

Hitachi Unified Storage Replication User Guide

FASTFIND LINKS

- [Changes in this revision](#)
- [Changes in this revision](#)
- [Document organization](#)
- [Contents](#)

© 2012-2015 Hitachi, Ltd. All rights reserved.

No part of this publication may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying and recording, or stored in a database or retrieval system for any purpose without the express written permission of Hitachi, Ltd. and Hitachi Data Systems Corporation (hereinafter referred to as "Hitachi").

Hitachi, Ltd. and Hitachi Data Systems reserve the right to make changes to this document at any time without notice and assume no responsibility for its use. Hitachi, Ltd. and Hitachi Data Systems products and services can only be ordered under the terms and conditions of Hitachi Data Systems' applicable agreements.

All of the features described in this document may not be currently available. Refer to the most recent product announcement or contact your local Hitachi Data Systems sales office for information on feature and product availability.

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

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

All other trademarks, service marks, and company names are properties of their respective owners.



Contents

Preface	xxiii
Intended audience	xxiv
Product version	xxiv
Product Abbreviations	xxiv
Document revision level	xxv
Changes in this revision	xxvi
Document organization	xxvii
Related documents	xxx
Document conventions	xxxii
Convention for storage capacity values	xxxii
Accessing product documentation	xxxiii
Getting help	xxxiii
Comments	xxxiii
1 Replication overview	1-1
ShadowImage® In-system Replication	1-2
Key features and benefits	1-2
Copy-on-Write Snapshot	1-4
Key features and benefits	1-4
TrueCopy® Remote Replication	1-6
Key features and benefits	1-6
TrueCopy® Extended Distance	1-8
Key features and benefits	1-8
TrueCopy® Modular Distributed	1-10
Differences between ShadowImage and Snapshot	1-12
Comparison of ShadowImage and Snapshot	1-13
Redundancy	1-14
2 ShadowImage In-system Replication theory of operation	2-1
ShadowImage In-system Replication software	2-2
Hardware and software configuration	2-2

How ShadowImage works	2-3
Volume pairs (P-VOLs and S-VOLs).	2-4
Creating pairs.	2-6
Initial copy operation	2-7
Automatically split the pair following pair creation	2-7
MU number	2-8
Splitting pairs.	2-8
Re-synchronizing pairs	2-9
Re-synchronizing normal pairs	2-11
Quick mode	2-12
Restore pairs	2-12
Re-synchronizing for split or split pending pair.	2-13
Re-synchronizing for suspended pair.	2-14
Suspending pairs	2-14
Deleting pairs.	2-14
Differential Management Logical Unit (DMLU)	2-15
Ownership of P-VOLs and S-VOLs.	2-16
Command devices.	2-17
Consistency group (CTG).	2-18
ShadowImage pair status.	2-19
Interfaces for performing ShadowImage operations	2-21
 3 Installing ShadowImage.	 3-1
System requirements	3-2
Supported platforms	3-3
Installing ShadowImage.	3-4
Enabling/disabling ShadowImage	3-6
Uninstalling ShadowImage	3-7
 4 ShadowImage setup	 4-1
Planning and design.	4-2
Plan and design workflow.	4-2
Copy frequency	4-2
Copy lifespan	4-3
Lifespan based on backup requirements	4-3
Lifespan based on business uses	4-4
Establishing the number of copies.	4-4
Ratio of S-VOLs to P-VOL	4-5
Requirements and recommendations for volumes.	4-6
RAID configuration for ShadowImage volumes	4-7
Operating system considerations and restrictions.	4-8
Identifying P-VOL and S-VOL in Windows	4-8
Volume mapping with CCI	4-9
AIX	4-9

Microsoft Cluster Server (MSCS)	4-9
Veritas Volume Manager (VxVM).	4-9
Windows 2000	4-9
Windows Server	4-9
Linux and LVM configuration	4-10
Concurrent use with Volume Migration	4-11
Concurrent use with Cache Partition Manager	4-12
Concurrent use of Dynamic Provisioning	4-12
Concurrent use of Dynamic Tiering	4-16
Windows Server and Dynamic Disk	4-16
UNMAP Short Length Mode	4-16
Limitations of Dirty Data Flush Number	4-16
VMware and ShadowImage configuration	4-17
Creating multiple pairs in the same P-VOL	4-19
Load balancing function.	4-19
Enabling Change Response for Replication Mode	4-19
Calculating maximum capacity.	4-19
Configuration.	4-22
Setting up primary, secondary volumes	4-22
Location of P-VOLs and S-VOLs	4-23
Locating multiple volumes within same drive column	4-23
Pair status differences when setting multiple pairs	4-24
Drive type P-VOLs and S-VOLs	4-24
Locating P-VOLs and DMLU	4-24
Setting up the DMLU	4-25
Removing the designated DMLU	4-27
Add the designated DMLU capacity	4-28
Setting the ShadowImage I/O switching mode	4-29
Setting the system tuning parameter	4-31

5 Using ShadowImage. 5-1

ShadowImage workflow	5-2
Prerequisites for creating the pair	5-2
Pair assignment	5-3
Confirming pair status	5-4
Setting the copy pace	5-5
Create a pair	5-6
Split the ShadowImage pair	5-9
Resync the pair	5-10
Delete a pair	5-12
Edit a pair	5-13
Restore the P-VOL	5-13
Use the S-VOL for tape backup, testing, reports	5-15

6	Monitoring and troubleshooting ShadowImage	6-1
	Monitor pair status	6-2
	Monitoring pair failure	6-4
	Monitoring of pair failure using a script	6-5
	Troubleshooting	6-6
	Pair failure	6-7
	Path failure	6-9
	Cases and solutions using the DP-VOLs	6-9
7	Copy-on-Write Snapshot theory of operation	7-1
	Copy-on-Write Snapshot software	7-2
	Hardware and software configuration	7-2
	How Snapshot works	7-3
	Volume pairs — P-VOLs and V-VOLs	7-4
	Creating pairs	7-5
	Creating pairs options	7-6
	Splitting pairs	7-6
	Re-synchronizing pairs	7-7
	Restoring pairs	7-7
	Deleting pairs	7-10
	DP pools	7-10
	Consistency Groups (CTG)	7-11
	Command devices	7-13
	Differential data management	7-14
	Snapshot pair status	7-15
	Interfaces for performing Snapshot operations	7-18
8	Installing Snapshot	8-1
	System requirements	8-2
	Supported platforms	8-3
	Installing or uninstalling Snapshot	8-4
	Installing Snapshot	8-4
	Uninstalling Snapshot	8-6
	Enabling or disabling Snapshot	8-7
9	Snapshot setup	9-1
	Planning and design	9-2
	Plan and design workflow	9-3
	Assessing business needs	9-3
	Copy frequency	9-4
	Selecting a reasonable time between Snapshots	9-4
	Establishing how long a copy is held (copy lifespan)	9-5
	Lifespan based on backup requirements	9-5

Lifespan based on business uses	9-5
Establishing the number of V-VOLs	9-6
DP pool capacity	9-6
DP pool consumption	9-7
Determining DP pool capacity	9-7
Replication data	9-7
Management information	9-8
Calculating DP pool size	9-12
Requirements and recommendations for Snapshot Volumes	9-14
Pair assignment	9-15
RAID configuration for volumes assigned to Snapshot	9-16
Pair resynchronization and releasing	9-16
Locating P-VOLS and DP pools	9-17
Command devices	9-19
Operating system host connections	9-20
Veritas Volume Manager (VxVM)	9-20
AIX	9-20
Linux and LVM configuration	9-20
Tru64 UNIX and Snapshot configuration	9-20
Cluster and path switching software	9-20
Windows Server and Snapshot configuration	9-20
Microsoft Cluster Server (MSCS)	9-21
Windows Server and Dynamic Disk	9-21
UNMAP Short Length Mode	9-22
Windows 2000	9-22
VMware and Snapshot configuration	9-23
Array functions	9-25
Identifying P-VOL and V-VOL volumes on Windows	9-25
Volume mapping	9-26
Concurrent use of Cache Partition Manager	9-26
Concurrent use of Dynamic Provisioning	9-26
Concurrent use of Dynamic Tiering	9-28
User data area of cache memory	9-29
Limitations of dirty data flush number	9-30
Load balancing function	9-30
Enabling Change Response for Replication Mode	9-30
Configuring Snapshot	9-31
Configuration workflow	9-31
Setting up the DP pool	9-31
Setting the replication threshold (optional)	9-32
Setting up the Virtual Volume (V-VOL) (manual method) (optional)	9-34
Deleting V-VOLs	9-34
Setting up the command device (optional)	9-35
Setting the system tuning parameter (optional)	9-36

10 Using Snapshot	10-1
Snapshot workflow	10-2
Confirming pair status	10-3
Create a Snapshot pair to back up your volume	10-6
Splitting pairs	10-10
Updating the V-VOL	10-11
Making the host recognize secondary volumes with no volume number	10-12
Restoring the P-VOL from the V-VOL	10-13
Deleting pairs and V-VOLs	10-14
Editing a pair	10-15
Use the V-VOL for tape backup, testing, and reports	10-16
Tape backup recommendations	10-17
Restoring data from a tape backup	10-18
Quick recovery backup	10-20
11 Monitoring and troubleshooting Snapshot	11-1
Monitoring Snapshot	11-2
Monitoring pair status	11-2
Monitoring pair failure	11-3
Monitoring pair failure using a script	11-4
Monitoring DP pool usage	11-5
Expanding DP pool capacity	11-7
Other methods for lowering DP pool usage	11-8
Troubleshooting	11-9
Pair failure	11-9
DP Pool capacity exceeds replication threshold value	11-14
Cases and solutions using DP-VOLs	11-15
Recovering from pair failure due to a hardware failure	11-15
Confirming the event log	11-16
Snapshot fails in a TCE S-VOL - Snapshot P-VOL cascade configuration	11-17
Message contents of the Event Log	11-18
12 TrueCopy Remote Replication theory of operation	12-1
TrueCopy Remote Replication	12-2
How TrueCopy works	12-2
Typical environment	12-3
Volume pairs	12-3
Remote Path	12-4
Differential Management LU (DMLU)	12-4
Command devices	12-5
Consistency group (CTG)	12-6
TrueCopy interfaces	12-6
Typical workflow	12-7
Operations overview	12-7

13	Installing TrueCopy Remote	13-1
	System requirements	13-2
	Installation procedures	13-3
	Installing TrueCopy Remote	13-3
	Enabling or disabling TrueCopy Remote	13-4
	Uninstalling TrueCopy Remote	13-5
	Prerequisites	13-5
	To uninstall TrueCopy	13-5
14	TrueCopy Remote setup	14-1
	Planning for TrueCopy	14-2
	The planning workflow	14-3
	Planning disk arrays	14-3
	Planning volumes	14-4
	Volume pair recommendations	14-4
	Volume expansion	14-5
	Operating system recommendations and restrictions	14-7
	Host time-out	14-7
	P-VOL, S-VOL recognition by same host on VxVM, AIX®, LVM	14-7
	Setting the Host Group options	14-8
	Windows 2000 Servers	14-11
	Windows Server	14-11
	Dynamic Disk in Windows Server	14-12
	UNMAP Short Length Mode	14-12
	Identifying P-VOL and S-VOL in Windows	14-12
	VMware and TrueCopy configuration	14-13
	Volumes recognized by the same host restrictions	14-14
	Concurrent use of Dynamic Provisioning	14-14
	Concurrent use of Dynamic Tiering	14-17
	Load balancing function	14-18
	Enabling Change Response for Replication Mode	14-18
	Calculating supported capacity	14-19
	Setup procedures	14-20
	Setting up the DMLU	14-20
	Adding or changing the remote port CHAP secret (iSCSI only)	14-23
	Setting the remote path	14-24
	Changing the port setting	14-26
	Remote path design	14-27
	Determining remote path bandwidth	14-28
	Measuring write-workload	14-28
	Optimal I/O performance versus data recovery	14-31
	Remote path requirements, supported configurations	14-32
	Management LAN requirements	14-33
	Remote path requirements	14-33

Remote path configurations	14-34
Remote path configurations for Fibre Channel	14-34
Direct connection	14-35
Fibre Channel switch connection 1	14-36
Fibre Channel switch connection 2	14-37
One-Path-Connection between Arrays	14-39
Fibre Channel extender	14-40
Path and switch performance	14-41
Port transfer rate for Fibre Channel	14-41
Remote path configurations for iSCSI	14-42
Direct iSCSI connection	14-43
Single LAN switch, WAN connection	14-44
Multiple LAN switch, WAN connection	14-45
Connecting the WAN Optimization Controller	14-46
Switches and WOCs connection (1)	14-47
Switches and WOCs connection (2)	14-48
Two sets of a pair connected via the switch and WOC (1)	14-49
Two sets of a pair connected via the switch and WOC (2)	14-51
Using the remote path — best practices	14-52
Remote processing	14-52
Supported connections between various models of arrays	14-54
Restrictions on supported connections	14-54

15 Using TrueCopy Remote 15-1

TrueCopy operations	15-2
Pair assignment	15-2
Checking pair status	15-3
Pairs Operations	15-3
Creating pairs	15-3
Prerequisite information and best practices	15-4
Copy pace	15-4
Fence level	15-6
Operation when the fence level is “never”	15-6
Creating pairs procedure	15-7
Splitting pairs	15-9
Resynchronizing pairs	15-10
Swapping pairs	15-11
Editing pairs	15-11
Deleting pairs	15-12
Deleting the remote path	15-13
Operations work flow	15-14
TrueCopy ordinary split operation	15-15
TrueCopy ordinary pair operation	15-17
Data migration use	15-18

TrueCopy disaster recovery	15-19
Resynchronizing the pair	15-20
Data path failure and recovery	15-20
Host server failure and recovery	15-21
Host timeout	15-22
Production site failure and recovery	15-22
Automatic switching using High Availability (HA) software	15-22
Manual switching	15-24
Special problems and recommendations	15-25
16 Monitoring and troubleshooting TrueCopy Remote	16-1
Monitoring and maintenance	16-2
Monitoring pair status	16-3
Monitoring pair failure	16-7
Monitoring of pair failure using a script	16-8
Monitoring the remote path	16-9
Troubleshooting	16-10
Pair failure	16-10
Restoring pairs after forcible release operation	16-10
Recovering from a pair failure	16-11
Cases and solutions Using the DP-VOLs	16-12
17 TrueCopy Extended Distance theory of operation	17-1
How TrueCopy Extended Distance works	17-2
Configuration overview	17-2
Operational overview	17-3
Typical environment	17-5
TCE Components	17-6
Remote path	17-7
Alternative path	17-7
Confirming the path condition	17-7
Port connection and topology for Fibre Channel Interface	17-8
Port transfer rate for Fibre Channel	17-8
DP pools	17-9
Guaranteed write order and the update cycle	17-10
Extended update cycles	17-11
Consistency Group (CTG)	17-11
Command Devices	17-14
TCE interfaces	17-14
18 Installing TrueCopy Extended	18-1
TCE system requirements	18-2
Installation procedures	18-3

Installing TCE	18-3
Enabling or disabling TCE	18-5
Uninstalling TCE	18-6
19 TrueCopy Extended Distance setup	19-1
Plan and design — sizing DP pools and bandwidth	19-2
Plan and design workflow	19-3
Assessing business needs — RPO and the update cycle	19-3
Measuring write-workload	19-4
Collecting write-workload data	19-4
DP pool size	19-5
DP pool consumption	19-6
How much capacity TCE consumes	19-6
Determining bandwidth	19-7
Performance design	19-8
Plan and design — remote path	19-11
Remote path requirements	19-12
Management LAN requirements	19-13
Remote data path requirements	19-13
Remote path configurations	19-14
Fibre Channel	19-14
Direct connection	19-15
Fibre Channel switch connection 1	19-16
Fibre Channel switch connection 2	19-17
One-Path-Connection between Arrays	19-19
Fibre Channel extender connection	19-20
Port transfer rate for Fibre Channel	19-21
iSCSI	19-22
Direct connection	19-23
Single LAN switch, WAN connection	19-24
Connecting Arrays via Switches	19-25
WAN optimization controller (WOC) requirements	19-26
Combining the Network between Arrays	19-27
Connections with multiple switches, WOCs, and WANs	19-28
Multiple array connections with LAN switch, WOC, and single WAN	19-29
Multiple array connections with LAN switch, WOC, and two WANs	19-30
Local and remote array connection by the switches and WOC	19-31
Using the remote path — best practices	19-32
Plan and design—disk arrays, volumes and operating systems	19-33
Planning workflow	19-34
Supported connections between various models of arrays	19-35
Connecting HUS with AMS500, AMS1000, or AMS2000	19-35
Planning volumes	19-36
Prerequisites and best practices for pair creation	19-36

Volume pair and DP pool recommendations	19-37
Operating system recommendations and restrictions	19-38
Host time-out	19-38
P-VOL, S-VOL recognition by same host on VxVM, AIX®, LVM.	19-38
HP server	19-38
Windows Server 2000	19-40
Windows Server 2003 or 2008	19-40
Windows Server and TCE configuration volume mount	19-41
Volumes to be recognized by the same host	19-41
Identifying P-VOL and S-VOL in Windows	19-41
Dynamic Disk in Windows Server	19-42
UNMAP Short Length Mode	19-42
VMware and TCE configuration	19-43
Changing the port setting	19-44
Concurrent use of Dynamic Provisioning	19-45
Concurrent use of Dynamic Tiering	19-49
Load balancing function	19-49
Enabling Change Response for Replication Mode	19-49
User data area of cache memory	19-50
Setup procedures	19-51
Pair Assignment	19-51
Setting up DP pools	19-54
Setting the replication threshold (optional)	19-54
Setting the cycle time	19-56
Adding or changing the remote port CHAP secret	19-57
Setting the remote path	19-58
Deleting the remote path	19-60
Operations work flow	19-61

20 Using TrueCopy Extended 20-1

TCE operations	20-2
Checking pair status	20-2
Creating the initial copy	20-2
Create pair procedure	20-3
Splitting a pair	20-6
Resynchronizing a pair	20-8
Swapping pairs	20-9
Editing pairs	20-10
Deleting pairs	20-11
Example scenarios and procedures	20-12
CLI scripting procedure for S-VOL backup	20-12
Scripted TCE, Snapshot procedure	20-14
Procedure for swapping I/O to S-VOL when maintaining local disk array.	20-18
Procedure for moving data to a remote disk array	20-19

Example procedure for moving data	20-21
Process for disaster recovery	20-21
Takeover processing	20-21
Swapping P-VOL and S-VOL.	20-22
Failback to the local disk array	20-22
21 Monitoring and troubleshooting TrueCopy Extended	21-1
Monitoring and maintenance.	21-2
Monitoring pair status	21-3
Monitoring DP pool capacity	21-10
Monitoring DP pool usage	21-10
Checking DP pool status or changing threshold value of the DP pool	21-10
Adding DP pool capacity	21-10
Processing when DP pool is Exceeded.	21-11
Monitoring the remote path	21-14
Changing remote path bandwidth.	21-14
Monitoring cycle time.	21-14
Changing cycle time	21-15
Changing copy pace.	21-16
Monitoring synchronization	21-16
Monitoring synchronization using CCI	21-17
Monitoring synchronization using Navigator 2	21-18
Routine maintenance	21-19
Deleting a volume pair	21-19
Deleting the remote path.	21-20
TCE tasks before a planned remote disk array shutdown	21-20
TCE tasks before updating firmware	21-20
Troubleshooting	21-21
Correcting DP pool shortage	21-23
Cycle copy does not progress	21-25
Message contents of event log	21-26
Correcting disk array problems	21-26
.	21-27
Deleting replication data on the remote array	21-28
Delays in settling of S-VOL Data.	21-28
DP-VOLs troubleshooting.	21-29
Correcting resynchronization errors	21-30
Using the event log	21-32
Miscellaneous troubleshooting.	21-33
22 TrueCopy Modular Distributed theory of operation	22-1
TrueCopy Modular Distributed overview	22-2
Distributed mode.	22-3

23	Installing TrueCopy Modular Distributed	23-1
	TCMD system requirements	23-2
	Installation procedures	23-3
	Installing TCMD	23-3
	Uninstalling TCMD	23-5
	Enabling or disabling TCMD	23-7
24	TrueCopy Modular Distributed setup	24-1
	Cautions and restrictions	24-2
	Precautions when writing from the host to the Hub array or Edge array	24-2
	Setting the remote paths for each HUS in which TCMD is installed	24-2
	Setting the remote path: HUS 100 (TCMD install) and AMS2000/500/1000	24-3
	Adding the Edge array in the configuration of the set TCMD	24-3
	Configuring TCMD adding an array to configuration (TCMD not used)	24-4
	Changing an Array (part of a TCMD configuration with a different array)	24-4
	Precautions when setting the remote port CHAP secret	24-5
	Recommendations	24-6
	Configuration guidelines	24-7
	Environmental conditions	24-9
	Setting the remote path	24-12
	Deleting the remote path	24-15
	Setting the remote port CHAP secret	24-16
25	Using TrueCopy Modular Distributed	25-1
	Configuration example: centralized backup using TCE	25-2
	Perform the aggregation backup	25-2
	Data delivery using TrueCopy Remote Replication	25-3
	Creating data delivery configuration	25-3
	Create a pair in data delivery configuration	25-7
	Executing the data delivery	25-10
	Setting the distributed mode	25-13
	Changing the Distributed mode to Hub from Edge	25-13
	Changing the Distributed Mode to Edge from Hub	25-14
26	Troubleshooting TrueCopy Modular Distributed	26-1
	Troubleshooting	26-2
27	Cascading replication products	27-1
	Cascading ShadowImage	27-2
	Cascading ShadowImage with Snapshot	27-2
	Restriction when performing restoration	27-3
	I/O switching function	27-3

