable it.

  If the freeze option is not specified to the RCU with CU units, the freeze option is disabled.

Examples

For the RCU for which the product number: 64034, the RAID type: R700 and the path ID: 1 are set, setting the options: the minimum number of paths <mpth>4, RIO MIH time<rto>15 seconds, and round trip time 20 milliseconds are set.

```
# raidcom modify rcu -cu_free 64034 R700 1 -rcu_option 4 15 20
```

raidcom add rcu_iscsi_port

This command registers the RCU-side iSCSI port of the remote storage system to the MCU-side iSCSI port of the local storage system. If the iSCSI port of the specified local storage system does not exist, the command is rejected with EX_ENOOBJ.

📄 **Note:** If a failure occurs in a processor on an MP blade or a controller, this command cannot be run even on an MP blade or a controller on which no failure occurs.

📄 **Note:**

- To create an iSCSI path between the local storage system and a remote storage system, you must register the iSCSI target port on the MCU-side iSCSI port of the local storage system first.

- If you register an RCU port connected via iSCSI to the MCU-side iSCSI port, you cannot perform TCz or URz remote copy using iSCSI path connection.

Syntax

```
raidcom add rcu_iscsi_port -port <port#> -rcu_port <port#>
  -rcu_id <serial#> <id> -rcu_address <IP address>
  [-tcp_port <value>]
```

Options and parameters

-port <port#>

Specifies the port number of the local storage system. For example:

- CL1-A

-rcu_port <port#>

Specifies the iSCSI port number of the remote storage system.

-rcu_id <serial#> <id>

<serial#> specifies the storage system serial number:

- When specifying the serial number for VSP 5000 series, add a "5" at the beginning of the serial number. For example, for serial number 12345, enter 512345.

- When specifying the serial number for VSP G1x00 and VSP F1500, add a "3" at the beginning of the serial number. For example, for serial number 12345, enter 312345.

<id> specifies the storage system model:

- R900: VSP 5000 series

- R800: VSP G1x00 and VSP F1500

- M800: VSP E series, VSP G/F350, G/F370, G/F700, G/F900, VSP G200, G400, G600, G800, and VSP F400, F600, F800

-rcu_address <IP address>

Specifies the IP address of the iSCSI target on the remote storage system. An IPv4 address or an IPv6 address can be specified.

Note that the following addresses cannot be specified when you specify the IPv4 address:

- Network address (for example: 192.168.10.0, 0.120.10.1)

- Broadcast address (for example: 255.255.255.255, 10.1.255.255)

- Loop back address (for example: 127.0.0.1)

Note that the following addresses cannot be specified when you specify the IPv6 address:

- Unspecified address (for example: ::)

- Multicast address (for example: ff:1024:1215::01)

- Loop back address (for example: ::1)

[-tcp_port <value>]
> Specifies the TCP port number of the iSCSI target on the remote storage system. If this option is omitted, the TCP port number of the iSCSI target port which is specified with -port option is configured.

Examples

Registers the iSCSI port: CL1-A (IP address: 158.214.135.100) of the remote storage system (product serial number: 400031, model: VSP G200, G400, G600, G800 and VSP F400, F600, F800) in the iSCSI port: CL4-E of the local storage system:

```
# raidcom add rcu_iscsi_port -port CL4-E -rcu_port CL1-A -rcu_id 400031 M800 -
rcu_address 158.214.135.100
```

Registers iSCSI ports of the remote strage system which is registered in the iSCSI port: CL1-E of the local storage system to the iSCSI port: CL4-E of the local storage system:

```
# raidcom get rcu_iscsi_port | rmawk @1-eq:CL1-E exe="raidcom add rcu_iscsi_port -
port CL4-E -rcu_port @4 -rcu_id @2 @3 -rcu_address @5"
```

raidcom delete rcu_iscsi_port

This command deletes the iSCSI port of the remote storage system that is registered as the RCU side port from MCU-side iSCSI port of the local storage system.

If the specified iSCSI port does not exist in the local storage system, the command is rejected with EX_ENOOBJ. If the specified port of the remote storage system does not exist, the command is ignored.

> **Note:** If a failure occurs in a processor on an MP blade or a controller, this command cannot be run even on an MP blade or a controller on which no failure occurs.

Syntax

```
raidcom delete rcu_iscsi_port -port <port#> -rcu_port <port#>
   -rcu_id <serial#> <id>
```

Options and parameters

-port <port#>
> Specifies the port number of the local storage system. For example:
> - CL1-A

-rcu_port <port#>
> Specifies the iSCSI port number of the remote storage system.

-rcu_id <serial#> <id>

<serial#> specifies the storage system serial number:

- When specifying the serial number for VSP 5000 series, add a "5" at the beginning of the serial number. For example, for serial number 12345, enter 512345.

- When specifying the serial number for VSP G1x00 and VSP F1500, add a "3" at the beginning of the serial number. For example, for serial number 12345, enter 312345.

<id> specifies the storage system model:

- R900: VSP 5000 series

- R800: VSP G1x00 and VSP F1500

- M800: VSP E series, VSP G/F350, G/F370, G/F700, G/F900, VSP G200, G400, G600, G800, and VSP F400, F600, F800

Examples

Deletes the iSCSI port: CL1-A (IP address: 158.214.135.100) of the remote storage system (product serial number: 400031, model: VSP G200, G400, G600, G800 and VSP F400, F600, F800) from the iSCSI port: CL4-E of the local storage system:

```
# raidcom delete rcu_iscsi_port -port CL4-E -rcu_port CL1-A -rcu_id 400031 M800
```

Deletes the iSCSI port of the remote storage system which is registered in the iSCSI port: CL1-E of the local storage system from the iSCSI port: CL4-E of the local storage system.

```
# raidcom get rcu_iscsi_port | rmawk @1-eq:CL1-E exe="raidcom delete rcu_iscsi_port -
port CL4-E -rcu_port @4 -rcu_id @2 @3"
```

raidcom get rcu_iscsi_port

This command displays the RCU-side iSCSI port of the remote storage system which is registered in the MCU-side iSCSI port of the local storage system.

If the specified iSCSI port does not exist in the local storage system, the command is rejected with EX_ENOOBJ.

Only the remote storage port registered in the port to which the user who executes the command can refer are output. For details about the port to which the user can refer, see the descriptions of the relationship of the resource group and the command operation in the *Command Control Interface User and Reference Guide*.

Syntax

```
raidcom get rcu_iscsi_port
```

Options and parameters

None.

Examples

Displays the port of the remote storage system which are registered in the iSCSI port of the local storage system:

```
#raidcom get rcu_iscsi_port
```

```
PORT    Serial# ID RPORT   IP_ADDR          IP_PORT#
CL4-E   400031 M8 CL1-A   158.214.135.100      3260
CL2-E   400031 M8 CL1-A   158.214.135.100      3260
CL1-E   400031 M8 CL1-A   158.214.135.100      3260
```

Description of each column in output example:

PORT
> Displays the port number.

Serial#
> Displays the product serial number of the remote storage system.

ID
> Displays the model of the remote storage system.
> - R8: VSP G1x00 and VSP F1500
> - R9: VSP 5000 series
> - M8: VSP E series, VSP G/F350, G/F370, G/F700, G/F900, VSP G200, G400, G600, G800 and VSP F400, F600, F800

RPORT
> Displays the port number of the remote storage system.

IP_ADDR
> Displays the IP address of the remote storage system.

IP_PORT#
> Displays the TCP port number of the port of the remote storage system.

raidcom add rcu_path

Adds logical paths to RCUs.

This command is executed asynchronously with the command input. Check the completion of this process on the raidcom get command_status command.

> **Note:** If a failure occurs in a processor on an MP blade or a controller, this command cannot be run even on an MP blade or a controller on which no failure occurs.

Syntax

```
raidcom add rcu_path {-rcu <serial#> <mcu#> <rcu#> -ssid <ssid>
    | -cu_free <serial#> <id> <pid>} -mcu_port <port#>
    -rcu_port <port#>
```

Options and parameters

-rcu <serial#> <mcu#> <rcu#>

Specifies the CU specified by serial number, <mcu#>, and <rcu#>. You can use hexadecimal numbers (add 0x prefix) or decimal numbers for the <mcu#> and <rcu#>.

<serial#> specifies the storage system serial number:

- When specifying the serial number for VSP 5000 series, add a "5" at the beginning of the serial number. For example, for serial number 12345, enter 512345.

- When specifying the serial number for VSP G1x00 and VSP F1500, add a "3" at the beginning of the serial number. For example, for serial number 12345, enter 312345.

-ssid <ssid>

Specifies the storage subsystem ID.

You can use hexadecimal numbers (add 0x prefix) or decimal numbers for the <ssid>.

Note: This option is for enterprise storage systems. You need not specify the -ssid option for the Unified Storage VM, but this option is enabled.

-cu_free <serial#><id><pid>

This parameter is used to specify CU free specified by serial number, ID, and PID.

<serial#> specifies the storage system serial number:

- When specifying the serial number for VSP 5000 series, add a "5" at the beginning of the serial number. For example, for serial number 12345, enter 512345.

- When specifying the serial number for VSP G1x00 and VSP F1500, add a "3" at the beginning of the serial number. For example, for serial number 12345, enter 312345.

<id> specifies the storage system:

- R900: VSP 5000 series

- R800: VSP G1x00 and VSP F1500

- R700: VSP

- M800: VSP E series, VSP G/F350, G/F370, G/F700, G/F900, VSP G200, G400, G600, G800 and VSP F400, F600, F800

- M700: HUS VM

<pid> specifies the path group ID (1-255).

📄 **Note:**

The RCU is displayed by default on the GUI of TrueCopy, Universal Replicator, or Universal Replicator for Mainframe when "0" is specified for <pid>.

-mcu_port <port#>

Specifies the port number on the MCU.

-rcu_port <port#>

Specifies the port number on the RCU (storage system port on the remote side).

Specify the port that the attribute is MCU initiator port (MCU) or RCU target port (RCU).

Examples

To add a path to the RCU for which the serial number is 64034, the storage ID is R700 and the path ID is 1 (from MCU port: CL1-A to RCU port: CL1-B).

```
# raidcom add rcu_path -cu_free 64034 R700 1 -mcu_port CL1-A
-rcu_port CL1-B
```

raidcom delete rcu_path

Deletes logical paths from a specified RCU.

This command is executed asynchronously with the command input. Check the completion of this process on the raidcom get command_status command.

📄 **Note:** If a failure occurs in a processor on an MP blade or a controller, this command cannot be run even on an MP blade or a controller on which no failure occurs.

Syntax

```
raidcom delete rcu_path {-rcu <serial#> <mcu#> <rcu#>
     -ssid <ssid>  | -cu_free <serial#> <id> <pid>} -mcu_port
     <port#> -rcu_port <port#>
```

Options and parameters

-rcu <serial#> <mcu#> <rcu#>

Specifies the CU specified by serial number, <mcu#>, and <rcu#>. You can use hexadecimal numbers (add 0x prefix) or decimal numbers for the <mcu#> and <rcu#>.

<serial#> specifies the storage system serial number:

- When specifying the serial number for VSP 5000 series, add a "5" at the beginning of the serial number. For example, for serial number 12345, enter 512345.

- When specifying the serial number for VSP G1x00 and VSP F1500, add a "3" at the beginning of the serial number. For example, for serial number 12345, enter 312345.

-ssid <ssid>

Specifies the storage subsystem ID.

You can use hexadecimal numbers (add 0x prefix) or decimal numbers for the <ssid>.

Note: This option is for enterprise storage systems. You need not specify the -ssid option for the Unified Storage VM, but this option is enabled.

-cu_free <serial#><id><pid>

This parameter is used to specify CU free specified by serial number, ID, and PID.

<serial#> specifies the storage system serial number:

- When specifying the serial number for VSP 5000 series, add a "5" at the beginning of the serial number. For example, for serial number 12345, enter 512345.

- When specifying the serial number for VSP G1x00 and VSP F1500, add a "3" at the beginning of the serial number. For example, for serial number 12345, enter 312345.

<id> specifies the storage system:

- R900: VSP 5000 series

- R800: VSP G1x00 and VSP F1500

- R700: VSP

- M800: VSP E series, VSP G/F350, G/F370, G/F700, G/F900, VSP G200, G400, G600, G800 and VSP F400, F600, F800

- M700: HUS VM

<pid> specifies the path group ID (1-255).

> 📑 **Note:**
>
> The RCU is displayed by default on the GUI of TrueCopy, Universal Replicator, or Universal Replicator for Mainframe when "0" is specified for <pid>.

-mcu_port <port#>

Specifies the port number on the MCU.

-rcu_port <port#>

Specifies the port number on the RCU (storage system port on the remote side).

Specifies the port that the attribute is MCU initiator port (MCU) or RCU target port (RCU).

Examples

From the RCU for which the product number: 64034, the RAID type: R700 and the path ID: 1 are set, deleting RCU path (MCU port: CL1-A and RCU port: CL1-B).

```
# raidcom delete rcu_path -cu_free 64034 R700 1 -mcu_port CL1-A
-rcu_port CL1-B
```

raidcom add ssid

Adds the specified SSID to the RCU that is specified by serial number, <mcu#>, and <rcu#>.

📄 **Note:** If a failure occurs in a processor on an MP blade or a controller, this command cannot be run even on an MP blade or a controller on which no failure occurs.

Syntax

```
raidcom add ssid -rcu <serial#> <mcu#> <rcu#> <id> -ssid <ssid>
```

Options and parameters

-rcu <serial#> <mcu#> <rcu#> <id>

Specifies the CU specified by serial number, <mcu#>, <rcu#>, and <id>. You can use hexadecimal (add the prefix 0x) or decimal for the <mcu#> and <rcu#> numbers.

<serial#> specifies the storage system serial number:

- When specifying the serial number for VSP 5000 series, add a "5" at the beginning of the serial number. For example, for serial number 12345, enter 512345.

- When specifying the serial number for VSP G1x00 and VSP F1500, add a "3" at the beginning of the serial number. For example, for serial number 12345, enter 312345.

Use the following strings for <id> to specify the RAID storage system type:

- R900: VSP 5000 series

- R800: VSP G1x00 and VSP F1500

- M800: VSP E series, VSP G/F350, G/F370, G/F700, G/F900, VSP G200, G400, G600, G800, and VSP F400, F600, F800

- R700: VSP

-ssid <ssid>

Specifies the storage subsystem ID (SSID) to add to the RCU.

You can use hexadecimal (add the 0x prefix) or decimal for the <ssid> number.

📄 **Note:**

> This option is for enterprise storage systems. You need not specify the -ssid option for the Unified Storage VM, but this option is enabled.

Examples

Adding SSID:345 to the RCU with serial number: 64034, RAID type: R700, MCU#:0, and RCU#:1 are set.

```
raidcom add ssid -rcu 64034 0 1 R700 -ssid 345
```

raidcom delete ssid

Deletes the specified SSID from the RCU that is specified by a serial number, <mcu#>, and <rcu#>.

📄 **Note:** If a failure occurs in a processor on an MP blade or a controller, this command cannot be run even on an MP blade or a controller on which no failure occurs.

Syntax

```
raidcom delete ssid -rcu <serial#> <mcu#> <rcu#> -ssid <ssid>
```

Options and parameters

-rcu <serial#> <mcu#> <rcu#>

This parameter is used to specify a CU that is specified by a serial number, <mcu#>, and <rcu#>. This option specifies <mcu#> and <rcu#> with hexadecimal numbers (adding 0x) or decimal numbers.

<serial#> specifies the storage system serial number:

- When specifying the serial number for VSP 5000 series, add a "5" at the beginning of the serial number. For example, for serial number 12345, enter 512345.

- When specifying the serial number for VSP G1x00 and VSP F1500, add a "3" at the beginning of the serial number. For example, for serial number 12345, enter 312345.

-ssid <ssid>

Specifies the storage subsystem ID to be deleted from the RCU.

You can specify the ssid as a hexadecimal number (add the 0x prefix) or a decimal number.

Note: This option is for enterprise storage systems. You need not specify the -ssid option for the Unified Storage VM, but this option is enabled.

Examples

Deleting SSID:345 from the RCU where the serial number: 64034, MCU#:0, and RCU#:1 are set.

```
raidcom delete ssid -rcu 64034 0 1 -ssid 345
```

raidcom get ssid

Supported storage systems:

- VSP G1x00 or VSP F1500

- VSP 5000 series

Displays the SSID allocation information for the boundary area of the CU number and LDEV number.

Syntax

```
raidcom get ssid
```

Examples

Displays the SSID allocation information for the boundary area of the CU number and LDEV number.

```
# radicom get ssid
CU LDEV_B SSID
00 00-FF 0004
01 00-FF 0005
02 00-FF     -
03 00-FF 0012
04 00-FF     -
...
FE 00-FF FEFF
```

Description of each column in output example:

CU

  Displays the CU number in hexadecimal notation.

LDEV_B

  Displays the boundary area of the LDEV number.

SSID

  Displays the SSID in hexadecimal notation.

  A hyphen (-) is displayed if the SSID is not specified

raidcom add resource

Creating resource groups. If you specify only the resource group name, an empty resource group is created. If you specify the resource group name and the information on the virtual storage machine, an empty resource group that corresponds to the virtual storage machine is created.

When you input a resource group name and a resource group ID, the current name of the resource group whose ID you specify is changed to the new resource group name.

When you specify resource group name, LDEV number, port number, host group number, parity group ID or an external group ID, the specified resource is registered to the specified resource group. If the specified resource group does not exist, an error occurs. When the resource group is already created, the specified resource is added to the resource group. You can specify a device group name instead of an LDEV number.

When the relevant LDEVs configure the pool, journal, and LUSE, all LDEVs must be added to the same resource group.

> 📄 **Note:** If a failure occurs in a processor on an MP blade or a controller, this command cannot be run even on an MP blade or a controller on which no failure occurs.

Syntax

When creating a blank resource group

```
raidcom add resource -resource_name <resource group name>
```

When creating a virtual storage machine and a blank resource group which correspond to the virtual storage machine

```
raidcom add resource -resource_name <resource group name> -virtual_type <serial#> <id>
```

When changing the resource group name of the specified resource group ID

```
raidcom add resource -resource_name <resource group name> -resource_id <resource group_id>
```

When registering a resource (LDEV, port, host group, parity group, external volume group, or NVM subsystem) to the resource group (An error occurs if the specified resource group is not available)

```
raidcom add resource -resource_name <resource group name> [-ldev_id <ldev#> | -port <port#> [<host group name>] | -parity_grp_id <gno-sgno>| -external_grp_id <gno-sgno>| -nvm_subsystem_id <nvm subsystem id>]
```

When registering an LDEV which belongs to the device group, to the resource group (An error occurs if the specified resource group is not available)

```
raidcom add resource -resource_name <resource group name> -grp_opt ldev -
device_grp_name <device group name> [<device name>]
```

> **Note:** If the corresponding LDEV is a volume that includes pool, journal, or LUSE, all LDEVs need to be assigned to the same resource group.

Options and parameters

-resource_name <resource group name>
Specifies the resource group name. Up to 32 characters can be specified.

[-virtual_type <serial#> <id>]
Specifies the serial number and the system name (type identifier) of the virtual storage machine.

- serial#: Serial number of the virtual storage machine.

  - When specifying the serial number for VSP 5000 series, add a "5" at the beginning of the serial number. For example, for serial number 12345, enter 512345.

  - When specifying the serial number for VSP G1x00 and VSP F1500, add a "3" at the beginning of the serial number. For example, for serial number 12345, enter 312345.

- id: Type identifier of the virtual storage machine.

  - RH10HF: VSP 5200, VSP 5600

  - RH10HG: VSP 5200H, VSP 5600H

  - R900F: VSP 5500, VSP 5100

  - R900G: VSP 5100H, VSP 5500H

  - R800: VSP G1x00, VSP F1500

  - R700: VSP

  - M900S: VSP E590

  - M900SG: VSP E590H

  - M900M: VSP E790

  - M900MG: VSP E790H

  - M900H: VSP E990

  - RH10MHF: VSP E1090

  - RH10MHG: VSP E1090H

  - M850S1: VSP G350

  - M850S1F: VSP F350

  - M850S2: VSP G370

  - M850S2F: VSP F370

  - M850M3: VSP G700

  - M850M3F: VSP F700

  - M850H: VSP G900

  - M850HF: VSP F900

  - M800S: VSP G200

  - M800M: VSP G400, VSP G600, VSP F400, VSP F600

  - M800H: VSP G800, VSP F800

  - M700: HUS VM

Chapter 5: Configuration setting commands

Storage System	R500 RK500 R600 RK600 R700 M700	R800	M800 M800 M M800 H	M850 XS M850 S1 M850 S1F M850 S2 M850 S2F M850 M3 M850 M3F M850 H M850 HF	RH10 HG RH10 MHG RH10 HF RH10 MHF	R900G R900F	M900H M900M M900S M900MG M900SG
Virtual storage machine that can be created in the storage system: Identifier of the virtual storage machine							
VSP G1000, VSP G1500, VSP F1500	Y	Y	Y	N	N	N	N
VSP G100, VSP G200, G400, G600, G800 and VSP F400, F600, F800	Y	Y	Y	N	N	N	N
VSP G130, G/F350, G/F370,	Y	Y	Y	Y	N	N	N

Storage System	Virtual storage machine that can be created in the storage system: Identifier of the virtual storage machine						
	R500 RK500 R600 RK600 R700 M700	R800	M800 M800 M M800 H	M850 XS M850 S1 M850 S1F M850 S2 M850 S2F M850 M3 M850 M3F M850 H M850 HF	RH10 HG RH10 MHG RH10 HF RH10 MHF	R900G R900F	M900H M900M M900S M900MG M900SG
G/F700, G/F900							
VSP 5100, VSP 5500, VSP 5100H, and VSP 5500H	Y	Y	Y	Y	Y	Y	N
VSP 5200, VSP 5600, VSP 5200H, and VSP 5600H	Y	Y	Y	Y	Y	Y	Y
VSP E590, VSP E790,	Y	Y	Y	Y	Y	Y	Y

Chapter 5: Configuration setting commands

Storage System	Virtual storage machine that can be created in the storage system: Identifier of the virtual storage machine						
	R500 RK500 R600 RK600 R700 M700	R800	M800 M800M M800H	M850XS M850S1 M850S1F M850S2 M850S2F M850M3 M850M3F M850H M850HF	RH10HG RH10MHG RH10HF RH10MHF	R900G R900F	M900H M900M M900S M900MG M900SG
VSP E990, VSP E1090, VSP E590H, and VSP E790H, VSP E1090H							
Legend Y: Supported, N: Not supported							

-resource_id <resource group_id>
> Specifies the resource group ID (1-1023) (for example, 5).

-ldev_id <ldev#>
> Specifies the LDEV number (0-65279) (for example, -ldev_id 200).

-port <port#> [<host group name>]
> Specifies the port number, host group ID, and host group name. If the host group name is more than 64 characters, use the host group ID instead of the host group name (for example, CL1-A or CL1-A-g, where g is from 0 to 255).

Chapter 5: Configuration setting commands

-parity_grp_id <gno-sgno>
> Specifies the parity group number (gno:1-52, sgno:1-32) (for example, 3-1).

-external_grp_id <gno-sgno>
> Specifies the external volume group number (gno:1-16384, sgno:1-4096) (for example, 52-11, "E" is not required).

-nvm_subsystem_id <nvm subsystem id>
> Specifies the NVM subsystem ID.

-grp_opt ldev -device_grp_name <device group name> [<device name>]
> Specifies the name of the device group (maximum 32 characters).
>
> To specify an LDEV in the device group, use the device name of the LDEV (maximum 32 characters).
>
> If the device name is omitted, this command is applied to all LDEVs in the specified device group.

Examples

Creating a resource group of resource group name: sql_srv.

```
# raidcom add resource -resource_name sql_srv
```

Creating a virtual storage machine: rsg_vir, and the serial number of the virtual storage machine: 1000.

```
# raidcom add resource -resource_name rsg_vir -virtual_type 1000 R700
```

Changing the resource group name of the resource group ID:5 to sql_srv.

```
# raidcom add resource -resource_name sql_srv -resource_id 5
```

Add LDEV: 400 to the resource group: sql_srv.

```
# raidcom add resource -resource_name sql_srv -ldev_id 400
```

Adding a port of CL1-A to a resource group: sql_srv.

```
# raidcom add resource -resource_name sql_srv -port CL1-A
```

Adding a host group of CL1-A-0 to a resource group: sql_srv.

```
# raidcom add resource -resource_name sql_srv -port CL1-A-0
```

Adding a parity group:5-2 to a resource group: sql_srv.

```
# raidcom add resource -resource_name sql_srv -parity_grp_id 5-2
```

Chapter 5: Configuration setting commands

Adding an external volume group:01-02 to a resource group: sql_srv.

```
# raidcom add resource -resource_name sql_srv -external_grp_id 01-02
```

Adding an NVM subsystem ID: 1 to a resource group: sql_srv.

```
# raidcom add resource -resource_name sql_srv -nvm_subsystem_id 1
```

raidcom modify resource

Supported storage systems:

- VSP

- HUS VM

Sets the virtual storage mode of a resource group. This is an asynchronous command.