Performance when cascading P-VOL of ShadowImage with Snapshot . . .	27-4
Cascading a ShadowImage S-VOL with Snapshot.	27-7
Restrictions	27-7
Cascading restrictions with ShadowImage P-VOL and S-VOL.	27-11
Cascading ShadowImage with TrueCopy.	27-11
Cascading a ShadowImage S-VOL	27-14
Cascading a ShadowImage P-VOL and S-VOL	27-16
Cascading restrictions on TrueCopy with ShadowImage and Snapshot .	27-18
Cascading restrictions on TCE with ShadowImage	27-18
Cascading Snapshot.	27-19
Cascading Snapshot with ShadowImage	27-19
Cascading restrictions with ShadowImage P-VOL and S-VOL.	27-19
Cascading Snapshot with TrueCopy Remote	27-21
Cascading a Snapshot P-VOL	27-22
Cascading a Snapshot V-VOL	27-22
Configuration restrictions on the Cascade of TrueCopy with Snapshot .	27-25
Cascade restrictions on TrueCopy with ShadowImage and Snapshot . .	27-26
Cascading Snapshot with True Copy Extended.	27-27
Restrictions on cascading TCE with Snapshot.	27-28
Cascading TrueCopy Remote	27-29
Cascading with ShadowImage	27-30
Cascade overview	27-30
Cascade configurations	27-30
Configurations with ShadowImage P-VOLs	27-31
Configurations with ShadowImage S-VOLs	27-35
Configurations with ShadowImage P-VOLs and S-VOLs.	27-37
Cascading a TrueCopy P-VOL with a ShadowImage P-VOL	27-39
Volume shared with P-VOL on ShadowImage and P-VOL on TrueCopy.	27-40
Pair Operation restrictions for cascading TrueCopy/ShadowImage . . .	27-42
Cascading a TrueCopy S-VOL with a ShadowImage P-VOL	27-43
Volume shared with P-VOL on ShadowImage and S-VOL on TrueCopy.	27-44
Volume shared with TrueCopy S-VOL and ShadowImage P-VOL	27-45
Cascading a TrueCopy P-VOL with a ShadowImage S-VOL	27-46
Volume shared with S-VOL on ShadowImage and P-VOL on TrueCopy.	27-48
Volume Shared with TrueCopy P-VOL and ShadowImage S-VOL.	27-49
Volume Shared with S-VOL on TrueCopy and ShadowImage.	27-50
Cascading TrueCopy with ShadowImage P-VOL and S-VOL 1:1	27-51
Simultaneous cascading of TrueCopy with ShadowImage	27-52
Cascading TrueCopy with ShadowImage P-VOL and S-VOL 1:3	27-53
Cascade with a ShadowImage S-VOL (P-VOL: S-VOL=1:3)	27-54
Simultaneous cascading of TrueCopy with ShadowImage	27-55
Swapping when cascading TrueCopy and ShadowImage Pairs.	27-56
Creating a backup with ShadowImage	27-57
Cascading with Snapshot.	27-59

Cascade overview	27-59
Cascade configurations	27-59
Configurations with Snapshot P-VOLs	27-60
Cascading with a Snapshot V-VOL	27-62
Cascading a TrueCopy P-VOL with a Snapshot P-VOL	27-63
Volume shared with P-VOL on Snapshot and P-VOL on TrueCopy.	27-64
V-VOLs number of Snapshot.	27-65
Cascading a TrueCopy S-VOL with a Snapshot P-VOL	27-66
Volume Shared with Snapshot P-VOL and TrueCopy S-VOL	27-67
V-VOLs number of Snapshot.	27-68
Cascading a TrueCopy P-VOL with a Snapshot V-VOL	27-69
Transition of statuses of TrueCopy and Snapshot pairs	27-70
Swapping when cascading a TrueCopy pair and a Snapshot pair	27-72
Creating a backup with Snapshot	27-74
When to create a backup	27-75
Cascading with ShadowImage and Snapshot	27-76
Cascade restrictions of TrueCopy with Snapshot and ShadowImage	27-76
Cascade restrictions of TrueCopy S-VOL with Snapshot V-VOL	27-76
Cascading restrictions	27-77
Concurrent use of TrueCopy and ShadowImage or Snapshot.	27-77
Cascading TCE.	27-78
Cascading with Snapshot	27-78
V-VOLs number of Snapshot.	27-78
DP pool	27-78
Cascading a TCE P-VOL with a Snapshot P-VOL	27-78
Cascading a TCE S-VOL with a Snapshot P-VOL	27-81
Snapshot cascade configuration local and remote backup operations	27-85
TCE with Snapshot cascade restrictions	27-90

A ShadowImage In-system Replication reference information . . . A-1

ShadowImage general specifications	A-2
Operations using CLI	A-6
Installing and uninstalling ShadowImage	A-7
Installing ShadowImage	A-7
Uninstalling ShadowImage	A-8
Enabling or disabling ShadowImage	A-8
Setting the DMLU	A-9
Setting the ShadowImage I/O switching mode.	A-11
Setting the system tuning parameter	A-11
ShadowImage operations	A-12
Confirming pairs status	A-12
Creating ShadowImage pairs	A-12
Splitting ShadowImage pairs	A-14

Re-synchronizing ShadowImage pairs	A-15
Restoring the P-VOL	A-15
Deleting ShadowImage pairs	A-16
Editing pair information	A-16
Creating ShadowImage pairs that belong to a group	A-17
Splitting ShadowImage pairs that belong to a group	A-18
Sample back up script for Windows	A-19
Operations using CCI	A-20
Setting up CCI	A-21
Setting the command device	A-21
Setting LU mapping	A-22
Defining the configuration definition file	A-23
Setting the environment variable	A-26
ShadowImage operations using CCI	A-28
Confirming pair status	A-29
Creating pairs (paircreate)	A-30
Pair creation using a consistency group	A-31
Splitting pairs (pairsplit)	A-32
Resynchronizing pairs (pairresync)	A-32
Releasing pairs (pairsplit -S)	A-33
Pair, group name differences in CCI and Navigator 2	A-33
I/O switching mode feature	A-34
I/O Switching Mode feature operating conditions	A-35
Specifications	A-36
Recommendations	A-37
Enabling I/O switching mode	A-38
Recovery from a drive failure	A-39
.	A-40

B Copy-on-Write Snapshot reference information B-1

Snapshot specifications	B-2
Operations using CLI	B-6
Installing and uninstalling Snapshot	B-7
Important prerequisite information	B-7
Installing Snapshot	B-7
Uninstalling Snapshot	B-9
Enabling or disabling Snapshot	B-10
Operations for Snapshot configuration	B-11
Setting the DP pool	B-11
Setting the replication threshold (optional)	B-11
Setting the V-VOL (optional)	B-13
Setting the system tuning parameter (optional)	B-14
Performing Snapshot operations	B-14
Creating Snapshot pairs using CLI	B-14
Splitting Snapshot Pairs	B-15

Re-synchronizing Snapshot Pairs	B-16
Restoring V-VOL to P-VOL using CLI	B-16
Deleting Snapshot pairs	B-17
Changing pair information	B-18
Creating multiple Snapshot pairs that belong to a group using CLI	B-18
Sample back up script for Windows	B-20
Operations using CCI	B-21
Setting up CCI	B-22
Setting the command device	B-22
Setting LU Mapping information	B-23
Defining the configuration definition file	B-24
Setting the environment variable	B-27
Performing Snapshot operations	B-29
Confirming pair status	B-30
Pair create operation	B-30
Pair creation using a consistency group	B-31
Pair Splitting	B-32
Re-synchronizing Snapshot pairs	B-32
Restoring a V-VOL to the P-VOL	B-33
Deleting Snapshot pairs	B-34
Pair and group name differences in CCI and Navigator 2	B-34
Performing Snapshot operations using raidcom	B-35
Setting the command device for raidcom command	B-35
Creating the configuration definition file for raidcom command	B-36
Setting the environment variable for raidcom command	B-36
Creating a snapshotset and registering a P-VOL	B-36
Creating a Snapshot data	B-36
Example of creating the Snapshot data of the multiple P-VOLs	B-37
Discarding Snapshot data	B-37
Restoring Snapshot data	B-38
Changing the Snapshotset name	B-38
Volume number mapping to the Snapshot data	B-38
Volume number un-mapping of the Snapshot data	B-39
Changing the volume assignment number of the Snapshot data	B-39
Deleting the snapshotset	B-39
Using Snapshot with Cache Partition Manager	B-40
.	B-42

C TrueCopy Remote Replication reference information C-1

TrueCopy specifications	C-2
Operations using CLI	C-5
Installation and setup	C-6
Installing	C-6
Enabling or disabling	C-6
Uninstalling	C-7

Setting the Differential Management Logical Unit	C-8
Release a DMLU	C-8
Adding a DMLU capacity	C-8
Setting the remote port CHAP secret	C-9
Setting the remote path	C-9
Deleting the remote path.	C-13
Pair operations	C-15
Displaying status for all pairs	C-15
Displaying detail for a specific pair	C-15
Creating a pair	C-16
Creating pairs belonging to a group	C-17
Splitting a pair	C-17
Resynchronizing a pair	C-17
Swapping a pair	C-18
Deleting a pair	C-18
Changing pair information	C-19
Sample scripts.	C-20
Backup script	C-20
Pair-monitoring script	C-21
Operations using CCI	C-22
Setting up CCI.	C-23
Preparing for CCI operations.	C-24
Setting the command device	C-24
Setting LU mapping	C-25
Defining the configuration definition file	C-26
Setting the environment variable	C-29
Pair operations	C-30
Multiple CCI requests and order of execution.	C-30
Operations and pair status.	C-30
Confirming pair status	C-32
Creating pairs (paircreate)	C-32
Splitting pairs (pairsplit)	C-34
Resynchronizing pairs (pairresync)	C-34
Suspending pairs (pairsplit -R)	C-35
Releasing pairs (pairsplit -S)	C-35
Mounting and unmounting a volume.	C-35

D TrueCopy Extended Distance reference information D-1

TCE system specifications.	D-2
Operations using CLI	D-7
Installation and setup	D-8
Installing	D-8
Enabling and disabling	D-9
Un-installing TCE	D-10
Setting the DP pool.	D-10

Setting the replication threshold	D-10
Setting the cycle time	D-12
Setting mapping information	D-12
Setting the remote port CHAP secret.	D-13
Setting the remote path.	D-14
Deleting the remote path.	D-17
Pair operations.	D-18
Displaying status for all pairs	D-18
Displaying detail for a specific pair	D-18
Creating a pair	D-19
Splitting a pair	D-20
Resynchronizing a pair.	D-20
Swapping a pair	D-20
Deleting a pair	D-21
Changing pair information	D-22
Monitoring pair status	D-22
Confirming consistency group (CTG) status	D-23
Procedures for failure recovery	D-24
Displaying the event log.	D-24
Reconstructing the remote path	D-24
Sample script.	D-25
Operations using CCI	D-26
Setup.	D-26
Setting the command device	D-26
Setting mapping information	D-27
Defining the configuration definition file	D-28
Setting the environment variable	D-30
Pair operations.	D-32
Checking pair status	D-32
Creating a pair (paircreate)	D-33
Splitting a pair (pairsplit)	D-33
Resynchronizing a pair (pairresync).	D-34
Suspending pairs (pairsplit -R)	D-34
Releasing pairs (pairsplit -S).	D-35
Splitting TCE S-VOL/Snapshot V-VOL pair (pairsplit -mscas)	D-35
Confirming data transfer when status is PAIR	D-37
Pair creation/resynchronization for each CTG.	D-37
Response time of pairsplit command.	D-39
Pair, group name differences in CCI and Navigator 2.	D-42
TCE and Snapshot differences	D-42
Initializing Cache Partition when TCE and Snapshot are installed	D-43
Wavelength Division Multiplexing (WDM) and dark fibre.	D-45

E TrueCopy Modular Distributed reference information E-1

TCMD system specifications	E-2

.....	E-4
Operations using CLI	E-5
Installation and uninstalling	E-6
Installing TCMD	E-6
Un-installingTCMD	E-8
Enabling and disabling	E-9
Setting the Distributed Mode	E-10
Changing the Distributed mode to Hub from Edge	E-11
Changing the Distributed Mode to Edge from Hub	E-12
Setting the remote port CHAP secret	E-13
Setting the remote path	E-14
Deleting the remote path	E-18

Glossary

Index



Preface

Welcome to the Hitachi Unified Storage Replication User Guide.
This document describes how to use the Hitachi Unified Storage Replication software.

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

This preface includes the following information:

- ☐ [Intended audience](#)
- ☐ [Product version](#)
- ☐ [Changes in this revision](#)
- ☐ [Changes in this revision](#)
- ☐ [Document organization](#)
- ☐ [Related documents](#)
- ☐ [Document conventions](#)
- ☐ [Convention for storage capacity values](#)
- ☐ [Accessing product documentation](#)
- ☐ [Getting help](#)
- ☐ [Comments](#)

Intended audience

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

This document assumes the user has a background in data processing and understands storage systems and their basic functions, Microsoft Windows and its basic functions, and Web browsers and their basic functions.

Product version

This document applies to Hitachi Unified Storage firmware version 0981/D or later and to HSNM2 version 28.12 or later.

Replication products require the following firmware and HSNM2 versions (or later versions).

Product	Firmware Version	HSNM2 Version
ShadowImage	0915/B	21.50
Snapshot	0915/B	21.50
TrueCopy Remote	0916/A	22.00
TrueCopy Extended Distance	0916/A	21.60
TrueCopy Modular Distributed	See TCMD system requirements (page 23-2)	

Product Abbreviations

Product Abbreviation	Product Full Name
ShadowImage	ShadowImage In-system Replication
Snapshot	Copy-on-Write Snapshot
TrueCopy Remote	TrueCopy Remote Replication
TCE	TrueCopy Extended Distance
TCMD	TrueCopy Modular Distributed
Windows Server	Windows Server 2003, Windows Server 2008, and Windows Server 2012.

Changes in this revision

This release includes updates or changes to the following:

- Under [Create a Snapshot pair to back up your volume \(page 10-6\)](#), [Confirming pair status \(page A-29\)](#), [Creating Snapshot pairs using CLI \(page B-14\)](#), [Pair create operation \(page B-30\)](#), and [Creating a snapshotset and registering a P-VOL \(page B-37\)](#), added note about the names that are assigned automatically if a pair name is not specified.

Document organization

Thumbnail descriptions of the chapters are provided in the following table. Click the [chapter title](#) in the first column to go to that chapter. The first page of every chapter or appendix contains links to the contents.

Chapter/Appendix Title	Description
Chapter 1, Replication overview	Provides short descriptions of the Replication software and describes how they differ from each other.
Chapter 2, ShadowImage In-system Replication theory of operation	Provides descriptions of ShadowImage components and how they work together.
Chapter 3, Installing ShadowImage	Provides ShadowImage requirements and instructions for enabling ShadowImage.
Chapter 4, ShadowImage setup	Provides detailed planning and design information and configuration information.
Chapter 5, Using ShadowImage	Provides directions for the common tasks performed with ShadowImage.
Chapter 6, Monitoring and troubleshooting ShadowImage	Provides directions on how to monitor and troubleshoot ShadowImage.
Chapter 7, Copy-on-Write Snapshot theory of operation	Provides descriptions of Snapshot components and how they work together.
Chapter 8, Installing Snapshot	Provides Snapshot requirements and instructions for enabling Snapshot.
Chapter 9, Snapshot setup	Provides detailed planning and design information and configuration information.
Chapter 10, Using Snapshot	Provides directions for the common tasks performed with Snapshot.
Chapter 11, Monitoring and troubleshooting Snapshot	Provides directions on how to monitor and troubleshoot Snapshot.
Chapter 12, TrueCopy Remote Replication theory of operation	Provides descriptions of TrueCopy Remote components and how they work together.
Chapter 13, Installing TrueCopy Remote	Provides TrueCopy Remote requirements and instructions for enabling TrueCopy Remote.

Chapter/Appendix Title	Description
Chapter 14, TrueCopy Remote setup	Provides detailed planning and design information and configuration information.
Chapter 15, Using TrueCopy Remote	Provides directions for the common tasks performed with TrueCopy Remote and for disaster recovery.
Chapter 16, Monitoring and troubleshooting TrueCopy Remote	Provides directions on how to monitor and troubleshoot TrueCopy Remote.
Chapter 17, TrueCopy Extended Distance theory of operation	Provides descriptions of TrueCopy Extended components and how they work together.
Chapter 18, Installing TrueCopy Extended	Provides TrueCopy Extended requirements and instructions for enabling TrueCopy Remote.
Chapter 19, TrueCopy Extended Distance setup	Provides detailed planning and design information and configuration information.
Chapter 20, Using TrueCopy Extended	Provides directions for the common tasks performed with TrueCopy Extended and for disaster recovery.
Chapter 21, Monitoring and troubleshooting TrueCopy Extended	Provides directions on how to monitor and troubleshoot TrueCopy Extended.
Chapter 22, TrueCopy Modular Distributed theory of operation	Provides descriptions of TrueCopy Modular Distributed components and how they work together
Chapter 23, Installing TrueCopy Modular Distributed	Provides TrueCopy Modular Distributed requirements and instructions for enabling TrueCopy Modular Distributed
Chapter 24, TrueCopy Modular Distributed setup	Provides detailed planning and design information and configuration information.
Chapter 25, Using TrueCopy Modular Distributed	Provides directions for the common tasks performed with TrueCopy Modular Distributed.
Chapter 26, Troubleshooting TrueCopy Modular Distributed	Provides directions on how to monitor and troubleshoot TrueCopy Modular Distributed.
Chapter 27, Cascading replication products	Provides information on how to cascade the replication products with each other.
Appendix A, ShadowImage In-system Replication reference information	Provides specifications, how to use CLI, how to use CCI, enabling I/O switching, cascading with Snapshot, and cascading with TrueCopy
Appendix B, Copy-on-Write Snapshot reference information	Provides specifications, how to use CLI, how to use CCI, cascading with Snapshot, and cascading with TrueCopy
Appendix C, TrueCopy Remote Replication reference information	Provides specifications, how to use CLI, how to use CCI, Cascading with ShadowImage, Cascading with Snapshot, and Cascading with ShadowImage and Snapshot

Chapter/Appendix Title	Description
Appendix D, TrueCopy Extended Distance reference information	Provides specifications, how to use CLI, how to use CCI, Cascading with Snapshot, Initializing Cache Partition when TCE and Snapshot are installed, and Wavelength Division Multiplexing (WDM) and dark fibre.
Appendix E, TrueCopy Modular Distributed reference information	Provides specifications, and how to use CLI.



Replication also provides a command-line interface that lets you perform operations by typing commands from a command line. For information about using the Replication command line, refer to the *Hitachi Unified Storage Command Line Interface Reference Guide*.

Related documents

This Hitachi Unified Storage documentation set consists of the following documents.

Hitachi Unified Storage Firmware Release Notes, RN-91DF8304

Contains late-breaking information about the storage system firmware.

Hitachi Storage Navigator Modular 2 Release Notes, RN-91DF8305

Contains late-breaking information about the Storage Navigator Modular 2 software.

Read the release notes before installing and using this product. They may contain requirements and restrictions not fully described in this document, along with updates and corrections to this document.

Hitachi Unified Storage Getting Started Guide, MK-91DF8303

Describes how to get Hitachi Unified Storage systems up and running in the shortest period of time. For detailed installation and configuration information, refer to the Hitachi Unified Storage Hardware Installation and Configuration Guide.

Hitachi Unified Storage Hardware Installation and Configuration Guide, MK-91DF8273

Contains initial site planning and pre-installation information, along with step-by-step procedures for installing and configuring Hitachi Unified Storage systems.

Hitachi Unified Storage Hardware Service Guide, MK-91DF8302

Provides removal and replacement procedures for the components in Hitachi Unified Storage systems.

Hitachi Unified Storage Operations Guide, MK-91DF8275

Describes the following topics:

- Adopting virtualization with Hitachi Unified Storage systems
- Enforcing security with Account Authentication and Audit Logging
- Creating DP-Vols, standard volumes, Host Groups, provisioning storage, and utilizing spares
- Tuning storage systems by monitoring performance and using cache partitioning
- Monitoring storage systems using email notifications and Hi-Track
- Using SNMP Agent and advanced functions such as data retention and power savings
- Using functions such as data migration, volume expansion and volume shrink, RAID Group expansion, DP pool expansion, and mega VOLS

Hitachi Unified Storage Replication User Guide, MK-91DF8274 — **this document**

Describes how to use the four types of Hitachi replication software to meet your needs for data recovery:

- ShadowImage In-system Replication
- Copy-on-Write Snapshot
- TrueCopy Remote Replication
- TrueCopy Extended Distance

Hitachi Unified Storage Command Control Interface Installation and Configuration Guide, MK-91DF8306

Describes Command Control Interface installation, operation, and troubleshooting.

Hitachi Unified Storage Provisioning Configuration Guide, MK-91DF8277

Describes how to use virtual storage capabilities to simplify storage additions and administration.

Hitachi Unified Storage Command Line Interface Reference Guide, MK-91DF8276

Describes how to perform management and replication activities from a command line.

Document conventions

The following typographic conventions are used in this document.

Convention	Description
Bold	Indicates text on a window, other than the window title, including menus, menu options, buttons, fields, and labels. Example: Click OK .
<i>Italic</i>	Indicates a variable, which is a placeholder for actual text provided by you or the system. Example: copy <i>source-file target-file</i> Angled brackets (< >) are also used to indicate variables.
screen or code	Indicates text that is displayed on screen or entered by you. Example: <code># pairedisplay -g oradb</code>
< > angled brackets	Indicates a variable, which is a placeholder for actual text provided by you or the system. Example: <code># pairedisplay -g <group></code> Italic font is also used to indicate variables.
[] 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.
underline	Indicates the default value. Example: [<u>a</u> b]

This document uses the following symbols to draw attention to important safety and operational information.

Symbol	Meaning	Description
	Tip	Tips provide helpful information, guidelines, or suggestions for performing tasks more effectively.
	Note	Notes emphasize or supplement important points of the main text.
	Caution	Cautions indicate that failure to take a specified action could result in damage to the software or hardware.
	WARNING:	Warns that failure to take or avoid a specified action could result in severe conditions or consequences (for example, loss of data).

Convention for storage capacity values

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

Physical capacity unit	Value
1 KB	1,000 bytes
1 MB	1,000 KB or 1,000 ² bytes
1 GB	1,000 MB or 1,000 ³ bytes
1 TB	1,000 GB or 1,000 ⁴ bytes
1 PB	1,000 TB or 1,000 ⁵ bytes
1 EB	1,000 PB or 1,000 ⁶ bytes

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

Logical capacity unit	Value
1 block	512 bytes
1 KB	1,024 (2 ¹⁰) bytes
1 MB	1,024 KB or 1024 ² bytes
1 GB	1,024 MB or 1024 ³ bytes
1 TB	1,024 GB or 1024 ⁴ bytes
1 PB	1,024 TB or 1024 ⁵ bytes
1 EB	1,024 PB or 1024 ⁶ bytes

Accessing product documentation

The Hitachi Unified Storage user documentation is available on the HDS Support Portal: <https://portal.hds.com>. 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 Data Systems customer support staff is available 24 hours a day, seven days a week. If you need technical support, please log on to the HDS Support Portal for contact information: <https://portal.hds.com>

Comments

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

Thank you!

Replication overview

There are five types of Hitachi replication software applications designed to meet your needs for data recovery. The key topics in this chapter are:

- ❑ [ShadowImage® In-system Replication](#)
- ❑ [Copy-on-Write Snapshot](#)
- ❑ [TrueCopy® Remote Replication](#)
- ❑ [TrueCopy® Extended Distance](#)
- ❑ [TrueCopy® Modular Distributed](#)
- ❑ [Differences between ShadowImage and Snapshot](#)

ShadowImage® In-system Replication

Hitachi ShadowImage In-system Replication software is an In-System, full volume mirror or clone solution that is common to all Hitachi storage platforms. ShadowImage clone operations are host/application independent and have zero impact on application/storage performance or throughput.

You can create up to 8 clones on Hitachi Unified Storage systems. These full volume clones can be used for recovery, testing, application development or non disruptive tape backup. The nondestructive operation of ShadowImage increases the availability of revenue producing business applications.

For flexibility, ShadowImage is bundled with Copy-on-write Snapshot in the Hitachi Base Operating System M software bundle.

Key features and benefits

- Immediately available copies (clones) of the volumes being replicated for concurrent use by other applications or backups
- No host processing cycles required,
- Immediate restore to production if needed
- Whole volume replication
- No impact to production volumes
- Support for consistency groups
- Can be combined with remote replication products for off-site protection of data

ShadowImage uses local mirroring technology to create full-volume copies within the array. In a ShadowImage pair operation, all data blocks in the original data volume are sequentially copied onto the secondary volume upon creation, subsequent updates are incremental changes only.

The original and secondary data volumes remain synchronized until they are split. While synchronized, updates to the original data volume are continually mirrored to the secondary volume. When the secondary volume is split from the original volume, it contains a mirror image of the original volume at that point in time. That point in time can be application consistent when combined with application quiescing abilities

After the pair is split, the secondary volume can be used for offline testing or analytical purposes since there is no common data sharing with the original volume. Since there are no dependencies between the original and secondary volumes, each can be written to by separate hosts. Changes to both volumes are tracked so they can be re-synchronized in either direction and as an incremental only.

ShadowImage is recommended to create a Gold Copy to be used to recover in the event of a rolling disaster. There should be at least one copy on the recovery side and on the production side.

Figure 1-1 shows a typical ShadowImage system configuration.




Figure 1-1: ShadowImage configuration

Copy-on-Write Snapshot

An essential component of business continuity is the ability to quickly replicate data. Hitachi Copy-on-Write Snapshot software provides logical snapshot data replication within Hitachi storage systems for immediate use in decision support, software testing and development, data backup, or rapid recovery operations.

Copy-on-Write Snapshot rapidly creates up to 1024 point-in-time snapshot copies of any data volume within Hitachi storage systems, without impacting host service or performance levels. Since these snapshots only store the changed data blocks in the DP pool, the amount of storage capacity required for each snapshot copy is substantially smaller than the source volume. As a result, a significant savings is realized when compared with full cloning methods.

For flexibility, Copy-on-write Snapshot is bundled with ShadowImage in the Hitachi Base Operating System M software bundle.



NOTE: Copy-on-write Snapshot can be used together with TrueCopy Extended.

Key features and benefits

- Point-in-time copies of only changed blocks in the DP pool, not full volume, reducing storage consumption
- Instantaneous restore of just the differential data you need, back to the source volume
- Versioning of backups for easy restore
- RAID protection of all Copy-on-Write Snapshot copies
- Near-instant copy creation and deletion
- Can be integrated with industry-leading backup software applications
- Supports consistency groups
- Can be combined with application quiescing capabilities to create application-aware point-in-times
- Can be combined with remote replication for off site protection

[Figure 1-2 on page 1-5](#) shows a typical Snapshot system configuration.




Figure 1-2: Snapshot configuration

TrueCopy® Remote Replication

For the most mission-critical data situations, replication urgency and backup certainty of already saved data are of the utmost importance. Hitachi TrueCopy Remote Replication addresses these challenges with immediate and robust replication capabilities. This software is built with the same engineering expertise used to develop Hitachi remote replication software for enterprise-level storage environments.

TrueCopy Remote can be combined with ShadowImage or Snapshot, on either or both local and remote sites. These in-system copy tools allow restoration from one or more additional copies of critical data at a specific point in time.

With TrueCopy Remote, you receive confirmation of replication and achieve the highest level of replication integrity as compared to asynchronous replication. You can also adopt best practices, such as disaster recovery plan testing with online data.

Besides disaster recovery, TrueCopy backup copies can be used for development, data warehousing and mining, or migration applications.

Key features and benefits

- Used for distances within the same metro area, where latency of network is not a concern
- Provides the highest level of replication integrity, because its real time copies are the same as the originals.
- Can be used with ShadowImage Replication or Copy-on-Write Snapshot software
- Minimal performance impact on primary system
- Essential when a recovery point objective of zero must be maintained for business for regulatory reasons




Figure 1-3: TrueCopy Remote backup system

TrueCopy® Extended Distance

When both fast performance and geographical distance capabilities are vital, Hitachi TrueCopy Extended Distance (TCE) software for Hitachi Unified Storage provides bi-directional, long-distance, remote data protection. TrueCopy Extended Distance supports data copy, failover and multi generational recovery without affecting your applications.

TrueCopy Extended Distance maximizes bandwidth utilization and reduces cost by providing write-consistent incremental changes. TrueCopy Extended Distance software enables simple, easy-to-manage business continuity that is independent from complexities of host-based operating systems and applications.

ShadowImage and Snapshot are Hitachi Unified Storage copy solutions for replication within an array. These are effective solution for creating backups within a local site, which can be used to recover from failures such as data corruption in a certain volume in an array.

On the other hand, the TCE and TrueCopy Remote copy solution extends over two arrays. TCE has functionality to copy data from an array to another. Therefore, TCE can be a solution to restart business off-site, using another target array, which has backup of the original data at a certain point in time. Even if a disaster such as earthquake, hurricane, or terrorist attack occurs, the business can be restarted by the target array, since it is located far away from the local site where the business is run.

Key features and benefits

- Copies data over any distance without interrupting the application
- Provides a bi-directional remote data protection solution
- Delivers premier data integrity with minimal performance impact on the primary system
- Provides failover and fail-back recovery capabilities
- Operates in midrange environments to maximize the use and limit the cost of bandwidth by only copying incremental changes
- Can be combined with disk-based, point-in-time copies to provide for online DR testing without disrupting replication




Figure 1-4: TrueCopy Extended Distance backup system

TrueCopy® Modular Distributed

TrueCopy Modular Distributed (TCMD) expands the capabilities of Hitachi TrueCopy Extended Distance (TCED) by allowing up to eight (8) HUS systems to remotely replicate to a single HUS system. This configuration is referred to as a “fan-in” configuration and is typically used to enable the storage systems at remote sites to backup their critical data to a single storage system at the customer’s primary data center. Replication from the storage system at the primary data center to a maximum of 8 storage systems at remote sites may also be implemented and this is referred to as a “fan-out” configuration. Note that in the “fan-out” configuration the storage system in the primary data center will be replicating separate volumes to each of the remote storage systems.

Based on the remote replication capabilities of TCE, TCMD supports data copy, failover and multi generational recovery without affecting your applications. TCMD maximizes bandwidth utilization and reduces cost by providing write-consistent incremental changes in the same way that TCE does.

TCMD software enables simple, easy-to-manage business continuity that is independent from complexities of host-based operating systems and applications.

You can also employ TCMD and Snapshot at the same time for further data protection. This allows you to have backup copies within the arrays themselves for the master data on the local arrays and for the backup data on the remote array.

Since TCMD is an add-on function for TCE, you cannot employ TCMD as a stand-alone. You need to install TCE for TCMD to work.

Key features and benefits

- Establishes TrueCopy Extended Distance connections between up to 9 arrays
- Consolidates backup copies on up to 8 local arrays to a remote array
- Copies data over any distance without interrupting the application
- Provides a bi-directional remote data protection solution
- Delivers premier data integrity with minimal performance impact on the primary system
- Provides failover and fail-back recovery capabilities
- Operates in midrange environments to maximize the use and limit the cost of bandwidth by only copying incremental changes
- Can be combined with disk-based, point-in-time copies to provide for online DR testing without disrupting replication




Figure 1-5: Remote backup system using TCE and TCMD

Differences between ShadowImage and Snapshot

ShadowImage and Snapshot both create a duplicate copies within an array; however, their uses are different because their specifications and data assurance measures are different. Advantages, limitations, and the functions of ShadowImage and Snapshot are shown in [Figure 1-6](#) and [Table 1-1 on page 1-13](#).




Figure 1-6: ShadowImage and Snapshot array copy solution

Comparison of ShadowImage and Snapshot

Table 1-1 shows advantages, limitations, and the functions of ShadowImage and Snapshot.

Table 1-1: ShadowImage and Snapshot functions

Contents	ShadowImage	Snapshot
Advantages	- When a hardware failure occurs in the P-VOL (source), it has no effect on the S-VOL (target). - When a failure occurs in the S-VOL, it has no effect on the generations of other S-VOLs. - Access performance is only slightly lowered in comparison with ordinary cases because the P-VOL and S-VOL are independent asynchronous volumes.	- Amount of physical data to be used for the V-VOL is small because only the differential data is copied. - Up to 1024 snapshots per P-VOL can be created, for a maximum of 100,000. - The Dynamic Provisioning (DP) pool can be used by the two or more P-VOLs and the same number of V-VOLs by sharing between them; single instancing of its capacity can be done. - A pair creation/resynchronization is completed in a moment.
Limitations	- Only eight S-VOLs can be created per P-VOL. - The S-VOL must have the same capacity as the P-VOL. - A pair creation/resynchronization requires time because it accompanies data copying from the P-VOL to S-VOL.	- If there is a hardware failure in the P-VOL, all the V-VOLs associated with the P-VOL in which the failure has occurred are placed in the Failure status. - If there is a hardware failure or a shortage of the DP pool capacity in the DP pool, all the V-VOLs that use the DP pool in which the failure has occurred are placed in the Failure status. - Careful management of write rates must be done to ensure that space savings are maintained - When the V-VOL is accessed, the performance of the P-VOL can be affected because the V-VOL data is shared among the P-VOL and DP pool.
Uses	- Not recommended for backup for quick recovery (instantaneous recovery from multiple point in times greater than 8.) - Recommended for online backup when many I/O operations are required at night or an amount of data to be backed up is too large to be disposed during the night.	- Backup for quick recovery from multiple point in times. - To make a restoration quickly when software failure occurs, managing multiple backups (for example, by making backups every several hours and managing them according to their generations). - It is important to backup onto a tape device due to low redundancy. - Online backup.

Redundancy

Snapshot and ShadowImage are identical functions from the viewpoint of producing a duplicate copy of data within a array. While both technologies provide equal levels of protection against logical corruptions in the application, consideration must be given to the unlikely event of physical failure in the array. The duplicated volume (S-VOL) of ShadowImage is a full copy of the entire P-VOL data to a single volume; the duplicated volume (V-VOL) of Snapshot consists of the P-VOL data and only changed data saved in the DP pool. Therefore, when a hardware failure, such as a double failure of drives occurs in the P-VOL, a similar failure also occurs in the V-VOL and the pair status is changed to Failure (see [Volume pairs — P-VOLs and V-VOLs on page 7-4](#)).

The DP pool can be used by two or more P-VOLs and V-VOLs who share them. However, when a hardware failure occurs in the DP pool (such as a double failure of drives), similar failures occur in all the V-VOLs that use the DP pool and their pair statuses are changed to Failure. When the DP pool capacity is insufficient, all the V-VOLs which use the DP pool are placed in Failure status because the replication data cannot be saved and the pair relationship cannot be maintained. If the V-VOL is placed in the Failure status, data retained in the V-VOL cannot be restored.

When hardware failures occur in the DP pool and S-VOL during a restoration for both Snapshot and ShadowImage, the P-VOL being restored accepts no Read/Write instruction. The difference between Snapshot and ShadowImage in redundancy is shown in [Figure 1-7](#), [Figure 1-8 on page 1-15](#), and [Figure 1-9 on page 1-16](#).




Figure 1-7: P-VOL failure




Figure 1-8: DP pool S-VOL failures




Figure 1-9: DP pool S-VOL failures during restore operation

ShadowImage In-system Replication theory of operation

Hitachi ShadowImage In-system Replication software uses local mirroring technology to create a copy of any volume in the array. During copying, host applications can continue to read/write to and from the primary production volume. Replicated data volumes can be split as soon as they are created for use with other applications.

The key topics in this chapter are:

- ❑ [ShadowImage In-system Replication software](#)
- ❑ [Hardware and software configuration](#)
- ❑ [How ShadowImage works](#)
- ❑ [ShadowImage pair status](#)
- ❑ [Interfaces for performing ShadowImage operations](#)

ShadowImage In-system Replication software

Hitachi's ShadowImage uses local mirroring technology to create full-volume copies within the array. In a ShadowImage pair operation, all data blocks in the original data volume are sequentially copied onto the secondary volume.

The original and secondary data volumes remain synchronized until they are split. While synchronized, updates to the original data volume are continually and asynchronously mirrored to the secondary volume. When the secondary volume is split from the original volume, it contains a mirror image of the original volume at that point in time.

After the pair is split, the secondary volume can be used for offline testing or analytical purposes since there is no common data sharing with the original volume. Since there are no dependencies between the original and secondary volumes, each can be written to by separate hosts. Changes to both volumes are tracked so they can be re-synchronized.

Hardware and software configuration

The typical replication configuration includes a array, a host connected to the array, and software to configure and operate ShadowImage (management software). The host is connected to the array using fibre channel or iSCSI connections. The management software is connected to arrays via a management LAN.

The logical configuration of the array includes a command device, a Differential Management Logical Unit (DMLU), and primary data volumes (P-VOLs) belonging to the same group.

ShadowImage employs a primary volume, a secondary volume or volumes, and the Hitachi Storage Navigator Modular2 (Navigator 2) graphical user interface (GUI). Additional user functionality is made available through Navigator 2 Command-Line Interface (CLI) or Hitachi Command Control Interface (CCI).

[Figure 2-1 on page 2-3](#) shows a typical ShadowImage environment.




Figure 2-1: ShadowImage environment

The following sections describe how these components work together.

How ShadowImage works

The array contains and manages both the original and copied ShadowImage data. ShadowImage supports a maximum of 1,023 pairs for HUS 110, or 2,047 pairs for HUS 130 or HUS 150.

ShadowImage creates a duplicate volume of another volume. This volume “pair” is created when you:

- Select a volume that you want to replicate
- Identify another volume that will contain the copy
- Associate the primary and secondary volumes
- Copy all primary volume data to the secondary volume

When the initial copy is made, all data on the P-VOL is copied to the S-VOL. The P-VOL remains available for read/write I/O during the operation. Write operations performed on the P-VOL are always duplicated to the S-VOL.

When the pair is split, the primary volume continues being updated, but data in the secondary volume remains as it was at the time of the split. At this time:

- The secondary volume becomes available for read/write access by secondary host applications.

- Changes to primary and secondary volumes are tracked by differential bitmaps.
- The pair can be made identical again by re-synchronizing changes from primary-to-secondary, or secondary-to-primary volumes.

Volume pairs (P-VOLs and S-VOLs)

The array contains and manages both the original and copied ShadowImage data.

Each ShadowImage pair can consist of up to 8 secondary volumes (S-VOL). The primary and secondary volumes are located in the same array.

ShadowImage P-VOLs are the primary volumes, which contain the original data. ShadowImage S-VOLs are the secondary or mirrored volumes, which contain the duplicated data. During ShadowImage operations, the P-VOLs remain available to all hosts for read and write I/O operations (except during reverse resync). The S-VOLs become available for host access only after the pair has been suspended or deleted.

The S-VOL only becomes accessible to a host after the pair is split. When a pair is split, the pair status becomes split. While a pair is split, the array keeps track of changes to the P-VOL and S-VOL in differential bitmaps.

When the pair is re-synchronized, differential (changed) data in the P-VOL is copied to the S-VOL; then the S-VOL is again identical to the P-VOL. A reverse-resync can also be performed when you want to update the P-VOL with the S-VOL data. In a reverse resync, the differential data in the S-VOL is copied to the P-VOL.

A pair name can be assigned to each pair to make identification easy. The pair name should be the maximum of 31 characters and unique in the group. This pair name can be assigned when creating the pair and can be changed later. Once a pair name is assigned, a target pair can be specified by the pair name at the time of the pair operation. For changing the pair information see [Editing pair information on page A-15](#) or CLI operation.

ShadowImage supports a maximum of 1,023 pairs (HUS 110), 2,047 pairs (HUS 130/HUS 150).

One P-VOL can be shared by up to eight pairs. That is, a configuration of linking the maximum of eight S-VOLs to one P-VOL can be set.

ShadowImage operations can be performed from the UNIX®/PC host using CCI software and/or Navigator 2.

[Figure 2-2 on page 2-5](#) shows basic operations. [Figure 2-3 on page 2-6](#) shows pair operations using Storage Navigator Modular 2 GUI.




Figure 2-2: Basic ShadowImage Operations

Initial copy operation

Select whether to make an initial copy from the P-VOL to the S-VOL. The default is the initial copy.

The ShadowImage initial copy operation takes place when you create a new ShadowImage pair, as shown in [Figure 2-4](#). The ShadowImage initial copy operation copies all data on the P-VOL to the associated S-VOL. The P-VOL remains available to all hosts for read and write I/Os throughout the initial copy operation. Write operations performed on the P-VOL during the initial copy operation will always be duplicated to the S-VOL. The status of the pair is Synchronizing while the initial copy operation is in progress. The pair status changes to Paired when the initial copy is complete.

When creating pairs, you can select the pace for the initial copy operation(s): Slow, Medium, and Fast. The default setting is Medium.

If you select not to make the initial copy, the pair status is immediately changed to Paired and you must ensure that the data of the volume specified in the P-VOL and the S-VOL is the same.




Figure 2-4: Adding a ShadowImage pair

Automatically split the pair following pair creation

When creating a new ShadowImage pair, you can execute an initial copy in the background and perform pair creation and pair split continuously by specifying Quick Mode. If you execute the command, you will be able to make Read/Write access for the S-VOL immediately. The S-VOL data accessed by the host becomes the same as the P-VOL data at the time of command execution. The pair status under the initial copy is Split Pending. The pair status changes to Split when the initial copy is completed.

MU number

The MU numbers used in CCI can be specified. The MU number is the management number that is used for configurations where a single volume is shared among multiple pairs. You can specify any value from 0 to 39 by selecting **Manual**. The MU numbers already used by other ShadowImage pairs or SnapShot pairs which share the P-VOL cannot be specified. The free MU numbers are assigned in ascending order from MU number 1 by selecting **Automatic** (the default). The MU number is attached to the P-VOL. The MU number for the S-VOL is fixed as 0.