📄 **Note:** If a failure occurs in a processor on an MP blade or a controller, this command cannot be run even on an MP blade or a controller on which no failure occurs.

Syntax

```
raidcom modify resource -resource_name <resource group name>
     -virtual_switch <y/n>
```

Options and parameters

-resource_name <resource group name>
Specifies the resource group name. Up to 32 characters can be specified.

-virtual_switch <y/n>
Sets the virtual storage mode of the resource group to ON/OFF.

- y: Enable the virtual storage mode

- n: Disable the virtual storage mode

Examples

Resource group: Enable the virtual storage mode of sql_srv.

```
# raidcom modify resource -resource_name sql_srv -virtual_switch y
```

raidcom delete resource

Deletes resource groups. You can delete a resource group only after all resources that are registered to that resource group have been deleted.

LDEV number, port number, host group number, parity group, and external group are deleted from the specified resource groups. The deleted resources are moved to resource group 0. Even if the specified resources do not exist in the specified resource group, the specified resources are moved to resource group 0, and then the command is complete. You can specify a device group name instead of a LDEV number.

If an LDEV is a volume that configures a pool, journal, or LUSE, all LDEVs of the pool, journal, or LUSE must be assigned to the same resource group.

> 📄 **Note:** If a failure occurs in a processor on an MP blade or a controller, this command cannot be run even on an MP blade or a controller on which no failure occurs.

> 📄 **Note:** Use Device Manager - Storage Navigator to delete resource groups that are assigned to user groups where All Resource Groups Assigned is set to No. Do not use the `raidcom delete resource` command by CCI.
>
> Procedure for deleting resource groups:
>
> 1. Release the assigned information to the user group where the All Resource Groups Assigned is set to No in the **Edit Resource Groups Assignment** window.
>
> 2. Delete the resource group.
>
>   If the resource group is deleted before the assigned information is released, the deleted resource group ID is displayed in the **Edit Resource Groups Assignment** window. The deleted resource group name is not displayed.

Syntax

```
raidcom delete resource -resource_name <resource group name>
    [-ldev_id <ldev#> | -port <port#> [<host group name>] | -parity_grp_id <gno-
sgno> | -external_grp_id
    <gno-sgno> | -grp_opt <group option> -device_grp_name <device group name>
[<device name>]| -nvm_subsystem_id <nvm subsystem id>]
```

Options and parameters

-resource_name <resource group name>
Specifies the resource group name. Up to 32 characters can be specified.

-ldev_id <ldev#>
Specifies the LDEV number (0-65279). For example:

- -ldev_id 200

-port <port#> [<host group name>]
Specifies the port number, host group ID, and host group name. If the host group name is more than 64 characters, use the host group ID instead of the host group name. For example:

- CL1-A or CL1-A-g (g is from 0 to 255)

Chapter 5: Configuration setting commands

-parity_grp_id <gno-sgno>
>   Specifies the parity group number (gno:1-52, sgno:1-32). For example:
>
>   ▪ 3-1

-external_grp_id <gno-sgno>
>   Specifies the external volume group number (gno:1-16384, sgno:1-4096). For example:
>
>   ▪ 52-11 ("E" is not required)

-grp_opt <group option>
>   Specifies the device information about the LDEV in the device group. Specify "ldev" (fixed). The information about the LDEV in the device group is used.

-device_grp_name <device group name> [<device name>]
>   Specifies the name of device group (maximum 32 characters) to be deleted.
>
>   To specify the specific LDEV in the device group, specify the device name of the LDEV (maximum 32 characters) in the device group.
>
>   If the device name is omitted, the command is applied to all LDEVs in the device group.

-nvm_subsystem_id <nvm subsystem id>
>   Specifies the NVM subsystem ID.

Examples

Deleting the LDEV: 400 from the resource group: sql_srv.

```
# raidcom delete resource -resource_name sql_srv -ldev_id 400
```

Deleting a port of CL1-A from the resource group: sql_srv.

```
# raidcom delete resource -resource_name sql_srv -port CL1-A
```

Deleting a host group of CL1-A-0 from the resource group: sql_srv.

```
# raidcom delete resource -resource_name sql_srv -port CL1-A-0
```

Deleting a parity group: 5-2 from the resource group: sql_srv.

```
# raidcom delete resource -resource_name sql_srv -parity_grp_id 5-2
```

Deleting an external volume group: 01-02 from the resource group: sql_srv.

```
# raidcom delete resource -resource_name sql_srv -external_grp_id 01-02
```

Deleting the resource group: sql_srv.

```
# raidcom delete resource -resource_name sql_srv
```

Chapter 5: Configuration setting commands

Deleting the NVM subsystem ID: 0 from the resource group: sql_srv.

```
# raidcom delete resource -resource_name sql_srv -nvm_subsystem_id 0
```

raidcom get resource

Displays resource group information.

Syntax

```
raidcom get resource -key <option>
```

Options and parameters

-key <option>
Specify "opt" as the option to display resource group information on the virtual mode.

Examples

Displaying a resource group and resource group lock information.

```
# raidcom get resource
```

```
RS_GROUP        RGID      stat Lock_owner  Lock_host  Serial#
meta_resource    0    Unlocked        -           -    64556
```

Displaying a resource group status on the virtual mode.

```
#raidcom get resource -key opt
```

```
RS_GROUP       RGID    V_Serial#   V_ID   V_IF   Serial#
meta_resource   0       302624     R8      Y     302624
USP_002         1        64035     R5      Y     302624
```

Description of each column in output example:

RS_GROUP
Displays the resource group name.

RGID
Displays the resource group ID. RGID=0 is used for meta resource group.

stat
Displays the locking status of the resource group name on HUS VM and VSP.

Chapter 5: Configuration setting commands

Lock_owner
> Displays the (authorized) user who locks the resource group. A hyphen (-) is displayed if the user who locks the resource group has already logged out, or the resource group is not locked.

Lock_host
> Displays the host name of a user who locks the resource group. A hyphen (-) is displayed if the user who locks the resource group has already logged out, or the resource group is not locked.
>
> If REST API locks the resource group, the IP address that is used internally in the storage system might be displayed.

Serial#
> Displays the product serial number.
>
> - The serial number for VSP 5000 series is displayed with a "5" at the beginning ("512345" = serial number 12345).
>
> - The serial number for VSP G1x00 and VSP F1500 is displayed with a "3" at the beginning ("312345" = serial number 12345).

V_Serial#
> Displays the product serial number on the virtual mode.
>
> - The serial number for VSP 5000 series is displayed with a "5" at the beginning ("512345" = serial number 12345).
>
> - The serial number for VSP G1x00 and VSP F1500 is displayed with a "3" at the beginning ("312345" = serial number 12345).

V_ID

Displays the storage system type:*

- RH10HF: VSP 5200, VSP 5600

- RH10HG: VSP 5200H, VSP 5600H

- R9F: VSP 5100, VSP 5500

- R9G: VSP 5100H, VSP 5500H

- R8: VSP F1500, VSP G1x00

- R7: VSP

- RH10MHF: VSP E1090

- RH10MHG: VSP E1090H

- M9H: VSP E990

- M9MG: VSP E790H

- M9M: VSP E790

- M9S: VSP E590

- M9SG: VSP E590H

- M850HF: VSP F900

- M850H: VSP G900

- M850M3F: VSP F700

- M850M3: VSP G700

- M850S2F: VSP F370

- M850S2: VSP G370

- M850S1F: VSP F350

- M850S1: VSP G350

- M850XS: VSP G130

- M8H: VSP G800, VSP F800

- M8M: VSP G400, G600, VSP F400, F600

- M8S: VSP G200

- M7: HUS VM

* The storage system uses the storage system type information to determine the model information to respond to the host. The storage system type information for a resource group that does not belong to a virtual storage machine matches the actual model information of the storage system. In addition, when the storage system controller has been upgraded and changed, the storage system type information before the controller upgrade is displayed and does not match the actual storage model. To see the actual model information of a storage system whose controller has been upgraded, refer to the "MODEL" information of the `raidcom get system` command.

V_IF

Displays the status of the virtual mode:

- Y: The virtual mode is enabled.

- N: The virtual mode is disabled.

raidcom lock resource

This locks the specified resource group name.

When you perform these commands, lock the resource group to which resource is allocated before executing the command.

- **add**

- **delete**

- **modify**

- **initialize**

- **check_ext_storage**

- **disconnect**

- **set**

- **reset**

- **reallocate**

- **monitor**

If the specified resource group name does not exist, this command is rejected with EX_ENOOBJ. If the resource group to be locked is deleted while executing the **raidcom lock resource** command, the command might fail with the error EX_CMDRJE (SSB1=2E30, SSB2=0026). If the command fails with this error, retry the command.

In authentication mode, the user executing this command must have a permission for the resource group name.

Syntax

```
raidcom lock resource [-resource_name <resource group name>]
      [-time <time(sec)>] [-automatic_unlock <time>]
```

Options and parameters

[-resource_name <resource group name>]

Specifies the name of resource group (maximum 32 characters).

Specify defined resource group names.

If this parameter is not specified, all resource groups that are assigned to the user are locked.

[-time <time(sec)>]

This parameter is used for specifying the latency until the specified resource is locked.

The TOV time of the lock instruction is specified.

When <time> is specified as "0", it is executed as "nowait (no waiting time)" mode.

If this parameter is not specified, the default waiting time (7200 seconds) is used.

[-automatic_unlock <time>]

Supported storage systems:

- VSP 5000 series
- VSP G1x00, VSP F1500 (firmware 80-06-0x or later)
- VSP E series
- VSP G130, G/F350, G/F370, G/F700, G/F900
- VSP G200, G400, G600, G800, VSP F400, F600, F800 (firmware 83-05-0x or later)

Other storage systems ignore this option even if it is specified.

Specifies the latency until resource lock is released automatically. If this option is specified, resource lock is obtained, and then released automatically if another raidcom command (except raidcom `-login` or raidcom `-h`) is not executed within the time period specified by <time>.

Configuration setting and reference commands raidcom command (except when the `-login` or `-h` option is specified)

<time> must be specified in seconds. When 0 is specified, resource lock is not released automatically. If the specified value is 10 or smaller, resource lock might be released automatically during command execution.

Examples

Resource group: Locking the resource of the meta_resource.

```
# raidcom lock resource -resource_name meta_resource
```

raidcom unlock resource

This unlocks the specified resource group name.

If the specified resource group name does not exist, this command is rejected with EX_ENOOBJ.

In authentication mode, a user executing this command must have a permission for the resource group name.

Syntax

```
raidcom unlock resource [-resource_name <resource group name>]
```

Options and parameters

[-resource_name <resource group name>]

Specifies the name of resource group (maximum 32 characters).

Specify defined resource group names.

If this parameter is not specified, all resource groups that are assigned to the user are unlocked.

Examples

Resource group: Unlocking the resource of the meta_resource.

```
# raidcom unlock resource -resource_name meta_resource
```

raidcom map resource

Arrange a resource to the virtual storage machine. This is a synchronous command.

📄 **Note:** If a failure occurs in a processor on an MP blade or a controller, this command cannot be run even on an MP blade or a controller on which no failure occurs.

If you specify the `-request_id` option for the asynchronous command, the processing is run asynchronously with the command input. Verify that the processing is complete by using the **`raidcom get command_status`** command. There is a limit to the number of asynchronous commands that can be accepted simultaneously. For details, see <u>Maximum number of acceptable asynchronous commands (on page 228)</u>.

Syntax

```
raidcom map resource {-ldev_id <ldev#> -virtual_ldev_id
    {<ldev#>|reserve} [-ssid<ssid>] [-emulation <emulation type>][-request_id auto]
    | -port <port#> -virtual_port <port#>}
```

Options and parameters

-ldev_id <ldev#>

Specify an LDEV ID (0-65279).

Example:

▪ -ldev_id 400

-virtual_ldev_id {<ldev#>| reserve}

Specify an LDEV ID (0-65279) to be used in the virtual storage machine.

If you specify "reserve" instead of the LDEV ID, the reserve attribute of global-active device is set.

Chapter 5: Configuration setting commands

Example:

- -virtual_ldev_id 100

- -virtual_ldev_id reserve

-ssid <ssid>

Specify an SSID related to an LDEV in the virtual storage machine.

[-emulation <emulation type>]

Specify the emulation type of a relevant LDEV on the virtual storage machine. This setting is reflected in the inquiry response.

Specify the emulation type by adding "*n" in the LUSE configuration or by adding "-CVS" in the CVS configuration. ("n" shows the number of LUSE components.)

Apply in order from "*n" (LUSE configuration) to "-CVS" (CVS configuration) when it is LUSE configuration and CVS configuration.

Example:

- -emulation OPEN-3-CVS

- -emulation OPEN-3*6

- -emulation OPEN-3*6-CVS

[-request_id auto]

Supported storage systems:

- VSP E series (DKCMAIN firmware version 93-06-41-XX/XX or later)

- VSP 5000 series (DKCMAIN firmware version 90-08-41-XX/XX or later)

It is recommended to specify this option when running the command. The command is run as an asynchronous command only when this option is specified.

<request #> is the Request ID assigned for each command run. For details, see Request ID function (on page 246).

-port <port#>

Supported storage systems:

- VSP

- HUS VM

Specify a port number. Specify the port number whose attribute is Target.

Example:

- CL1-A

-virtual_port <port#>

Supported storage systems:

- VSP

- HUS VM

Specify a port number to be used in the virtual storage machine.

Example:

- CL3-B

Examples

Create the virtual LDEV100 in the LDEV400 (run as a synchronous command).

```
# raidcom map resource -ldev_id 400 -virtual_ldev_id 100
```

Create the virtual LDEV100 in the LDEV400 (run as an asynchronous command).

```
raidcom map resource -ldev_id 400 -virtual_ldev_id 100 -request_id auto REQID : 1
```

Set the global-active device reserve attribute to the LDEV400.

```
# raidcom map resource -ldev_id 400 -virtual_ldev_id reserve
```

Create the virtual port CL2-B in port CL1-A.

```
# raidcom map resource -port CL1-A -virtual_port CL2-B
```

raidcom unmap resource

Cancel the resource arrangement in the virtual storage machine. This is a synchronous command.

📄 **Note:** If a failure occurs in a processor on an MP blade or a controller, this command cannot be run even on an MP blade or a controller on which no failure occurs.

If you specify the `-request_id` option for the asynchronous command, the processing is run asynchronously with the command input. Verify that the processing is complete by using the **`raidcom get command_status`** command. There is a limit to the number of asynchronous commands that can be accepted simultaneously. For details, see .

Syntax

```
raidcom unmap resource {-ldev_id <ldev#> -virtual_ldev_id
     {<ldev#>|reserve}[-request_id auto] | -port <port#> -virtual_port <port#>}
```

Options and parameters

-ldev_id <ldev#>
        Specify an LDEV ID (0-65279).

Example:

- -ldev_id 400

-virtual_ldev_id {<ldev#>| reserve}

Specify an LDEV ID (0-65279) to be used in the virtual storage machine.

If you specify "reserve" instead of the LDEV ID, the reserve attribute of global-active device is released.

Example:

- -virtual_ldev_id 100

- -virtual_ldev_id reserve

[-request_id auto]

Supported storage systems:

- VSP E series (DKCMAIN firmware version 93-06-41-XX/XX or later)

- VSP 5000 series (DKCMAIN firmware version 90-08-41-XX/XX or later)

It is recommended to specify this option when running the command. The command is run as an asynchronous command only when this option is specified.

<request #> is the Request ID assigned for each command run. For details, see .

-port <port#>

Specify a port number. Specify the port number whose attribute is Target.

Example:

- CL1-A

-virtual_port <port#>

Specify a port number to be used in the virtual storage machine.

Example:

- CL3-B

Examples

Cancel a virtual LDEV100 in an LDEV400 (run as a synchronous command).

```
# raidcom unmap resource -ldev_id 400 -virtual_ldev_id 100
```

Cancel a virtual LDEV100 in an LDEV400 (run as an asynchronous command).

```
# raidcom unmap resource -ldev_id 400 -virtual_ldev_id 100 -request_id auto
REQID : 1
```

Release the global-active device reserve attribute to the LDEV400.

```
# raidcom unmap resource -ldev_id 400 -virtual_ldev_id reserve
```

Chapter 5: Configuration setting commands

Cancel a virtual port CL2-B in a port CL1-A.

```
# raidcom unmap resource -port CL1-A -virtual_port CL2-B
```

raidcom add snap_pool

Creates pools and adds pool VOLs for Thin Image and Copy-on-Write Snapshot by the specified resource.

When specifying a pool that is already created for Thin Image or Copy-on-Write Snapshot, the specified resource is added as a pool volume. A parity group, an LDEV, and a device group can be specified as a resource.

> **Note:** If a failure occurs in a processor on an MP blade or a controller, this command cannot be run even on an MP blade or a controller on which no failure occurs.

> **Tip:** For VSP 5000 series, VSP E series, VSP G130, G/F350, G/F370, G/F700, G/F900, specifying a parity group is recommended. When a parity group is specified, an LDEV is created and a pool volume is added. It is not necessary to create an LDEV in advance. In addition, you cannot add an existing LDEV as a pool volume.

Specify either one of Pool ID or Pool Name certainly. If both the Pool ID and Pool Name options are omitted, this command is rejected with EX_REQARG.

This command is executed asynchronously with the command input. Check the completion of this process on the `raidcom get command_status` command.

Syntax

```
raidcom add snap_pool {-pool_id <pool ID#> | -pool_name <pool
    naming> | -pool_id <pool ID#>
    -pool_name <pool naming>}
     {-parity_grp_id <gno-sgno> [-resource_id <resource group id >] |
    -ldev_id <ldev#> …[-cnt<count>] | -grp_opt <group option>
     -device_grp_name <device group name> [<device name>]}
     [-user_threshold <%> ][-thinsnap][-auto_add_poolvol {enable|disable}]
```

Options and parameters

-pool_id <pool ID#>

Specifies the Pool ID (0-127) of a Thin Image or Copy-on-Write Snapshot pool.

If a `-pool_id` option is non-numeric, the specified value is recognized as a pool name to identify the pool ID.

When you omit specifying the `-pool_id` option, you need to specify `-pool_name` option.

Chapter 5: Configuration setting commands

When specifying the `-pool_name` option without a `-pool_id` option, the specified pool is searched first. If the pool exists, the pool volume is added to the corresponding pool. If the pool does not exist, a new pool is created and is assigned a pool ID. However, if the pool does not exist and multiple commands without the `-pool_id` option are executed simultaneously, the same pool ID might be assigned to multiple newly created pools.

To avoid this behavior, use the **raidcom lock** resource command in advance to lock all resource groups. If all the resource groups are locked, commands executed by other users cannot take out unused pool IDs from all the resource groups. If you create multiple pools, execute the next **raidcom add snap_pool** command after the execution of the current **raidcom add snap_pool** command is complete.

After the execution of all the **raidcom add snap_pool** commands are complete, use the **raidcom unlock resource** command to unlock all the resource groups.

-pool_name <pool naming>

Specifies the pool name of a pool for Thin Image or Copy-on-Write Snapshot. Up to 32 characters can be specified.

When specifying a pool ID or a pool name, if a pool name exists in the specified pool ID, the pool name is overwritten. If the pool volume is added by specifying only a pool ID, the pool name is not changed. When the specification of Pool Name is omitted, a Pool ID must be specified. When the Pool ID is specified and the Pool Name is omitted, a pool name is allocated automatically in the form of "New Pool<number>".

The `-pool_name` option cannot be specified by configuring only in numeric because numeric values are preferentially identified as a pool ID. Specifies a pool ID with the `-pool_id<pool ID#>` option.

-parity_grp_id <gno-sgno> [-resource_id <resource group_id >]

Supported storage systems:

- VSP E series

- VSP G130, G/F350, G/F370, G/F700, G/F900

Specifies the parity group number (gno: 1 to 52, sgno: 1 to 32).

For example: `3-1`

When `-resource_id <resource grp id>` is specified, select the LDEV that is not installed in the specified resource group and has the largest LDEV number to create a pool volume. This option is valid only when a pool is created. If you omit this option, the smallest ID of a resource group for which the user has operation permissions is set.

-ldev_id <ldev#>

Specifies the LDEV number (0 to 65279). Up to 64 of LDEVs can be specified at a time. For example:

- -ldev_id 100

- -ldev_id 100-110

- -ldev_id 100 -cnt 10

Chapter 5: Configuration setting commands

[-cnt <count>]

Specifies the count (2 to 64).

The count becomes singular if not specified.

Up to 64 of LDEVs can be specified at a time.

-grp_opt <group option>

Specifies the device information extracted from the LDEV belonging to the device group. Specify "ldev" (fixed). The information of LDEV belonging to the device group is used.

-device_grp_name <device group name> [<device name>]

Specifies the name of device group (maximum 32 characters) to be operated.

To operate an LDEV in the device group, specify the device name of LDEV (maximum 32 characters) within the device group.

If the device name is omitted, all the LDEVs belonging to the device group are operated.

When the `-pool_name` option is omitted, the device group name changes into the pool name.

[-user_threshold <%>]

Specifies the user defined threshold value (20 to 95) (%).

If this option is omitted, "80" is used.

When you add a pool volume, this option is ignored even if it is specified. If you want to change the user defined threshold value of the additional pool volume, execute the **raidcom modify pool** command.

[-thinsnap]

Supported storage systems:

▪ VSP

▪ HUS VM

When this option is specified, a pool for Thin Image is created.

[-auto_add_poolvol {enable|disable}]

Specifies whether to automatically add pool volumes for which accelerated compression is enabled. This option is valid only for creating pools. This option can be specified for VSP G1x00 or VSP F1500.

▪ enable: Automatically adds pool volumes for which accelerated compression is enabled according to the compression ratio of a parity group for which accelerated compression is enabled.

▪ disable: The pool volumes for which accelerated compression is enabled are not added automatically.

When these options are not specified, the default values will be applied as follows:

- enable: When the storage system is VSP 5000 series, VSP E series, VSP G/F350, G/F370, G/F700, G/F900.

- disable: When the storage system is VSP G1x00 or VSP F1500.

Examples

Using LDEVs:400, 401, and 402, creating a pool of Pool ID:1, Pool Name: my_ss_pool for Copy-on-Write Snapshot.

```
# raidcom add snap_pool -pool_id 1 -pool_name my_ss_pool -ldev_id 400 401 402
```

Using LDEVs:410, 411, and 412, creating a pool of Pool ID:3, Pool Name: my_ss_pool for Thin Image.

```
# raidcom add snap_pool -pool_id 3 -pool_name my_ss_pool -ldev_id 410 411 412 -
thinsnap
```

Using LDEVs:500, 501, and 502, creating a pool of Pool ID: Allocated automatically, Pool Name: my_ss_pool for Copy-on-Write Snapshot.

```
# raidcom add snap_pool -pool_name my_ss_pool -ldev_id 500 501 502
```

Using LDEVs:600, 601, and 602, creating a pool of Pool ID: 2, Pool Name: Allocated automatically for Copy-on-Write Snapshot.

```
# raidcom add snap_pool -pool_id 2 -ldev_id 600 601 602
```

Using LDEV belonging to the device group: grp1, creating a pool of Pool ID: 1, Pool Name: Allocated automatically for Copy-on-Write Snapshot.

```
# raidcom add snap_pool -pool_id 1 -grp_opt ldev -device_grp_name grp1
```

Using an LDEV ID of resource group ID: 0, creating an LDEV for parity group: 1-1, and then creating a pool of Pool ID: 1, Pool Name: Allocated automatically for Thin Image.

```
# raidcom add snap_pool -pool_id 1 -parity_grp_id 1-1 -resource_id 0 -thinsnap
```

Using LDEVs: 410, 411, and 412, creating a pool of Pool ID: 3 and Pool Name: my_ss_pool for Thin Image for which automatic addition settings of pool volumes whose accelerated compression is enabled: enable.

```
# raidcom add snap_pool -pool_id 3 -pool_name my_ss_pool -ldev_id 410 411
412 -thinsnap -auto_add_poolvol enable
```

raidcom get snap_pool

Displays pool information for Thin Image or Copy-on-Write Snapshot.

Syntax

```
raidcom get snap_pool
```

Options and parameters

None.

Examples

Displaying pool information for Thin Image or Copy-on-Write Snapshot.

```
# raidcom get snap_pool
```

```
PID POLS U(%) SSCNT Available(MB)  Capacity(MB)  Seq# Num LDEV# H(%)  FMT_CAP(MB)
003 POLS 100 10000           100   1000000000   62500  1   375  70    100
```

Description of each column in output example:

PID
>   pool ID

POLS
>   Displays status of the pool.
>
>   - POLN: "Pool Normal" (The pool is in the normal status.)
>
>   - POLF: "Pool Full" (The pool is in the overflow status exceeding the threshold.)
>
>   - POLS: "Pool Suspend" (The pool is in the overflow status exceeding the threshold and is blocked.)
>
>   - POLE: "Pool failure" (The pool is suspended in the failure status.) In this status, the pool information cannot be displayed.