NOTE: If the MU numbers from 0 to 39 are already used, no more ShadowImage pairs can be created. When creating SnapShot pairs, specify the MU numbers from 40 and more. When creating SnapShot pairs, if you select **Automatic**, the MU numbers are assigned in descending order from 1032.

Splitting pairs

Split the pair to retain the backup data in the S-VOL. The ShadowImage splitting pairs operation splits the paired P-VOL and S-VOL, and changes the pair state of the P-VOL and S-VOL to Split. Once a pair is split, the subsequent operation to reflect the update for the P-VOL in the S-VOL stops and the backup data at the time of the split instruction is retained in the S-VOL.

When splitting pairs is performed the S-VOL becomes identical to the P-VOL and then provides full Read/Write access to the S-VOL.

Pair splitting options include:

- **Suspending Pairs Operation:** Split pair with Suspend operation in progress and force the pair into a failure state operation. You can suspend the pairing and change it to the Failure status. The copy processing of ShadowImage is a process to give a load to the array when the copy pace is Fast or Medium. This option is used when the copy processing of ShadowImage is forcibly suspended. Since the copy processing is suspended, the S-VOL data is incomplete. Furthermore, since the pair status is Failure, the Write access to the S-VOL cannot be performed. Once the ShadowImage pair is suspended, the entire P-VOL differential map is marked as the differential data. If the resynchronization operation is executed in the Failure status pair, the entire P-VOL is copied to the S-VOL. Since the resynchronization operation for the ShadowImage pair in the Split or Split Pending status only copies the difference, the required time is significantly shortened. However, the resynchronization operation for the pair in the Failure status takes as much time as the initial copy of ShadowImage.
- **Attach description to identify:** The character string of the maximum of 31 characters can be added to the split pair. You can also check this character string on the pair list. This is useful for indicating the information of when and for what the backup data retained in the S-

VOL was backed up. This character string is only retained while splitting.

- **Quick mode:** The pair split has the quick mode and the normal mode. Specify the quick mode for the pair whose status is Synchronizing. By specifying the quick mode, even the P-VOL and the S-VOL are during synchronization, the P-VOL data at the point can be immediately retained in the S-VOL. In this case, the pair status is changed to Split Pending and changed to Split after the copy processing is completed. The Read/Write access to the S-VOL can be possible immediately after the split instruction.
- **Normal mode:** Specify the normal mode for the pair whose status is Paired. If it is split, the pair status becomes Split and the data at the point is retained in the S-VOL as the backup data. In the Split status, since the Read/Write access to the S-VOL is possible, the backup data can be read/written.

You can split the pair whose status is Synchronizing or Paired Internally Synchronizing by executing the split operation in Quick Mode. To perform the split operation in Quick Mode for the pair whose status is Synchronizing, specify the option at the time of command execution. Moreover, the split operation for the pair whose status is Paired Internally Synchronizing is executed in Quick Mode regardless of the Quick Mode specification option. In the split operation in Quick Mode, if you execute the command, you will be able to make Read/Write access for the S-VOL immediately, and the S-VOL data accessed by the host becomes the same as the P-VOL data at the time of command execution. The data to make the S-VOL data same as the P-VOL data at the time of command execution is copied in the background, and the status becomes Split Pending until the copy is completed. The status changes to Split when the copy is completed.

This feature provides point-in-time backup of your data, and also facilitates real data testing by making the ShadowImage copies (S-VOLs) available for host access.

When the split operation is complete, the pair status changes to Split or Split Pending, and you have full Read/Write access to the split S-VOL. While the pair is split, the array establishes a track map for the split P-VOL and S-VOL and records all updates to both volumes. The P-VOL remains fully accessible during splitting pairs operation. Splitting pairs operations cannot be performed on suspended (Failure) pairs. Also, when the P-VOL that will be an operation target configures a ShadowImage pair in the Split Pending status with another S-VOL, the split operation in the Quick Mode cannot be executed.

Re-synchronizing pairs

When discarding the backup data retained in the S-VOL by split or recovering the suspended pair (Failure status), perform the pair resynchronization to resynchronize the S-VOL and the P-VOL. When the resynchronization copy starts, the pair status becomes Synchronizing or Paired Internally Synchronizing. When the resynchronization copy is completed, the pair status becomes Paired.

When the resynchronization is executed, the Write access from the host to the S-VOL becomes impossible. The Read/Write access from the host to the P-VOL continues.

ShadowImage allows you to perform two types of re-synchronizing pairs operations:

- Re-synchronizing normal pairs
- Quick mode

Re-synchronizing normal pairs

The normal re-synchronizing pairs operation (see [Figure 2-5 on page 2-11](#)) resynchronizes the S-VOL with the P-VOL. ShadowImage allows you to perform re-synchronizing operations on split, Split Pending, and Failure pairs. The following operation is performed in the resynchronization operation and the required time differs depending on the pair status at the time of resynchronization

- Re-synchronizing for Split or Split Pending pair. When a re-synchronizing pairs operation is performed on a split or split pending pair (status = Split or Split Pending), the array merges the S-VOL differential track map into the P-VOL track differential map and copies all flagged data from the P-VOL to the S-VOL. This ensures that the P-VOL and S-VOL are properly resynchronized in the desired direction, and also greatly reduces the time needed to resynchronize the pair.
- Re-synchronizing for suspended pair. When a re-synchronizing pairs operation is performed on a suspended pair (status = Failure), the array copies all data on the P-VOL to the S-VOL, since all P-VOL tracks were flagged as differential data when the pair was suspended. It takes the same time as the initial copy of ShadowImage until the copy for resynchronization is completed and the status changes to Paired.



NOTE: The P-VOL to be the operation target configures another S-VOL, two pairs which are in the Paired status, Paired Internally Synchronizing status, Synchronizing status, or Split Pending status. The re-synchronizing pairs operation cannot be executed. Or, if one of two pairs is in the Split Pending status and the other pair status is any of the Paired status, Paired Internally Synchronizing status, or Synchronizing status, re-synchronizing pairs operation cannot be executed.




Figure 2-5: Re-synchronizing pairs operation

Quick mode

If the quick mode is specified, the pair status becomes Paired Internally Synchronizing, the split operation is executed without specifying the quick mode option and the new backup data can be retained in the S-VOL immediately.

In the Paired Internally Synchronizing status, the P-VOL data is being copied in the S-VOL as well as the pair in the Synchronizing status. Note that, even if the copy pace of the background copy at this time is specified as Fast, it is executed at Medium. For operating the data copy by resynchronization at Fast, execute it without specifying the quick mode. Quick mode

If you use the Quick mode for creating or updating the copying volume (SVOL) with ShadowImage, the Read/Write access for the S-VOL becomes available immediately. Since the S-VOL data accessed from the host becomes the same as the P-VOL data at the time when the command was executed, you can start the backup from the S-VOL without waiting for the completion of the data copy.s

Table 2-1: Quick mode characteristics

Mode	Advantages	Considerations
With Quick mode	The Read/Write Access from the host to the S-VOL becomes possible without waiting for data copy completion while creating or splitting the copying volume.	Since the P-VOL data is used when accessing the S-VOL, the load for the S-VOL affects the I/O performance of the P-VOL. The I/O performance of the S-VOL is affected by the load of the P-VOL and is more deteriorated than the case where Quick Mode is not used.
Without Quick mode	Since access from the host to the S-VOL is independent of the P-VOL, the I/O performance is less affected.	Access to the S-VOL cannot begin unless the data copy is completed while creating or updating the copying volume. Therefore, wait for data copy completion before accessing the S-VOL.

Restore pairs

When the P-VOL data is in the unusable status and returned to the backup data retained in the S-VOL, execute pair restoration.

Restore pairs operation (see [Figure 2-6](#)) synchronizes the P-VOL with the S-VOL. However, when the P-VOL that will be an operation target configures a ShadowImage pair in the Paired status, the Paired Internally Synchronizing status, the Synchronizing status, the Reverse Synchronizing status, the Split Pending status or the Failure (Restore) status with another S-VOL, the restore operation cannot be executed. The copy direction for a restore pairs operation is S-VOL to P-VOL. The pair status during a restore operation is Reverse Synchronizing, and the S-VOL becomes inaccessible to

all hosts for write operations during a restore pairs operation. The P-VOL remains accessible for both read and write operations, and the write operation on P-VOL will always be reflected to S-VOL (see [Figure 2-7](#)). When operating restore, you cannot specify Quick Mode.

ShadowImage allows you to perform re-synchronizing operations on Split, Split Pending, and Failure pairs.




Figure 2-6: Restore Pairs Operation




Figure 2-7: Reflecting write data to S-VOL during reverse re-synchronizing pairs operation (restore)

Re-synchronizing for split or split pending pair

When a re-synchronizing pairs operation is performed on a split pair or split pending (status = Split or Split Pending), the array merges the S-VOL differential track map into the P-VOL track differential map and copies all flagged data from the P-VOL to the S-VOL. When a reverse re-synchronizing pairs operation is performed on a split pair, the array merges the P-VOL differential track map into the S-VOL differential track map and then copies all flagged tracks from the S-VOL to the P-VOL. This ensures that the P-VOL and S-VOL are properly resynchronized in the desired direction, and also greatly reduces the time needed to resynchronize the pair.

Re-synchronizing for suspended pair

When a re-synchronizing pairs operation is performed on a suspended pair (status = Failure), the array copies all data on the P-VOL to the S-VOL, since all P-VOL tracks were flagged as differential data when the pair was suspended. It takes the same time as the initial copy of ShadowImage until the copy for resynchronization is completed and the status changes to Paired.

Suspending pairs

The ShadowImage suspending pairs operation (split pair with Suspend operation in progress and force the pair in failure state) immediately suspends the ShadowImage copy operations to the S-VOL of the pair. The user can suspend a ShadowImage pair at any time. When a ShadowImage pair is suspended on error (status = Failure), the array stops performing ShadowImage copy operations to the S-VOL, continues accepting write I/O operations to the P-VOL, and marks the entire P-VOL track map as differential data. When a re-synchronizing pairs operation is performed on a suspended pair, the entire P-VOL is copied to the S-VOL (when a restore operation is performed, the entire S-VOL is copied to the P-VOL). While re-synchronizing pairs operation for a split or split pending ShadowImage pair greatly reduces the time needed to resynchronize the pair, re-synchronizing pairs operation for a suspended on error pair will take as long as the initial copy operation.

The array will automatically suspend a ShadowImage pair when copy operation cannot be continued or cannot keep the pair mirrored for any reason. When the array suspends a pair, a file is output to the system log or event log to notify the host (CCI only). The array will automatically suspend a pair under the following conditions:

- When the ShadowImage volume pair has been suspended or deleted from the UNIX®/PC host using the CCI.
- When the array detects an error condition related to an initial copy operation. When a volume pair with Synchronizing status is suspended on error, the array aborts the initial copy operation, changes the status of the P-VOL and S-VOL to Failure and accepts all subsequent write I/Os to the P-VOL.

Deleting pairs

The ShadowImage deleting pairs operation stops the ShadowImage copy operations to the S-VOL of the pair and deletes the volume in paired status. The user can delete a ShadowImage pair at any time except when the volumes are already in Simplex or Split Pending status. In both ShadowImage volumes, the status will change to Simplex.

Differential Management Logical Unit (DMLU)

The Differential Management Logical Unit (DMLU) is an exclusive volume used for storing ShadowImage information. The DMLU is treated the same as other volumes in the storage array, but is hidden from a host. See [Setting up the DMLU on page 4-25](#) to configure.

To create a ShadowImage pair, it is necessary to prepare one DMLU in the array. The differential information of all ShadowImage pairs is managed by this one DMLU.

The DMLU size must be greater than or equal to 10 GB. The recommended size is 64 GB. Since the DMLU capacity affects the capacity of creatable pairs, refer to the [Calculating maximum capacity on page 4-19](#) and determine the DMLU capacity.

As shown in [DMLU on page 2-16](#), the array accesses the differential information stored in the DMLU and refers to/updates it in the copy processing to synchronize the P-VOL and the S-VOL and the processing to manage the difference of the P-VOL and the S-VOL.

DMLU precautions:

- -The volume belonging to RAID 0 cannot be set as a DMLU.
- -When setting the unified volume as a DMLU, it cannot be set if the capacity of each unified volume becomes less than 1 GB on an average. For example, when setting a volume of 10 GB as a DMLU, if the volume consists of 11 sub-volumes, it cannot be set as a DMLU.
- -The volume assigned to the host cannot be set as a DMLU.
- In the DMLU expansion, select a RAID group which meets the following conditions:
 - The drive type and the combination are the same as the DMLU
 - A new volume can be created
 - A sequential free area for the capacity to be expanded exists
- When either pair of ShadowImage, TrueCopy, or Volume Migration exists, the DMLU cannot be removed.
- Notes on the combination and the drive types in the RAID group to which the DMLU is located:
 - When a failure occurs in the DMLU, all the pairs of ShadowImage, TrueCopy, and/or Volume Migration are changed to Failure. Therefore, secure sufficient redundancy for the RAID group in which the DMLU is located.
 - In the status where the pair status is Split, Split Pending, or Reverse Synchronizing, the I/O performance of the DMLU may effect the host I/O performance on the volume that configures the pair. Using RAID 1+0 or SSD/FMD drives can decrease the effect on host I/O performance.




Figure 2-8: DMLU

Ownership of P-VOLs and S-VOLs

The ownership of the volume specified in the S-VOL of the ShadowImage pair is the same as the ownership of the volume specified in the P-VOL. This ownership change operates regardless of the setting status of load balancing.

For example, if creating a ShadowImage pair by specifying the volume whose ownership is controller 0 as a P-VOL and specifying the volume whose ownership is controller 1 as an S-VOL, the ownership of the volume specified in the S-VOL is changed to controller 0.




Figure 2-9: Ownership of P-VOLs and S-VOLs

When the same controller has the P-VOL ownerships of two or more ShadowImage pairs, the ownerships of all the pairs are biased toward the same controller and the load is concentrated. To diversify the load, specify the ownership to be equal when creating a ShadowImage pair.

If the ownership of a volume has been changed at pair creation, the ownership is not changed at pair deletion. After deleting a pair, set ownership again considering load balance.

Command devices

The command device is a user-selected, dedicated logical volume on the disk array, which functions as the interface to the CCI software. ShadowImage commands are issued by CCI (HORCM) to the disk array command device.

A command device must be designated in order to issue ShadowImage commands. The command device must be defined in the HORCM_CMD section of the configuration definition file for the CCI instance on the attached host. 128 command devices can be designated for the array. You can designate command devices using Navigator 2.



NOTE: Volumes set for command devices must be recognized by the host. The command device volume size must be greater than or equal to 33 MBs.

Consistency group (CTG)

Application data often spans more than one volume. With ShadowImage, it is possible to manage operations spanning multiple volumes as a single group. In a group, all primary logical volumes are treated as a single entity.

Managing ShadowImage primary volumes as a group allows multiple operations to be performed on grouped volumes concurrently. Write order is guaranteed across application logical volumes, since pairs can be split at the same time.

By making multiple pairs belong to the same group, the pair operation is possible in units of groups. In the group whose Point-in-time attribute of the group is enabled, the backup data of the S-VOL created in units of groups is the data of the same time.

For setting a group, specify a new group number for a group to be assigned after pair creation when creating a ShadowImage pair. The maximum of 1,024 groups can be created in ShadowImage.

A group name can be assigned to a group. You can select one pair belonging to the created group and assign a group name arbitrarily by using the pair edit function.



NOTE: Group restrictions:

- The Point-in-time attribute of a group created in ShadowImage is always enabled. It cannot be disabled.
 - You cannot change the group specified at the time of the pair creation. To change it, delete the pair once, and specify another group when creating a pair again.
-

Splitting the group without specifying the quick mode option is possible only when the pairs of **Paired and Paired Internally Synchronizing** are included in the group. Moreover, splitting the group by specifying the quick mode option is possible only when the pairs of **Paired, Paired Internally Synchronizing**, and **Synchronizing** are included in the group.

ShadowImage pair status

ShadowImage displays the pair status of all ShadowImage volumes. [Figure 2-10](#) shows the ShadowImage pair status transitions and the relationship between the pair status and the ShadowImage operations.




Figure 2-10: ShadowImage pair status transitions

[Table 2-2](#) lists and describes the ShadowImage pair status conditions and accessibility. If a volume is not assigned to a ShadowImage pair, its status is Simplex.

When you create a ShadowImage pair not specifying Quick Mode, the status of the P-VOL and S-VOL changes to Synchronizing. When the initial copy operation is complete, the pair status becomes Paired. When specifying Quick Mode, if the ShadowImage pair creation starts in the pair creation operation, the pair status becomes Split Pending. In this status, the initial copy is in progress in the background. When the initial copy operation is completed, the pair status changes to Split.

If the array cannot maintain the data copy for any reason or if you suspend the pair, the pair status changes to Failure. When you split a pair, the pair status changes to Synchronizing.

When splitting pairs operation is complete, the pair status changes to Split or Split Pending to enable you to access the split S-VOL. When you start a re-synchronizing pairs operation, the pair status changes to Synchronizing or Paired Internally Synchronizing. When you specify reverse mode for a re-synchronizing pairs operation (restore), the pair status changes to Reverse Synchronizing (data is copied in the reverse direction from the S-VOL to the P-VOL). When re-synchronizing pairs operation is complete, the pair status changes to Paired. When you delete a pair, the pair status changes to Simplex.

Table 2-2: ShadowImage pair status

Pair status	Description	P-VOL access	S-VOL access
Simplex	The volume is not assigned to a ShadowImage pair. If the created pair is deleted, the pair status becomes Simplex. Note that the Simplex volume is not displayed on the list of the ShadowImage pair. The array accepts Read and Write I/Os for all Simplex volumes.	Read and write	Read and write
Synchronizing	Creating a pair or re-synchronizing a pair, the copy operation is in progress. The array continues to accept read and write operations for the P-VOL but does not accept write operations for the S-VOL. When a split pair is resynchronized in normal mode, the array copies only the P-VOL differential data to the S-VOL. When creating a pair or a Failure pair is resynchronized, the array copies the entire P-VOL to the S-VOL.	Read and write	Read only
Paired	The copy operation is complete, and the array starts copying the write operation taken to the P-VOL data onto the S-VOL. The P-VOL and S-VOL of a duplex pair (Paired status) is identical. The array rejects all write I/Os for S-VOLs with the status Paired.	Read and write	Read only
Paired Internally Synchronizing	The copy operation in progress is the same as Synchronizing. The P-VOL and the S-VOL are not yet the same. The pair split in the Paired Internally Synchronizing status operates in the Quick Mode even without specifying the option and changes to Split Pending.	Read and write	Read only
Split	The array starts accepting write I/Os for Split S-VOLs. The array keeps track of all updates to the split P-VOL and S-VOL so that the pair can be resynchronized quickly.	Read and write	Read and write. The S-VOL can be mounted.

Table 2-2: ShadowImage pair status (Continued)

Pair status	Description	P-VOL access	S-VOL access
Split Pending	Although the array starts accepting the I/O operation of Write for the S-VOL in the Split Pending status, the data copy from the P-VOL to the S-VOL is in progress in the background. The array records the positions of all updates to the split P-VOL and S-VOL. You cannot delete the pair in the Split Pending status.	Read and write	Read and write. The S-VOL can be mounted.
Reverse Synchronizing	The array does not accept write I/Os for Reverse Synchronizing S-VOLs. When a split pair is resynchronized in reverse mode, the array copies only the S-VOL differential data to the P-VOL.	Read and write	Read only
Failure	The array continues accepting read and write I/Os for a Failure (suspended under error) P-VOL (however, if the status transits from Reverse Synchronizing, all access to P-VOL is disabled). The array marks the entire P-VOL track map as differential data, so that the entire P-VOL is copied to the S-VOL when the Failure pair is resumed. Use re-synchronizing pairs operation to resume a Failure pair.	Read and write	Read only
Failure (S-VOL Switch)	This is a state in which a double failure (triple failures for RAID 6) of drives occurred in a P-VOL and the P-VOL was switched to an S-VOL internally. This state is displayed as PSUE with CCI. For details, see Setting up CCI on page A-20	Read and write	Read/write is not available
Failure (R)	This is a state in which the P-VOL data becomes unjustified due to a Failure during restoration (in Reverse Synchronizing status).	Read/write is not available.	Read/write is not available

Interfaces for performing ShadowImage operations

ShadowImage can be operated using of the following interfaces:

- The Hitachi Storage Navigator Modular 2 Graphical User Interface is a browser-based interface from which ShadowImage can be set up, operated, and monitored. The GUI provides the simplest method for performing operations, requiring no previous experience. Scripting is not available.
- CLI (Hitachi Storage Navigator Modular 2 Command Line Interface), from which ShadowImage can be setup and all basic pair operations can be performed—create, split, resynchronize, restore, swap, and delete. The GUI also provides these functionalities. CLI also has scripting capability.
- CCI (Hitachi Command Control Interface), used to display volume information and perform all copying and pair-managing operations. CCI provides a full scripting capability which can be used to automate replication operations. CCI requires more experience than the GUI or

CLI. CCI is required on Windows 2000 Server for performing mount/unmount operations.

HDS recommends using the GUI to begin operations for new users with no experience with CLI or CCI. Users who are new to replication software but have CLI experience in managing arrays may want to continue using CLI, though the GUI is an option. The same recommendation applies to CCI users.



NOTE: Hitachi Replication Manager can be used to manage and integrate ShadowImage. It provides a GUI representation of the ShadowImage system, with monitoring, scheduling, and alert functions. For more information, visit the Hitachi Data Systems website, <https://portal.hds.com>. .



CAUTION! Storage Navigator 2 CLI is provided for users with significant storage management expertise. Improper use of this CLI could void your Hitachi warranty. Please consult with your reseller before using the CLI.

Installing ShadowImage

This chapter provides instructions for installing and enabling ShadowImage.

- ❑ [System requirements](#)
- ❑ [Installing ShadowImage](#)
- ❑ [Enabling/disabling ShadowImage](#)
- ❑ [Uninstalling ShadowImage](#)

System requirements

Table 3-1 shows the minimum requirements for ShadowImage. See [Installing ShadowImage](#) for more information.

Table 3-1: ShadowImage requirements

Item	Minimum requirements
Firmware	Version 0915/B or later is required for the array.
Storage Navigator 2	version 21.50 or later is required for the management PC.
CCI	Version 01-27-03/02 or later is required for the host when CCI is used for the ShadowImage operation.
License key	Required for ShadowImage
Number of controllers	2 (dual configuration).
Command devices	Maximum of 128. The command device is required only when CCI is used for ShadowImage operations. CCI is provided for advanced users only. The command device volume size must be greater than or equal to 33 MB.
DMLU	Maximum of 1. The Differential Management volume size must be greater than or equal to 10 GB and less than 128 GB.
Volume size	The S-VOL block count must be equal to the P-VOL block count.

Supported platforms

Table 3-2 shows the supported platforms and operating system versions required for ShadowImage.

Table 3-2: Supported platforms

Platforms	Operating system version
SUN	Solaris 8 (SPARC)
	Solaris 9 (SPARC)
	Solaris 10 (SPARC)
	Solaris 10 (x86)
	Solaris 10 (x64)
PC Server (Microsoft)	Windows 2000
	Windows Server 2003 (IA32)
	Windows Server 2008 (IA32)
	Windows Server 2003 (x64)
	Windows Server 2008 (x64)
	Windows Server 2003 (IA64)
	Windows Server 2008 (IA64)
Red Hat	Red Hat Linux AS2.1 (IA32)
	Red Hat Linux AS/ES 3.0 (IA32)
	Red Hat Linux AS/ES 4.0 (IA32)
	Red Hat Linux AS/ES 3.0 (AMD64/EM64T)
	Red Hat Linux AS/ES 4.0 (AMD64/EM64T)
	Red Hat Linux AS/ES 3.0 (IA64)
	Red Hat Linux AS/ES 4.0 (IA64)
HP	HP-UX 11i V1.0 (PA-RISC)
	HP-UX 11i V2.0 (PA-RISC)
	HP-UX 11i V3.0 (PA-RISC)
	HP-UX 11i V2.0 (IPF)
	HP-UX 11i V3.0 (IPF)
	Tru64 UNIX 5.1
IBM®	AIX 5.1
	AIX 5.2
	AIX 5.3
SGI	IRIX 6.5.x

Installing ShadowImage

If ShadowImage was purchased at the same time as the order for the Hitachi Unified Storage was placed, then ShadowImage is bundled with the array and no installation is necessary. Proceed to [Enabling/disabling ShadowImage on page 3-6](#).

If ShadowImage was purchased on an order separate from Adaptable, it must be installed before enabling.



NOTE: A key code or key file is required to install or uninstall. If you do not have the key file or code, you can obtain it from the download page on the HDS Support Portal, <https://portal.hds.com>

- For CLI instructions, see [Installing and uninstalling ShadowImage on page A-6](#) (advanced users only).

Before installing or uninstalling ShadowImage, verify that the array is operating in a normal state. Installation/Un-installation cannot be performed if a failure has occurred.

To install ShadowImage

1. Start Navigator 2.
2. Log in as a registered user.
3. In the Navigator 2 GUI, click the check box for the array where you want to install ShadowImage.
4. Click **Show & Configure array**. The tree view appears.
5. Select the **Install Licenses** icon in the **Common array Task**.



6. The **Install License** screen appears.

Install License Help

License Property

Enter the information for the license to be installed.

*Install with: ☐ Key File: Browse...
Input the Key File Name.

☒ Key Code:
Input the Key Code.

* Required field

OK Cancel

7. Select the **Key File** or **Key Code** option, then enter the file name or key code. You may browse for the Key File.
8. Click **OK**.
9. Click **Confirm** on the screen requesting confirmation to install ShadowImage.
10. Click **Close** on the confirmation screen.

Enabling / disabling ShadowImage

Enable or disable ShadowImage using the following procedure.



NOTE: All ShadowImage pairs must be deleted and their volume status returned to Simplex before enabling or disabling ShadowImage.

To enable or disable ShadowImage

1. Start Navigator 2.
2. Log in as a registered user to Navigator 2.
3. Select the array where you want to enable or disable ShadowImage
4. Click the **Show & Configure array** button.
5. Click **Settings** in the tree view, then click **Licenses**.
6. Select **SHADOWIMAGE** in the Licenses list.
7. Click **Change Status**. The Change License screen appears.

8. To enable, click the **Enable: Yes** check box.
To disable, clear the **Enable: Yes** check box.
9. Click **OK**.
10. A message appears confirming that ShadowImage is enabled or disabled. Click **Close**.




Uninstalling ShadowImage

To uninstall ShadowImage, the key code or key file provided with the optional feature is required. Once uninstalled, ShadowImage cannot be used again until it is installed using the key code or key file.

All ShadowImage pairs must be deleted and their volume status returned to Simplex before uninstalling.

To uninstall ShadowImage

1. Start Navigator 2.
2. Log in as a registered user to Navigator 2.
3. In the Navigator 2 GUI, click the check box for the array where you want to uninstall ShadowImage.
4. Click the **Show & Configure disk array** button.
5. In the tree view, click **Settings**, then click **Licenses**.



The **Licenses** list appears.

6. Click **De-Install License**.

The **De-Install License** screen appears.



The image shows a 'De-Install License' dialog box. It has a title bar with 'De-Install License' and a 'Help' button. Below the title bar is a green header with 'License Property'. The main area contains the text 'Enter the information for the license to be de-installed.' and a section titled '*De-install with:'. There are two radio buttons: 'Key File:' and 'Key Code:'. The 'Key File:' option is selected, and it has a text box and a 'Browse...' button. Below it is the instruction 'Input the Key File Name.' The 'Key Code:' option is also present with a text box and the instruction 'Input the Key Code.' At the bottom left is the text '* Required field'. At the bottom right are 'OK' and 'Cancel' buttons, with a mouse cursor pointing at the 'OK' button.

7. When you uninstall the option using the key code, click the **Key Code** option, and then set up the key code. When you uninstall the option using the key file, click the **Key File** option, and then set up the path for the key file name. Click **OK**.



NOTE: Browse is used to set the path to a key file

8. On the confirmation screen, click **Close** to confirm.

ShadowImage setup

This chapter provides information for setting up your system for ShadowImage. It includes:

Planning and design

- ☐ [Planning and design](#)
- ☐ [Plan and design workflow](#)
- ☐ [Copy frequency](#)
- ☐ [Copy lifespan](#)
- ☐ [Establishing the number of copies](#)
- ☐ [Requirements and recommendations for volumes](#)
- ☐ [Calculating maximum capacity](#)

Configuration

- ☐ [Configuration](#)
- ☐ [Setting up primary, secondary volumes](#)
- ☐ [Setting up the DMLU](#)
- ☐ [Removing the designated DMLU](#)
- ☐ [Add the designated DMLU capacity](#)
- ☐ [Setting the ShadowImage I/O switching mode](#)
- ☐ [Setting the system tuning parameter](#)

Planning and design

With ShadowImage, you create copies of your production data so that it can be used to restore the P-VOL or used for tape backup, development, data warehousing, and so on.

This topic guides you in planning a design that meets your organizational business requirements.

Plan and design workflow

A ShadowImage system can only be successful when your business needs are assessed. Business needs determine your ShadowImage copy frequency, lifespan, and number of copies.

- Copy frequency means how often a P-VOL is copied.
- Copy lifespan means how long the copy is held before it is updated.
- Knowing the frequency and lifespan help you determine the number of copies that are required.

These objectives are addressed in detail in this chapter. Three additional tasks are required before your design can be implemented, which are also addressed in this chapter.

- The primary and secondary logical volumes must be set up. Recommendations and supported configurations are provided.
- The ShadowImage maximum capacity must be calculated and compared to the disk array maximum supported capacity. This has to do with how the disk array manages storage segments.
- Equally important in the planning process are the ways that various host operating systems interact with ShadowImage. Make sure to review the information at the end of the chapter.

Copy frequency

How often copies are made is determined by how much data could be lost in a disaster before business is significantly impacted.

Ideally, a business desires no data loss. In the real world, disasters occur and data is lost. You or your organization's decision makers must decide the number of business transactions that could be lost, the number of hours required to key in lost data, and so on to decide how often copies must be made.

For example, if losing 4 hours of business transaction could be tolerated, but not more, then copies should be planned for every 4 hours. If 24 hours of business transaction could be lost, copies should be planned every 24 hours.

[Figure 4-1 on page 4-3](#) shows copy frequency.




Figure 4-1: Copy frequency

Copy lifespan

Copy lifespan is the length of time a copy (S-VOL) is held, before a new backup is made to the volume. Lifespan is determined by two factors:

- Your organization's data retention policy for holding onto backup copies
- Secondary business uses of the backup data

Lifespan based on backup requirements

Copy lifespan is based on backup requirements.

- If the copy is to be used for tape backups, the minimum lifespan must be greater than the copy time from S-VOL to tape. For example:

Hours to copy an S-VOL to tape = 4

Therefore, S-VOL lifespan = 4 hours

- If the copy is to be used as a disk-based backup available for online recovery, you can determine the lifespan by multiplying the number of copies you will require to keep online times the copy frequency. For example:

Copies held = 4

Copy frequency = 4 hours

$4 \times 4 = 16$ hours

S-VOL lifespan = 16 hours

Figure 4-2 on page 4-4 shows copy lifespan.




Figure 4-2: Copy lifespan

Lifespan based on business uses

Copy life span is also based on business requirements

- If copy data is used for testing an application, the testing requirements determine the amount of time a copy is held.
- If copy data is used for development purposes, development requirements determine the time the copy is held.
- If copy data is used for business reports, the reporting requirements determine the backup's lifespan.

Establishing the number of copies

Data retention and business-use requirements for the secondary volume determine a copy's lifespan. They also determine the number of S-VOLs needed per P-VOL.

For example: If your data must be backed up every 12 hours, and business-use of secondary volume data requires holding it 36 hours, then your ShadowImage system requires 3 S-VOLs. This is illustrated in [Figure 4-3](#).




Figure 4-3: Number of S-VOLs required

Ratio of S-VOLs to P-VOL

Hitachi recommends setting up at least two S-VOLs per P-VOL. When an S-VOL is re-synchronizing with the P-VOL, it is in an inconsistent state and therefore not usable. Thus, if at least two S-VOLs exist, one is always available for restoring the P-VOL in an emergency.

A workaround when employing one S-VOL only is to backup the S-VOL to tape. However, this operation can be lengthy, and recovery time from tape is more time-consuming than from an S-VOL. Also, if a failure occurs during the updating of the copy, both the P-VOL and the single S-VOL are invalid.

Requirements and recommendations for volumes

This section relates mostly to primary and secondary volumes. However, recommendations for the DMLU and command device are also included.

Please review the following key requirements and recommendations.

Also, review the information on setting up volumes for ShadowImage volumes, DMLUs, and command devices in:

- [System requirements on page 3-2](#)
- [ShadowImage general specifications on page A-2](#)

When preparing for ShadowImage, please observe the following regarding the P-VOL and S-VOL:

- They must be the same in size, with identical block counts. You can verify block size. In the Navigator 2 GUI, navigate to **Groups>RAID Groups>volumes** tab. Click the desired volume. On the popup window that appears, review the Capacity field. This shows block size.
- Use SAS drives, SAS7.2K drives, or SSD/FMD drives to increase performance.
- Assign four or more disks to the data disks.
- Volumes used for other purposes should not be assigned as a primary volume. If such a volume must be assigned, move as much of the existing write workload to non-ShadowImage volumes as possible.
- When locating multiple P-VOLs in the same parity group, performance is best when the status of their pairs are the same (Split, Paired, Resync, and so on).

RAID configuration for ShadowImage volumes

Please observe the following regarding RAID levels when setting up ShadowImage pairs and Differential Management LUs.

- Volumes should be assigned to different RAID groups on the disk array to reduce I/O impact.
- If assigned to the same RAID group, limit the number of pairs in the group to reduce impact on performance.
- Avoid locating P-VOLs and S-VOLs within the same ECC group of the same RAID group for the following reasons:
 - A single drive failure causes status degeneration in the P-VOL and S-VOL.
 - Initial copy, coupling, and resync processes incurs a drive bottleneck which decreases performance.
- A RAID level with redundancy is recommended for both P-VOLs and S-VOLs.
- Redundancy for the P-VOL should be the same as the redundancy for the S-VOL.
- The recommended RAID configuration for P-VOLs and S-VOLs is RAID 5 (4D+1).
- When the DMLU or two or more command devices (when using CCI) are set within the one disk array, assign them to the respective RAID groups for redundancy.

Operating system considerations and restrictions

This section describes the system considerations and restrictions that apply to ShadowImage volumes.

Identifying P-VOL and S-VOL in Windows

In Navigator 2, the P-VOL and S-VOL are identified by their volume number. In Windows, volumes are identified by HLUN. These instructions provide procedures for the fibre channel and iSCSI interfaces. To confirm the H-LUN:

1. From the Windows Server 2003 Control Panel, select Computer Management/Disk Administrator.
2. Right-click the disk whose HLUN you want to know, then select **Properties**. The number displayed to the right of "LUN" in the dialog window is the HLUN.

For Fibre Channel interface:

Identify HLUN-to-VOL Mapping for the Fibre Channel interface, as follows:

1. In the Navigator 2 GUI, select the desired disk array.
2. In the array tree, click the **Group** icon, then click **Host Groups**.



WARNING! Your host group changes will be applied to multiple ports. This change will delete existing host group mappings and corresponding Host Group IDs, corrupting or removing data associated with the host groups. To keep specified host groups you do not want to remove, please cancel this operation and make changes to only one host group at a time.

3. Click the **Host Group** to which the volume is mapped.
4. On the screen for the host group, click the **Volumes** tab. The volumes mapped to the Host Group display. You can confirm the VOL that is mapped to the H-LUN.

For iSCSI interface:

Identify HLUN-to-VOL Mapping for the iSCSI interface as follows.

1. In the Navigator 2 GUI, select the desired array.
2. In the array tree that displays, click the **Group** icon, then click the **iSCSI Targets** icon in the Groups tree.
3. On the iSCSI Target screen, select an iSCSI target.
4. On the target screen, select the **Volumes** tab. Find the identified HLUN. The LUN displays in the next column.
5. If the HLUN is not present on a target screen, on the iSCSI Target screen, select another iSCSI target and repeat Step 4.

Volume mapping with CCI

In order to operate pair operation using CCI, you need to map the P-VOL and S-VOL to ports that are described in the configuration definition file for CCI. When you want to operate the P-VOL and S-VOL, but do not want the host to recognize them, you can map them to a host group that is not assigned to the host using Volume Manager.

If you use HSNM2 instead of CCI for pair operation, there is no need to map the P-VOL or S-VOL to a port or a host group.

AIX

To ensure that the same host recognizes both a P-VOL and an S-VOL, version 04-00-/B or later of HDLM (JP1/HiCommand Dynamic Link Manager) is required.

Microsoft Cluster Server (MSCS)

To create an S-VOL that is recognized by a host, observe the following:

- Use the CCI mount command. Do not use Disk Administrator.
- Do not place the MSCS Quorum Disk in CCI.
- The command device cannot be shared between the different hosts in the cluster.
- Assign the exclusive command device to each host.

Veritas Volume Manager (VxVM)

A host cannot recognize both a P-VOL and its S-VOL at the same time. Map the P-VOL and S-VOL to separate hosts.

Windows 2000

- A host cannot recognize both a P-VOL and its S-VOL at the same time. Map the P-VOL and S-VOL to separate hosts.
- When mounting a volume, you must use the CCI mount command, even if you are operating the pairs using Navigator 2 GUI or CLI. Do not use the Windows mountvol command because the data residing in server memory is not flushed. The CCI mount command flushes data in server memory, which is necessary for ShadowImage operations. For more information, see the *Hitachi Unified Storage Command Control Interface Installation and Configuration Guide*.

Windows Server

Volume mount:

- In order to make a consistent backup using a storage-based replication such as ShadowImage, you must have a way to flush the data residing on the server memory to the array, so that the source volume of the

replication has the complete data. You can flush the data on the server memory using the umount command of CCI to un-mount the volume. When using the umount command of CCI for un-mount, use the mount command of CCI for mount. (For more detail about mount/umount command, see the *Hitachi Unified Storage Command Control Interface Installation and Configuration Guide*.

- If you are using Windows Server 2003, mountvol /P to flush data on the server memory when un-mounting the volume is supported. Please understand the specification of the command and run sufficient test before you use it for your operation.
- In Windows Server 2008, use the umount command of CCI to flush the data on the memory of the server at the time of the unmount. Do not use the mountvol command of Windows standard. Refer to the *Hitachi Unified Storage Command Control Interface Installation and Configuration Guide* for the detail of the restrictions of Windows Server 2008 when using the mount/umount command
- Windows Server may write for the un-mounted volume. If a pair is resynchronized while remaining the data to the S-VOL on the memory of the server, the compatible backup cannot be collected. Therefore, execute the sync command of CCI immediately before re-synchronizing the pair for the un-mounted S-VOL.

Volumes recognized by the host:

- If you recognize the P-VOL and S-VOL on Windows Server 2008 at the same time, it may cause an error because the P-VOL and S-VOL have the same disk signature. You can use the uniqueid command to rewrite a disk signature. See the *Hitachi Unified Storage Command Control Interface Installation and Configuration Guide* for the detail.
- Multiple S-VOLs per P-VOL cannot be recognized from one host. Make S-VOL recognized to limit recognition from a host to only one S-VOL per P-VOL.

Command devices

- -When a path detachment, which is caused by a controller detachment or interface failure, continues for longer than one minute, the command device may be unable to be recognized at the time when recovery from the path detachment is made. To make the recovery, execute the "re-scanning of the disks" of Windows. When Windows cannot access the command device although CCI is able to recognize the command device, restart CCI.

Linux and LVM configuration

A host cannot recognize both a P-VOL and its S-VOL at the same time. Map the P-VOL and S-VOL to separate hosts.

Concurrent use with Volume Migration

The array limits the maximum number of ShadowImage pairs and Volume Migration pairs to 1,023 (HUS 110), or 2,047 (HUS 130/HUS 150) together. The numbers of ShadowImage pairs that can be executed are calculated by subtracting the number of migration pairs from the maximum number of pairs.

The number of copying operations that can be performed at the same time in the background is called *copying multiplicity*. The copying multiplicity limits of both the Volume Migration pairs and ShadowImage pairs together to four (HUS 110) per controller (HUS 130/HUS 150: eight per controller). When ShadowImage is used together with Volume Migration, the copying multiplicity of ShadowImage becomes smaller than the maximum number because ShadowImage and Volume Migration share the copying multiplicity

Because the disk array basically executes copying operations for ShadowImage and Volume Migration in sequential order. There may be times when copying is not started immediately when previously issued copy operations are being performed. See [Figure 4-4](#) and [Figure 4-5 on page 4-12](#).




Figure 4-4: The copying operation of ShadowImage is made to wait (four copying multiplicity)




Figure 4-5: The copying operation of volume migration is made to wait (four copying multiplicity)

Concurrent use with Cache Partition Manager

ShadowImage can be used with Cache Partition Manager. See the section on restrictions in the *Hitachi Storage Navigator Modular 2 Storage Features Reference Guide* for more information.

Concurrent use of Dynamic Provisioning

ShadowImage and Dynamic Provisioning can be used together. Refer to the *Hitachi Unified Storage Dynamic Provisioning Configuration Guide* for detailed information regarding Dynamic Provisioning. The volume created in the RAID group is called a normal volume and the volume created in the DP pool is called a DP-VOL.

- When using a DP-VOL as a DMLU

Check that the free capacity (formatted) of the DP pool to which the DP-VOL belongs is more than or equal to the capacity of the DP-VOL which can be used as the DMLU, and then set the DP-VOL as a DMLU. If the free capacity of the DP pool is less than the capacity of the DP-VOL which can be used as the DMLU, the DP-VOL cannot be set as a DMLU.