U(%)
>   Usage rate of the pool, including the mapped capacity and the capacity for Full Allocation

SSCNT
>   number of volumes in the pool

Available (MB)
>   Available capacity for the data volumes in the pool

Capacity (MB)
>   Total capacity of the pool.

Seq#

Serial number.

Num

Number of LDEVs in the pool

LDEV#

Number of the first LDEV in the pool

H(%):

threshold value for the pool

FMT_CAP(MB)

Formatted capacity of the pool

- (hyphen): This information is not available for this pool.

raidcom add snapshot

Add a combination of the specified LDEV number and Pool ID to a snapshot group. If there is no name of specified snapshot group, create a new snapshot group.

Syntax

```
raidcom add snapshot -ldev_id <ldev#(P)> <ldev#(S)> -pool
{<pool ID#> | <pool naming>}
-snapshotgroup <name> [-snap_mode <mode>] [-mirror_id <mu#>]
```

Options and parameters

-ldev_id <ldev#(P)> <ldev#(S)>

Specifies the LDEV number with the snapshot data created. LDEV numbers for P-VOL must be included. If VOL_ATTR of the specified LDEV is a DRS-VOL, a Thin Image Advanced pair is created.

-pool {<pool ID#> | <pool naming>}

Specifies the pool ID or the pool name created for Snapshot.

-snapshotgroup <name>

Specifies a name to be given for snapshot group.

[-mirror_id <mu#>]

Supported storage systems:

- VSP 5000 series

- VSP G1x00 and VSP F1500

- VSP E series

- VSP G130, G/F350, G/F370, G/F700, G/F900

Specifies the mirror ID.

If this option is omitted, an unused mirror ID is allocated automatically.

[-snap_mode <mode>]

Specifies the mode to create a snapshot group and the mode to add a P-VOL or S-VOL to the snapshot group. If you specify multiple modes, use a space as a separator.

The following modes can be specified:

- <mode>= CTG: For creating in CTG mode. The consistency group number is allocated internally by itself.

  If this option is omitted, the snapshot group is created in normal mode.

  This option is effective only when a new snapshot group is to be created.

  This option is ignored if this is specified for the existing snapshot group.

- <mode> = clone: For adding a P-VOL or S-VOL as a volume for which you can create a pair with the clone attribute (clone mode).

- <mode>= cascade: For adding a P-VOL or S-VOL as a volume in which you can create a cascade pair (cascade mode).

Example:

- `-snap_mode CTG clone`

- `-snap_mode CTG`

Examples

Adding a combination of the P-VOL (LDEV number 10:10), the S-VOL (LDEV number 20:20), and the Pool (SnapPool00) to the snapshot group (db1).

```
# raidcom add snapshot -ldev_id 0x1010 0x2020 -pool SnapPool00 -snapshotgroup db1
```

raidcom map snapshot

Maps the specified snapshot data to the S-VOL. The S-VOL to be mapped snapshot data must be created before it is specified.

For a Thin Image Advanced pair, the pair status changes to CPYP while the specified operation is in process, and then changes to PSUS when the process is complete. To verify that the operation is complete, run the raidcom get snapshot command to verify that STAT has changed from CPYP to other than CPYP.

Syntax

```
raidcom map snapshot -ldev_id <ldev#(P)> <ldev#(S)> {-snapshotgroup <name> | -
mirror_id <mu#>}
```

Options and parameters

-ldev_id <ldev#(P)> <ldev#(S)>
Specifies the LDEV number of P-VOL and S-VOL.

-snapshotgroup <name>
Specifies the name of the snapshot group in which the snapshot data as the operation target is included.

One MU that corresponds to the specified P-VOL is mapped from the specified snapshot group.

> **Note:**
>
> Because the MU is selected automatically, an unexpected MU could be mapped. So if you intend to specify the specific snapshot data, specify the snapshot data by an MU number instead of a snapshot group.

-mirror_id <mu#>
Specifies the mirror ID of a snapshot data to be a target.

Examples

Mapping the snapshot data of the LDEV number 10:10 that is included in the snapshot group (db1) to the LDEV number 20:00.

```
# raidcom map snapshot -ldev_id 0x1010 0x2000 -snapshotgroup db1
```

Mapping the snapshot data of the LDEV number 10:10 and the Mirror ID 10 to the LDEV number 20:00.

```
# raidcom map snapshot -ldev_id 0x1010 0x2000 -mirror_id 10
```

raidcom unmap snapshot

Unmaps the S-VOL which is mapping the snapshot data.

For a Thin Image Advanced pair, the pair status changes to CPYP while the specified operation is in process, and then changes to PSUS when the process is complete. To verify that the operation is complete, run the raidcom get snapshot command to verify that STAT has changed from CPYP to other than CPYP.

Syntax

```
raidcom unmap snapshot -ldev_id <ldev#> [-snapshotgroup <name> | -mirror_id <mu#>]
```

Options and parameters

-ldev_id <ldev#>
Specifies the LDEV number to be unmapped.

Snapshot group name or MU number must be specified when you specify the LDEV number of P-VOL to identify the snapshot data.

Do not specify the snapshot group name and MU number when you specify the LDEV number of S-VOL.

[-snapshotgroup <name>]

Specifies the name of the snapshot group in which the snapshot data as the operation target is included.

One MU which corresponds to the specified P-VOL is unmapped from the specified snapshot group.

> **📄 Note:**
>
> Because the MU is selected automatically, an unexpected MU could be unmapped. So if you intend to specify the specific snapshot data, specify the snapshot data by an MU number instead of a snapshot group.

[-mirror_id <mu#>]

Specifies the mirror ID of the snapshot data to be a target of unmapping when you specify the LDEV number of P-VOL.

Examples

Unmapping the S-VOL (LDEV number 20:00).

```
# raidcom unmap snapshot -ldev_id 0x2000
```

Unmapping the snapshot data of the LDEV number 10:10 that is included in the snapshot group (db1).

```
# raidcom unmap snapshot -ldev_id 0x1010 -snapshotgroup db1
```

Unmapping the P-VOL (LDEV number 10:10) and the mirror ID 10.

```
# raidcom unmap snapshot -ldev_id 0x1010 -mirror_id 10
```

raidcom delete snapshot

Deletes the snapshot data and the snapshot group. The relevant snapshot data of the LDEV is deleted from the snapshot group by specifying LDEV number. When no snapshot data left in the snapshot group, the snapshot group is deleted.

Syntax

```
raidcom delete snapshot {-snapshotgroup <name> | -ldev_id <ldev#>
{-mirror_id <mu#> |-snapshotgroup <name> | -range tree}}
```

Options and parameters

-snapshotgroup <name>

    Specifies the snapshot group in which the target data to be deleted is included.

    If the snapshot group is specified as the target, all the snapshot data and the snapshot group are deleted.

-ldev_id <ldev#>

    Specifies the LDEV number of P-VOL or S-VOL for the snapshot data to be deleted.

    When P-VOL is specified, specify the snapshot data by specifying the MU number or the snapshot group (Specifying the MU number or the snapshot group is mandatory).

    When you specify the S-VOL, do not specify a MU number or a Snapshot group. If you specify the MU number or the Snapshot group, the P-VOL of specified LDEV number becomes the subject of deletion.

-mirror_id <mu#>

    Specifies the Mirror ID of the snapshot data to be deleted.

-snapshotgroup <name>

    Specifies the snapshot group to be deleted.

    The smallest number of MU in the snapshot group becomes the subject to be deleted.

-range tree

    Specify this option to delete snapshot data in all layers of a snapshot tree for which the root volume is specified by -ldev_id *<ldev#>*. This option is not supported by Thin Image Advanced. To specify this option, the following conditions must be met:

- Root volume is an LDEV specified by -ldev_id *<ldev#>*.

- The pair to be deleted was created in the cascade or clone mode.

- The user has the operation authority for volumes of all pairs to be deleted.

> 📄 **Note:**
>
> When multiple snapshot trees are cloned, specify the root volume of the first snapshot tree for -ldev_id *<ldev#>*, and execute the command. Only the first snapshot tree is deleted, and the status of the pairs under the first snapshot tree changes to PSUE. Specify the root volume of a snapshot tree under the first snapshot tree for -ldev_id *<ldev#>*, and execute the command again. When this option is specified, the SIM code (4b3xxx) that indicates abnormal termination of the Thin Image option might be output. However, the snapshot tree has been deleted normally and there is no problem.

Examples

Deleting the snapshot data of the snapshot group (db1).

```
# raidcom delete snapshot -snapshotgroup db1
```

Deleting the snapshot data of the P-VOL (LDEV number 10:10) and the Mirror ID (10).

```
# raidcom delete snapshot -ldev_id 0x1010 -mirror_id 10
```

Deleting the snapshot data of the LDEV number 10:10 that is included in the snapshot group (db1).

```
# raidcom delete snapshot -ldev_id 0x1010 -snapshotgroup db1
```

```
# raidcom delete snapshot -snapshotgroup db1 -ldev_id 0x1010
```

Deleting the snapshot data of the S-VOL (LDEV number 20:10).

```
# raidcom delete snapshot -ldev_id 0x2010
```

Deleting the snapshot data in all layers of a snapshot tree whose root volume has LDEV number 20:10.

```
# raidcom delete snapshot -ldev_id 0x2010 -range tree
```

raidcom modify snapshot

Operate the specified snapshot group.

Syntax

```
raidcom modify snapshot -ldev_id <ldev#> {-snapshotgroup <name>
| -mirror_id <mu#>} -snapshot_data <op> [-copy_pace <copy pace>]

raidcom modify snapshot -snapshotgroup <name> -snapshot_data <op>
[-copy_pace <copy pace>]

raidcom modify snapshot -ldev_id <ldev#> -snapshot_data <op>
[-copy_pace <copy pace>]

raidcom modify snapshot -snapshotgroup <current name>
<new name> -snapshot_data rename
```

Options and parameters

-ldev_id <ldev#>
Specifies the LDEV number of P-VOL (or S-VOL) to be performed.

-snapshotgroup <name>
Specifies the snapshot group name in which the snapshot data is included.

-mirror_id <mu#>

        Specifies the mirror ID of a snapshot data.

-snapshot_data <op>

        Specifies the operation to be performed for the specified snapshot group. The parameter of the operation to be specified is the following. (Note that parameters clone, delete_garbage, and stop_deleting_garbage cannot be used in Thin Image Advanced.):

- create: Takes snapshot data.[1, 4, 6]

- split: Takes snapshot data.[1, 4, 6]

- resync: Deletes the created snapshot data.[1]

- restore: Restores the Thin Image pair.[1]

  For a Thin Image Advanced pair, perform the following steps:

  1. Verify that the status of the pair to be restored is PSUS or SSUS.
  2. Run the `raidcom modify snapshot -snapshot_data` restore, and then verify that the command is complete normally.
  3. Wait until the status of the pair to be restored migrates to a status that is not RCPY.
  4. If the pair status changes to PSUS or SSUS in step 3, the pair restore operation is successfully complete.

- clone: Clones pairs.[2, 5]

- delete_garbage: Starts the snapshot garbage data deletion processing (defragmentation)[3] of the specified Thin Image Advanced root volume.[3, 5]

- stop_deleting_garbage: Stops the snapshot garbage data deletion (defragmentation)[3] of the specified Thin Image Advanced root volume.[3, 5]

> 📄 **Note:**
>
> 1. Clone mode pairs do not support this operation.
> 2. Cascade mode pairs do not support this operation.
> 3. For the details about the defragmentation, see the Hitachi Thin Image User Guide.
> 4. For a Thin Image Advanced pair, the pair status changes to PSUP while the specified operation is in process, and then changes to PSUS when the processing is complete. To verify that the operation is complete, run the raidcom get snapshot command to verify that STAT has changed from PSUP to the status other than PSUP.
> 5. This operation is not available for a Thin Image Advanced pair.
> 6. For a Thin Image Advanced pair, the group operation with the `-snapshotgroup` option specified might not be performed if all of the following conditions are met. If the following conditions are met, perform the operation on one pair at a time by specifying the `-ldev_id` and `-mirror_id` options.
>
>   - The snapshot group is not in the CTG mode.
>
>   - The snapshot group contains two or more pairs created using the same P-VOL.

[-copy_pace <copy pace>]

Specifies the copy speed. Enabled only when `clone` is specified for `-snapshot_data`. The following speeds can be specified:

- slower

- medium

- faster

When `clone` is specified for `-snapshot_data`, if you did not specify the copy speed, `medium` is set. When you specify a value other than `clone`, if you specify the copy speed, this option is not enabled.

-snapshotgroup <current name> <new name> -snapshot_data rename

Specifies when you change the snapshot group name. The snapshot group name specified in <current name> is changed to the name specified in <new name>.

The following shows the operation of the snapshot data with the combination of options and parameters:

When creating a snapshot data (when specifying create/split)

#	The ways to specify the parameter	CTG mode	normal mode
1	LDEV number and snapshot group.	P-VOL in the snapshot group.	P-VOL in the snapshot group and the smallest number of MU.
2	LDEV number and MU number.	Specified P-VOL and MU number.	Specified P-VOL and MU number.
3	LDEV number only (Specifying S-VOL).	P-VOL and MU number that are mapped by the specified S-VOL.	P-VOL and MU number that are mapped by the specified S-VOL.
4	Snapshot group.	All the P-VOLs related to the snapshot group. The consistency is endured.	All the P-VOLs related to the snapshot group. The consistency is not endured.

When discarding or restoring the snapshot data (when specifying resync/restore)

#	The ways to specify the parameter	CTG mode	normal mode
1	LDEV number and snapshot group.	P-VOL in the snapshot group.	P-VOL in the snapshot group and the smallest number of MU.
2	LDEV number and MU number.	Specified P-VOL and MU number.	Specified P-VOL and MU number.
3	LDEV number only (Specifying S-VOL).	P-VOL and MU number that are mapped by the specified S-VOL.	P-VOL and MU number that are mapped by the specified S-VOL.
4	Snapshot group.	All the P-VOLs related to the snapshot group.	All the P-VOLs related to the snapshot group.

When deleting (defragmenting) the snapshot garbage data(when specifying delete_garbage/stop_deleting_garbage) (Only for VSP 5000 series)

#	The ways to specify the parameter	CTG mode	normal mode
1	LDEV number and snapshot group.	Cannot be specified.	Cannot be specified.
2	LDEV number and MU number.	Cannot be specified.	Cannot be specified.
3	LDEV number only (Specifying root volume).	Snapshot data area of the specified root volume.	Snapshot data area of the specified root volume.
4	Snapshot group.	Cannot be specified.	Cannot be specified.

Examples

Creating a snapshot data for the P-VOL (LDEV number 10:10) that is included in the snapshot group (db1).

```
# raidcom modify snapshot -ldev_id 0x1010 -snapshotgroup db1 -snapshot_data create
```

Creating a snapshot data for the P-VOL (LDEV number 10:10) and the mirror ID 10.

```
# raidcom modify snapshot -ldev_id 0x1010 -mirror_id 10 -snapshot_data create
```

Creating a snapshot data for the S-VOL (LDEV number 20:10).

```
# raidcom modify snapshot -ldev_id 0x2010 -snapshot_data create
```

Creating a snapshot data for all the P-VOLs that are included in the snapshot group (db1).

```
# raidcom modify snapshot -snapshotgroup db1 -snapshot_data create
```

Discarding the snapshot data for the P-VOL (LDEV number 10:10) and the mirror ID 10.

```
# raidcom modify snapshot -ldev_id 0x1010 -mirror_id 10 -snapshot_data resync
```

Restoring the snapshot data for the S-VOL (LDEV number 20:10).

```
# raidcom modify snapshot -ldev_id 0x2010 -snapshot_data restore
```

Changing the snapshot group name (db1) to db2.

```
# raidcom modify snapshot -snapshotgroup db1 db2 -snapshot_data rename
```

Chapter 5: Configuration setting commands

Starting the snapshot garbage data deletion (defragmentation) of the root volume (LDEV number 10:10) (only for VSP 5000 series).

```
# raidcom modify snapshot -ldev_id 0x1010 -snapshot_data delete_garbage
```

Stopping the snapshot garbage data deletion (defragmentation) of the root volume (LDEV number 10:10) (only for VSP 5000 series).

```
# raidcom modify snapshot -ldev_id 0x1010 -snapshot_data stop_
deleting_garbage
```

raidcom get snapshot

Displays the information of snapshot group and snapshot data that are defined in the device. If this option is omitted, the list of snapshot group is displayed.

Syntax

```
raidcom get snapshot [-ldev_id <ldev#> [-key opt] | -snapshotgroup <name>
[-key {opt|detail}]] [-format_time] [{-check_status | -check_status_not} <string>...
[-time <time>]]
```

Options and parameters

[-ldev_id <ldev#>]

Specifies the LDEV number to be displayed the snapshot data information. Specifies either one of P-VOL or S-VOL for the LDEV number.

[-snapshotgroup <name>]

Specifies the snapshot group in which you want to display the snapshot data information.

[-key opt]

Specify the `-key opt` option together with the `-ldev_id` or the `-snapshotgroup` option to display the following:

- SMPP as the snapshot data status

- (VSP 5000 series, VSP G130, G/F350, G/F370, G/F700, G/F900) SLU attribute and information of the Diff Clone pair that is being used in the VASA environment

- (VSP 5000 series, VSP G1x00 and VSP F1500, VSP E series, VSP G130, G/F350, G/F370, G/F700, G/F900) Status of expansion processing of DP-VOL

- (VSP 5000 series, VSP E series, VSP G130, G/F350, G/F370, G/F700, G/F900) Status of the Read Only attribute of the Snapshot data.

[-key detail]

Specify the `-key detail` option together with the `-ldev_id` or the `-snapshotgroup` option to display the following:

- SMPP as the snapshot data status

- (VSP 5000 series, VSP G130, G/F350, G/F370, G/F700, G/F900) SLU attribute and information of the Diff Clone pair that is being used in the VASA environment

- Status of expansion processing of DP-VOL

- (VSP 5000 series, VSP G1x00 and VSP F1500, VSP E series, VSP G130, G/F350, G/F370, G/F700, G/F900) Status of the Read Only attribute of the Snapshot data

- (VSP 5000 series) MODE (R) indicating snapshot data of a Thin Image Advanced pair in the snapshot group mode

Specifying this option expands the display area of MODE.

The following lists the differences in the display format depending on whether to specify this option:

- When the `-key detail` option is not specified:

  MODE

  -N--

  (Displayed in a format available to display up to 4 modes.)

- When the `-key detail` option is specified:

  MODE

  --A-R-------

  (Displayed in a format available to display up to 12 modes.)

[–format_time]

Specify this option to display the time when a snapshot data (SPLT-TIME) is created, with the format[1] below. The time zone used in the storage system[2] is displayed. Zero "0" is added to the beginning if the value of each element is shorter than the regulated length.

Format: YYYY-MM-DDThh:mm:ss

Where:

- YYYY-MM-DD = year-month-date

- hh = time in 24 hours, range 0-23

- mm:ss = minute:second

For example, December 2, 2016, one o'clock p.m. is displayed as 2016-12-02T13:00:00

Notes:

1. "T" in the format separates the date and the time.

Chapter 5: Configuration setting commands

2. For VSP 5000 series, VSP E series, VSP G130, G/F350, G/F370, G/F700, G/F900, VSP G200, G400, G600, G800 and VSP F400, F600, F800, the time zone set in the device is used.

For other models, UTC+0 is used.

[-check_status <string> [-time <time>]]

Check if the snapshot group or the snapshot data is in the same state as specified in <string>. If the option contains multiple states, the OR condition check is performed and verifies that the snapshot group or snapshot data is in one of the states contained in the option.

The following strings are specified in the <string>:

- COPY: Copy status.

- PAIR: Pair status.

- PSUS: Suspend status.

- PSUE: Suspend failure status.

- PFUL: Indicates that the Thin Image or Copy-on-Write Snapshot pool exceeds the threshold in the PAIR status.

- PFUS: Indicates that the Thin Image or Copy-on-Write Snapshot pool exceeds the threshold in the PSUS status.

- RCPY: Shows that the copying is in progress by resynchronization.

- SMPP: Indicates that the pair is being deleted. When specifying SMPP, specify -key opt as well.

- PSUP: Indicates that the pair with the clone attribute is split.

If "-time" is specified, the status of the snapshot group or the snapshot data is checked every three seconds until the end of the specified <time> (seconds).

When this option is specified, the returned values are as follows:

- The snapshot group or the snapshot data is in one of the specified states: 0

- The snapshot group or the snapshot data is in none of the specified states (without -time option): 1

- The snapshot group or the snapshot data is in none of the specified states (when the specified <time> passed): EX_EWSTOT

Note:

When you specify -snapshotgroup *<name>*, and the pair is deleted, EX_ENOOBJ is returned.

Chapter 5: Configuration setting commands

[-check_status_not <string> [-time <time>]]

Check if the snapshot group or the snapshot data is not in the same state as specified in <string>. If the option contains multiple states, the OR condition check is performed and verifies that the snapshot group or snapshot data is not in any of the states contained in the option. The following strings are specified in the <string>.

- COPY: Copy status.

- PAIR: Pair status.

- PSUS: Suspend status.

- PSUE: Suspend failure status.

- PFUL: Indicates that the Thin Image or Copy-on-Write Snapshot pool exceeds the threshold in the PAIR status.

- PFUS: Indicates that the Thin Image or Copy-on-Write Snapshot pool exceeds the threshold in the PSUS status.

- RCPY: Shows that the copying is in progress by resynchronization.

- SMPP: Indicates that the pair is being deleted. When specifying SMPP, specify -key opt as well.

- PSUP: Indicates that the pair with the clone attribute is split.

If "-time" is specified, the status of the snapshot group or the snapshot data is checked every three seconds until the end of the specified <time> (seconds).

When this option is specified, the returned values are as follows:

- The snapshot group or the snapshot data is not in any of the specified states: 0

- The snapshot group or the snapshot data is in one of the specified states (without -time option): 1

- The snapshot group or the snapshot data is in one of the specified states (when the specified <time> passed): EX_EWSTOT

> 📄 **Note:**
>
> When you specify -snapshotgroup *<name>*, and the pair is deleted, EX_ENOOBJ is returned.

Examples

Displaying the list of snapshot groups.

```
# raidcom get snapshot
```

```
SnapShot_name P/S STAT  Serial# LDEV# MU# P-LDEV# PID % MODE SPLT-TIME
snap1          -   -     85000098 -    -   -       -  - ----     -
snap2          -   -     85000098 -    -   -       -  - ----     -
snap3          -   -     85000098 -    -   -       -  - ----     -
```

Displaying the snapshot data related to the specific P-VOL (LDEV number: 14536).

```
# raidcom get snapshot -ldev_id 14536
```

```
SnapShot_name P/S STAT Serial# LDEV# MU# P-LDEV# PID % MODE SPLT-TIME
snap1       P-VOL PAIR 85000098 14536 1010    -    2 100 ----   -
snap2       P-VOL PAIR 85000098 14536 1011  13000  2 100 G---   -
snap3       P-VOL PAIR 85000098 14536 1012    -    2 100 ----   -
```

Displaying the snapshot data related to the specific P-VOL (LDEV number: 10).

```
# raidcom get snapshot -ldev_id 10 -key detail
SnapShot_name P/S STAT Serial# LDEV# MU# P-LDEV#
PID % MODE SPLT-TIME SLU C_LDEV# P R
snap1 P-VOL PAIR 85000098 10 0
20 0 - --A-R------- - Y - - D
```

Displaying the snapshot data related to the specific S-VOL (LDEV number: 13000).

```
# raidcom get snapshot -ldev_id 13000
```

```
SnapShot_name P/S STAT Serial# LDEV# MU# P-LDEV# PID % MODE SPLT-TIME
snap2       S-VOL PAIR 85000098 13000 1011 14536   2 100 G---   -
```

Displaying the snapshot data included in the specific snapshot group.

```
# raidcom get snapshot -snapshotgroup snap2
```

```
SnapShot_name P/S STAT Serial# LDEV# MU# P-LDEV# PID % MODE SPLT-TIME
snap2       P-VOL PAIR 85000098 14536 1011 13000   2 100 G---   -
snap2       P-VOL PAIR 85000098 14537 1011 13001   2 100 G---   -
snap2       P-VOL PAIR 85000098 14538 1011 13002   2 100 G---   -
```

Displaying SPLT-TIME in YYYY-MM-DDThh:mm:ss format.

```
# raidcom get snapshot -ldev_id 14356  -format_time
```

```
SnapShot_name   P/S STAT   Serial# LDEV#   MU# P-LDEV# PID    % MODE SPLTTIME
       snap1 P-VOL PSUS 85000098 14536 1010        -     2 100 ---- 2016-07-
22T10:18:20
        snap2 P-VOL PSUS 85000098 14536 1011   13000  2 100 G--- 2016-07-22T10:18:20
        snap3 P-VOL PSUS 85000098 14536 1012       -    2 100 ---- 2016-07-22T10:18:20
```

Chapter 5: Configuration setting commands

Displays the snapshot data with the SLU attribute related to the specific S-VOL (LDEV # 40960) by specifying the `-key opt` option.

```
# raidcom get snapshot -ldev_id 40960 -key opt
SnapShot_name P/S STAT Serial# LDEV# MU# P-LDEV# PID %
MODE SPLT-TIME SLU C_LDEV#  P  R
Snap2          S-VOL PAIR 85000098 40960 1012  258  2 100
A---           -      Y         - - D
```

Description of each column in output example:

SnapShot_name

      Displays the name of snapshot group defined in the device.