- Volume type that can be set for a P-VOL or an S-VOL of ShadowImage
The DP-VOL can be used for a P-VOL or an S-VOL of ShadowImage. [Table 4-1](#) shows a combination of a DP-VOL and a normal volume that can be used for a P-VOL or an S-VOL of ShadowImage. When using the DP-VOL, which is already used as an S-VOL at the time of the ShadowImage pair creation, a pair can be created by using the used DP-VOL. In that case, however, the initial copy time may be long. Therefore, create a pair after initializing the DP-VOL.

Table 4-1: Combination of a DP-VOL and a normal volume

ShadowImage P-VOL	ShadowImage S-VOL	Comments
DP-VOL	DP-VOL	Available. The P-VOL and S-VOL capacity can be reduced compared to the normal volume*
DP-VOL	Normal volume	Available. In this combination, copying after pair creation takes about the same time it takes when the normal volume is a P-VOL. When executing the restore, the DP pool of the same capacity as the normal volume (S-VOL) is used.
Normal volume	DP-VOL	Available. In this combination, the DP pool of the same capacity as the normal volume (P-VOL) is used. Therefore, this combination is not recommended.
*When both the P-VOL and the S-VOL use DP-VOLs, a pair cannot be created by combining the DP-VOLs which have different setting of Enabled/Disabled for Full Capacity Mode.		

Depending on the usage condition of the volume, the consumed capacity of the P-VOL and the S-VOL may differ even in the Paired status. Execute the DP Optimization and zero page reclaim as needed.

- Volume type that can be set for a DMLU.
The DP-VOL created by Dynamic Provisioning can be set for a DMLU. Set the normal volume for the DMLU.
- Volume type that can be set for a command device
The DP-VOL created by Dynamic Provisioning can be set for a command device. Set the normal volume for a command device.
- Assigning the controlled processor core of a P-VOL or an S-VOL that uses the DP-VOL
When the controlled processor core of the DP-VOL used for a ShadowImage P-VOL or S-VOL differs from the normal volume, switch the S-VOL controlled processor core assignment to the P-VOL controlled processor core automatically, and create a pair. This applies to HUS 130/ HUS 150.
- DP pool designation of a P-VOL or S-VOL which uses the DP-VOL
When using the DP-VOL for a ShadowImage P-VOL or S-VOL, using the DP-VOL designated in a separate DP pool of a P-VOL or S-VOL is recommended considering the performance implications.
- Pair status at the time of DP pool capacity depletion

When the DP pool is depleted after operating the ShadowImage pair which uses the DP-VOL, the pair status of the pair concerned may be a Failure. [Table 4-2](#) shows the pair statuses before and after the DP pool capacity depletion. When the pair status becomes a Failure caused by the DP pool capacity depletion, add the DP pool capacity whose capacity is depleted, and execute the pair operation again.

Table 4-2: Pair Statuses before and after DP Pool Capacity Depletion

Pair statuses before the DP pool capacity depletion	Pair statuses after the DP pool capacity depletion belonging to P-VOL	Pair statuses after the DP pool capacity depletion belonging to S-VOL
Simplex	Simplex	Simplex
Synchronizing	Synchronizing Failure*	Failure
Reverse Synchronizing	Failure	Reverse Synchronizing Failure*
Paired	Paired Failure*	Failure
Paired Internally Synchronizing	Paired Internally Synchronizing Failure*	Failure
Split	Split	Split
Split Pending	Split Pending Failure*	Failure
Failure	Failure	Failure
* When write is performed to the P-VOL to which the capacity depletion DP pool belongs, the copy cannot be continued and the pair status changes to Failure.		

- DP pool status and availability of pair operation

When using the DP-VOL for a P-VOL or S-VOL of the ShadowImage pair, the pair operation may not be executed depending on the status of the DP pool to which the DP-VOL belongs. [Table 4-3 on page 4-15](#) shows the DP pool status and availability of the ShadowImage pair operation. When the pair operation fails due to the DP pool status, correct the DP pool status and execute the pair operation again.

Table 4-3: DP pool statuses and availability of ShadowImage pair operation

Pair operation	DP pool, DP pool Capacity, and DP pool optimization statuses					
	Normal	Capacity in growth	Capacity depletion	Regressed	Blocked	DP in optimization
Create pair	YES ¹	YES ¹	YES ¹ YES ²	YES	NO	YES
Create pair (split option)	YES	YES	YES	YES	NO	YES
Split pair	YES	YES	YES	YES	NO	YES
Resync pair	YES	YES	YES	YES	NO	YES
Restore pair	YES	YES	YES	YES	NO	YES
Delete pair	YES	YES	YES	YES	YES	YES
Notes: 1. Refer to the status of the DP pool to which the DP-VOL of the S-VOL belongs. If the status exceeds the DP pool capacity belonging to the S-VOL by the pair operation, the pair operation cannot be executed. 2. Refer to the status of the DP pool to which the DP-VOL of the P-VOL belongs. If the status exceeds the DP pool capacity belonging to the P-VOL by the pair operation, the pair operation cannot be executed.						

YES indicates a possible case

NO indicates a case not supported



NOTE: When the DP pool was created or the capacity was increased, the DP pool underwent formatting. If pair creation, pair resynchronization, or restoration is performed during formatting, depletion of usable capacity may occur. Since the formatting progress is displayed when checking the DP pool status, check if sufficient usable capacity is secured according to the formatting progress, and then start the operation.

- Operation of the DP-VOL while using ShadowImage
When using the DP-VOL for a ShadowImage P-VOL or S-VOL, any of the operations among the capacity growing, capacity shrinking, and volume deletion, and Full Capacity Mode changing of the DP-VOL in use cannot be executed. To execute the operation, delete the ShadowImage pair of which the DP-VOL to be operated is in use, and then execute it again. The attribute edit and capacity addition of the DP pool can be executed regardless of the ShadowImage pair.
- Operation of the DP pool while using ShadowImage
When using the DP-VOL for a ShadowImage P-VOL or S-VOL, the DP pool to which the DP-VOL in use belongs cannot be deleted. To execute the operation, delete the ShadowImage pair of which the DP-VOL is in

SSD/FMD use belonging to the DP pool to be operated, and then execute it again. The attribute edit and capacity addition of the DP pool can be executed usually regardless of the ShadowImage pair.

- Volume write during Split Pending

When using the DP-VOL for a ShadowImage P-VOL or S-VOL, if writing to a P-VOL or an S-VOL when the pair status is Split Pending, the capacity of the DP pool to which both volumes belong may be consumed.

Concurrent use of Dynamic Tiering

The considerations for using the DP pool or the DP-VOL whose tier mode is enabled by using Dynamic Tiering are described. For the detailed information related to Dynamic Tiering, refer to *Hitachi Unified Storage 100 Dynamic Tiering User's Guide*. Other considerations are common with Dynamic Provisioning.

- When using a DP-VOL whose tier mode is enabled as a DMLU

When using the DP-VOL whose tier mode is enabled as DMLU, check that the free capacity (formatted) of the Tier other than SSD/FMD of the DP pool to which the DP-VOL belongs is more than or equal to the DP-VOL used as DMLU, and then set it. At the time of the setting, the entire capacity of DMLU is assigned from the 1st Tier. However, the Tier configured by SSD/FMD is not assigned to the DMLU. Furthermore, the area assigned to the DMLU is out of the relocation target.

Windows Server and Dynamic Disk

In an environment of Windows Server, you cannot use ShadowImage pair volumes as dynamic disks. The reason for this restriction is because if you restart Windows or use the Rescan Disks command after creating or re-synchronizing a ShadowImage pair, there are cases where the S-VOL is displayed as Foreign in Disk Management and becomes inaccessible.

UNMAP Short Length Mode

Enable UNMAP Short Length Mode when connecting to Windows 2012. If you do not enable it, UNMAP commands may not be completed due to a time-out.

Limitations of Dirty Data Flush Number

This setting determines the number of times processing is executed for flushing the dirty data in the cache to the drive at the same time. This setting is effective when ShadowImage is enabled. When all the volumes in the disk array are created in the RAID group of RAID 1 or RAID 1+0, the SAS drives are configured and in the DP pool. If this setting is enabled, the dirty data flush number is limited even though ShadowImage is enabled. When the dirty data flush number is limited, the response time in I/O, which has a low load and high Read rate, shortens. Note that, when TrueCopy or TCE is unlocked at the same time, this setting is not effective.

See [Setting the system tuning parameter on page 4-31](#) or for CLI [Setting the system tuning parameter on page A-10](#) for the setting method about the Duty Data Flush Number Limit.

VMware and ShadowImage configuration

When creating a backup of the virtual disk in the vmfs format using ShadowImage, shutdown the virtual machine that accesses the virtual disk, and then split the pair.

If one volume is shared by multiple virtual machines, it is required to shutdown all the virtual machines that share the volume when creating a backup. Therefore, it is not recommended to share one volume by multiple virtual machines in the configuration that creates a backup using ShadowImage.

The VMware ESX has a function to clone the virtual machine. Although the ESX clone function and ShadowImage can be linked, cautions are required for the performance at the time of execution.

When the volume which becomes the ESX clone destination is a ShadowImage P-VOL pair whose pair status is Paired, Synchronizing, Paired Internally Synchronizing, Reverse Synchronizing or Split Pending, the data may be written to the S-VOL for writing to the P-VOL. And when the volume which becomes the ESX clone destination is a ShadowImage P-VOL pair whose pair status is Synchronizing, Paired Internally Synchronizing, Reverse Synchronizing or Split Pending, since background copy is executed for re-synchronizing the P-VOL and S-VOL, the load on the drive becomes large. Therefore, the time required to clone may become longer and the clone may be terminated abnormally in some cases.

To avoid this, make the ShadowImage pair status Split or Simplex and resynchronize or create the pair after executing the ESX clone. Also, if you execute the ESX clone while ShadowImage is in a pair state, such as Synchronizing, where background copy is being executed, set the copy pace to Slow. Do the same for executing the functions such as migrating the virtual machine, deploying from the template, inflating the virtual disk and Space Reclamation.




Figure 4-6: VMware ESX

UNMAP Short Length Mode

It is recommended that you enable UNMAP Short Length Mode when connecting to VMware. If you do not enable it, UNMAP commands may not be completed due to a time-out.

Creating multiple pairs in the same P-VOL

Consider the following when creating multiple pairs in the same P-VOL:

- Copy operation order when creating multiple pair configurations in the same P-VOL

In the configuration where multiple pairs are created in the same P-VOL, only one physical copy (background copy) operates in the configuration. It can be changed to the status that the background copy operates in the maximum of two pairs at the same time for the same P-VOL.

Therefore, when the background copy of either pair is operating, the other pair is waiting for the background copy. When the background copy of either pair is completed, the other pair that was waiting for the background copy starts the operation.

- Performance when creating multiple pairs in the same P-VOL

Among the ShadowImage pairs in the Paired, Synchronizing, Paired Internally Synchronizing, or Split Pending status, the data copy processing is operating from the P-VOL to the S-VOL (differential copy or background copy). The pairs that the data copy processing operates up to two pairs at the same time can be created in the same P-VOL.

Therefore, the host I/O performance for the P-VOL, compared with the pair configuration of P-VOL:S-VOL=1:1, deteriorates at the maximum of 40% in the pair configuration of P-VOL:S-VOL=1:2.

Load balancing function

The load balancing function applies to a ShadowImage pair. When the load balancing function is activated for a ShadowImage pair, the ownership of the P-VOL and S-VOL changes to the same controller. When the pair state is **Synchronizing** or **Reverse Synchronizing**, the ownership of the pair will change across the cores but not across the controllers.

Enabling Change Response for Replication Mode

When write commands are being executed on the P-VOL or S-VOL in Split Pending state, if background copy is timed-out for some reason, the array returns Medium Error (03) to the host. Some hosts receiving Medium Error (03) may determine that the P-VOL or S-VOL is inaccessible, and stop accessing it. In such cases, enabling the Change Response for Replication Mode makes the array return Aborted Command (0B) to the host. When the host receives Aborted Command (0B), it will retry the command to the P-VOL or S-VOL and the operation will continue.

Calculating maximum capacity


Table 4-4 shows the maximum capacity of the S-VOL by the DMLU capacity in TB. The maximum capacity of the S-VOL is the total value of the S-VOL capacity of ShadowImage, TrueCopy, and Volume Migration.

The maximum capacity shown in [Table 4-4](#) is the value smaller than the pair creatable capacity displayed in Navigator 2. That's because the pair creatable capacity in Navigator 2 is treated not as the real capacity but as the value rounded up by 1.5 TB unit, not as the actual capacity when calculating the S-VOL capacity. The maximum capacity (the capacity of which the pair can be surely created) reduced by the capacity capable of rounding up by the number of S-VOLs becomes the capacity shown in [Table 4-4](#).


Table 4-4: Maximum S-VOL capacity / DMLU capacity in TB

S-VOL number	DMLU capacity				
	10 GB	32GB	64GB	96 GB	128 GB
2	256				
32	1,031	3,411	4,096		
64	983	3,363	6,327	7,200	
128	887	3,267	6,731		
512	311	2,691	6,155		
1,024	N/A	1,923	5,387		
4,096	N/A	N/A	779	4,241	7,200

ShadowImage supported capacity is calculated not based on the P-VOL capacity but based on the S-VOL capacity only. The total sum of the P-VOL and S-VOL capacities varies depending on whether the pair configuration (correspondence between the P-VOL and S-VOL) is one-to-one or not. An example of the pair configuration, which can be constructed when the maximum S-VOL capacity that is supported is 3 TB, is shown below.




① Pair configuration is one-to-one:
The sum of the P-VOL and S-VOL capacities is 6 TB.



② Pair configuration is one-to-three:
The sum of the P-VOL and S-VOL capacities is 4 TB.

When considering the capacity of the S-VOL whose pair status is Split Pending, the capacity is assumed to be twice as the actual capacity. The example that can be configured when the maximum S-VOL support capacity is 3 TB is shown below.



Configuration

This topic provides required information for setting up your system for ShadowImage.

Setup for ShadowImage consists of making certain that primary and secondary volumes are set up correctly.

Setting up primary, secondary volumes

The primary and secondary volumes must be set up prior to making ShadowImage copies. When doing so, adhere to the following:

- The P-VOL and S-VOL must have identical block counts.
- Verify block size in the Navigator 2 GUI by navigating to Groups/Groups/Volumes tab.
- Click the volume whose block size you want to check.
- On the popup window that appears, review the Capacity field. This shows block size.

Refer to [Appendix A, ShadowImage In-system Replication reference information](#) for all key requirements and recommendations.

Location of P-VOLs and S-VOLs

DO NOT locate P-VOLs and S-VOLs within the same ECC group of the same RAID group, because:

- A single drive failure causes status regression in the P-VOL and S-VOL.
- Initial copy, coupling, and resync processes incurs a drive bottleneck which decreases performance.




Table 4-5: Locations for P-VOLs and S-VOLs (not recommended and recommended)

Locating multiple volumes within same drive column

If multiple volumes are set within the same parity group and if each pair state differs, it is difficult to estimate the performance and to design the system operational settings. An example is as follows: the VOL0 and VOL1 are both P-VOLs and exist within the same group in the same drive (S-VOLs are located in a different parity group) and when the VOL0 is in Paired status and the VOL1 is in Reverse Synchronizing status




Figure 4-7: Locating multiple volumes within the same drive column

Pair status differences when setting multiple pairs

Even with a single volume per parity group, it is recommended that you keep the status of pairs the same (such as Simplex, Paired, and Split) when setting multiple ShadowImage pairs. If each ShadowImage pair status differs, it is difficult to estimate the performance when designing the system operational settings

Drive type P-VOLs and S-VOLs

SAS or SSD/FMD drive performance exceeds SAS7.2K drive performance; therefore, when a P-VOL or S-VOL is located in a RAID group consisting of SAS7.2K drives performance is lower than when a P-VOL or S-VOL is located in a RAID group consisting of SAS or SSD/FMD drives. We recommend:

- Locate a P-VOL in a RAID group consisting of SAS or SSD/FMD drives.
- When locating an S-VOL in a RAID group consisting of the SAS7.2K drives, conduct a thorough investigation beforehand.

Locating P-VOLs and DMLU

Locate the P-VOL and the DMLU in the different RAID groups. If a drive dual failure (triple failure in case of RAID 6) occurs in the RAID group to which the DMLU belongs, the differential data is lost. Therefore, the pair status becomes Failure and the ShadowImage I/O switching function does not operate.

Setting up the DMLU

A DMLU is an abbreviation of a Differential Management Logical Unit and a volume exclusive for storing differential information of a P-VOL and an S-VOL of a ShadowImage pair. To create a ShadowImage pair, it is necessary to prepare one DMLU in the array. The differential information of all ShadowImage pairs is managed by this one DMLU. The DMLU in the array is treated in the same way as the other logical units. However, a logical unit that is set as the DMLU is not recognized by a host (it is hidden).

When the DMLU is not set, it must be created.

Prerequisites


- The DMLU size must be greater than or equal to 10 GB. The recommended size is 64 GB. The minimum DMLU size is 10 GB, the maximum size is 128GB.
- The stripe size is 64KB minimum, 256KB maximum.
- If you are using a merged volume for DMLU, it is necessary that each sub volume capacity is more than 1GB in average.
- There is only one DMLU. Redundancy is necessary because a secondary DMLU is not available.
- A SAS drive and RAID 1+0 is recommended for performance.
- When a failure occurs in the DMLU, all the pairs of ShadowImage, TrueCopy, and/or Volume Migration are changed to Failure. Therefore, secure sufficient redundancy for the RAID group to which the DMLU is located.
- When the pair status is Split, Split Pending, or Reverse Synchronizing, the I/O performance of the DMLU may effect the host I/O performance of the volume which configures the pair. Using RAID 1+0 or SSD/FMD drives can decrease the effect to the host I/O performance.
- Also see DMLU items in [ShadowImage general specifications on page A-2](#).



NOTE: When either pair of ShadowImage, TrueCopy, or Volume Migration exist and when only one DMLU is set, the DMLU cannot be removed.


To set up a DMLU

1. Select the **DMLU** icon in the Setup tree view of the Replication tree view. The Differential Management Logical Units screen displays.
2. Click **Add DMLU**. The Add DMLU screen displays.



3. Select the LUN you want to set as the DMLU and click OK. A confirmation message displays.
4. Select the **Yes, I have read...** check box, then click **Confirm**. When the success message displays, click **Close**.


Add DMLU - 1003



Removing the designated DMLU

When either pair of ShadowImage, TrueCopy, or Volume Migration exist there is this restriction: When only one DMLU is set, the DMLU cannot be removed.

1. Select the **DMLU** icon in the Setup tree view of the Replication tree view. The **Differential Management Logical Units** list appears.
2. Select the LUN you want to remove, and click **Remove DMLU**.
3. A message displays. Click Close.



Add the designated DMLU capacity

To add the designated DMLU capacity:

1. Select the **DMLU** icon in the Setup tree view of the Replication tree view. The Differential Management Logical Units list appears.
2. Select the LUN you want to add, and click **Add DMLU Capacity**. The **Add DMLU Capacity** screen appears.

Add DMLU Capacity - 1003 Help

Differential Management Logical Unit Capacity Property

Enter the information for logical unit added to a DMLU.

LUN : 1003
Total Capacity: 64.0GB
RAID Group: 000
RAID Level: RAID5(4D+1P)
Drive Type: SAS

*New Capacity: 96 GB
To 128GB

Use RAID Group:

Assignable RAID Groups

Rows/Page: 25 | Page 1 of 1

RAID Group	RAID Level	Capacity	
		Total	Free
000	RAID5(4D+1P)	4.0TB	4.0TB

* Required field

OK Cancel

3. Enter a capacity after the expansion in units of GB to the **New Capacity** and click **OK**. Select the RAID group which can acquire the capacity to be expanded in the sequential free area (selection is not necessary when using the DMLU in the DP pool).
4. A message displays. Click Close.

Add DMLU Capacity - 1003


i DMLU capacity added successfully.

Close

Setting the ShadowImage I/O switching mode


To set the ShadowImage I/O Switching Mode:

1. Start Navigator 2.
2. Log in as registered user to Navigator 2.
3. Select the disk array in which you will set ShadowImage.
4. Click **Show & Configure disk array**.
5. Select the **System Parameters** icon in the **Settings** tree view.




6. Click **Edit System Parameters**.

The **Edit System Parameters** screen appears.



7. Select the **ShadowImage I/O Switch Mode** in the **Options** and click **OK**.
8. A message displays. Click **Close**.



NOTE: When turning off the ShadowImage I/O Switching mode, it is required to make pair statuses of all ShadowImage pairs to those other than Failure (S-VOL Switch) and Synchronizing (S-VOL Switch).


Setting the system tuning parameter

This setting determines whether to limit the number of times processing is executed for flushing the dirty data in the cache to the drive at the same time.

To set the Dirty Data Flush Number Limit of a system tuning parameters:

1. Select the **System Tuning** icon in the **Tuning Parameter** of the **Performance** tree view.

The **System Tuning** screen appears.



2. Click **Edit System Tuning Parameters**.

The **Edit System Tuning Parameters** screen appears.

Edit System Tuning Parameters Help

System Tuning Parameters

Specify system tuning parameters.