P/S

      Displays the attribute of the target LDEV. It displays P-VOL for the P-VOL and S-VOL for the S-VOL. In the list of snapshot, "-" is displayed.

STAT

      Displays the following status of each snapshot data:

- SMPP: The pair is being deleted.

- PSUP: The pair is being suspended.

- CPYD: The differential data is being copied.

- Other statuses: See **pairdisplay** command.

Serial#

      Displays the product serial number.

LDEV#

      Displays the LDEV number related to the snapshot.

MU#

      Displays the Mirror ID of the P-VOL for the snapshot.

P-LDEV#

      Displays the LDEV number of the volume (P-VOL or S-VOL) of the pair associated with the snapshot data. If the LDEV that is paired belongs to a different virtual storage machine, "----" is displayed.

PID

      Displays the pool ID.

%

When MODE is clone or cascade mode, displays one of the following according to the status of STAT:

- When STAT is COPY, CPYP, RCPY, SMPP, or PSUP: Displays the progress rate of each processing.

- For other statuses: Displays the concordance rate between the P-VOL and the S-VOL of a pair.

  When MODE is other than clone or cascade mode, the concordance rate between the P-VOL and the S-VOL is always displayed regardless of the STAT status.

  When MODE is R that indicates snapshot data of a Thin Image Advanced pair, a hyphen (-) is always displayed in the concordance rate.

The following shows the concordance rate for each pair status.

- PSUS: 0% to 100%
- RCPY: 0% to 100%[1]
- PAIR: 100%
- COPY: 0% to 100%[2]
- CPYD: 0% to 100%

> 📄 **Note:**
>
> 1. Displays the concordance rate before changing to RCPY status. The displayed value might decrease during RCPY status because the copy works.
> 2. Displays the concordance rate before changing to COPY status. The displayed value does not change during COPY status.

MODE

Displays the status of snapshot data:

- G: The snapshot data created in CTG mode.

- W: The status when the data are written in the secondary volume from the host in the PSUS/PFUS status.

- C: The snapshot data created in clone mode.

- A: The snapshot data created in cascade mode.

- D: The snapshot data created in Diff Clone mode that is used in the VASA environment.

- R: The snapshot data of a Thin Image Advanced pair.

SPLT-TIME

Displays the time when a snapshot data is created. When `-format_time` option is specified, displays in YYYY-MM-DDThh:mm:ss format. When `-format_time` option is not specified, indicates accumulated time in seconds from January 1, 1970 (GMT).

Chapter 5: Configuration setting commands

SLU

Displays whether the SLU attribute is set for the snapshot:

- Y: The SLU attribute is set.
- N: Non-SLU attribute is set.

    A hyphen (-) is displayed for unsupported DKCMAIN microcode version.

C_LDEV#

Displays the LDEV number of the diff compare volume (base volume), which is used for comparing the difference between the Diff Clone pair volumes in the VASA environment.

A hyphen (-) is displayed for unsupported microcode version.

P

Displays the status of expansion processing of DP-VOL (VSP 5000 series, VSP G1x00 and VSP F1500, VSP E series, and VSP G130, G/F350, G/F370, G/F700, G/F900 only).

- E: Expansion is in progress.
- N: Not processed.
- - (hyphen): This information is not available.

A hyphen (-) is displayed if the microcode version is not supported.

R

Displays the status of the Read Only attribute of the snapshot data (VSP 5000 series, VSP E series, and VSP G130, G/F350, G/F370, G/F700, G/F900 only).

- E: Read Only attribute is enabled.
- D: Read Only attribute is disabled.
- - (hyphen): This information is not available.

A hyphen (-) is displayed if the microcode version is not supported.

raidcom replace snapshot

Replaces the snapshot data that is mapped to the S-VOL.

For a Thin Image Advanced pair, the pair status changes to CPYP while the specified operation is in process, and then changes to PSUS when the processing is complete. To verify that the operation is complete, run the raidcom get snapshot command to verify that STAT has changed from CPYP to the status other than CPYP.

Syntax

```
raidcom replace snapshot -ldev_id <ldev#> {-snapshotgroup <name> | -mirror_id <mu#> }
```

Options and parameters

-ldev_id <ldev#>

Specifies the LDEV number of the S-VOL to be replaced.

-snapshotgroup <name>

Specifies the name of the snapshot group in which the snapshot data as the operation target is included.

One MU which corresponds to the snapshot group which is specified by the P-VOL, corresponding to the specified S-VOL, is mapped.

> 📄 **Note:**
>
> Because the MU is selected automatically, an unexpected MU could be mapped. So if you intend to specify the specific snapshot data, specify the snapshot data by an MU number instead of a snapshot group.

-mirror_id <mu#>

Specifies the mirror ID of the specified snapshot data.

The specified snapshot data is mapped to the S-VOL.

Examples

Replacing the snapshot data of S-VOL (LDEV number 20:00) to the snapshot group snap3.

```
# raidcom replace snapshot -ldev_id 0x2000 -snapshotgroup snap3
```

raidcom add spm_wwn

Specifies the Server Priority Manager name for preferred/non-preferred WWNs.

Syntax

```
raidcom add spm_wwn -port <port#> -spm_name <nick_name> -hba_wwn <wwn_strings>
```

Options and parameters

-port <port#>

Specifies the port number with the Target attribute, for example:

- CL1-A

-spm_name <nick_name>

Specifies the SPM name.

Up to 64 characters can be specified by CLI.

SPM names are managed uniquely in the entire system.

-hba_wwn <wwn_strings>

Specifies the monitored WWN.

Chapter 5: Configuration setting commands

Before setting the SPM name, the WWN must be registered as preferred or non-preferred.

The WWN is specified as a 16-digit hexadecimal value. The 17th digit and later are ignored. The value can be split in units of 4 bytes by "," (comma). For example:

- 210000e08b0256f8
- 210000e0,8b0256f8

Examples

Specifies the SPM name (WWN_NICK_LINUX) to WWN (50060e8005fa0f36).

```
# raidcom add spm_wwn -port CL4-E -spm_name WWN_NICK_LINUX -hba_wwn 50060e80,05fa0f36
```

raidcom delete spm_wwn

Deletes WWN from the Server Priority Manager targets.

Syntax

```
raidcom delete spm_wwn -port <port#> [-hba_wwn <wwn_string> | -spm_name <nick_name>]
```

Options and parameters

-port <port#>

Specifies the port number whose attribute is Target. For example:

- CL1-A

-hba_wwn <wwn_string>

Specifies the WWN to be deleted.

The WWN is specified as a 16-digit hexadecimal value. The 17th digit and later are ignored. The value can be split in units of 4 bytes by "," (comma). For example:

- 210000e08b0256f8
- 210000e0,8b0256f8

-spm_name <nick_name>

Specifies the SPM name to be deleted.

Up to 64 characters can be specified by CLI.

Examples

Deletes the SPM name (WWN_NICK_LINUX) from the SPM targets.

```
# raidcom delete spm_wwn -port CL4-E -spm_name WWN_NICK_LINUX
```

Deletes the WWN (50060e8005fa0f36) from the SPM targets.

```
# raidcom delete spm_wwn -port CL4-E -hba_wwn 50060e80,05fa0f36
```

raidcom modify spm_wwn

Specifies the Server Priority Manager information to the Server Priority Manager target WWN.

> 📄 **Note:** For VSP 5000 series, VSP G1x00, VSP F1500, VSP E series, VSP G130, G/F350, G/F370, G/F700, G/F900, VSP G200, G400, G600, G800 and VSP F400, F600, F800, this command cannot be executed if SPM information is already set for the storage system by the **raidcom modify spm_ldev command**.

Syntax

```
raidcom modify spm_wwn -port <port#> [-spm_priority <y/n>]
{-limit_io | -limit_kb | -limit_mb } <value>
{-hba_wwn <wwn_strings> | -spm_name <nick_name>}
```

Options and parameters

-port <port#>

Specifies the port number whose attribute is Target. For example:

- CL1-A

-spm_priority <y/n>

Specifies preferred/non-preferred WWN.

- y: preferred WWN

- n: non-preferred WWN

{-limit_io | -limit_kb | -limit_mb } <value>

Specifies maximum value/threshold value by the I/O rate or the transmission rate.

- -limit_io: I/O rate, maximum value: 2,147,483,647 [IOPS]

- -limit_kb: transmission rate (unit of KB), maximum value: 2,147,483,647 [KB]

- -limit_mb: transmission rate (unit of MB), maximum value: 2,097,151 [MB]

  (If the value is specified by MB, 1 MB is calculated as 1024 KB.)

The threshold value for the prioritized WWN is the same as the threshold value for the entire system.

-hba_wwn <wwn_strings>

Specifies WWN that you set the SPM information.

The WWN is specified as a 16-digit hexadecimal value. The 17th digit and later are ignored. The value can be split in units of 4 bytes by "," (comma). For example:

- 210000e08b0256f8
- 210000e0,8b0256f8

-spm_name <nick_name>

Specifies the SPM name.

Up to 64 characters can be specified by CLI.

Examples

Specifies WWN (50060e8005fa0f36) as non-preferred WWN, and specifies 5000/[I/O] (I/O rate) as the maximum value.

```
# raidcom modify spm_wwn -port CL4-E -spm_priority n -limit_io 5000 -hba_wwn 50060e80,
05fa0f36
```

Specifies the SPM name (WWN_NICK_LINUX) as non-preferred WWN, and specifies 500[MB/s] (transmission rate) as the maximum value.

```
# raidcom modify spm_wwn -port CL4-E -spm_priority n -limit_mb 500 -spm_name
WWN_NICK_LINUX
```

Specifies WWN (50060e8005fa0f36) as preferred WWN, and specifies 5000/[I/O] (I/O rate) as the threshold value.

```
# raidcom modify spm_wwn -port CL4-E -spm_priority y -limit_io 5000 -hba_wwn 50060e80,
05fa0f36
```

Specifies the SPM name (WWN_NICK_LINUX) as preferred WWN, and specifies 500000[KB/s] (transmission rate) as the maximum value.

```
# raidcom modify spm_wwn -port CL4-E -spm_priority y -limit_kb 500000 -spm_name
WWN_NICK_LINUX
```

Deletes the SPM information of WWN (50060e8005fa0f36).

```
# raidcom modify spm_wwn -port CL4-E -hba_wwn 50060e80,05fa0f36
```

Deletes the SPM information of the SPM name (WWN_NICK_LINUX).

```
# raidcom modify spm_wwn -port CL4-E -spm_name WWN_NICK_LINUX
```

raidcom get spm_wwn

Gets the Server Priority Manager information of the Server Priority Manager target WWN.

Chapter 5: Configuration setting commands

Syntax

```
raidcom get spm_wwn -port <port#> [-hba_wwn <wwn_strings>
     | -spm_name <nick_name>
```

Options and parameters

-port <port#>

Specifies the port number whose attribute is Target. For example:

- CL1-A

-hba_wwn <wwn_strings>

Specifies WWN that you get the SPM information.

The WWN is specified as a 16-digit hexadecimal value. The 17th digit and later are ignored. The value can be split in units of 4 bytes by "," (comma). For example:

- 210000e08b0256f8

- 210000e0,8b0256f8

-spm_name <nick_name>

Specifies the SPM name that you get the SPM information.

Up to 64 characters can be specified by CLI.

Examples

Gets the SPM information of WWN that belongs to the specified port (CL4-E).

```
#raidcom get spm_wwn -port CL4-E
PORT   SPM_MD              SPM_WWN NICK_NAME  GRP_NAME    Serial#
CL4-E   WWN 210000e08b0256f8 WWN_NICK_LINUX_0 OLA_NODE0_CTL 63528
CL4-E   WWN 210000e08b0256f7 WWN_NICK_LINUX_1 OLA_NODE0_CTL 63528
```

Specifies WWN (50060e8005fa0f36) to get the SPM information.

```
# raidcom get spm_wwn -port CL4-E -hba_wwn 50060e80,05fa0f36
```

```
PORT  SPM_MD PRI IOps KBps Serial#
CL4-E WWN    Y 5000    - 63528
```

Specifies the SPM name (WWN_NICK_LINUX) to get the SPM information.

```
# raidcom get spm_wwn -port CL4-E -spm_name WWN_NICK_LINUX
```

```
PORT  SPM_MD PRI IOps KBps Serial#
CL4-E WWN    Y    - 5000 63528
```

Chapter 5: Configuration setting commands

Description of each column in output example:

PORT

Displays the port to which the WWN is set.

SPM_MD

Displays the Control mode of the SPM.

- WWN: WWN control

- PORT: Port control

SPM_WWN

Displays the set WWN.

NICK_NAME

Displays the SPM name set to the WWN. If the SPM name is not set, a hyphen (-) is displayed.

GRP_NAME

Displays the SPM group name to which the WWN belongs. If the WWN does not belong to the group, a hyphen (-) is displayed.

Serial#

Displays the Seq#.

PRI

Displays the settings of the priority.

- Y: Preferred

- N: Non-preferred

IOps

If the PRI setting is Y, it displays the threshold. If the PRI setting is N, it displays the maximum value. If the value is set by the transferred rate (KBps), it displays a hyphen (-).

KBps

If the PRI setting is Y, it displays the threshold. If the PRI setting is N, it displays the maximum value. If the value is set by the IO rate (IOPS), it displays a hyphen (-). If the value is set by MB, 1 MB is calculated as 1024 KB, and displays the value by KB.

raidcom monitor spm_wwn

Gets the monitoring information of Server Priority Manager target WWN.

Syntax

```
raidcom monitor spm_wwn {-hba_wwn <wwn_strings> | -spm_name <nick_name>}
```

Options and parameters

-hba_wwn <wwn_strings>

Specifies WWN that you get the monitoring information.

The WWN is specified as a 16-digit hexadecimal value. The 17th digit and later are ignored. The value can be split in units of 4 bytes by "," (comma). For example:

- 210000e08b0256f8

- 210000e0,8b0256f8

-spm_name <nick_name>

Specifies the SPM name that you get the monitoring information.

Up to 64 characters can be specified by CLI.

Examples

Specifies WWN (50060e8005fa0f36) to get the monitoring information.

```
# raidcom monitor spm_wwn -hba_wwn 50060e80,05fa0f36
```

```
PORT  SPM_MD IOps    KBps Serial#
CL4-E PORT   5000 5000000 63528
```

Specifies the SPM name (WWN_NICK_LINUX) to get the monitoring information.

```
# raidcom monitor spm_wwn -spm_name WWN_NICK_LINUX
```

```
PORT  SPM_MD IOps    KBps Serial#
CL4-E PORT   5000 5000000 63528
```

Description of each column in output example:

PORT

Displays the port to which the WWN is set.

SPM_MD

Displays the Control mode of the SPM.

- WWN: WWN control

- PORT: Port control

IOps

Displays the current IO rate (IOPS) of the specified WWN of or the specified SPM name.

Chapter 5: Configuration setting commands

KBps

Displays the current transferred rate (KBps) of the specified WWN or the specified SPM name.

Serial#

Displays the Seq#.

raidcom add spm_group

Registers Server Priority Manager target WWN to Server Priority Manager group. For VSP 5000 series, VSP E series, VSP G130, G/F350, G/F370, G/F700, G/F900, VSP G200, G400, G600, G800 and VSP F400, F600, F800, this command cannot be executed if SPM information is already set for the storage system by the `raidcom modify spm_ldev` command.

Syntax

For registering the WWN to the SPM group:

```
raidcom add spm_group -port <port#> -spm_group <group_name>
   {<nick_name> | -hba_wwn <wwn_strings>}
```

(VSP and VSP G1x00 and VSP F1500 only) For associating the host group with the SPM group:

```
raidcom add spm_group -spm_group <group name> -port <port#>
   [<host group name>] -spm_host_grp
```

Options and parameters

-port <port#>

Specifies the port number whose attribute is Target. For example:

- CL1-A

-spm_group <group_name> [<nick_name>]

Supported storage systems:

- VSP 5000 series

- VSP E series

- VSP G130, G/F350, G/F370, G/F700, G/F900

- VSP G200, G400, G600, G800 and VSP F400, F600, F800

Specifies the SPM group name and the SPM name.

Up to 64 characters can be specified by CLI. If <nick_name> is omitted, -hba_wwn is required. The SPM group name must be unique across the system.

-spm_group <group_name>

Specifies the SPM group name.

Up to 64 characters can be specified by CLI.

<nick_name>

Specifies the SPM target WWN using the SPM name.

Up to 64 characters can be specified by CLI.

-hba_wwn <wwn_strings>

Specifies SPM target WWN.

The WWN is specified as a 16-digit hexadecimal value. The 17th digit and later are ignored. The value can be split in units of 4 bytes by "," (comma). For example:

- 210000e08b0256f8
- 210000e0,8b0256f8

-port <port#> [<host group name>] -spm_host_grp

VSP and VSP G1x00 and VSP F1500 only. When registering the WWN registered in the host group to the SPM group, specifies the port number and the host group ID, or the host group name. If you register the WWN by this method, the host group is associated with the SPM group. When adding or deleting the WWN to or from the host group, the WWN is also added or deleted to or from the SPM group. Also, if you delete a host group, all WWNs registered in the host group are deleted from the SPM group.

Examples:

- CL1-A-*g*, where g is a value from 0 to 254.
- CL1-A Linux_X86

Examples

Registers WWN (50060e8005fa0f36) as the SPM group name (WWN_GRP_LINUX).

```
# raidcom add spm_group -port CL4-E -spm_group WWN_GRP_LINUX
-hba_wwn 50060e80,05fa0f36
```

Registers the SPM name (WWN_NICK_LINUX) as the SPM group name (WWN_GRP_LINUX).

```
# raidcom add spm_group -port CL4-E -spm_group WWN_GRP_LINUX WWN_NICK_LINUX
```

(VSP and VSP G1x00 and VSP F1500 only) Registers the host group (CL4-E-1) as the SPM group name (WWN_GRP_LINUX).

```
# raidcom add spm_group -port CL4-E-1 -spm_host_grp
-spm_group WWN_GRP_LINUX
```

Command Control Interface Command Reference 602

raidcom delete spm_group

Deletes the WWN of the specified port or the host group from the Server Priority Manager group to delete the SPM information of the WWN. If no WWN is registered in the SPM group, the SPM group itself is also deleted.

Syntax

```
raidcom delete spm_group -port <port#> -spm_group <group_name>
```

VSP and VSP G1x00 and VSP F1500 only:

```
raidcom delete spm_group {-port <port#> -spm_group <group_name> |
   -port <port#> [<host group name>] -spm_host_grp}
```

Options and parameters

-port <port#>
Specifies the port number whose attribute is Target. For example:

- CL1-A

-spm_group <group_name>
Specifies the SPM group name.

Up to 64 characters can be specified by CLI.

-port <port#> [<host group name>] -spm_host_grp
VSP and VSP G1x00 and VSP F1500 only. When setting SPM information for the SPM group associated with the host group, specifies the port number and the host group ID, or the host group name.

Examples:

- CL1-A-*g*, where g is a value from 0 to 254.
- CL1-A Linux_X86

Examples

Deletes the SPM group (WWN_GRP_LINUX).

```
# raidcom delete spm_group -port CL4-E -spm_group WWN_GRP_LINUX
```

(VSP and VSP G1x00 and VSP F1500 only) Deletes SPM associated with the host group (CL4-E-1).

```
# raidcom delete spm_group -port CL4-E-1 -spm_host_grp
```

Chapter 5: Configuration setting commands

raidcom modify spm_group

Specifies the Server Priority Manager information to the Server Priority Manager target group.

Syntax

```
raidcom modify spm_group -port <port#> -spm_group <group_name> [-spm_priority <y/n>]
   {-limit_io | -limit_kb | -limit_mb } <value>
```

VSP and VSP G1x00 and VSP F1500 only:

```
raidcom modify spm_group {-port <port#> -spm_group <group_name> |
   -port <port#> [<host group name>] -spm_host_grp} [-spm_priority <y/n>]
   {-limit_io | -limit_kb | -limit_mb} <value>
```

Options and parameters

-port <port#>

Specifies the port number whose attribute is Target. For example:

- CL1-A

-port <port#> [<host group name>] -spm_host_grp

VSP and VSP G1x00 and VSP F1500 only. When you specify SPM information for the SPM group associated with the host group, specify the port number, host group ID, or host group name.

Examples:

- CL1-A-*g*, where g is a value from 0 to 254.

- CL1-A Linux_X86

-spm_priority <y/n>

Specifies preferred/non-preferred WWN.

- y: preferred WWN

- n: non-preferred WWN

{-limit_io | -limit_kb | -limit_mb } <value>

Specifies maximum value/threshold value by the I/O rate or the transmission rate.

- -limit_io: I/O rate, maximum value: 2,147,483,647 [IOPS]

- -limit_kb: transmission rate (unit of KB), maximum value: 2,147,483,647 [KB]

- -limit_mb: transmission rate (unit of MB), maximum value: 2,097,151 [MB]

  If you specify the value per MB, 1 MB is calculated as 1024 KB.

The threshold value for the prioritized WWN is the same as the threshold value for the entire system.

-spm_group <group_name>

Specifies the SPM group name that you set the SPM information.

Up to 64 characters can be specified by CLI.

Examples

Specifies the SPM group name (WWN_GRP_LINUX) as non-preferred WWN, and specifies 5000[I/O] (I/O rate) as the maximum value.

```
# raidcom modify spm_group -port CL4-E -spm_priority n -limit_io 5000 -spm_group
WWN_GRP_LINUX
```

Specifies the SPM group name (WWN_GRP_LINUX) as non-preferred WWN, and specifies 500[MB/s] (transmission rate) as the maximum value.

```
# raidcom modify spm_group -port CL4-E -spm_priority n -limit_mb 500 -spm_group
WWN_GRP_LINUX
```

Specifies the SPM group name (WWN_GRP_LINUX) as preferred WWN, and specifies 5000[I/O] (I/O rate) as the threshold value.

```
# raidcom modify spm_group -port CL4-E -spm_priority y -limit_io 5000 -spm_group
WWN_GRP_LINUX
```

Specifies the SPM group name (WWN_GRP_LINUX) as preferred WWN, and specifies 500000[KB/s] (transmission rate) as the threshold value.

```
# raidcom modify spm_group -port CL4-E -spm_priority y -limit_kb 500000 -spm_group
WWN_GRP_LINUX
```

Deletes the SPM information of the SPM group name (WWN_GRP_LINUX).

```
# raidcom modify spm_group -port CL4-E -spm_group WWN_GRP_LINUX
```

(VSP and VSP G1x00 and VSP F1500 only) Specifies the SPM group associated with the host group (CL4-E-1) as non-preferred WWN, and specifies 5000[I/O] (I/O rate) as the maximum value.

```
# raidcom modify spm_group -port CL4-E-1 -spm_host_grp -spm_priority n
-limit_io 5000
```

raidcom get spm_group

Gets the Server Priority Manager information of the Server Priority Manager target WWN in the specified port by the Server Priority Manager group unit.

Syntax

```
raidcom get spm_group -port <port#> -spm_group <group_name>
```

VSP and VSP G1x00 and VSP F1500 only:

```
raidcom get spm_group {-port <port#> -spm_group <group_name> |
   -port <port#> [<host group name>] -spm_host_grp}
```

Options and parameters

-port <port#>

Specifies the port number whose attribute is Target. For example:

- CL1-A

-spm_group <group_name>

Specifies the SPM group name.

Up to 64 characters can be specified by CLI.

-port <port#> [<host group name>] -spm_host_grp

VSP and VSP G1x00 and VSP F1500 only. When obtaining SPM information of the SPM group associated with the host group, specifies the port number and the host group ID, or the host group name.

Examples:

- CL1-A-*g*, where g is a value from 0 to 254.
- CL1-A Linux_X86

Examples

Specifies the port (CL4-E) and the SPM group name (WWN_GRP_LINUX) to get the SPM information.

```
# raidcom get spm_group -port CL4-E -spm_group WWN_GRP_LINUX
```

```
PORT  SPM_MD PRI IOps KBps Serial#
CL4-E PORT     Y 5000    - 63528
```

(VSP and VSP G1x00 and VSP F1500 only) Specifies the host group (CL4-E-1) to obtain SPM information.

```
# raidcom get spm_group -port CL4-E-1 -spm_host_grp
```

```
PORT  SPM_MD PRI      IOps KBps  Serial#
CL4-E   WWN   Y 1000000000   -    63528
```

Chapter 5: Configuration setting commands

Description of each column in output example:

PORT

Displays the port to which the WWN is set.

SPM_MD

Displays the Control mode of the SPM.

- WWN: WWN control

- PORT: Port control

PRI

Displays the settings of the priority.

- Y: Preferred

- N: Non-preferred

IOps

If the PRI setting is Y, it displays the threshold. If the PRI setting is N, it displays the maximum value. If the value is set by the transferred rate (KBps), it displays a hyphen (-).

KBps

If the PRI setting is Y, it displays the threshold. If the PRI setting is N, it displays the maximum value. If the value is set by the IO rate (IOPS), it displays a hyphen (-). If the value is set by MB, 1 MB is calculated as 1024 KB, and displays the value by KB.

Serial#

Displays the Seq#.

raidcom monitor spm_group

Gets the monitoring information of Server Priority Manager target WWN by the Server Priority Manager group unit.

Syntax

```
raidcom monitor spm_group -spm_group <group_name>
```

Options and parameters

-spm_group <group_name>

Specifies the SPM group name.

Up to 64 characters can be specified by CLI.

Examples

Specifies the SPM group name (WWN_GRP_LINUX) to get the monitoring information.

```
# raidcom monitor spm_group -spm_group WWN_GRP_LINUX
```

```
PORT  SPM_MD IOps KBps    Serial#
CL4-E PORT   5000 5000000 63528
```

Description of each column in output example:

PORT

Displays the port to which the WWN is set.

SPM_MD

Displays the Control mode of the SPM.

- WWN: WWN control

- PORT: Port control

IOps

Displays the current IO rate (IOPS) of the specified WWN or the specified SPM name.

KBps

Displays the current transferred rate (KBps) of the specified WWN or the specified SPM name.

Serial#

Displays the Seq#.

raidcom modify spm_ldev

Supported storage systems:

- VSP 5000 series

- VSP G1x00 and VSP F1500

- VSP E series

- VSP G130, G/F350, G/F370, G/F700, G/F900

- VSP G200, G400, G600, G800 and VSP F400, F600, F800

Sets the Server Priority Manager information for a combination of an LDEV and WWN, or an LDEV and iSCSI name.

> 📄 **Note:**
>
> - This command cannot be executed if SPM information is already set for the storage system by the raidcom modify spm_wwn or raidcom add spm_group command.
>
> - When the number of LDEVs with SPM configured exceeds 4,096, I/O response for the 4,097th and subsequent LDEVs is slower than for the other 4,096 LDEVs. To improve host I/O response, delete the existing SPM information so that the total number of LDEVS with SPM configured will be 4,096 or fewer.

Syntax

```
raidcom modify spm_ldev -ldev_id
   <ldev#>  {-hba_wwn <wwn strings>
   | -hba_iscsi_name <initiator iscsi name>}
   [-spm_priority {y/n}] [{-limit_io | -limit_mb} <value>]
```

Options and parameters

-ldev_id <ldev#>

Specifies the LDEV number (0 to 65279) whose SPM information you want to set.

Example: -ldev_id 200

-hba_wwn <wwn_strings>

Specifies the WWN of the host bus adapter (initiator) for which you set the SPM information.

The WWN is specified as a 16-digit hexadecimal value. The 17th digit and later are ignored. The value can be split in units of 4 bytes by "," (comma). For example:

- 210000e08b0256f8

- 210000e0,8b0256f8

-hba_iscsi_name <initiator iscsi name>

Specifies the iSCSI name of the host bus adapter (initiator) for which you set the SPM information. A maximum of 223 characters can be specified.

Example:

- iqn.win2k8.example.of.iqn.form

- eui.0123456789ABCDEF

-spm_priority <y/n>

Specifies the priority (prioritized or non-prioritized) for the combination of an LDEV and WWN, or an LDEV and iSCSI name.

- y: prioritized. Obtains monitoring information, and the upper limit is not set. This combination is prioritized.

- n: non-prioritized. Obtains monitoring information, but the upper limit is set. This combination is not prioritized.

Non-prioritized is set if omitted. If y is set, only monitoring is performed.

For details about operations for setting SPM with LDEVs and WWNs, or LDEVs and iSCSI names, see the *Command Control Interface User and Reference Guide*.

{-limit_io | -limit_mb } <value>

Specifies the upper limit for the I/O rate or the transmission rate. If omitted, the I/O rate or the transmission rate is not suppressed by the upper limit.

- -limit_io: I/O rate

  Maximum value: 65,535 [IOPS]

- -limit_mb: Transmission rate (in MB)

  Maximum value: 31 [MB]

Examples

The following shows an example of setting n (non-prioritized) for the combination of LDEV: 1024 and WWN: 50060e8005fa0f36, and I/O rate: 5000 [IOPS] as the upper limit:

```
# raidcom modify spm_ldev -ldev_id 1024 -hba_wwn 50060e80,05fa0f36 -spm_priority n -
limit_io 5000
```

The following shows an example of setting n (non-prioritized) for the combination of LDEV: 1024 and iSCSI name: iqn.z1, and I/O rate: 5000 [IOPS] as the upper limit:

```
# raidcom modify spm_ldev -ldev_id 1024 -hba_iscsi_name iqn.z1 - spm_priority n -
limit_io 5000
```

raidcom delete spm_ldev

Supported storage systems:

- VSP 5000 series
- VSP G1x00 and VSP F1500
- VSP E series
- VSP G130, G/F350, G/F370, G/F700, G/F900
- VSP G200, G400, G600, G800 and VSP F400, F600, F800

Deletes a combination of the SPM target LDEV and WWN, or the SPM target LDEV and iSCSI name from the SPM target.

Syntax

```
raidcom delete spm_ldev -ldev_id <ldev#>  {-hba_wwn <wwn strings>
   | - hba_iscsi_name <initiator iscsi name>}
```

Options and parameters

-ldev_id <ldev#>

Specifies the LDEV number (0 to 65279) whose SPM information you want to delete.

Example: -ldev_id 200

-hba_wwn <wwn_strings>

Specifies the WWN of the host bus adapter (initiator) whose SPM information you want to delete.

The WWN is specified as a 16-digit hexadecimal value. The 17th digit and later are ignored. The value can be split in units of 4 bytes by "," (comma). For example:

- 210000e08b0256f8

- 210000e0,8b0256f8

-hba_iscsi_name <initiator iscsi name>

Specifies the iSCSI name of the host bus adapter (initiator) whose SPM information you want to delete. A maximum of 223 characters can be specified.

Example:

- iqn.win2k8.example.of.iqn.form

- eui.0123456789ABCDEF

Examples

The following shows an example of deleting the combination of LDEV: 1024 and WWN: 50060e8005fa0f36 from the SPM target:

```
# raidcom delete spm_ldev -ldev_id 1024 -hba_wwn 50060e80,05fa0f36
```

The following shows an example of deleting the combination of LDEV: 1024 and iSCSI name: iqn.z1 from the SPM target:

```
# raidcom delete spm_ldev -ldev_id 1024 -hba_iscsi_name iqn.z1
```

raidcom monitor spm_ldev

Supported storage systems:

- VSP 5000 series

- VSP G1x00 and VSP F1500

- VSP E series

- VSP G130, G/F350, G/F370, G/F700, G/F900

- VSP G200, G400, G600, G800 and VSP F400, F600, F800

Chapter 5: Configuration setting commands

Obtains the monitoring information for the combination of an LDEV and WWN, or an LDEV and iSCSI name. This command displays information when both the following conditions are met:

- SPM information is set for the specified LDEV.

- The current I/O rate or transfer rate for the LDEV with the specified WWN or iSCSI name is not 0.

Syntax

```
raidcom monitor spm_ldev -ldev_id <ldev#> {-hba_wwn <wwn strings>
   | - hba_iscsi_name <initiator iscsi name>}
```

Options and parameters

-ldev_id <ldev#>

Specifies the LDEV number (0 to 65279) whose monitoring information you want to obtain.

Example: -ldev_id 200

-hba_wwn <wwn_strings>

Specifies the WWN of the host bus adapter (initiator) whose monitoring information you want to obtain.

The WWN is specified as a 16-digit hexadecimal value. The 17th digit and later are ignored. The value can be split in units of 4 bytes by "," (comma). For example:

- 210000e08b0256f8

- 210000e0,8b0256f8

-hba_iscsi_name <initiator iscsi name>

Specifies the iSCSI name of the host bus adapter (initiator) whose monitoring information you want to obtain. A maximum of 223 characters can be specified.

Example:

- iqn.win2k8.example.of.iqn.form

- eui.0123456789ABCDEF

Examples

The following shows an example of obtaining monitoring information for the combination of LDEV: 1024 and WWN: 50060e8005fa0f36:

```
# raidcom monitor spm_ldev -ldev_id 1024 -hba_wwn 50060e80,05fa0f36
```

```
Serial#  LDEV  IOps     KBps    WWN
  63528  1024  5000  5000000  50060e8005fa0f36
```

Chapter 5: Configuration setting commands

The following shows an example of obtaining monitoring information for the combination of LDEV: 1024 and iSCSI name: iqn.z1:

```
# raidcom monitor spm_ldev -ldev_id 1024 -hba_iscsi_name iqn.z1
Serial#  LDEV  IOps     KBps   IQN
  63528  1024  5000  5000000  iqn.z1
```

Description of each column in output example:

LDEV
>   Displays the LDEV number.

IOps
>   Displays the current I/O rate (IOPS) of the LDEV identified by the specified WWN or iSCSI name.

KBps
>   Displays the current transfer rate (KBps) for the LDEV identified by the specified WWN or iSCSI name.

WWN
>   Displays the WWN of the host bus adapter (initiator).

IQN
>   Displays the iSCSI name of the host bus adapter (initiator).

raidcom get spm_ldev

Supported storage systems:

- VSP 5000 series
- VSP G1x00 and VSP F1500
- VSP E series
- VSP G130, G/F350, G/F370, G/F700, G/F900
- VSP G200, G400, G600, G800 and VSP F400, F600, F800

Displays the Server Priority Manager information for a combination of an LDEV and WWN or an LDEV and iSCSI name.

Syntax

```
raidcom get spm_ldev -ldev_id <ldev#>  {-hba_wwn <wwn strings>
   | -hba_iscsi_name <initiator iscsi name>}
```

Options and parameters

-ldev_id <ldev#>
>   Specifies the LDEV number (0 to 65279) whose SPM information you want to obtain.

Chapter 5: Configuration setting commands

Example: -ldev_id 200

This command terminates without displaying information if SPM information is not set for the combination of an WWN and the specified LDEV, or an iSCSI name and the specified LDEV.

-hba_wwn <wwn_strings> |

Specifies the WWN of the host bus adapter (initiator) whose SPM information you want to obtain.

The WWN is specified as a 16-digit hexadecimal value. The 17th digit and later are ignored. The value can be split in units of 4 bytes by "," (comma). For example:

- 210000e08b0256f8

- 210000e0,8b0256f8

This command terminates without displaying information if SPM information is not set for the combination of an LDEV and the specified WWN.

-hba_iscsi_name <initiator iscsi name>

Specifies the iSCSI name of the host bus adapter (initiator) whose SPM information you want to obtain. A maximum of 223 characters can be specified.

Example:

- iqn.win2k8.example.of.iqn.form

- eui.0123456789ABCDEF

This command terminates without displaying information if SPM information is not set for the combination of an LDEV and the specified iSCSI name.

Examples

The following shows an example of displaying SPM information for the combination of LDEV: 1024 and WWN, and LDEV: 1024 and iSCSI name:

```
# raidcom get spm_ldev -ldev_id 1024
```

```
Serial# LDEV PRI IOps  KBps    T Name
63528  1024    Y      -        -   W 50060e8005fa0f3
63528  1024    N      - 31744   I iqn.z1
```

The following shows an example of displaying SPM information for the combination of an LDEV and iSCSI name: iqn.z1:

```
# raidcom get spm_ldev -hba_iscsi_name iqn.z1
Serial# LDEV PRI IOps KBps  T Name
```

```
63528  1024    Y      -     -    I iqn.z1
63528  1025    N      - 31744 I iqn.z1
```

Chapter 5: Configuration setting commands

The following shows an example of displaying SPM information for the combination of an LDEV and WWN: 50060e8005fa0f36:

```
# raidcom get spm_ldev -hba_wwn 50060e8005fa0f3
```

```
Serial# LDEV PRI IOps KBps  T Name
63528  1024   Y     -     -   W 50060e8005fa0f3
63528  1025   N     - 31744 W 50060e8005fa0f3
```

The following shows an example of displaying SPM information for the combination of an LDEV and WWN, and an LDEV and iSCSI name:

```
# raidcom get spm_ldev
```

```
Serial# LDEV PRI IOps  KBps     T      Name
63528 1024    Y     -      -       W     50060e8005fa0f3
63528 1024    N     - 31744      I      iqn.z1
63528 1025    N 5000      -      I      iqn.z1
```

Description of each column in output example:

Serial#
> Displays the product serial number.

LDEV
> Displays the LDEV number.

PRI
> Displays the priority.
> - Y: Prioritized
> - N: Non-prioritized

IOps
> For the non-prioritized setting, displays the upper limit. For the prioritized setting, displays a hyphen (-) if the transfer rate is set.

KBps
> Displays the upper limit for the non-prioritized setting. For the prioritized setting, displays a hyphen (-) if the I/O rate is set. In addition, if the value is set in megabytes, it is converted to kilobytes by using the following formula: 1 MB = 1024 KB

T
> Displays the type of name to be displayed in the Name column:
> - W: WWN
> - I: iSCSI name

Chapter 5: Configuration setting commands

Name
Displays the WWN or iSCSI name.

raidcom add dp_pool

Create a pool for Dynamic Provisioning or Dynamic Provisioning for Mainframe by the specified resource. Or, add pool volumes to the pools for Dynamic Provisioning, Dynamic Provisioning for Mainframe, Dynamic Tiering, Dynamic Tiering for Mainframe, active flash, or active flash for mainframe. A parity group, an LDEV, and a device group can be specified as a resource.

> 📄 **Note:** If a failure occurs in a processor on an MP blade or a controller, this command cannot be run even on an MP blade or a controller on which no failure occurs.
>
> If a cache memory, HIE, or a logical path inside the storage is blocked, the operation cannot be performed.

> 💡 **Tip:** For VSP E series, VSP G130, G/F350, G/F370, G/F700, G/F900, specifying a parity group is recommended. When a parity group is specified, an LDEV is created and a pool volume is added. It is not necessary to create an LDEV in advance. In addition, you cannot add an existing LDEV as a pool volume.

When specifying a pool that is already created for Dynamic Provisioning, Dynamic Provisioning for Mainframe, Dynamic Tiering, Dynamic Tiering for Mainframe, active flash, or active flash for mainframe, the specified LDEV is added as a pool volume.

Specify either one of Pool ID or Pool Name certainly. If both the Pool ID and Pool Name options are omitted, this command is rejected with EX_REQARG.

This command is executed asynchronously with the command input. Check the completion of this process on the **`raidcom get command_status`** command.

Syntax

```
raidcom add dp_pool {-pool_id <pool ID#> | -pool_name <pool
   naming> | -pool_id <pool ID#>
   -pool_name <pool naming>}
  {-parity_grp_id <gnosgno> [-resource_id <resource group id >]
  | -ldev_id <ldev#> …[-cnt <count>] | -grp_opt <group option>
  -device_grp_name <device group name> [<device name>]}
  [ -user_threshold <threshold_1>[<threshold_2>] ]
  [-suspend_tipair {yes | no}]
  [-auto_add_poolvol {enable|disable}]
```

Options and parameters

-pool_id <pool ID#>

Specifies the Pool ID (0-127) for Dynamic Provisioning, Dynamic Provisioning for Mainframe, Dynamic Tiering, Dynamic Tiering for Mainframe, active flash, or active flash for mainframe. When the specification of Pool ID is omitted, a Pool Name must be specified.

If a `-pool_id` option is non-numeric, the specified value is recognized as a pool name to identify the pool ID.

When you omit specifying `-pool_id` option, you need to specify `-pool_name` option.

When specifying the `-pool_name` option without a `-pool_id` option, the specified pool is searched first. If the pool exists, the pool volume is added to the corresponding pool. If the pool does not exist, a new pool is created and is assigned a pool ID. However, if the pool does not exist and multiple commands without the `-pool_id` option are executed simultaneously, the same pool ID might be assigned to multiple newly created pools.

To avoid this behavior, use the **raidcom lock** resource command in advance to lock all resource groups. If all the resource groups are locked, commands executed by other users cannot take out unused pool IDs from all the resource groups. If you create multiple pools, execute the next **raidcom add snap_pool** command after the execution of the current **raidcom add snap_pool** command is complete.

After the execution of all the **raidcom add snap_pool** commands are complete, use the **raidcom unlock resource** command to unlock all the resource groups.

-pool_name <pool naming>

Specifies a pool name of a pool for Dynamic Provisioning, Dynamic Provisioning for Mainframe, Dynamic Tiering, Dynamic Tiering for Mainframe, active flash, or active flash for mainframe. Up to 32 characters can be specified.

When specifying a pool ID or a pool name, if the pool name exists in the specified pool ID, the pool name is overwritten. If the pool volume is added by specifying only a pool ID for already existing pool, the pool name is not changed. When the specification of Pool Name is omitted, a Pool ID must be specified. When the Pool ID is specified and the Pool Name and `-device_grp_name` option are omitted, a pool name is allocated automatically in the form of "New Pool<number>".

The `-pool_name` option cannot be specified by configuring only in numeric because numeric values are preferentially identified as a pool ID. Specifies a pool ID with the '-pool_id<pool ID#>' option.

-parity_grp_id <gno-sgno> [-resource_id <resource group id >]

Supported storage systems:

- VSP E series
- VSP G130, G/F350, G/F370, G/F700, G/F900

Specifies the parity group number (gno: 1 to 52, sgno: 1 to 32).

For example: `3-1`

When `-resource_id <resource grp id>` is specified, select the LDEV which is not installed in the specified resource group and has the largest LDEV number to create a pool volume. This option is valid only when a pool is created. If you omit this option, the smallest ID of a resource group for which the user has operation permissions is set.

-ldev_id <ldev#> …

Specifies the LDEV number (0 to 65279). Up to 64 of LDEVs can be specified at a time. For example:

- -ldev_id 100

- -ldev_id 100 - 110

- -ldev_id 100 -cnt 10

[-cnt <count>]

Specifies the count (2 to 64).

If this specificaion is omitted, "1" is used.

Up to 64 of LDEVs can be specified at a time.

-grp_opt <group option>

Specifies the device information extracted from the LDEV belonging to the device group. Specify "ldev" (fixed). The information of LDEV belonging to the device group is used.

-device_grp_name <device group name> [<device name>]

Specifies the name of device group (maximum 32 characters) to be operated.

To operate an LDEV in the device group, specify the device name of LDEV (maximum 32 characters) within the device group.

If the device name is omitted, all the LDEVs belonging in the device group are operated.

When the -pool_name option is omitted, the device group name changes into the pool name.

[-user_threshold <threshold_1> [<threshold_2>]]

You may specify two user-defined thresholds, and their valid ranges are 1-100%, respectively. If you specify <threshold_1> and <threshold_2>:

the value of <threshold_1> is set as the threshold for WARNING specified to a pool.

the value of <threshold_2> is set as the threshold for High water mark specified to a pool.

If you specify only <threshold_1>, your specified value and the system default value (80%) are applied. If you omit to specify the value, 70% and 80% are applied, automatically.

When you add a pool volume, this option is ignored even if it is specified. If you want to change the user defined threshold value of the additional pool volume, execute the raidcom modify pool command.

[-suspend_tipair {yes | no}]

Specifies whether to suspend Thin Image pairs when the threshold for high water mark is exceeded. This option is valid only for creating pairs.

- yes: Thin Image pairs are suspended.

- no: Thin Image pairs are not suspended.

If this option is omitted , "yes" is set.

[-auto_add_poolvol {enable|disable}]

Specifies whether to automatically add pool volumes for which accelerated compression is enabled. This option is valid only for creating pools. This option can be specified for VSP G1x00 or VSP F1500.

- enable: Automatically adds pool volumes for which accelerated compression is enabled according to the compression ratio of a parity group for which accelerated compression is enabled.

- disable: The pool volumes for which accelerated compression is enabled are not added automatically.

When these options are not specified, the default values will be applied as follows:

- enable: When the storage system is VSP 5000 series, VSP E series or VSP G/F350, G/F370, G/F700, G/F900.

- disable: When the storage system is VSP G1x00 or VSP F1500.

Examples

By using LDEVs:400, 401, and 402, creating a pool of Pool ID:1, Pool Name: my_pool for Dynamic Provisioning/Dynamic Provisioning for Mainframe.

```
# raidcom add dp_pool -pool_id 1 -pool_name my_pool -ldev_id 400 401 402
```

By using LDEVs:500, 501, and 502, creating a pool of Pool ID: Allocated automatically, Pool Name: my_pool for Dynamic Provisioning/Dynamic Provisioning for Mainframe.

```
# raidcom add dp_pool -pool_name my_pool -ldev_id 500 501 502
```

By using LDEVs:600, 601, and 602, creating a pool of Pool ID: 2, Pool Name: Allocated automatically for Dynamic Provisioning/Dynamic Provisioning for Mainframe.

```
# raidcom add dp_pool -pool_id 2 -ldev_id 600 601 602
```

By using LDEVs: 700, 701, and 702, creating a pool of Pool ID: 3, Pool Name: my_pool for Dynamic Provisioning, Dynamic Provisioning for Mainframe, with the user-defined thresholds WARNING specified to 70% and High water mark to 80%.

```
# raidcom add dp_pool -pool_id 3 -pool_name my_pool
-ldev_id 700 701 702 -user_threshold 70 80
```

Add LDEV: 368 to the pool ID: 10 for Dynamic Provisioning, Dynamic Provisioning for Mainframe, Dynamic Tiering, Dynamic Tiering for Mainframe, active flash, or active flash for mainframe.

```
# raidcom add dp_pool -pool_id 10 -ldev_id 368
```

By using LDEV belonging to the device group: grp1, creating a pool of Pool ID:1, pool name: Allocated automatically for Dynamic Provisioning/Dynamic Provisioning for Mainframe.

```
# raidcom add dp_pool -pool_id 1 -grp_opt ldev -device_grp_name grp1
```

By using LDEVs: 800, 801, and 802, creating a pool of Pool ID: 4, Pool Name: my_pool for Dynamic Provisioning, with the user-defined thresholds WARNING specified to 70% and High water mark to 80%, and Thin Image pairs are suspended when the high water mark threshold is exceeded.

```
# raidcom add dp_pool -pool_id 4 -pool_name my_pool -ldev_id 800 801 802 -
user_threshold 70 80 -suspend_tipair yes
```

Using an LDEV ID of resource group ID: 0, creating an LDEV for parity group: 1-1, and then creating a pool of Pool ID: 1, Pool Name: Allocated automatically for Dynamic Provisioning or Dynamic Provisioning for Mainframe.

```
# raidcom add dp_pool -pool_id 1 -parity_grp_id 1-1 -resource_id 0
```

Using LDEVs: 500, 501, and 502, creating a pool of Pool ID: Allocated automatically and Pool Name: my_pool for Dynamic Provisioning for which automatic addition settings of pool volumes whose accelerated compression is enabled: enabled.

```
# raidcom add dp_pool -pool_name my_pool -ldev_id 500 501 502 -auto_add_poolvol enable
```

raidcom get dp_pool

Displays pool information for Dynamic Provisioning, Dynamic Provisioning for Mainframe, Dynamic Tiering, active flash, Dynamic Tiering for Mainframe, and active flash for mainframe.

Syntax

```
raidcom get dp_pool [-key <keyword> [-fx]
    [-pool {<pool_id> | <pool naming>} ]] [-pcap]
```

Options and parameters

[-key <keyword> [-fx]]

Specifies the display keyword. Specify opt as <keyword>. If you specify opt, information about Dynamic Tiering, Dynamic Tiering for Mainframe, active flash, or active flash for mainframe pool is displayed.

Chapter 5: Configuration setting commands

If you specify the -fx option, TL_RANGE and TD_RANGE are displayed in hexadecimal notation.

[-pool {<pool_id> | <pool naming>}]

Specifies the pool ID or the pool name of Dynamic Tiering, Dynamic Tiering for Mainframe, active flash, or active flash for mainframe when you want to display the range for each tiering level of Dynamic Tiering, Dynamic Tiering for Mainframe, active flash, or active flash for mainframe pools. When you specify only digits, it is recognized as a pool ID. When the pool name consists of digits, specify the pool ID instead of the pool name. This option is effective only when the -key option is specified.