* Dirty Data Opportunity: %
From 0 to 50

* Dirty Data Stop Opportunity: %
From 0 to 50

Cache Control Mode:

☒ First In First Out
☐ Least Recently Used

Detailed Trace Mode:

☒ Enable
☐ Disable

Load Balancing:

☒ Enable
☐ Disable

Load Balancing Monitoring Time: minute(s)

Dirty Data Flush Number Limit:

☒ Enable
☐ Disable

Load Reduction for Changing Configuration Mode:

☐ Enable
☒ Disable


* Required field

OK Cancel

OK

3. Select the **Enable** option of the **Dirty Data Flush Number Limit**.
4. Click **OK**.
5. A message appears. Click **Close**.

Edit System Tuning Parameters



Using ShadowImage

This chapter describes ShadowImage operations.

- ❑ [ShadowImage workflow](#)
- ❑ [Prerequisites for creating the pair](#)
- ❑ [Create a pair](#)
- ❑ [Split the ShadowImage pair](#)
- ❑ [Resync the pair](#)
- ❑ [Delete a pair](#)
- ❑ [Edit a pair](#)
- ❑ [Restore the P-VOL](#)
- ❑ [Use the S-VOL for tape backup, testing, reports](#)

ShadowImage workflow

A typical ShadowImage workflow consists of the following:

- Check pair status. Each operation requires a pair to have a specific status.
- Create the pair, in which the S-VOL becomes a duplicate of the P-VOL.
- Split the pair, which separates the primary and secondary volumes and allows use of the data in the S-VOL by secondary applications.
- Re-synchronize the pair, in which the S-VOL again mirrors the on-going, current data in the P-VOL.
- Restore the P-VOL from the S-VOL.
- Delete a pair.
- Edit pair information.

For an illustration of basic ShadowImage operations, see [Figure 2-1 on page 2-3](#).

Prerequisites for creating the pair

Please review the following before creating a pair.

- When you want to perform a specific ShadowImage operation the pair must be in a state that allows the operation. For instructions on checking pair status plus status definitions, see [Monitor pair status on page 6-2](#)
- When the primary volume *is not* part of another ShadowImage pair, both primary and secondary volumes must be in the SMPL (simplex) state.
- When the primary volume *is* part of another ShadowImage pair or pairs, only one of those pairs can be Paired status, Paired Internally Synchronizing status, or Synchronizing status.
- The primary and secondary volumes must have identical block counts and must be assigned to the same controller.
- Because pair creation affects the performance on the host, observe the following:
 - Create a pair when I/O load is light.
 - Limit the number of pairs that you create simultaneously.
- During the Create Pair operation, the following takes place:
 - All data in the P-VOL is copied to the S-VOL.
 - The P-VOL remains available to the host for read/write throughout the copy operation.
 - Writes to the P-VOL during pair creation are copied to the S-VOL.
 - Pair status is Synchronizing while the initial copy operation is in progress.
 - Status changes to Paired when the initial copy is complete.

- New writes to the P-VOL continue to be copied to the S-VOL in the Paired status.

Pair assignment

- Do not assign a volume (required for a quick response to a host) to a pair.

When volumes are Paired, data written to a P-VOL is also written to an S-VOL. This occurs particularly when the writing load become heavier due to a large number of write operations, writing data with a large block size, frequent write I/O operations, and continuous writing. Select the ShadowImage pair carefully. When applying ShadowImage to a volume with a heavy writing load, make the loads on other volumes lighter

- Assign two different RAID groups to each of P-VOL and S-VOL.

When an S-VOL is assigned to a RAID group in which the P-VOL has been assigned, the reliability of data is lowered because a failure that occurs in a single drive affects both of the P-VOL and S-VOLs. The performance becomes limited because the load of writing applied on a drive is doubled. Therefore, it is recommended to assign P-VOL and S-VOLs to respective RAID groups.

- Assign a small number of volumes within the same RAID group.

When volumes are assigned to the same RAID group and used as pair volumes, there may be a case where a pair creation or resynchronization for one of the volumes causes a restriction to be placed on performance of a host I/O, pair creation, resynchronization, etc., for the other volume(s) because of contention between drives. It is recommended that you assign a small number of (one or two) volumes to be paired to the same RAID group. When creating two or more pairs within the same RAID group, standardize the controllers that control volumes in the same RAID group and pay attention to make the pair creation or resynchronization timely.

- For a P-VOL, use the SAS drives or the SSD/FMD drives


When a P-VOL is located in a RAID group consisting of the SAS7.2K drives, performance of a host I/O, pair creation, and pair resynchronization, etc., is lowered because of the lower performance of the SAS7.2K drives. Therefore, it is recommended to assign a P-VOL to a RAID group consisting of the SAS drives or the SSD/FMD drives.

- Assign four or more disks to the data disks.


When the data disks that compose a RAID group are not sufficient, it affects the host performance and/or copying performance adversely because reading/writing from/to the drives is restricted. Therefore, when operating pairs with ShadowImage, it is recommended that you use a volume consisting of four or more data disks.

Confirming pair status


1. Select the **Local Replication** icon in the Replication tree view.



2. The Pairs list appears. The pair with the secondary volume without the volume number is not displayed. To display the pair with the secondary volume without the volume number, open the Primary Volumes tab and select the primary volume of the target pair.



3. The list of the primary volumes is displayed in the Primary Volumes tab.



4. When the primary volume is selected, all the pairs of the selected primary volume including the secondary volume without the volume number are displayed.
 - **Pair Name:** The pair name displays.
 - **Primary Volume:** The primary volume number displays
 - **Secondary Volume:** The secondary volume number displays. The secondary volume without the volume number is displayed as N/A.

- **Status:** The pair status displays.
 - **Simplex:** A pair is not created.
 - **Reverse Synchronizing:** Update copy (reverse) is in progress.
 - **Paired:** Initial copy or update copy is completed.
 - **Split:** A pair is split.
 - **Failure:** A failure occurs.
 - **Failure(R):** A failure occurs in restoration.
 - **---**: Other than above.
- **DP Pool:**
 - **Replication Data:** A Replication Data DP pool number displays.
 - **Management Area:** A Management Area DP pool number displays. Since this is the information used in Snapshot, N/A is displayed for the ShadowImage pair.
- **CopyType:** Snapshot or ShadowImage displays.
- **Group Number:** A group number, group name, or ---:{Ungrouped} displays.
- **GroupName**
- **Point-in-Time:** A point-in-time attribute displays. Enable is always displayed for the pair belonging to the group. N/A is displayed for the pair not belonging to the group.
- **Backup Time:** Acquired backup time or N/A displays.
- **Split Description:** A character strings appears when you specify **Attach description to identify the pair upon split**. If this is not specified **Attach description to identify the pair upon split, N/A** displays.
- **MU Number:** An MU number used in CCI displays.

Setting the copy pace

Copy pace is the speed at which a pair is created or re-synchronized. You select the copy pace when you create or resynchronize a pair (if using CCI, you enter a copy pace parameter).

Copy pace impacts host I/O performance. A slow copy pace has less impact than a medium or fast pace. The pace is divided on a scale of 1 to 15 (as in the CCI command option of -c), as follows:

- **Slow** — between 1-5. The process takes longer when host I/O activity is heavy. The amount of time to complete an initial copy or resync cannot be guaranteed.
- **Medium** — between 6-10. (Recommended) The process is performed continuously, but the amount of time to complete the initial copy or resync cannot be guaranteed. Actual pace varies according on host I/O activity.

- Fast — between 11-15. The copy/resync process is performed continuously and takes priority. Host I/O performance is restricted. The amount of time to complete an initial copy or resync is guaranteed.

Create a pair

To create a ShadowImage pair:

To use CLI, see [Creating ShadowImage pairs on page A-11](#).

1. In Navigator 2 GUI, select the desired disk array, then click the **Show & Configure disk array** button.
2. From the Replication tree, select the **Local Replication** icon. The **Create Pair** screen appears.

3. Select **ShadowImage** in the **CopyType**.
4. Enter a Pair Name if necessary.
5. Select a primary volume and secondary volume.



NOTE: The LUN may be different from H-LUN, which is recognized by the host.

6. After making selections on the **Basic** tab, further customize the pair by clicking the **Advanced** tab.

Local Pair Property

Select a copy type of local replication and enter property information as desired.

* Copy Type : ☒ ShadowImage ☐ Snapshot

Basic **Advanced**

Enter any additional properties for the pair.

Copy Pace : Medium

* After pair creation, add the pair to a group.- Group Assignment:

☒ (Ungrouped)

☐ New or existing Group Number : From 0 to 1023

☐ Existing Group Name : 31 characters or less (alphanumeric characters, special symbols, %, #, *, -, ., /, \, |, =, >, <, <N>, <P>, <S>, <T>, or *)

Do initial copy from the primary volume to the secondary volume : ☒ Yes

Allow read access to the secondary volume after the pair is created : ☒ Yes

Automatically split the pair immediately after they are created: ☐ Yes

* MU Number :

☒ Automatic

☐ Manual : From 0 to 39

* Required field

OK Cancel

7. From the **Copy Pace** dropdown list, select the speed at which copies will be made. Select Slow, Medium, or Fast. (See [Setting the copy pace on page 5-5](#) for more information.)
8. In the **Group Assignment** area, you have the option of assigning the new pair to a consistency group. (For a description, see [Consistency group \(CTG\) on page 2-18](#).) Do one of the following:
 - If you do not want to assign the pair to a consistency group, leave the **Ungrouped** button selected.
 - To create a group and assign the new pair to it, click the **New or existing Group Number** button and enter a new number for the group in the box. Specify a group number from 0 to 255.
 - To assign the pair to an existing group, enter its number in the **Group Number** box, or enter the group name in the **Existing Group Name** box.




NOTE: You can also add a Group Name for a consistency group as follows:

- a. After completing the create pair procedure, on the **Pairs** screen, check the box for the pair belonging to the group.
- b. Click the **Edit Pair** button.
- c. On the Edit Pair screen, enter the **Group Name** then click **OK**.


9. In the **Do initial copy from the primary volume...** field, leave **Yes** checked to copy the primary to the secondary volume.
Clear the check box to create a pair without copying the P-VOL at this time, and thus reduce the time it takes to create the pair. The system treats the two volumes as a pair.
10. In the **Allow read access to the secondary volume after the pair is created** field, leave **Yes** checked to allow access to the secondary volume after the pair is created. Clear the check box to prevent read/write access to the S-VOL from a host after the pair is created. This option (un-checking) insures that the S-VOL is protected and can be used as a backup.
11. Add a check mark to the box **Automatically split the pair immediately after they are created** when you want to automatically split the pair after creation.
12. When specifying a specific MU number, select **Manual** and specify the MU number in the range **0 - 39**.
13. Click **OK**.
14. A confirmation message displays. Check the **Yes, I have read the above warning and want to create the pair** check box, and click **Confirm**.

Create Pair



15. A confirmation message displays. Click **Close**.

Create Pair



Split the ShadowImage pair

When a primary and secondary volume are in Pair status, all data that is written to the primary volume is copied to the secondary volume. This continues until the pair is split.


When it is split, updates continue to be written to the primary volume, but not to the secondary volume. Data in the S-VOL is frozen at the time of the split. After the Split Pair operation:

- The secondary volume becomes available for read/write access by secondary host applications.
- Separate track tables record updates to the P-VOL and to the S-VOL.
- The pair can be made identical again by re-synchronizing from primary-to-secondary or secondary-to-primary.

To split the pair

To use CLI, see [Splitting ShadowImage pairs on page A-13](#).

1. In Navigator 2 GUI, select the desired disk array, then click the **Show & Configure disk array** button.
2. From the **Replication** tree, select the **Local Replication** icon. The **Pairs** screen displays.
3. Select the pair you want to split in the **Pairs** list.
4. Click the **Split Pair** button at the bottom of the screen. View further instructions by clicking the Help button, as needed.



5. Mark the check box of the **Suspend operation in progress and force the pair into a failure state** if necessary.
6. Enter a character strings to the **Attach description to identify the pair upon split** if necessary.
7. When you want to split Quick Mode, add a check mark to **Quick Mode**.
8. Click **OK**.
9. A confirmation message displays. Click **Close**.

Resync the pair

Re-synchronizing a pair that has been split updates the S-VOL so that it is again identical with the P-VOL. A reverse resync updates the P-VOL so that it is identical with the S-VOL.

- The pair must be in Split status.
- Pair status during a normal re-synchronizing is Synchronizing.
- Status changes to Paired when the resync is complete.
- When the pair is re-synchronized, it can then be split for tape backup or other uses of the updated S-VOL.



NOTE: Because updating the S-VOL affects performance in the RAID group to which the pair belongs, best results are realized by performing the operation when I/O load is light. Priority should be given to the Resync process.

To resync the pair

To use CLI, see [Re-synchronizing ShadowImage pairs on page A-14](#).

1. In Navigator 2 GUI, select the desired disk array, then click the **Show & Configure disk array** button.
2. From the Replication tree, select the **Local Replication** icon. The **Pairs** screen displays.
3. Select the pair you want to resync.
4. Click the **Resync Pair** button. The Resync Pair screen appears as shown below. View further instructions by clicking the **Help** button, as needed.



5. When you want to re-synchronize Quick Mode, place a check mark in the Yes box for **Quick Mode**.
6. Click **OK**.

A confirmation message displays.

Resync Pair - SI_LU1000_LU1001



After re-synchronizing the pair, you cannot use existing data in the secondary volume. Are you sure you want to re-synchronize the pair?

YOU CANNOT UNDO THIS OPERATION.



Yes, I have read the above warning and want to re-synchronize selected pairs.

Confirm

Cancel

Confirm

7. For **Yes, I have read the above warning and want to re-synchronize selected pairs**. Place a check in the box, and click **Confirm**.
8. A confirmation message displays. Click **Close**.

Delete a pair

You can delete a pair when you no longer need it. When you delete a pair, the primary and secondary volumes return to their SIMPLEX state. Both are available for use in another pair. You can delete a ShadowImage pair at any time except when the volumes are already in Simplex or Split Pending status. When the status is Split Pending, delete it after the status becomes Split.

To delete a ShadowImage pair


To use CLI, see [Deleting ShadowImage pairs on page A-15](#).

When executing the pair deletion sequentially in the batch file or the script, insert a five-second delay before executing the next step. An example of inserting a five-second delay in the batch file is shown below:

```
Ping 127.0.0.1 -n 5 > nul
```

1. In Navigator 2 GUI, select the desired disk array, then click the **Show & Configure disk array** button.
2. From the Replication tree, select the **Local Replication** icon. The **Pairs** screen displays.
3. Select the pair you want to delete.
4. Click **Delete Pair**. A confirmation message displays.

Delete Pair - SI_LU1000_LU1001



5. Check the **Yes, I have read the above warning and agree to delete selected pairs.** check box, and click **Confirm**.
6. A confirmation message displays. Click **Close**.

Edit a pair

You can edit the name, group name, and copy pace for a pair.

To edit pairs

To use CLI, refer to *Hitachi Unified Storage Command Line Interface Reference Guide (CLI) Reference Guide*. In Navigator 2 GUI, select the desired disk array, then click the **Show & Configure disk array** button.

1. From the Replication tree, select the **Local Replication** icon. The **Pairs** screen displays.
2. Select the pair that you want to edit.
3. Click **Edit Pair**. The Edit Pair screen appears.

4. Change the **Pair Name**, **Group Name**, or **Copy Pace** if necessary.
5. Click **OK**.
6. A confirmation message displays. Click **Close**.

Restore the P-VOL

ShadowImage enables you to restore your P-VOL to a previous point in time. You can restore from any S-VOL paired with the P-VOL.


The amount of time it takes to restore your data is dependent on the size of the P-VOL and the amount of data that has changed.

To restore the P-VOL from the S-VOL

1. Shut down the host application.
2. Un-mount the P-VOL from the production server.
3. In the Storage Navigator 2 GUI, select the **Local Replication** icon in the Replication tree view.

Advanced users using the Navigator 2 CLI, please refer to [Restoring the P-VOL on page A-14](#).

4. In the GUI, select the pair to be restored in the **Pairs** list.
5. Click **Restore Pair**.



6. A confirmation message displays.
7. Check the Yes, I have read the above warning and want to restore selected pairs. check box, and click **Confirm**.
8. A confirmation message displays. Click **Close**.
9. Mount the P-VOL.
10. Re-start the application.

Use the S-VOL for tape backup, testing, reports

Your ShadowImage copies can be used on a secondary server to fulfill a number of data management tasks. These might include backing up production data to tape, using the data to develop or test an application, generating reports, populating a data warehouse, and so on.

Whatever the task, the process for preparing and making your data available is the same. The following process can be performed using the Navigator 2 GUI or CLI, in combination with an operating system scheduler. The process should be performed during non-peak hours for the host application.

To use the S-VOL for secondary functions

1. Un-mount the S-VOL if it is being used by a host.
2. Resync the pair before stopping or quiescing the host application. This is done to minimize the down time of the production application.
 - Navigator 2 GUI users, please see [Resync the pair on page 5-10](#).
 - Advanced users using CLI, please see [Re-synchronizing ShadowImage pairs on page A-14](#).



NOTE: Some applications can continue to run during a backup operation, while others must be shut down. For those that continue running (placed in backup mode or quiesced rather than shut down), there may be a host performance slowdown.

3. When pair status becomes Paired, shut down or quiesce (quiet) the production application, if possible.
4. Split the pair. Doing this insures that the backup will contain the latest mirror image of the P-VOL.
 - GUI users please see [Split the ShadowImage pair on page 5-9](#).
 - Advanced users using CLI, please see [Splitting ShadowImage pairs on page A-13](#).
5. Un-quiesce or start up the production application so that it is back in normal operation mode.
6. Mount the S-VOL on the server, if needed.
7. Run the backup program using the S-VOL.

Monitoring and troubleshooting ShadowImage

This chapter provides information and instructions for monitoring and troubleshooting the ShadowImage system.

- ❑ [Monitor pair status](#)
- ❑ [Monitoring pair failure](#)
- ❑ [Troubleshooting](#)

Monitor pair status


Monitoring pair status insures the following:

- A pair is in the correct status for the ShadowImage operation you wish to perform.
- Pairs are operating correctly and status is changing to the appropriate state during and after an operation.
- Data is being updated from P-VOL to S-VOL in a pair resync, and from S-VOL to P-VOL in a pair reverse resync.
- Differential data management is being performed in the Split status. The Status column on the **Pairs** screen shows the percentage of synchronization. This can be used to estimate the amount of time a resync will take.

To check pair status

To use CLI, see [Confirming pairs status on page A-11](#).

1. In Navigator 2 GUI, select the desired disk array, then click the **Show & Configure disk array** button.
2. From the Replication tree, select the **Local Replication** icon.



3. The **Pairs** screen displays.



Pair Name	Primary Volume	Secondary Volume	Status	Copy Type	Group Number: GroupName	Backup Time	Split Description
LU1000_LU1001	1000	1001	Paired(10%)	ShadowImage	---:{Ungrouped}	N/A	N/A

4. Locate the pair and review the **Status** field.

Table 6-1 shows Navigator2 GUI status and descriptions. For CCI statuses see [Confirming pair status on page A-28](#).

Table 6-1: Pair statuses

GUI pair status	Description
Simplex	If a volume is not assigned to a ShadowImage pair, its status is Simplex. If the created pair is deleted, the pair status becomes Simplex. Note that the Simplex volume is not displayed on the list of the ShadowImage pair. The array accepts Read and Write I/Os for all Simplex volumes.
Paired	S-VOL is a duplicate of the P-VOL. Updates to the P-VOL are copied to the S-VOL.
Paired Internally Synchronizing	Re-synchronizing pairs specifying the Quick Mode.
Synchronizing	Initial or re-synchronization copy is in progress. The disk array continues to accept read and write operations for the P-VOL but does not accept write operations for the S-VOL. When a split pair is resynchronized in normal mode, the disk array copies only the P-VOL differential data to the S-VOL. When creating a pair or when a Failure pair is resynchronized, the disk array copies the entire P-VOL to the S-VOL.
Split	Updates stop from P-VOL to S-VOL. The S-VOL remains a copy of the P-VOL at the time of the split. P-VOL continues being updated by the host application.
Split Pending	A pair split specifying the Quick Mode.
Resynchronizing	The S-VOL is updated from the P-VOL. When this operation is completed, the status changes to Paired.
Reverse Synchronizing	P-VOL restoration from S-VOL is in progress.
Failure	Copying is suspended due to a failure occurrence. The disk array marks the entire P-VOL as differential data; thus, it must be copied in its entirety to the S-VOL when a Resync is performed.
Failure(R)	This is the status where the copy from the S-VOL to the P-VOL cannot be continued due to a failure during Reverse Synchronizing and the P-VOL data is in the unjustified status. The P-VOL cannot perform either the Read or Write access. To make it access, it is necessary to delete the pair.



NOTE: The identical rate displayed with the pair status shows the identical ratio of the P-VOL and S-VOL data that can be accessed from the host. When the pair status is Split Pending, even though the background copy is performed, if the P-VOL and S-VOL data viewed from the host is matched, the identical rate becomes 100%. The ratio that the background copy is completed is indicated by the Progress. You can check the Progress by the detail information of each pair.

Monitoring pair failure

It is necessary to check the pair statuses regularly to monitor that ShadowImage pairs are operating correctly and data is updated from P-VOLs to S-VOLs in the Paired status or that differential data management is performed in the Split status. When a hardware failure occurs, the failure may cause a pair failure and may change the pair status to Failure. Check that the pair status is other than Failure. When the pair status is Failure, it is required to restore the status. See [Pair failure on page 6-7](#).