[-pcap]

Supported storage systems:

- VSP 5000 series

- VSP G1x00 and VSP F1500

- VSP E series

- VSP G130, G/F350, G/F370, G/F700, G/F900

Specifies the actual capacity. This option is valid only when the -key opt option is specified. For the pool for which automatic addition settings of pool volumes whose accelerated compression is enabled are enabled, check the pool capacity by specifying this option.

If the microcode version does not support this option, nothing is displayed when this option is specified.

Example 1

Displaying the pool information of the pool for Dynamic Provisioning, Dynamic Provisioning for Mainframe, Dynamic Tiering, active flash, Dynamic Tiering for Mainframe, and active flash for mainframe.

```
# raidcom get dp_pool
```

```
PID POLS U(%) AV_CAP(MB) TP_CAP(MB) W(%) H(%) Num LDEV# LCNT TL_CAP(MB) BM TR_CAP(MB)
RCNT
001 POLN   10   45000000   50000000   50   80   2   265   33    65000000 PF
4000000     1
002 POLF   95      10000  100000000   50   80   3   270  900   100000000 PF
0     0
004 POLN    0   10000000  100000000   80   90   2   280    0   200000000 PF
0     0
```

Description of each column in output example:

PID

> Displays the pool ID of Dynamic Provisioning/Dynamic Provisioning for Mainframe, Dynamic Tiering, active flash, Dynamic Tiering for Mainframe, or active flash for mainframe.

POLS

> Displays the status of the pool.

> - POLN: "Pool Normal" (The pool is in the normal status.)
>
> - POLF: "Pool Full" (The pool is in the overflow status exceeding the threshold.)
>
> - POLS: "Pool Suspend" (The pool is in the overflow status exceeding the threshold and is blocked.)
>
> - POLE: "Pool failure" (The pool is suspended in the failure status.) In this status, the pool information cannot be displayed.

U(%)

> Displays the usage rate of the pool (including the mapped capacity and the capacity for Full Allocation).

AV_CAP(MB)

> Displays the available capacity for the volumes of Dynamic Provisioning, Dynamic Provisioning for Mainframe, Dynamic Tiering, active flash, Dynamic Tiering for Mainframe, or active flash for mainframe mapped to this pool.

TP_CAP(MB)

> Displays the total capacity of the pool.

W(%)

> Displays the threshold value for WARNING set for the pool.

H(%)

> Displays the threshold value set for the pool as the high water mark.

Num

> Displays the number of LDEVs composing the pool.

LDEV#

> Displays the LDEV number of a pool-VOL that includes the pool management area. 65535(ffff) is displayed if the pool is being created.

LCNT

> Displays the total number of Dynamic Provisioning, Dynamic Provisioning for Mainframe, Dynamic Tiering, active flash, Dynamic Tiering for Mainframe, or active flash for mainframe volumes mapped to the pool.

TL_CAP(MB)

> Displays the total capacity of all Dynamic Provisioning, Dynamic Provisioning for Mainframe, Dynamic Tiering, active flash, Dynamic Tiering for Mainframe, or active flash for mainframe virtual volumes and Thin Image pairs mapped to the pool.

BM

Displays the I/O Blocking Mode of the pool.

- PF (Pool Full): If the pool is full, you cannot read from the target DP-VOL or write to the target DP-VOL. If the pool VOL is blocked, you can read from the target DP-VOL or write to the target DP-VOL.

- PB (Pool vol Blockade): If the pool VOL is blocked, you cannot read from the target DP-VOL or write to the target DP-VOL. If the pool is full, you can read from the target DP-VOL or write to the target DP-VOL.

- FB (Full or Blockade): You cannot read from the target DP-VOL or write to the target DP-VOL if the pool is full and/or pool VOL is blocked.

- NB (No Blocking): You can read from the target DP-VOL or write to the target DP-VOL even if the pool is full or pool VOL is blocked.

- - (Not supported): The configuration that does not support the I/O Blocking Mode.
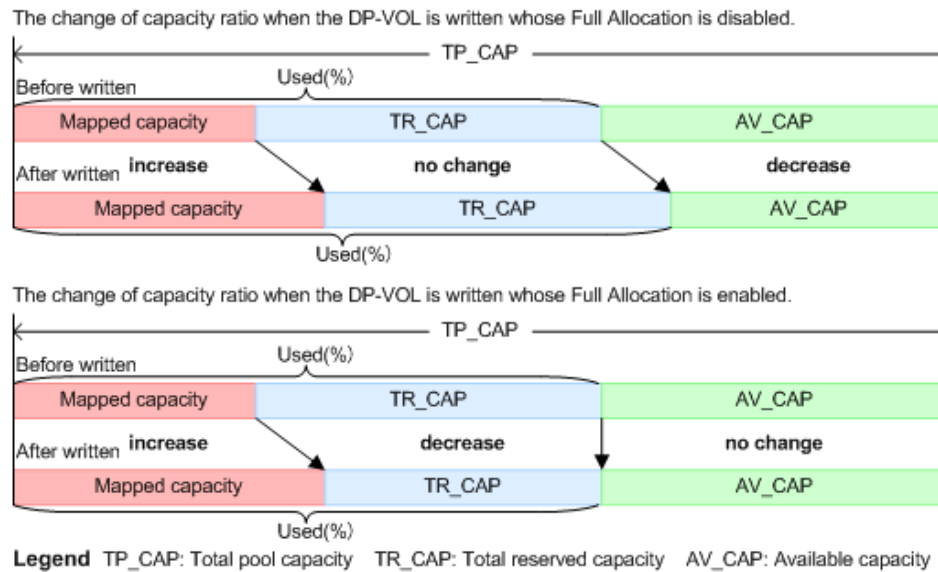
TR_CAP(MB)

Shows the sum of the pool capacities which are reserved for the volumes on which Full Allocation or Proprietary Anchor is enabled. For a configuration that does not support Full Allocation, a hyphen (-) is displayed.

RCNT

Shows the number of volumes for which Full Allocation is enabled that are mapped to a pool. In the configuration that does not support Full Allocation, a hyphen (-) is displayed.

The following figures show the difference between the capacity ratio changing according to whether Full Allocation is enabled or disabled when DP-VOL is written. The "Mapped capacity" in the figure is the total capacity of the user data in each virtual volumes and the page capacity which is storing the control information.

The change of capacity ratio when the DP-VOL is written whose Full Allocation is disabled.

The change of capacity ratio when the DP-VOL is written whose Full Allocation is enabled.

Legend TP_CAP: Total pool capacity TR_CAP: Total reserved capacity AV_CAP: Available capacity

Example 2

Displays the tier information of the pool for Dynamic Tiering, active flash, Dynamic Tiering for Mainframe, or active flash for mainframe. For details about the parameters, see the *Provisioning Guide* for the storage system.

```
# raidcom get dp_pool -key opt
```

```
PID POLS MODE STS DAT TNO TL_RANGE  TD_RANGE TU_CAP(MB)  TT_CAP(MB)  T(%)  P(%)  R(%)  B(%)
MM MC(h)
001 POLN DEF  STP VAL 1  00005000  00003000    200000     1000000   80   54   98    40
PM 24
001 POLN DEF  STP VAL 2  00003000  00002000    400000     1000000   80   54   98    30
PM 24
001 POLN DEF  STP VAL 3  00002000  00002000    600000     1000000   80   54   98    40
PM 24
002 POLF AUT  MON PND 1       -         -      500000     1000000   80   54  100     2
PM  8
```

Example 3

When you specify -pool, displays the range for each tiering level of corresponding pools.

```
# raidcom get dp_pool  -pool 1 -key opt
```

```
PID POLS MODE STS DAT TNO TL_RANGE  TD_RANGE   TU_CAP(MB)   TT_CAP(MB)  T(%)  P(%)  R(%)
B(%) MM MC(h)
001 POLN DEF  STP VAL 0 4294967294 00000000   8064 8190 10    28   100 2 PM  24
001 POLN DEF  STP VAL 0 4294967294 00000000   3948 3990 10   100   100 2 PM  24
001 POLN DEF  STP VAL 0   00000000 00000000    504 8190 10    33   100 2 PM  24
001 POLN DEF  STP VAL 1   00000000 00000000   8064 8190 10    28   100 2 PM  24
001 POLN DEF  STP VAL 1   00000000 00000000   3948 3990 10   100   100 2 PM  24
001 POLN DEF  STP VAL 1   00000000 00000000    504 8190 10    33   100 2 PM  24
001 POLN DEF  STP VAL 2   00000039 00000000   8064 8190 10    28   100 2 PM  24
001 POLN DEF  STP VAL 2   00000000 00000000   3948 3990 10   100   100 2 PM  24
001 POLN DEF  STP VAL 2   00000000 00000000    504 8190 10    33   100 2 PM  24
001 POLN DEF  STP VAL 3   00000100 00000000   8064 8190 10    28   100 2 PM  24
001 POLN DEF  STP VAL 3   00000041 00000000   3948 3990 10   100   100 2 PM  24
001 POLN DEF  STP VAL 3   00000000 00000000    504 8190 10    33   100 2 PM  24
001 POLN DEF  STP VAL 4   00000001 00000000   8064 8190 10    28   100 2 PM  24
001 POLN DEF  STP VAL 4   00000000 00000000   3948 3990 10   100   100 2 PM  24
001 POLN DEF  STP VAL 4   00000000 00000000    504 8190 10    33   100 2 PM  24
001 POLN DEF  STP VAL 5   00000001 00000000   8064 8190 10    28   100 2 PM  24
001 POLN DEF  STP VAL 5   00000001 00000000   3948 3990 10   100   100 2 PM  24
001 POLN DEF  STP VAL 5   00000000 00000000    504 8190 10    33   100 2 PM  24
```

Chapter 5: Configuration setting commands

Description of each column in output example:

PID

Displays the pool ID for Dynamic Provisioning, Dynamic Provisioning for Mainframe, Dynamic Tiering, active flash, Dynamic Tiering for Mainframe, or active flash for mainframe.

POLS

Displays the status of the pool.

- POLN: "Pool Normal" (The pool is in the normal status.)

- POLF: "Pool Full" (The pool is in the overflow status exceeding the threshold.)

- POLS: "Pool Suspend" (The pool is in the overflow status exceeding the threshold and is blocked.)

- POLE: "Pool failure" (The pool is suspended in the failure status.) In this status, the pool information cannot be displayed.

MODE

Displays the execution mode.

- DEF: The start/termination of the monitor is performed by the instruction of the CCI, and the Tier range setting is performed by the automatic calculation of the storage system.

- AUT: The start/termination of the monitor is performed by time specification, and the Tier range setting is performed by the automatic calculation of the storage system.

> 📄 **Note:**
>
> AUT cannot be instructed from the CCI. AUT can only performed if displayed when set from the SN2.

STS

Displays the operational status of the performance monitor and the tier relocation.

- STP: The performance monitor and the tier relocation are stopped.

- RLC: The performance monitor is stopped. The tier relocation is waiting or operating.

- MON: The performance monitor is operating. The tier relocation is stopped.

- RLM: The performance monitor is operating. The tier relocation is waiting or operating.

DAT

Displays the status of the monitor information.

- VAL: Valid.

- INV: Invalid.

- PND: Being calculated.

TNO

Tiering number. When you specify the pool, displays the tiering level (0 means "all" is set to the level of tiering policy).

TL_RANGE

Lowest limit value of the Tier in IOPH.

0 (0x00000000) to 4294967294 (0xFFFFFFFE)

When "all" is set to the level of tiering policy, the lowest limit value of the tier is displayed. When there is no V-VOL where "all" is set to the level of tiering policy, or when there is no V-VOL related to the pool, the invalid value "-" (0xFFFFFFFF) is displayed. If the -fix option is specified, this displays the value in parentheses given in hexadecimal.

TD_RANGE

Delta value of the Tier in IOPH. 0 (0x00000000) to 4294967294 (0xFFFFFFFE): When "all" is set to the level of tiering policy, the lowest limit value of the tier is displayed. When there is no V-VOL where "all" is set to the level of tiering policy, or when there is no V-VOL related to the pool, the invalid value "-" (0xFFFFFFFF) is displayed. If the -fix option is specified, this displays the value in parentheses given in hexadecimal.

TU_CAP(MB)

Tier capacity (Usage amount). Displays the actual capacity of the tier (usage amount) when the `-pcap` option is specified.

TT_CAP(MB)

Tier capacity (Total capacity). Displays the actual capacity of the tier (total capacity) when the `-pcap` option is specified.

T(%)

The free space percentage for the new allocation

P(%)

Performance working ratio

R(%)

Progress percentage of relocation.

- 0 to 99: Shows one of the following statuses.

  - When the value of STS is RLC or RLM: relocation is waiting or in progress.

  - When the value of STS is STP or MON: relocation is suspended.

- 100: Shows if the relocation operation is not in progress, or the relocation is complete.

B(%)

Displays the amount of buffer for the tier reallocation.

MM

Displays the mode of performance monitoring and the availability of active flash:

- PM: Periodical mode.

- CM: Continuous mode.

- RPM: Periodical mode and active flash is enabled.

- RCM: Continuous mode and active flash is enabled.

MC(h)

Displays the execution cycle of performance monitoring and tier reallocation. A hyphen (-) is displayed if MODE is not AUT.

For VSP G1x00, VSP F1500, VSP, HUS VM, VSP G200, G400, G600, G800, and VSP F400, F600, F800, a hyphen (-) is always displayed.

raidcom send ping

Sends a ping from the specified port to the specified host, and then displays the result.

Syntax

```
raidcom send ping -port <port#> -address <IP address>
   [-iscsi_virtual_port_id <iSCSI virtual port ID>]
```

Options and parameters

-port <port#>

Specifies the port number. For example:

- CL1-A

-address <IP address>

Specifies the IP address of the host that is the destination of the ping. You can specify IPv4 address or IPv6 address.

[-iscsi_virtual_port_id <iSCSI virtual port ID>]

Specifies the iSCSI virtual port ID when the specified port is iSCSI and the iSCSI virtual port mode is enabled. If this option is omitted, 0 is assumed.

An error occurs if you specify this option in any of the following cases:

- The virtual port of the specified iSCSI virtual port ID is not found.

- The iSCSI virtual port mode is disabled.

- The device does not support the iSCSI virtual port mode.

Returned values

The **raidcom send ping** command sets one of the following returned values in exit(), which allows users to check the execution results using a user program or script.

- **Normal termination:**

  - **0:** One or more packets were received.

  - **1:** A ping could not be sent, or a ping was sent, but no response packets have been received within the timeout period.

- **Abnormal termination:**

  - **EX_XXX:** The operation terminated abnormally. For details, see the descriptions of the command error messages in the *Command Control Interface User and Reference Guide*.

Examples

To send a ping from the port: CL4-E to the host whose IP address is 10.213.60.111, and then display the result (Normal).

```
# raidcom send ping -port CL4-E -address 10.213.60.111
raidcom : 5 packets transmitted, 5 packets received.
```

To send a ping from the port: CL4-E, iSCSI virtual port ID: 0 to the host whose IP address is 10.213.60.111, and then display the result (Normal).

```
# raidcom send ping -port CL4-E -address 10.213.60.111 -iscsi_virtual_port_id 0
raidcom : 5 packets transmitted, 5 packets received.
```

To send a ping from the port: CL4-E to the host whose IP address is 10.213.60.111, and then display the result (No response from the host).

```
# raidcom send ping -port CL4-E -address 10.213.60.111
raidcom : 5 packets transmitted, 0 packets received.
```

To send a ping from the port: CL4-E to the host whose IP address is 10.213.60.111, and then display the result (Fail to send a ping by an internal error etc).

```
# raidcom send ping -port CL4-E -address 10.213.60.111
raidcom : 0 packets transmitted.
```

Description of output example:

N **packets transmitted**

> Number of packets sent to the host.

N **packets received**

> Number of packets responded from the host normally.

raidcom add server

Supported storage systems:

- VSP E series
- VSP G130, G/F350, G/F370, G/F700, G/F900

Creates a server managed by Storage Advisor Embedded. To execute this command, you need to log in to the storage system with the user account belonging to the user group to which "all resource groups" are set. In addition, you must be able to perform operations with all implemented resource groups. Therefore, this command cannot be executed if any of the following conditions is met:

- The command device belonging to a resource group other than meta_resource is used in the in-band management.
- The virtual storage machine that contains only some of the implemented resource groups is specified as HORCM_VCMD in the configuration file.

> **Note:** If a failure occurs in a processor on an MP blade or a controller, this command cannot be run even on an MP blade or a controller on which no failure occurs.

When you run this command, the processing is run asynchronously with the command input. Verify that the processing is complete by using the `raidcom get command_status` command. There is a limit to the number of asynchronous commands that can be accepted simultaneously. For details, see .

Syntax

```
raidcom add server -server_name <name> -request_id auto
```

Options and parameters

-server_name <name> -request_id auto

Specifies the server nickname with one character through 229 characters.

The following are the characters that can be used:

0 through 9, A through Z, a through z, -. /: @ \ _ single-byte space

This command creates a server with no OS type set. 255 servers are created in the storage system.

The server ID is automatically selected. You can check the selected server ID using the `raidcom get command_status` command.

The request ID is output in the following format when the execution of the **raidcom add server** command ends:

▪ REQID : <request#>

<request#> is a request ID assigned each time the command is executed. For details, see .

Example

Creating a server with the server name: server0.

```
#raidcom add server -server_name server0 -request_id auto
 REQID:0
```

raidcom delete server

Supported storage systems:

▪ VSP E series

▪ VSP G130, G/F350, G/F370, G/F700, G/F900

Deletes a server managed by Storage Advisor Embedded. To execute this command, you need to log in to the storage system with the user account belonging to the user group to which "all resource groups" are set. In addition, you must be able to perform operations with all implemented resource groups. Therefore, this command cannot be executed if any of the following conditions is met:

▪ The command device belonging to a resource group other than meta_resource is used in the in-band management.

▪ The virtual storage machine that contains only some of the implemented resource groups is specified as HORCM_VCMD in the configuration file.

📄 **Note:** If a failure occurs in a processor on an MP blade or a controller, this command cannot be run even on an MP blade or a controller on which no failure occurs.

When you run this command, the processing is run asynchronously with the command input. Verify that the processing is complete by using the **raidcom get command_status** command. There is a limit to the number of asynchronous commands that can be accepted simultaneously. For details, see .

Syntax

```
raidcom delete server {-server_id <id>|-server_name <name>} -request_id auto
```

Options and parameters

{-server_id <id>|-server_name <name>} -request_id auto

Specifies the ID or nickname of the server to be deleted.

You need to specify the server ID (0 through 4095) in decimal or hexadecimal in -server_id <id>. To specify in hexadecimal, add "0x" at the beginning of the server ID.

For example:

- -server_id 128

- -server_id 0x80

To specify the server nickname, specify the server nickname in -server_name <name> with one character through 229 characters.

The following are the characters that can be used:

0 through 9, A through Z, a through z, -. /: @ \ _ single-byte space

The specified server is deleted by executing this command. The registered host groups or iSCSI targets are not deleted.

The request ID is output in the following format when the execution of the **raidcom delete server** command ends:

- REQID : <request#>

<request#> is a request ID assigned each time the command is executed. For details, see Request ID function (on page 246).

Example

Deleting a server with the server ID: 0.

```
#raidcom delete server -server_id 0 -request_id auto
REQID:0
```

raidcom modify server

Supported storage systems:

- VSP E series

- VSP G130, G/F350, G/F370, G/F700, G/F900

Modifies a server managed by Storage Advisor Embedded. To execute this command, you need to log in to the storage system with the user account belonging to the user group to which "all resource groups" are set. In addition, you must be able to perform operations with all implemented resource groups. Therefore, this command cannot be executed if any of the following conditions is met:

- The command device belonging to a resource group other than meta_resource is used in the in-band management.

- The virtual storage machine that contains only some of the implemented resource groups is specified as HORCM_VCMD in the configuration file.

Chapter 5: Configuration setting commands

📄 **Note:** If a failure occurs in a processor on an MP blade or a controller, this command cannot be run even on an MP blade or a controller on which no failure occurs.

When you run this command, the processing is run asynchronously with the command input. Verify that the processing is complete by using the `raidcom get command_status` command. There is a limit to the number of asynchronous commands that can be accepted simultaneously. For details, see Maximum number of acceptable asynchronous commands (on page 228).

Syntax

```
raidcom modify server {-server_id <id>|-server_name <name>} -request_id auto {-
server_operation {add_host_grp|delete_host_grp} -port <port> [<host group name>]}
```

Options and parameters

{-server_id <id>|-server_name <name>} -request_id auto

Specifies the ID or nickname of the server to be modified.

You need to specify the server ID (0 through 4095) in decimal or hexadecimal in -server_id <id>. To specify in hexadecimal, add "0x" at the beginning of the server ID. For example:

- -server_id 128

- -server_id 0x80

To specify the server nickname, specify the server nickname in -server_name <name> with one character through 229 characters.

The following are the characters that can be used:

0 through 9, A through Z, a through z, -. /: @ \ _ single-byte space

The specified server is deleted by executing this command. The registered host groups or iSCSI targets are not deleted.

The request ID is output in the following format when the execution of the `raidcom modify server` command ends:

- REQID : <request#>

<request#> is a request ID assigned each time the command is executed. For details, see Request ID function (on page 246).

-server_operation add_host_grp -port <port> [<host group name>]

Registers the host groups or the iSCSI targets specified with the -port <port> [<host group name>] on the server specified for the -server_id <id> or the -server_name <name>.

If the WWN or iSCSI name of the host is registered in the specified host groups or iSCSI targets, the WWN or iSCSI name of the host is registered on the server. In addition, if an LU path is defined for a specified host groups or iSCSI targets, the LDEV associated with the LU path is registered on the server.

For -port <port> [<host group name>], specify the port number and host groups (iSCSI target ID for iSCSI), or host group name (target alias for iSCSI).

If the host group name or the target alias is more than 64 characters, the host group name or the target alias cannot be specified. Use the host group ID or the target ID. For example:

- CL1-A -g (g is from 1 to 254)

- CL1-A Linux_X86

- CL1-A Target00

To check if host groups or iSCSI targets are registered on the server, execute the **raidcom get host_grp** command with the `-key server` option specified.

If any of the following conditions is met, the command execution fails:

- The specified server and host group/iSCSI target are not installed.

- The IDs of the specified host groups or iSCSI targets are 0.

- The host group/iSCSI target, LU path, or LDEV information on the specified server is being updated.

- The iSCSI target is specified even if the specified server protocol is Fibre. Alternatively, a host group is specified even if the specified server protocol is iSCSI.

- The host group/iSCSI target which has already been registered on other servers is specified.

- Another host group/iSCSI target on the same port as the one used for the specified host group/iSCSI target has already been registered on the specified server.

- If 32 WWNs or iSCSI names of the host have been registered on the specified server, the host group/iSCSI target including WWNs or iSCSI names that are not registered on the server is specified.

- The specified host group/iSCSI target, the port to which the host group/iSCSI target belongs, or the resource group to which the LDEV associated with the LU path of the host group/iSCSI target belongs is locked by another user.

- The LU path to an LDEV other than that for Dynamic Provisioning, Dynamic Tiering, or active flash exists in the specified host group/iSCSI target.

- The LDEV for a Volume Migration pair including LDEV other than that for Dynamic Provisioning, Dynamic Tiering, or active flash is specified.

- The host group/iSCSI target on the port with the LU security set to OFF is specified.

-server_operation delete_host_grp -port <port> [<host group name>]

Specifies the iSCSI virtual port ID when the specified port is iSCSI and the iSCSI virtual port mode is enabled. If this option is omitted, 0 is assumed.

From the server specified for -server_id <id> or -server_name <name>, delete the host groups or iSCSI targets specified with -port <port> [<host group name>].

The LDEV is deleted from the server if an LU path is defined for the specified host groups or iSCSI targets, and if the LU path is the last one of the LDEV associated with the LU path on the specified server.

Chapter 5: Configuration setting commands

For -port <port> [<host group name>], specify the port number and host groups (iSCSI target ID for iSCSI), or the host group name (target alias for iSCSI). If the host group name or the target alias is more than 64 characters, the host group name or the target alias cannot be specified. Use the host group ID or the target ID.

For example:

- CL1-A -g (g is from 1 to 254)

- CL1-A Linux_X86

- CL1-A Target00

To check if host groups or an iSCSI targets are registered on the server, execute the **raidcom get host_grp** command with the `-key server` option specified.

If any of the following conditions is met, the command execution fails:

- The specified server and host group/iSCSI target are not installed.

- The IDs of the specified host groups or iSCSI targets are 0.

- The host group/iSCSI target, LU path, or LDEV information on the specified server is being updated.

Examples

Adding CL1-A-1 to server ID: 0.

```
#raidcom modify server -server_id 0 -request_id auto -server_operation add_host_grp -
port CL1-A-1
REQID:0
```

Deleting CL1-A-1 from server ID: 0.

```
#raidcom modify server -server_id 0 -request_id auto -server_operation
delete_host_grp -port CL1-A-1
REQID:0
```

raidcom get server

Supported storage systems:

- VSP E series

- VSP G130, G/F350, G/F370, G/F700, G/F900

Displays a server managed by Storage Advisor Embedded. To execute this command, you need to log in to the storage system with the user account belonging to the user group to which "all resource groups" are set. In addition, you must be able to perform operations with all implemented resource groups. Therefore, this command cannot be executed if any of the following conditions is met:

- The command device belonging to a resource group other than meta_resource is used in the in-band management.

- The virtual storage machine that contains only some of the implemented resource groups is specified as HORCM_VCMD in the configuration file.

Syntax

```
raidcom get server [-key <keyword> [-server_id <id>|-server_name <name>]]
```

Options and parameters

[-key <keyword> [-server_id <id>|-server_name <name>]]
Specifies the server information to be displayed.

- opt: Specifies the detailed information about the server. To specify this parameter, make sure to specify the server ID or nickname that displays the detailed information in -server_id <id> or -server_name <name>.

You need to specify the server ID (0 through 4095) in decimal or hexadecimal in -server_id <id>. To specify in hexadecimal, add "0x" at the beginning of the server ID.

For example:

- -server_id 128

- -server_id 0x80

To specify the server nickname, specify the server nickname in -server_name <name> with one character through 229 characters.

The following are the characters that can be used:

0 through 9, A through Z, a through z, -. /: @ \ _ single-byte space

Example

Displaying server information.

```
#raidcom get server
SRVID NAME
0 "server0"
1 "server1"
155 "server2
```

Description of each column in the output example:

SRVID
Displays the server ID in decimal.

NAME

Displays the server nickname. A double quotation (") is added at the beginning and end of the server nickname if the -key opt option is not specified.

Specifying the server information to be displayed.

```
#raidcom get server -key opt -server_id 0
SRVID : 0
NAME : server0
STS : RCV DGG1
PROTOCOL : FIBRE
OS : WIN
OS_OPT : 1 3
```

Description of each column in the output example:

STS

Status of the server. The status of the server is not displayed if it is not in any of the following statuses:

- RSV: Protocol, OS type, and OS type option are not set.

  When setting the host group/iSCSI target to the server with the **raidcom modify server** command, the OS type is set based on the host mode of the host group/ iSCSI target.

- ILU: No LU path is defined between any of the LDEV registered on the server and any of the host group/iSCSI target registered on the server.

- IOT: The OS type of the server does not correspond to the host mode of any of the host group/iSCSI target registered on the server.

- IOP: The OS type of the server does not correspond to the host mode option of any of the host group/iSCSI target registered on the server.

- RCV: The host group/iSCSI target, LU path, or LDEV information on the server is being updated.

- DGG1: For VSP E series, VSP G130, G/F350, G/F370, G/F700, G/F900, the microcode version cannot be downgraded to 88-01-xx-xx/xx or earlier.

PROTOCOL

Displays the protocol of the server.

- FIBRE: The protocol is Fiber Channel.

- ISCSI: The protocol is iSCSI.

- - (hyphen): The protocol is not either Fiber Channel or iSCSI.

Chapter 5: Configuration setting commands

OS

Displays the OS type of the server. If other OS types are displayed, the OS type that is not supported by Storage Advisor Embedded is set.

- LINUX/IRIX: The OS type is "Linux".

- VMWARE: The OS type is "(Deprecated) VMware".

- HP-UX: The OS type is "HP-UX".

- OVMS: The OS type is "OpenVMS".

- TRU 64: The OS type is "Tru64".

- SOLARIS: The OS type is "Solaris".

- NETWARE: The OS type is "NetWare".

- WIN: The OS type is "(Deprecated) Windows".

- AIX: The OS type is "AIX".

- VMWARE_EX: The OS type is "VMWare".

- WIN_EX: The OS type is "Windows".

- - (hyphen): No OS type is set.

OS_OPT

OS type options that are set on the server. A hyphen (-) is displayed if the OS type is not set.

raidcom get apn

Supported storage systems:

- VSP 5000 series

- VSP G1x00 and VSP F1500

- VSP E series

- VSP G130, G/F350, G/F370, G/F700, G/F900

- VSP G200, G400, G600, G800 and VSP F400, F600, F800

Displays the AP numbers used in the storage system. These commands are used for troubleshooting. Use them only if you get any directions.

Syntax

```
raidcom get apn [-fx]
```

Options and parameters

-fx

Specifies the LDEV number of the command device in hexadecimal.

Displays the AP numbers used in the storage system.

```
#raidcom get apn
AP#    LDEV#   WWN
0001      1    5060000000000001
0002      2    5060000000000002
0eff     12    -
```

Description of each column in output example:

AP#

Displays the AP# being used in hexadecimal.

LDEV

Displays the LDEV number of the command device using AP#.

WWN

If the command device is used via Fibre Channel, the HBA WWN is displayed.

If a command device is used via iSCSI, a hyphen(-) is displayed.

raidcom add qos_grp

Supported storage systems:

▪ VSP 5000 series

Creates a QoS group or adds an LDEV to the specified QoS group.

This command is executed asynchronously with the command input. Check the completion of this processing using the `raidcom get command_status` command.

📄 **Note:** If a failure occurs in a processor on an MP blade or a controller, this command cannot be run even on an MP blade or a controller on which no failure occurs.

Syntax

```
raidcom add qos_grp -qos_grp_id <qos group#> [-ldev_id <ldev#>] -request_id auto
```

Options and parameters

-qos_grp_id <qos-group#>

Specifies a QoS group ID.

[-ldev_id <ldev#>]

Specifies an LDEV number (0-65279) to be set in the QoS group.

This option is available only when an LDEV is added to the QoS group.

Chapter 5: Configuration setting commands

Example:

- -ldev_id 200

If the command is being run continuously by specifying this option, run it within the allowable numbers of commands that the storage system can accept. For more information, see Maximum number of acceptable asynchronous commands (on page 228). The command ends normally even if an LDEV specified with -ldev exists in a QoS group specified with **-qos_grp_id**.

-request_id auto

<request#> is a request ID assigned each time the command is run. For more information, see Request ID function (on page 246).

Creating the QoS group whose QoS group ID is 1.

```
# raidcom add qos_grp -qos_grp_id 1 -request_id auto
```

Adding LDEV ID: 101 to the QoS group whose QoS group ID is 1.

```
# raidcom add qos_grp -qos_grp_id 1 -ldev_id 101 -request_id auto
```

Adding LDEV IDs: 1, 2, and 3 to the QoS group whose QoS group ID is 1.

```
# raidcom get ldev -ldev_id 1 -cnt 3 -key front_end | rmawk @L-ne:1
exe="raidcom add qos_grp -qos_grp_id 1 -ldev_id @2 -request_id auto"
REQID : 1
REQID : 2
REQID : 3
# raidcom get command_command status
# raidcom reset command_status -request_id all
# raidcom reset command_status
```

raidcom delete qos_grp

Supported storage systems:

- VSP 5000 series

Deletes an LDEV from the QoS group or deletes a QoS group.

This command is executed asynchronously with the command input. Check the completion of this processing using the **raidcom get command_status** command.

> **Note:** If a failure occurs in a processor on an MP blade or a controller, this command cannot be run even on an MP blade or a controller on which no failure occurs.

Syntax

```
raidcom delete qos_grp -qos_grp_id <qos group#> [-ldev_id <ldev#>] -request_id auto
```

Options and parameters

-qos_grp_id <qos-group#>

Specifies a QoS group ID.

[-ldev_id <ldev#>]

Specifies an LDEV number (0-65279) to be deleted from the QoS group.

Example:

▪ -ldev_id 200

If the command is being run continuously by specifying this option, run it within the allowable numbers of commands that the storage system can accept. For more information, see Maximum number of acceptable asynchronous commands (on page 228). The command ends normally even if an LDEV specified with -ldev does not exist in a QoS group specified with `-qos_grp_id`.

-request_id auto

<request#> is a request ID assigned each time the command is run. For more information, see Request ID function (on page 246).

Deleting LDEV ID: 101 from the QoS group whose QoS group ID is 1.

```
# raidcom delete qos_grp -qos_grp_id 1 -ldev_id 101 -request_id auto
```

Deleting the QoS group whose QoS group ID is 1.

```
# raidcom delete qos_grp -qos_grp_id 1 -request_id auto
```

Deleting up to 2000 LDEVs from the QoS group whose QoS group ID is 1.

```
#raidcom get qos_grp -qos_grp_id 1 -key resource | rmawk @L-ne:1 -a @L-le:2001
exe="raidcom delete qos_grp -qos_grp_id 1 -ldev_id @2 -request_id auto"
REQID : 1
REQID : 2
REQID : 3
…
# raidcom get command_command status
# raidcom reset command_status -request_id all
# raidcom reset command_status
```

raidcom modify qos_grp

Supported storage systems:

▪ VSP 5000 series

Changes the QoS settings for each QoS group.

This command is executed asynchronously with the command input. Check the completion of this processing using the `raidcom get command_status` command.

> 📄 **Note:** If a failure occurs in a processor on an MP blade or a controller, this command cannot be run even on an MP blade or a controller on which no failure occurs.

Syntax

```
raidcom modify qos_grp -qos_grp_id <qos group#> {-upper_throughput_io <upper
throughput io> | -upper_data_trans_mb <upper data trans mb> | -upper_alert_time
<upper alert time>| -response_priority <#priority> | -response_alert_time <response
alert time>} -request_id auto
```

Options and parameters

-qos_grp_id <qos-group#>
Specifies a QoS group ID.

[-upper_throughput_io <upper throughput io>]
Specifies the upper limit of the throughput per second. The upper limit is the total throughput for all LDEVs in the group. To disable it, specify zero.

[-upper_data_trans_mb <upper data trans mb>]
Specifies the upper limit of the data transfer volume (MB) per second. The upper limit is the total data transfer volumes for all LDEVs in the group. To disable it, specify zero.

[-upper_alert_time <upper alert time>]
Specifies the alert notification time (second) when the upper limit is exceeded. To disable it, specify zero.

An alert is issued if either of the following conditions continues for the alert notification time or longer:

▪ The number of received commands continuously exceeds the upper limit of the throughput per second.

▪ The transfer volume of received commands continuously exceeds the upper limit of the data transfer volume per second.

[-response_priority <#priority>]
Specifies the I/O processing priority with 1 to 3. To disable it, specify 0. 1 represents the lowest priority and a greater number represents a higher priority.

For details about the I/O processing priority, see the *Performance Guide*.

[-response_alert_time <response alert time>]

Specifies the alert notification time (second) when the target response time is not reached. If the target response time has not been continuously reached for more than the alert notification time, an alert is issued. To disable it, specify 0.

-request_id auto

<request#> is a request ID assigned each time the command is run. For more information, see Request ID function (on page 246).

Changing the upper limit of the data transfer volume per second to 100 MB/s in the QoS group whose QoS group ID is 1.

```
# raidcom modify qos_grp -qos_grp_id 1 -upper_data_trans_mb 100 -request_id auto
```

Specifying the I/O processing priority to 1 in the QoS group whose QoS group ID is 1.

```
# raidcom modify qos_grp -qos_grp_id 1 -response_priority 1 -request_id auto
```

raidcom get qos_grp

Supported storage systems:

- VSP 5000 series

Displays the QoS group information.

Syntax

```
raidcom get qos_grp [-qos_grp_id <qos group#>] [-key <resource | monitor>][-time_zone
<time zone>]
```

Options and parameters

-qos_grp_id <qos-group#>

Specifies a QoS group ID. This option displays only the specified QoS group information.

[-key <resource | monitor>]

Specifies the QoS group information to be displayed.

- resource: Specifies the information on LDEVs included in the QoS group. To specify this parameter, make sure to specify the QoS group ID.

- monitor: Specifies the performance monitor information in the QoS group. If the command is run immediately after creating the QoS group, the QoS group might not be displayed in the command output. In this case, rerun the command after a few seconds.

[-time_zone <time zone>]

> Specifies the time zone to be displayed. If this option is not specified, the time in the time zone set for the storage system is displayed.

> utc: displays the time in Coordinated Universal Time (UTC).

Example 1

Displaying the list of QoS groups. A hyphen (-) is displayed if the QoS information is not set.

```
# raidcom get qos_grp
QID UPPER_THROUGHPUT(IOps) UPPER_DATA_TRANS(MBps) UPPER_ALERT(s) UPPER_ALERT_NOTICE
PRIORITY RESPONSE_TARGET(ms) RESPONSE_ALERT(s) RESPONSE_ALERT_NOTICE
1 80 100 30 2019-07-31T10:15:20 1 80 100 2019-07-31T10:15:20
2 80 100 30 2019-07-31T10:15:20 1 80 100 2019-07-31T10:15:20
3 80 100 30 2019-07-31T10:15:20 1 80 100 2019-07-31T10:15:20
```

Description of each column in output example:

QID

> Displays the QoS group ID.

UPPER_THROUGHPUT(IOps)

> Displays the upper limit of the throughput per second. The upper limit is the total throughput of all LDEVs in the group.

UPPER_DATA_TRANS(MBps)

> Displays the upper limit of the data transfer volume (MB) per second. The upper limit is the total data transfer volumes of all LDEVs in the group.

UPPER_ALERT(s)

> Displays the alert notification time (second) when the upper limit is exceeded.

> An alert is issued if either of the following conditions continues for the alert notification time or longer:

> - The number of received commands continuously exceeds the upper limit of the throughput per second.

> - The transfer volume of received commands continuously exceeds the upper limit of the data transfer volume per second.

UPPER_ALERT_NOTICE

> Displays the last alert time of the following alerts:

> - The alert issued when the number of received commands continuously exceeds the upper limit of the throughput per second.

> - The alert issued when the transfer volume of received commands continuously exceeds the upper limit of the data transfer volume per second.

> If the last alert time does not exist, a hyphen (-) is displayed.

Chapter 5: Configuration setting commands

PRIORITY

Displays the I/O processing priority. 1 represents the lowest priority and a greater number represents a higher priority.

If the DKCMAIN microcode version is not supported, a hyphen (-) is displayed.

RESPONSE_TARGET(ms)

Displays the target response time.

If the DKCMAIN microcode version is not supported, a hyphen (-) is displayed.

RESPONSE_ALERT(s)

Displays the alert notification time (second) when the target response time is not reached. If the target response time has not been continuously reached for more than the alert notification time, an alert is issued.

If the DKCMAIN microcode version is not supported, a hyphen (-) is displayed.

RESPONSE_ALERT_NOTICE

Displays the last alert time when the target response time has not been continuously reached.

If the last alert time does not exist, or the DKCMAIN microcode version is not supported, a hyphen (-) is displayed.

Example 2

Displaying the information on the QoS group whose QoS group ID is 1. A hyphen (-) is displayed if the QoS information is not set.

```
# raidcom get qos_grp -qos_grp_id 1
QID UPPER_THROUGHPUT(IOps) UPPER_DATA_TRANS(MBps) UPPER_ALERT(s) UPPER_ALERT_NOTICE
PRIORITY RESPONSE_TARGET(ms) RESPONSE_ALERT(s) RESPONSE_ALERT_NOTICE
1 80 100 30 2019-07-31T10:15:20 1 80 100 2019-07-31T10:15:20
```

Displaying the information on LDEVs included in the QoS group whose QoS group ID is 1.

```
# raidcom get qos_grp -qos_grp_id 1 -key resource
QID LDEV#
1 100
1 101
```

Description of each column in output example:

LDEV#

Displays the numbers of LDEVs included in the QoS group.

Example 3

Displaying the performance monitor information on all LDEVs in the QoS group. If all performance monitor information is not valid, nothing is displayed.

```
# raidcom get qos_grp -key monitor
QID RCV_CMD(IOps) RCV_CMD_TRANS(KBps) THROUGHPUT(IOps) DATA_TRANS(KBps) RESPONSE(us)
```

Chapter 5: Configuration setting commands

```
MONITOR_TIME
1 3840 15 2650 10 5 2019-07-31T10:15:20
```

Description of each column in output example:

RCV_CMD(IOps)
> Displays the total numbers of received commands for all LDEVs in the QoS group.

RCV_CMD_TRANS(KBps)
> Displays the total data transfer volumes of received commands for all LDEVs in the QoS group.

THROUGHPUT(IOps)
> Displays the total throughput per second for all LDEVs in the QoS group.

DATA_TRANS(KBps)
> Displays the total data transfer volumes per second for all LDEVs in the QoS group.

RESPONSE(us)
> Displays the average response time for all LDEVs in the QoS group.

MONITOR_TIME
> Displays the monitoring collection time.

raidcom monitor resource

Supported storage systems:

- VSP 5000 series

Displays the performance information on the storage system. To run this command, you need to log in to the storage system with the user account belonging to the user group to which all resource groups are set. In addition, you must be able to perform the operations with all implemented resource groups. Therefore, this command cannot be run if either of the following conditions is met:

- The command device belonging to a resource group other than meta_resource is used in the in-band management.

- The virtual storage machine that contains only some of the implemented resource groups is specified as HORCM_VCMD in the configuration file.

Syntax

```
raidcom monitor resource -monitor_type qos_ldev [-qos_grp_id <qos group id>] [-start
_id <start id#>] [-cnt <count>] -monitor_interval <interval> [<repetition count>] [-
time_zone <time zone>]
raidcom monitor resource -monitor_type qos_grp [-start _id <start id#>][-cnt <count>]
-monitor_interval <interval> [<repetition count>] [-time_zone <time zone>]
```

Options and parameters

-monitor_type qos_ldev
> Displays the QoS target LDEV performance information that is collected every one second on the storage system. If the command is run immediately after setting the LDEV as a QoS target, the LDEV might not be displayed in the command output. In this case, rerun the command after a few seconds.

-monitor_type qos_grp
> Displays the QoS group performance information that is collected every one second on the storage system. If the command is run immediately after creating the QoS group, the QoS group might not be displayed in the command output. In this case, rerun the command after a few seconds.

[-qos_grp_id <qos group id>]
> Specifies a QoS group ID. This option displays the performance information on the LDEV in the QoS group specified with <qos group id>. This option is available only when `-monitor_type qos_ldev` is specified.

[-start _id <start id#>]
> Specifies a minimum ID of the resource with which the display is started in decimal.
>
> If `-monitor_type qos_ldev` is specified, specify a minimum LDEV ID with which the display is started.
>
> If `-monitor_type qos_grp` is specified, specify a minimum QoS group ID with which the display is started.

[-cnt <count>]
> Specifies a maximum number (1-65535) of the displayed resource in decimal. This option displays the resource information by the number specified with <count> in increasing order from the smallest ID of the resources that meet the display conditions. If the number of resources that meet the display conditions is less than specified with <count>, the resources that meet all requirements are displayed.
>
> If `-monitor_type qos_ldev` is specified, a maximum number of the displayed LDEVs is specified.
>
> If `-monitor_type qos_grp` is specified, a maximum number of the displayed QoS groups is specified.

-monitor_interval <interval> [<repetition count>]
> Specifies the repetition interval for monitoring the performance information in 100 ms as a decimal number (1-36000). (0.1[s]-1[hour])
>
> If <repetition count> (1-200000) is specified, the monitoring is performed the number of times specified with <repetition count>.
>
> If <repetition count> is not specified, the monitoring is performed until the user terminates the command with CTL+C/SIGTERM.
>
> <repetition count> specifies a decimal number.

If the number of displayed resources is large relative to the command execution rate, all information might not be collected within the repetition interval specified with <interval>. In that case, reduce the number of displayed resources or increase the value of <interval>.

[-time_zone <time zone>]

Specifies the time zone to be displayed in MONITOR_TIME. If this option is not specified, the time in the time zone set for the storage system is displayed.

Example 1

Of the LDEVs belonging to the QoS group ID: 1, displaying the performance information on at least two LDEVs whose LDEV number is 200 or more every one second. The number of times to display is 3 times.

```
# raidcom monitor resource -monitor_type qos_ldev -qos_grp_id 1 -start_id 200 -cnt 2 -
monitor_interval 10 3
LDEV# RCV_CMD(IOps) RCV_CMD_TRANS(KBps) THROUGHPUT(IOps)DATA_TRANS(KBps) RESPONSE(us)
MONITOR_TIME
200 3840 15 2650 10 5 2019-07-31T10:15:20
201 3840 15 2650 10 5 2019-07-31T10:15:20
200 3840 15 2650 10 5 2019-07-31T10:15:21
201 3840 15 2650 10 5 2019-07-31T10:15:21
200 3840 15 2650 10 5 2019-07-31T10:15:22
201 3840 15 2650 10 5 2019-07-31T10:15:22
```

Description of each column in output example:

LDEV#

Displays the numbers of LDEVs included in the QoS group.

RCV_CMD(IOps)

Displays the numbers of received commands.

RCV_CMD_TRANS(KBps)

Displays the data transfer volumes of received commands.

THROUGHPUT(IOps)

Displays the throughput per second.

DATA_TRANS(KBps)

Displays the data transfer volumes per second.

RESPONSE(us)

Displays the response time.

MONITOR_TIME

Displays the monitoring collection time.

Chapter 5: Configuration setting commands

Example 2

Displaying the performance information on at least two QoS groups whose QoS group ID is 200 or more every one second. The number of times to display is 3 times.

```
# raidcom monitor resource -monitor_type qos_grp -start_id 200 -cnt 2 -
monitor_interval 10 3
QID RCV_CMD(IOps) RCV_CMD_TRANS(KBps) THROUGHPUT(IOps) DATA_TRANS(KBps)RESPONSE(us)
MONITOR_TIME
200 3840 15 2650 10 5 2019-07-31T10:15:20
201 3840 15 2650 10 5 2019-07-31T10:15:20
200 3840 15 2650 10 5 2019-07-31T10:15:21
201 3840 15 2650 10 5 2019-07-31T10:15:21
200 3840 15 2650 10 5 2019-07-31T10:15:22
201 3840 15 2650 10 5 2019-07-31T10:15:22
```

Description of each column in output example:

QID

Displays the QoS group ID in decimal.

raidcom add nvm_subsystem

Supported storage systems:

- VSP 5000 series

- VSP E1090

Creates an NVM subsystem.

> **Note:** If a failure occurs in a processor on an MP blade or a controller, this command cannot be run even on an MP blade or a controller on which no failure occurs.

When you run this command, the processing is run asynchronously with the command input. Verify that the processing is complete by using the `raidcom get command_status` command. There is a limit to the number of asynchronous commands that can be accepted simultaneously. For details, see .

Syntax

```
raidcom add nvm_subsystem -nvm_subsystem_id <nvm subsystem id> [-
nvm_subsystem_name <nvm subsystem name>] [-host_mode <host mode> [-
host_mode_opt <host mode option> …]] [-namespace_security <enable |
disable>] -request_id auto
```

Options and parameters

-nvm_subsystem_id <nvm subsystem id>
>   Specifies the NVM subsystem ID.

[-nvm_subsystem_name <nvm subsystem name>]
>   Specifies the NVM subsystem name.
>
>   If this option is omitted, the name is automatically set by the system.

[-host_mode <host mode>]
>   Specifies the host mode.
>
>   Specify the following character strings. For more information, see the *Provisioning Guide for Open Systems*.
>
>   - LINUX
>   - VMWARE
>   - VMWARE_EX
>   - AIX

[-host_mode_opt <host mode option> …]
>   Specifies the host mode option. For more information, see the *Provisioning Guide for Open Systems*.
>
>   If this option is omitted, the host mode option is disabled.

[-namespace_security <enable | disable>]
>   Specifies the setting to disable or enable use of the namespace security.
>
>   - enable: Enable the namespace security.
>   - disable: Disable the namespace security.
>
>   If this option is omitted, the namespace security is enabled.

-request_id auto
>   <request#> is a request ID assigned each time the command is run. For more information, see Request ID function (on page 246).

Examples

Creating an NVM subsystem whose NVM subsystem ID is 1 and whose NVM subsystem name is `my_nvm_subsystem`.

```
# raidcom add nvm_subsystem -nvm_subsystem_id 1 -nvm_subsystem_name
my_nvm_subsystem -request_id auto
```

Creating an NVM subsystem whose NVM subsystem ID is 1 and for which the namespace security is disabled.

```
# raidcom add nvm_subsystem -nvm_subsystem_id 1 -namespace_security
disable -request_id auto
```

raidcom modify nvm_subsystem

Supported storage systems:

- VSP 5000 series

- VSP E1090

Changes an NVM subsystem.

The NVM subsystem name, the host mode, the host mode option, and the namespace security can be changed.

> 📄 **Note:** If a failure occurs in a processor on an MP blade or a controller, this command cannot be run even on an MP blade or a controller on which no failure occurs.

When you run this command, the processing is run asynchronously with the command input. Verify that the processing is complete by using the `raidcom get command_status` command. There is a limit to the number of asynchronous commands that can be accepted simultaneously. For details, see Maximum number of acceptable asynchronous commands (on page 228).

Syntax

```
raidcom modify nvm_subsystem {-nvm_subsystem_id <nvm subsystem id>} {-
rename_nvm_subsystem <nvm subsystem nick name> | -host_mode <host mode>
| -set_host_mode_opt <host mode option> … | -reset_host_mode_opt | -
namespace_security <enable | disable>} -request_id auto
```

Options and parameters

-nvm_subsystem_id <nvm subsystem id>
 Specifies the NVM subsystem ID.

-rename_nvm_subsystem <nvm subsystem nickname>
 Specifies the NVM subsystem name to which it is changed.

[-host_mode <host mode>]
 Specifies the host mode.

 Specify the following character strings. For more information, see the *Provisioning Guide for Open Systems*.

- LINUX

- VMWARE

- VMWARE_EX

- AIX

-set host_mode_opt <host mode option> …
 Specifies the host mode option.

Specify the host mode option for `<host mode option>`. If you specify this option, all host mode options that you do not specify are cleared. For more information, see the *Provisioning Guide for Open Systems*.

-reset_host_mode_opt

Clears all host mode options.

For more information about the host mode option, see the *Provisioning Guide for Open Systems*.

-namespace_security <enable | disable>

Specifies the setting to disable or enable use of the namespace security.

- enable: Enable the namespace security.

- disable: Disable the namespace security.

-request_id auto

<request#> is a request ID assigned each time the command is run. For more information, see Request ID function (on page 246).

Examples

Changing the NVM subsystem name with the NVM subsystem ID 1 to `my_subsystem`.

```
# raidcom modify nvm_subsystem -nvm_subsystem_id 1 -rename_nvm_subsystem
my_subsystem -request_id auto
```

Changing the namespace security for the NVM subsystem ID 1 to disable.

```
# raidcom modify nvm_subsystem -nvm_subsystem_id 1 -namespace_security
disable -request_id auto
```

raidcom delete nvm_subsystem

Supported storage systems:

- VSP 5000 series

- VSP E1090

Deletes an NVM subsystem.

> 📄 **Note:** If a failure occurs in a processor on an MP blade or a controller, this command cannot be run even on an MP blade or a controller on which no failure occurs.

When you run this command, the processing is run asynchronously with the command input. Verify that the processing is complete by using the `raidcom get command_status` command. There is a limit to the number of asynchronous commands that can be accepted simultaneously. For details, see Maximum number of acceptable asynchronous commands (on page 228).

Syntax

```
raidcom delete nvm_subsystem {-nvm_subsystem_id <nvm subsystem id>} -
request_id auto
```

Options and parameters

-nvm_subsystem_id <nvm subsystem id>
> Specifies the NVM subsystem ID.

-reset_host_mode_opt
> Clears all host mode options.

> For more information about the host mode option, see the *Provisioning Guide for Open Systems*.

Example

Deleting the NVM subsystem whose NVM subsystem ID is 1.

```
# raidcom delete nvm_subsystem -nvm_subsystem_id 1 -request_id auto
```

raidcom get nvm_subsystem

Supported storage systems:

- VSP 5000 series

- VSP E1090

Displays the NVM subsystem information.

Syntax

```
raidcom get nvm_subsystem [-nvm_subsystem_id <nvm subsystem id>] [-key
<keyword>]
```

Options and parameters

-nvm_subsystem_id <nvm subsystem id>
> Specifies the NVM subsystem ID.

[-key <keyword>]

> Specifies one of the following values:

- opt: Displays the information about the subsystem NQN of the specified NVM subsystem.

- namespace: Displays the information about the namespace created on the specified NVM subsystem.

- port: Displays the information about the port registered on the specified NVM subsystem.

- detail: Displays the information about the implemented and the unimplemented NVM subsystems.

- undefined: Displays the unimplemented NVM subsystem ID.

Example

Displaying the NVM subsystem list.

```
# raidcom get nvm_subsystem
NVMSS_ID RGID NVMSS_NAME SECURITY T10PI
HMD HMO_BITs
1 0 nvmss_id_00001(default_name) ENABLE DISABLE
VMWARE -
```

Description of each column in output example:

NVMSS_ID

> Displays the NVM subsystem ID.

RGID

> Displays the resource group ID.

NVMSS_NAME

> Displays the NVM subsystem name.

SECURITY

> Displays the setting to disable or enable use of the namespace security.

- ENABLE: Enable the namespace security.

- DISABLE: Disable the namespace security.

- – (hyphen): This information is not available.

T10PI

> Displays the setting to disable or enable the T10 PI mode.

- ENABLE: Enable the T10 PI mode.

- DISABLE: Disable the T10 PI mode.

- – (hyphen): This information is not available.

HMD

> Displays the data transfer volumes per second.

> - ▪ – (hyphen): This information is not available.

HMO_BITs

> Displays BIT to set the host mode option for the host group. If multiple host mode options are specified, list the specified host mode options separated by a colon (:). If no host mode options are specified, a hyphen (-) is displayed. For more information about BIT, see the *Provisioning Guide for Open Systems*.

Example

Displaying the information about a subsystem NQN of the NVM subsystem ID 1.

```
# raidcom get nvm_subsystem -nvm_subsystem_id 1 -key opt
NVMSS_ID NVMSS_NAME NVMSS_NQN
1 nvmss_id_00001(default_name) nqn.xxxxx
```

Description of each column in output example:

NVMSS_NQN

> Displays the subsystem NQN.

Example

Displaying the information about a namespace registered on the NVM subsystem with the NVM subsystem ID 1.

```
# raidcom get nvm_subsystem -nvm_subsystem_id 1 -key namespace
NVMSS_ID NVMSS_NAME NSID LDEVID
1 nvmss_id_00001(default_name) 1 1
1 nvmss_id_00001(default_name) 2 2
```

Description of each column in output example:

NSID

> Displays the namespace ID.

LDEVID

> Displays the LDEV ID.

Example

Displaying the information about an NVM subsystem port registered on the NVM subsystem with the NVM subsystem ID 1.

```
# raidcom get nvm_subsystem -nvm_subsystem_id 1 -key port
NVMSS_ID NVMSS_NAME PORT
1 nvmss_id_00001(default_name) CL1-A
1 nvmss_id_00001(default_name) CL1-B
```

Description of each column in output example:

PORT

Displays the port number.

Examples

Displays the list of the implemented and the unimplemented NVM subsystems.

```
# raidcom get nvm_subsystem -key detail
NVMSS_ID RGID NVMSS_NAME SECURITY T10PI
HMD HMO_BITs
0 0 nvmss_id_00001(default_name) ENABLE DISABLE
VMWARE 64
1 0 - - -
- -
. . .
```

Displays the list of the unimplemented NVM subsystem ID.

```
# raidcom get nvm_subsystem -key undefined
NVMSS_ID
2
3
5
. . .
```

raidcom add nvm_subsystem_port

Supported storage systems:

- VSP 5000 series

- VSP E1090

Adds an NVM subsystem port to the NVM subsystem.

📄 **Note:** If a failure occurs in a processor on an MP blade or a controller, this command cannot be run even on an MP blade or a controller on which no failure occurs.

When you run this command, the processing is run asynchronously with the command input. Verify that the processing is complete by using the `raidcom get command_status` command. There is a limit to the number of asynchronous commands that can be accepted simultaneously. For details, see Maximum number of acceptable asynchronous commands (on page 228).

Syntax

```
raidcom add nvm_subsystem_port {-nvm_subsystem_id <nvm subsystem id>} -
port <port#> -request_id auto
```

Options and parameters

-nvm_subsystem_id <nvm subsystem id>

Specifies the NVM subsystem ID.

-port <port#>

Specifies the port number.

-request_id auto

<request#> is a request ID assigned each time the command is run. For more information, see Request ID function (on page 246).

Example

Adding the port: CL1-A to the NVM subsystem with the NVM subsystem ID 1.

```
# raidcom add nvm_subsystem_port -nvm_subsystem_id 1 -port CL1-A -
request_id auto
```

raidcom delete nvm_subsystem_port

Supported storage systems:

- VSP 5000 series
- VSP E1090

Deletes an NVM subsystem port from the NVM subsystem.

📄 **Note:** If a failure occurs in a processor on an MP blade or a controller, this command cannot be run even on an MP blade or a controller on which no failure occurs.

When you run this command, the processing is run asynchronously with the command input. Verify that the processing is complete by using the `raidcom get command_status` command. There is a limit to the number of asynchronous commands that can be accepted simultaneously. For details, see Maximum number of acceptable asynchronous commands (on page 228).

Syntax

```
raidcom delete nvm_subsystem_port {-nvm_subsystem_id <nvm subsystem id>}
-port <port#> -request_id auto
```

Options and parameters

-nvm_subsystem_id \<nvm subsystem id>
>   Specifies the NVM subsystem ID.

-port \<port#>
>   Specifies the port number.

-request_id auto
>   \<request#> is a request ID assigned each time the command is run. For more
>   information, see Request ID function (on page 246).

Example

Deleting the port: CL1-A from the NVM subsystem with the NVM subsystem ID 1.

```
# raidcom delete nvm_subsystem_port -nvm_subsystem_id 1 -port CL1-A -
request_id auto
```

raidcom get nvm_subsystem_port

Supported storage systems:

- VSP 5000 series
- VSP E1090

Displays the NVM subsystem port information.

Syntax

```
raidcom get nvm_subsystem_port {-port <port> | -nvm_subsystem_id <nvm
subsystem id>}
```

Options and parameters

-nvm_subsystem_id \<nvm subsystem id>
>   Specifies the NVM subsystem ID.

-port \<port>
>   Specifies the port number.

Example

Displaying the information about an NVM subsystem port registered on the NVM subsystem
ID 1.

```
# raidcom get nvm_subsystem_port -nvm_subsystem_id 1
PORT NVMSS_ID NVMSS_NAME
CL1-A 1 my_nvm_subsystem
CL1-C 1 my_nvm_subsystem
```

Description of each column in output example:

PORT
Displays the port number.

NVMSS_ID
Displays the NVM subsystem ID.

NVMSS_NAME
Displays the NVM subsystem name.

Example

Displaying the information about an NVM subsystem on which the port: CL1-A is registered as an NVM subsystem port.

```
# raidcom get nvm_subsystem_port -port CL1-A
PORT NVMSS_ID NVMSS_NAME
CL1-A 1 my_nvm_subsystem
CL1-A 2 NVMSS_2
```

raidcom add host_nqn

Supported storage systems:

- VSP 5000 series
- VSP E1090

Registers a host NQN on the NVM subsystem.

> 📄 **Note:** If a failure occurs in a processor on an MP blade or a controller, this command cannot be run even on an MP blade or a controller on which no failure occurs.

When you run this command, the processing is run asynchronously with the command input. Verify that the processing is complete by using the `raidcom get command_status` command. There is a limit to the number of asynchronous commands that can be accepted simultaneously. For details, see Maximum number of acceptable asynchronous commands (on page 228).

Syntax

```
raidcom add host_nqn {-nvm_subsystem_id <nvm subsystem id>} -host_nqn
<host_nqn> -request_id auto
```

Options and parameters

-nvm_subsystem_id <nvm subsystem id>
Specifies the NVM subsystem ID.

-host_nqn <host_nqn>
    Specifies the host NQN.

-request_id auto
    <request#> is a request ID assigned each time the command is run. For more information, see Request ID function (on page 246).

Example

Adding a host NQN: `nqn.xxx` to the NVM subsystem with the NVM subsystem ID 1.

```
# raidcom add host_nqn -nvm_subsystem_id 1 -host_nqn nqn.xxx -request_id
auto
```

raidcom modify host_nqn

Supported storage systems:

- VSP 5000 series
- VSP E1090

Specifies a host NQN nickname.

If the specified host NQN is registered on multiple NVM subsystems, the setting is applied to all NVM subsystems.

> 📄 **Note:** If a failure occurs in a processor on an MP blade or a controller, this command cannot be run even on an MP blade or a controller on which no failure occurs.

When you run this command, the processing is run asynchronously with the command input. Verify that the processing is complete by using the `raidcom get command_status` command. There is a limit to the number of asynchronous commands that can be accepted simultaneously. For details, see Maximum number of acceptable asynchronous commands (on page 228).

Syntax

```
raidcom modify host_nqn -host_nqn <host_nqn> -nvm_subsystem_id <nvm
subsystem id> {-set_host_name <host nick name> | -reset_host_name} -
request_id auto
```

Options and parameters

-nvm_subsystem_id <nvm subsystem id>
    Specifies the NVM subsystem ID.

-host_nqn <host_nqn>
    Specifies the host NQN.

-set_host_name <host nick name>
> Specifies the host NQN nickname.

-reset_host_name
> Deletes the host NQN nickname.

-request_id auto
> <request#> is a request ID assigned each time the command is run. For more information, see Request ID function (on page 246).

Example

Changing the host NQN nickname from `nqn.xxx` to `my_host`.

```
# raidcom modify host_nqn -host_nqn nqn.xxx -nvm_subsystem_id 1 -
set_host_name my_host -request_id auto
```

raidcom delete host_nqn

Supported storage systems:

- VSP 5000 series
- VSP E1090

Deletes a host NQN from the NVM subsystem.

> 📄 **Note:** If a failure occurs in a processor on an MP blade or a controller, this command cannot be run even on an MP blade or a controller on which no failure occurs.

When you run this command, the processing is run asynchronously with the command input. Verify that the processing is complete by using the **raidcom get command_status** command. There is a limit to the number of asynchronous commands that can be accepted simultaneously. For details, see Maximum number of acceptable asynchronous commands (on page 228).

Syntax

```
raidcom delete host_nqn {-nvm_subsystem_id <nvm subsystem id>} -host_nqn
<host_nqn> -request_id auto
```

Options and parameters

-nvm_subsystem_id <nvm subsystem id>
> Specifies the NVM subsystem ID.

-host_nqn <host_nqn>
> Specifies the host NQN.

-request_id auto

        <request#> is a request ID assigned each time the command is run. For more information, see <u>Request ID function (on page 246)</u>.

Example

Deleting the host NQN: `nqn.xxx` from the NVM subsystem ID 1.

```
# raidcom delete host_nqn -nvm_subsystem_id 1 -host_nqn nqn.xxx -
request_id auto
```

raidcom get host_nqn

Supported storage systems:

- VSP 5000 series
- VSP E1090

Displays the host NQN information.

Syntax

```
raidcom get host_nqn {-nvm_subsystem_id <nvm subsystem id>} [-key
<keyword>]
```

Options and parameters

-nvm_subsystem_id <nvm subsystem id>

        Specifies the NVM subsystem ID.

[-key <keyword>]

        Specifies a display keyword:

- opt: Displays the detailed information about the host NQN.

Examples

Displaying the information about a host NQN registered on the NVM subsystem ID 1.

```
# raidcom get host_nqn -nvm_subsystem_id 1
NVMSS_ID NVMSS_NAME HOST_NQN
1 my_nvm_subsystem nqn.xxx
1 my_nvm_subsystem nqn.yyy
```

Displaying the detailed information about a host NQN registered on the NVM subsystem ID 1.

```
# raidcom get host_nqn -nvm_subsystem_id 1 -key opt
NVMSS_ID NVMSS_NAME NICK_NAME
HOST_NQN
1 my_nvm_subsystem my_host
```

Chapter 5: Configuration setting commands

```
nqn.xxx
1 my_nvm_subsystem host2 nqn.yyy
```

Description of each column in output example:

NVMSS_ID
Displays the NVM subsystem ID.

NVMSS_NAME
Displays the NVM subsystem name.

HOST_NQN
Displays the host NQN.

NICK_NAME
Displays the nickname of the host NQN. If no nicknames are specified, a hyphen (-) is displayed.

raidcom add namespace_path

Supported storage systems:

▪ VSP 5000 series

▪ VSP E1090

Adds a host-namespace path by specifying a host NQN and a namespace registered on the NVM subsystem.

📄 **Note:** If a failure occurs in a processor on an MP blade or a controller, this command cannot be run even on an MP blade or a controller on which no failure occurs.

When you run this command, the processing is run asynchronously with the command input. Verify that the processing is complete by using the **raidcom get command_status** command. There is a limit to the number of asynchronous commands that can be accepted simultaneously. For details, see Maximum number of acceptable asynchronous commands (on page 228).

Syntax

```
raidcom add namespace_path {-nvm_subsystem_id <nvm subsystem id>} -ns_id
<ns#> -host_nqn <host_nqn> -request_id auto
```

Options and parameters

-nvm_subsystem_id <nvm subsystem id>
Specifies the NVM subsystem ID.

--ns_id <ns#>

      Specifies the namespace ID.

-host_nqn <host_nqn>

      Specifies the host NQN.

-request_id auto

      <request#> is a request ID assigned each time the command is run. For more information, see Request ID function (on page 246).

Example

Adding a host-namespace path between the host NQN: `nqn.xxx` and the namespace with the namespace ID 1 on the NVM subsystem with the NVM subsystem ID 1.

```
# raidcom add namespace_path -nvm_subsystem_id 1 -ns_id 1 -host_nqn
nqn.xxx -request_id auto
```

raidcom delete namespace_path

Supported storage systems:

- VSP 5000 series
- VSP E1090

Deletes a host-namespace path.

> **Note:** If a failure occurs in a processor on an MP blade or a controller, this command cannot be run even on an MP blade or a controller on which no failure occurs.

When you run this command, the processing is run asynchronously with the command input. Verify that the processing is complete by using the `raidcom get command_status` command. There is a limit to the number of asynchronous commands that can be accepted simultaneously. For details, see Maximum number of acceptable asynchronous commands (on page 228).

Syntax

```
raidcom delete namespace_path {-nvm_subsystem_id <nvm subsystem id>} -
ns_id <ns#> -host_nqn <host_nqn> -request_id auto
```

Options and parameters

-nvm_subsystem_id <nvm subsystem id>

      Specifies the NVM subsystem ID.

--ns_id <ns#>

      Specifies the namespace ID.

-host_nqn <host_nqn>

>   Specifies the host NQN.

-request_id auto

>   <request#> is a request ID assigned each time the command is run. For more information, see Request ID function (on page 246).

Example

Deleting a host-namespace path between the host NQN: `nqn.xxx` and the namespace with the namespace ID 1 on the NVM subsystem with the NVM subsystem ID 1.

```
# raidcom delete namespace_path -nvm_subsystem_id 1 -ns_id 1 -host_nqn
nqn.xxxx -request_id auto
```

raidcom get namespace_path

Supported storage systems:

▪ VSP 5000 series

▪ VSP E1090

Displays a host-namespace path.

Syntax

```
raidcom get namespace_path {-nvm_subsystem_id <nvm subsystem id> [-ns_id
<ns#>]}
```

Options and parameters

-nvm_subsystem_id <nvm subsystem id>

>   Specifies the NVM subsystem ID.

--ns_id <ns#>

>   Specifies the namespace ID.

Examples

Displaying the information about a host-namespace path for the namespace with the namespace ID 1 on the NVM subsystem with the NVM subsystem ID 1.

```
# raidcom get namespace_path -nvm_subsystem_id 1 -ns_id 1
NVMSS_ID NVMSS_NAME NSID LDEV# HOST_NQN
1 nvmss_id_00001(default_name) 1 10 nqn.xxx
```

Displaying the information about a host-namespace path for the NVM subsystem with the NVM subsystem ID 1.

```
# raidcom get namespace_path -nvm_subsystem_id 1
NVMSS_ID NVMSS_NAME NSID LDEV# HOST_NQN
1 nvmss_id_00001(default_name) 1 10 nqn.xxx
1 nvmss_id_00001(default_name) 2 11 nqn.xxx
```

Description of each column in output example:

NVMSS_ID
> Displays the NVM subsystem ID.

NVMSS_NAME
> Displays the NVM subsystem name.

NSID
> Displays the namespace ID.

LDEV#
> Displays the LDEV ID.

HOST_NQN
> Displays the host NQN.

raidcom add namespace

Supported storage systems:

- VSP 5000 series
- VSP E1090

Assigns an LDEV to a namespace that is created on the specified NVM subsystem.

This command is run asynchronously with the command input. Verify that the operation is complete by using the **raidcom get command_status** command.

> **Note:** If a failure occurs in a processor on an MP blade or a controller, this command cannot be run even on an MP blade or a controller on which no failure occurs.

Syntax

```
raidcom add namespace {-nvm_subsystem_id <nvm subsystem id>} {-ns_id
<ns#> | -ns_id auto} -ldev_id <ldev#> -request_id auto
```

Options and parameters

-nvm_subsystem_id <nvm subsystem id>

> Specifies the NVM subsystem ID.

{-ns_id <ns#> | -ns_id auto}

> To create a namespace, specifies whether to assign the namespace ID manually or automatically.

-ldev_id <ldev#>

> Specifies the LDEV number (0 to 65279).

-request_id auto

> <request#> is a request ID assigned each time the command is run. For more information, see Request ID function (on page 246).

Example

Assigning the LDEV with the LDEV number 10 to the namespace with the namespace ID automatically assigned that is created on the NVM subsystem with the NVM subsystem ID 1.

```
# raidcom add namespace -nvm_subsystem_id 1 -ns_id auto -ldev_id 10 -
request_id auto
```

raidcom delete namespace

Supported storage systems:

- VSP 5000 series
- VSP E1090

Deletes a namespace.

This command is run asynchronously with the command input. Verify that the operation is complete by using the `raidcom get command_status` command.

> 📄 **Note:** If a failure occurs in a processor on an MP blade or a controller, this command cannot be run even on an MP blade or a controller on which no failure occurs.

Syntax

```
raidcom delete namespace {-nvm_subsystem_id <nvm subsystem id>} -ns_id
<ns#> -request_id auto
```

Options and parameters

-nvm_subsystem_id <nvm subsystem id>

> Specifies the NVM subsystem ID.

-ns_id <ns#>
Specifies the namespace ID.

-request_id auto
<request#> is a request ID assigned each time the command is run. For more information, see Request ID function (on page 246).

Example

Deleting the namespace with the namespace ID 1 on the NVM subsystem with the NVM subsystem ID 1.

```
# raidcom delete namespace -nvm_subsystem_id 1 -ns_id 1 -request_id auto
```

raidcom modify namespace

Supported storage systems:

- VSP 5000 series
- VSP E1090

Changes a namespace. The namespace nickname can be changed.

This command is run asynchronously with the command input. Verify that the operation is complete by using the `raidcom get command_status` command.

> **Note:** If a failure occurs in a processor on an MP blade or a controller, this command cannot be run even on an MP blade or a controller on which no failure occurs.

Syntax

```
raidcom modify namespace {-nvm_subsystem_id <nvm subsystem id>} -ns_id
<ns#> {-set_namespace_name <namespace name> | -reset_namespace_name} -
request_id auto
```

Options and parameters

-nvm_subsystem_id <nvm subsystem id>
Specifies the NVM subsystem ID.

-ns_id <ns#>
Specifies the namespace ID.

-set_namespace_name <namespace name>
Specifies the namespace nickname to which it is changed.

-reset_namespace_name
Deletes the namespace nickname.

Chapter 5: Configuration setting commands

-request_id auto

> <request#> is a request ID assigned each time the command is run. For more information, see Request ID function (on page 246).

Example

Adding a nickname: `my_namespace` to the namespace with the namespace ID 1 on the NVM subsystem with the NVM subsystem ID 1.

```
# raidcom modify namespace -nvm_subsystem_id 1 -ns_id 1 -
set_namespace_name my_namespace -request_id auto
```

raidcom get namespace

Supported storage systems:

- VSP 5000 series
- VSP E1090

Displays the namespace information.

Syntax

```
raidcom get namespace {-nvm_subsystem_id <nvm subsystem id>} [-ns_id
<ns#>] [-key <keyword>]
```

Options and parameters

-nvm_subsystem_id <nvm subsystem id>

> Specifies the NVM subsystem ID.

-ns_id <ns#>

> Specifies the namespace ID.

[-key <keyword>]

> Specifies a display keyword:
>
> - opt: Displays the detailed information about the namespace.

Example

Displaying the information about the namespace with the NVM subsystem ID 1.

```
# raidcom get namespace -nvm_subsystem_id 1
NVMSS_ID NVMSS_NAME NSID LDEVID CAPACITY(BLK)
1 nvmss_id_00001(default_name) 1 10 16777216
1 nvmss_id_00001(default_name) 2 11 16777216
```

Description of each column in output example:

NVMSS_ID

Displays the NVM subsystem ID.

NVMSS_NAME

Displays the NVM subsystem name.

NSID

Displays the namespace ID.

LDEVID

Displays the LDEV ID.

CAPACITY(BLK)

Displays the namespace capacity in block sized increments.

Example

Displaying the detailed information about the namespace with the NVM subsystem ID 1.

```
# raidcom get namespace -nvm_subsystem_id 1 -key opt
NVMSS_ID NVMSS_NAME NSID LDEVID CAPACITY(BLK)
NS_NAME
1 nvmss_id_00001(default_name) 1 10 16777216
my_namespace
1 nvmss_id_00001(default_name) 2 11 16777216
namespace2
```

NS_NAME

Displays the namespace nickname.

If the namespace nickname is not assigned, a hyphen (-) is displayed.

Hitachi Vantara