For ShadowImage, the following processes are executed when the pair failure occurs.

Table 6-2: Pair failure results

Management software	Results
Navigator 2	A message is displayed in the event log
	The pair status is changed to Failure or Failure(R)
CCI	The pair status is changed to PSUE.
	An error message is output to the system log file. (For UNIX® system and the Windows Server, the syslog file and eventlog file are shown respectively.)

When the pair status is changed to Failure or Failure(R), a trap is reported with SNMP Agent Support Function

When using CCI, the following message is output to the event log. For the details, refer to *Hitachi Unified Storage Command Control Interface Installation and Configuration Guide*.

Table 6-3: CCI system log message

Message ID	Condition	Cause
HORCM_102	The volume is suspended in code 0006	The pair status was suspended due to code 0006.

Monitoring of pair failure using a script

When SNMP Agent Support Function not used, it is necessary to monitor the pair failure using a Windows Server script that can be performed using Navigator 2 CLI commands.

The following is a script for monitoring the two pairs (SI_LU0001_LU0002 and SI_LU0003_LU0004) and informing the user when pair failure occurs. The following script is activated every several minutes. The disk array must be registered to beforehand.

```
echo OFF
REM Specify the registered name of the arrays
set UNITNAME=Array1
REM Specify the name of target group (Specify "Ungroup" if the
pair doesn't belong to any group)
set G_NAME=Ungrouped
REM Specify the name of target pair
set P1_NAME=SI_LU0001_LU0002
set P2_NAME=SI_LU0003_LU0004
REM Specify the value to inform "Failure"
set FAILURE=14

REM Checking the first pair
:pair1
aureplicationmon -unit %UNITNAME% -evwait -si -pairname
%P1_NAME% -gname %G_NAME% -nowait
if errorlevel %FAILURE% goto pair1_failure
goto pair2
:pair1_failure
<The procedure for informing a user>*

REM Checking the second pair
aureplicationmon -unit %UNITNAME% -evwait -si -pairname
%P2_NAME% -gname %G_NAME% -nowait
if errorlevel %FAILURE% goto pair2_failure
goto end
:pair2_failure
<The procedure for informing a user>*

:end
%
```

Troubleshooting

In the case of the ShadowImage pair, a pair failure may be caused by occurrence of a hardware failure, so restoring the pair status is required. When you perform the pair delete, the pair status is changed to Failure or Failure(R) in the same way that a pair failure occurs, so restoring the pair status is required. Furthermore, when the DP-VOLs are used for the volumes configuring a pair, a pair failure may occur depending on the consumed capacity of the DP pool, and the pair status may become Failure.

When a pair failure occurs because of a hardware failure, maintaining the array must be done first. In the maintenance work, the pair operation of ShadowImage may be required. Since you must perform the pair operation of ShadowImage, please cooperate with service personnel in the maintenance work.

Pair failure

A pair failure occurs when one of the following takes place:

- A hardware failure occurs.
- Forcible delete is performed by the user. This occurs when you halt a Pair Split operation. The array places the pair in Failure status.

If the pair was not forcibly suspended, the cause is hardware failure.

To restore pairs after a hardware failure

1. If the volumes were re-created after the failure, the pairs must be re-created.
2. If the volumes were recovered and it is possible to resync the pair, then do so. If resync is not possible, delete and then re-create the pairs.
3. If a P-VOL restore was in progress during a hardware failure, delete the pair, restore the P-VOL if possible, and create a new pair

Table 6-4: Data assurance and method for recovering the pair

State before failure	Data assurance	Action taken after pair failure
Failure or PSUE from other than RCPY	P-VOL: Assured S-VOL: Not assured	Resynchronize in the direction of P-VOL to S-VOL. Note that the pair may have been split due to the drive's multiple-malfunction in either or both volumes. In such case, confirm that the data exists in the P-VOL, and then recreate the pair.
Failure(R) or PSUE from RCPY	P-VOL: Not assured S-VOL: Not assured	Delete the pair, restore the backup data to P-VOL, and then create a pair. Note that the pair may have been split due to the drive's multiple-malfunction in either or both volumes. In such case, confirm that the backup data restoration has been completed to the P-VOL, and then recreate the pair

To restore pairs after forcible delete operation

Create or re-synchronize the pair. When an existing pair is re-synchronized, the entire P-VOL is re-copied to the S-VOL.

To recover from a pair failure

Figure 6-1 shows a workflow to be done when a pair failure occurs from determining the factor to restore of the pair status by pair operation.

Table 6-5 on page 6-9 shows the work responsibility schedule for the service personnel and a user.

Figure 6-1: Recovery from a pair failure




Table 6-5: Operational notes for ShadowImage operations

Action	Action taken by
Monitoring pair failure.	User
Verify that operation of suspends the pair by the user.	User
Verify the status of the array.	User
Call maintenance personnel when the array malfunctions.	User
For other reasons, call the Hitachi support center.	User (only for users that are registered in order to receive a support)
Hardware maintenance.	Hitachi Customer Service
Reconfigure and recover the pair.	User

Path failure

When using CCI, and a path fails for more than one minute, the command device may not be recognized when the path is recovered. Execute Windows' "re-scan the disks" to recovery. Restart CCI when Windows is able to recognize the command device but CCI cannot access the command device.

Cases and solutions using the DP-VOLs

When configuring a ShadowImage pair using the DP-VOL as a pair target volume, the ShadowImage pair status may become Failure depending on the combination of the pair status and the DP pool status shown in [Table 6-6 on page 6-10](#). Check the pair status and the DP pool status, and perform the countermeasure according to the conditions. For checking the DP pool status, check all the DP pools to which the P-VOLs and the S-VOLs of the pairs where pair failures have occurred belong. Refer to the *Dynamic Provisioning User's Guide* for how to check the DP pool status. When the DP pool tier mode is enabled, refer to the *Dynamic Tiering User's Guide*.

Table 6-6: Cases and solutions using DP-VOLs

Pair status	DP pool status	Cases	Solutions
Paired Paired Internally Synchronizing Synchronizing Reverse Synchronizing Split Pending	Formatting	Although the DP pool capacity is being added, the format progress is slow and the required area cannot be allocated.	Wait until the formatting of the DP pool for total capacity of the DP-VOLs created in the DP pool is completed.
	Capacity Depleted	The DP pool capacity is depleted and the required area cannot be allocated.	To make the DP pool status normal, perform the DP pool capacity growing and DP pool optimization, and increase the DP pool free capacity.

Copy-on-Write Snapshot theory of operation

Hitachi Copy-on-Write Snapshot creates virtual copies of data volumes within the Hitachi Unified Storage disk array. These copies can be used for recovery from logical errors. They are identical to the original volume at the point in time they were taken.

The key topics in this chapter are:

- ❑ [Copy-on-Write Snapshot software](#)
- ❑ [Hardware and software configuration](#)
- ❑ [How Snapshot works](#)
- ❑ [Snapshot pair status](#)
- ❑ [Interfaces for performing Snapshot operations](#)



NOTE: “Snapshot” refers to Copy-on-Write Snapshot software. A “snapshot” refers to a copy of the primary volume (P-VOL).

Copy-on-Write Snapshot software

Hitachi's Copy-on-Write Snapshot software creates virtual backup copies of any data volume within the disk array with minimal impact to host service or performance levels. These snapshots are suitable for immediate use in decision support, software testing and development, data backup, or rapid recovery operations.

Snapshot minimizes disruption of planned or unplanned outages for any application that cannot tolerate downtime for any reason or that requires non-disruptive sharing of data. Since each snapshot captures only the changes to the original data volume, the amount of storage space required for each Copy-on-Write Snapshot is significantly smaller than the original data volume.

The most probable types of target applications for Copy-on-Write Snapshot are:

- Database copies for decision support/database inquiries
- Non-disruptive backups from a Snapshot secondary volume
- Periodic point-in-time disk copies for rapid restores in the event of a corrupted data volume

Hardware and software configuration

A typical Snapshot hardware configuration includes a disk array, a host connected to the storage system, and management software to configure and manage Snapshot. The host is connected to the storage system using fibre channel or iSCSI connections. The management software is connected to the storage system via a management LAN.

The logical configuration of the disk array includes, primary data volumes (P-VOLs) belonging to the same group, virtual volumes (V-VOLs,) the DP pool and a command device (optional). Snapshot creates a volume pair from a primary volume (P-VOL), which contains the original data, and a Snapshot Image (V-VOL), which contains the snapshot data. Snapshot uses the V-VOL as the secondary volume (S-VOL) of the volume pair. Since each P-VOL is paired with its V-VOL independently, each volume can be maintained as an independent copy set.

The Snapshot system is operated using Hitachi Storage Navigator Modular 2 (Navigator 2) graphical user interface (GUI), Navigator 2 Command-Line interface (CLI), and Hitachi Command Control Interface (CCI).

[Figure 7-1 on page 7-3](#) shows the Snapshot configuration.




Figure 7-1: Snapshot functional component

The following sections describe how these components work together.

How Snapshot works

Snapshot creates a virtual duplicate volume of another volume. This volume "pair" is created when you:

- Select a volume that you want to replicate
- Identify another volume that will contain the copy
- Associate the primary and secondary volumes
- Create a snapshot of primary volume data in the virtual (secondary) volume.

Once a snapshot is made, it remains unchanged until a new snapshot instruction is issued. At that time, the new image replaces the previous image.

Volume pairs — P-VOLs and V-VOLs

A volume pair is a relationship established by Snapshot between two volumes. A pair consists of a production volume, which contains the original data and is called the primary volume (P-VOL), and from 1 to 1024 virtual volumes (V-VOLs), which contain virtual copies of the PVOL.

One P-VOL can pair with up to 1,024 V-VOLs and when one P-VOL pairs with 1,024 V-VOLs, the number of pairs is 1,024. The disk array supports up to 100,000 Snapshot pairs.

The V-VOL is created automatically when creating a pair by specifying the P-VOL. When this is done, the V-VOL can have a volume number at the same time as the pair creation by specifying the volume number for the V-VOL. If the volume number of the V-VOL is not specified, the automatically created V-VOL does not have a volume number.

It is also possible to create the V-VOL before creating the pair by the Snapshot volume creation function. The V-VOL created by the Snapshot volume creation function has a volume number.

There are V-VOLs which have volume numbers and which do not have volume numbers. If volume numbers are assigned to all V-VOLs, the number of creatable Snapshot pairs is restricted to the maximum number of volumes of the array. However, by utilizing the V-VOLs which do not have the volume numbers, the number of Snapshot pairs can be increased to the maximum of 100,000 pairs. The V-VOLs which do not have the volume numbers cannot be recognized by the host. To check the data in the V-VOLs, assign unused volume numbers to the V-VOLs, and further map the volumes and the H-LUNs to be recognized by the host.

To maintain the Snapshot image of the P-VOL when new data is written to the P-VOL, Snapshot copies data that is being replaced to the DP pool. V-VOL pointers in cache memory are updated to reference the original data's new location in the DP pool.

A V-VOL provides a virtual image of the P-VOL at the time of the snapshot. Unlike the P-VOL, which contains actual data, the V-VOL is made up of pointers to the data in the P-VOL, and to original data that has been changed in the P-VOL since the last snapshot and which has been copied to the DP pool.

V-VOL's are set up with volumes that are the same size as the related P-VOL. This capacity is not actually used and remains available as free storage capacity. This V-VOL sizing requirement (must be equal to the P-VOL), is necessary for Snapshot and disk array logic.

Creating pairs

The Snapshot creating pairs operation establishes two newly specified volumes (see [Figure 7-2 on page 7-5](#)). Once a pair is created, the P-VOL and the V-VOL are synchronized. However, since the data copy from the P-VOL to the V-VOL is unnecessary, the pair creation is immediately completed and the pair status becomes Paired. When the pair creation with the **Split the pair immediately after creation is completed** option is specified, the pair status becomes Split.




Figure 7-2: Creating a Snapshot pair

You must specify the P-VOL when creating a pair. There are three patterns for how to specify the V-VOL.

- Specify the existing V-VOL: Create a Snapshot pair using the V-VOL with the specified volume number. Create the V-VOL by the Snapshot volume creation in advance.
- Specify an unused volume number: Create a Snapshot pair using the V-VOL with the specified volume number. At this time, the V-VOL with the specified volume number is added. By using this specification method, a pair can be created without making the automatically created V-VOL by the Snapshot volume creation in advance.
- Omit the volume number of the V-VOL: Create a Snapshot pair by using the V-VOL without the volume number. Since there is no volume number for the V-VOL, it is necessary to use a pair name for identifying a pair in the subsequent pair operation. For checking the backup data in the V-VOL, it is necessary to add a volume number to the V-VOL by the pair edit function and map the volume number and the H-LUN.

Creating pairs options

The following can be specified as an option at the time of the pair creation.

- **Group:** You can select whether making a pair to be created belongs to the group, and if doing so, whether creating a new group to belong to or making it belong to the existing group. When creating a new group, specify a group number to belong to. When making it belong to the existing group, you can specify either group number or group name. The default value does not belong to the group. The group name of the pair not belonging to the group becomes Ungrouped.
- **Copy pace:** Select the copy pace at the pair restoration from Slow, Medium, and Fast. The default value is the Medium. You can change the copy pace which was once set later by using the pair edit function. You may change it when you feel that the creation time takes a long time in the pace specified at the time of the restoration or the effect on the host I/O is significant because the copy processing is given priority.
- **Split the pair immediately after creation is completed:** If specified, the pair status changes to Split and the Read/Write access for the V-VOL becomes possible immediately and the pair cannot belong to the group. The default is not to perform split the pair immediately after creation.
- **MU number:** The MU numbers used in CCI can be specified. The MU number is the management number that is used for configurations where a single volume is shared among multiple pairs. You can specify any value from 0 to 39 by selecting **Manual**. The MU numbers already used by other ShadowImage pairs or SnapShot pairs which share the P-VOL cannot be specified. The free MU numbers are assigned in ascending order from MU number 1 by selecting **Automatic** (the default). The MU number is attached to the P-VOL. The MU number for the S-VOL is fixed as 0.



NOTE: If the MU numbers from 0 to 39 are already used, no more ShadowImage pairs can be created. When creating SnapShot pairs, specify the MU numbers from 40 and more. When creating SnapShot pairs, if you select **Automatic**, the MU numbers are assigned in descending order from 1032.

Splitting pairs

Split the pair for retaining the backup data in the V-VOL. Split the pair for the pair in the Paired status. If the pair is split, the pair status is changed to Split and the backup data at the time of split instruction to the V-VOL can be retained.

The following can be specified as an option at the time of the pair splitting.

The following can be specified as an option at the time of the pair splitting.

- **Attach description to identify:** The character string of the maximum of 31 characters can be added to the split pair. You can also check this character string on the pair list. This is useful for indicating the information of when and for what the backup data retained in the V-VOL was backed up. This character string is only retained while splitting.

Re-synchronizing pairs

When discarding the backup data retained in the V-VOL by split, perform the pair resynchronization. Since the resynchronized pair is immediately changed to Paired, a new backup can be created by split again.

The replication data stored in the DP pool is deleted when re-synchronizing or deleting all the Snapshot pairs which used the same P-VOL. The deletion of replication data is not completed immediately after the pair status is changed to Paired or Simplex; it is completed after a few moments. Time required for the deletion process is proportional to the P-VOL capacity. As a standard, it takes about five minutes with the pair configuration of 1:1 and about 15 minutes with the pair configuration of 1:32 for a 100 GB P-VOL.

Restoring pairs

When the P-VOL data is in the unusable status and returned to the backup data retained in the V-VOL, execute pair restoration.

When the copy processing starts from the V-VOL to the P-VOL by restoration, the pair status becomes Reverse Synchronizing, the copy processing is completed, and the pair status becomes Paired if the P-VOL and the V-VOL are synchronized. If executing restoration, the Read/Write access from the host to the P-VOL can be continued immediately after the restore operation is executed even while the pairs are reverse resynchronizing. Even if the P-VOL and the V-VOL are not synchronized, since it appears to the host that there is the V-VOL data in the P-VOL immediately after restoration, the operation can restart immediately.



- The pair in the Reverse Synchronizing status cannot be split.
 - Read/Write access cannot be performed for the V-VOL in the Reverse Synchronizing status. Furthermore, in the configuration where multiple pairs are created in one P-VOL, Read/Write to the other V-VOLs also becomes impossible. Once the restoration is completed, Read/Write to the other V-VOL in the Split status becomes possible again.
 - The pair can be deleting while its status is Reverse Synchronizing, however, the P-VOL data being restored cannot be used logically. The V-VOLs correlated to the P-VOL and with a status other than Simplex are placed in Failure status. Do not delete a pair while the pair status is Reverse Synchronizing, except when an emergency exists.
-

When the restoration instruction of the V-VOL data is issued to the P-VOL, the pair status is not changed to Paired immediately but is changed to Reverse Synchronizing. The data of the P-VOL, however, is replaced promptly with the backup data retained in the V-VOL. When the pair is split to the other V-VOL after the issue of the restoration instruction, the V-VOL retains the data it had at the time of the pair splitting based on the P-VOL data. That data is replaced with the backup data, even before the restoration completes.

Here is a rough estimate on how much time is required for the search process to complete. Note that the actual time can depend on the configuration.

Test conditions: With a total of 100 GB of P-VOL, restoration runs on 4 P-VOLs without I/O at the same time.

- 1 P-VOL with 1 V-VOL: About 6 minutes
- 1 P-VOL with 8 V-VOLs: About 22 minutes
- 1 P-VOL with 32 V-VOLs: About 36 minutes





Figure 7-3: Snapshot operation performed to the other V-VOL during the restoration

When no differential exists between a P-VOL and V-VOL to be restored, restoration is not completed immediately. It takes time to examine the differential between the P-VOL and V-VOL. The method of search for the differentials is "Search All".





- We recommend the copy pace be "Medium". However, if you specify "Medium", the time to complete the copying may differ according to the host I/O load. If you specify the copy pace to "Fast", the host I/O performance deteriorates. When you want to suppress the deterioration of the host I/O performance further from the case you specify the copy pace to "Medium", specify "Slow".
- The restoration command can be issued up to 128 P-VOLs at the same time. However, the number of P-VOLs to which the physical copying (background copying) from a V-VOL can be done at the same time is up to four (HUS 110) per controller (HUS 130/HUS 150: eight per controller). When the background copying can be executing, background copies are completed in the order the command was issued. The other background copies are completed in the ascending order of volume numbers, after the preceding restoration is completed.

Deleting pairs

When the **Delete Pair** button is pushed, any pair in the Paired, Split, Reverse Synchronizing, Failure, or Failure(R) status can be deleted at any time after it is placed in the Simplex status.

When the **Delete Pair** button is pushed, the V-VOL data is annulled immediately and invalidated. Therefore, if you access the V-VOL after the pair is deleted, the data retained before the pair is deleted is not available.

The V-VOL without the volume number is automatically deleted with the pair deletion.

Unnecessary replication data is removed from the DP pool when a pair is deleted. Removing unnecessary replication data does not finish shortly after the pair status changes to Simplex and will take a while to complete. The time required for this process increases with the P-VOL capacity.

DP pools

A V-VOL is a virtual volume that does not actually have disk capacity. In order to make the V-VOL retain data at the time when the pair splitting instruction is issued, it is required to save the P-VOL data as the differential data before it is overwritten by the Write command. The saved differential data is called replication data.

The information to manage the Snapshot pair configuration and its replication data is called management information.

The replication data and the management information are stored in the DP pool. The DP pool storing the replication data is called the replication data DP pool and the DP pool storing the management information is called the management area DP pool. The replication data and the management information can be stored in the different DP pools separately or can be stored in the same DP pool. When they

are stored in the same DP pool, the replication data DP pool and the management area DP pool indicate the same DP pool. Since a Snapshot pair needs a DP pool, it is necessary to validate Hitachi Dynamic Provisioning (HDP).

Up to 64 DP pools (HUS 130/HUS 150) or up to 50 DP pools (HUS 110) can be created per disk array and a DP pool to be used by a certain P-VOL is specified when a pair is created. A DP pool to be used can be specified for each P-VOL, and V-VOLs that pair with the same P_VOL must use a common DP pool. Two or more Snapshot pairs can share a single DP pool.

It is not necessary to make the DP pool used for the Snapshot pair the DP pool specific for the Snapshot pair. The DP-VOL can be created in the DP pool used by the Snapshot pair

The Replication threshold value can be set in the DP pool. In the Replication threshold value, the replication Depletion Alert threshold value and the Replication Data Released threshold value can be set. The threshold value to be set is the ratio of the usage of the DP pool for the entire capacity of the DP pool. Setting the Replication threshold value helps prevent the DP pool from becoming depleted by Snapshot. Always set the larger Replication Data Released threshold value than the replication Depletion Alert threshold value.

When the usage rate of the replication data DP pool or management area DP pool reaches the Replication Depletion Alert threshold value, the pair status of the Split status pair changes to Threshold Over and notices that the usable amount of the DP pool is reduced. At the point when the usage rate of the DP pool recovers to over -5% of the replication Depletion Alert threshold value, it is returned to the Split status. The replication Data Released threshold value cannot be set within the range of -5% of the Replication Depletion Alert threshold value. When the usage rate of the replication data DP pool or management area DP pool reaches the Replication Data Released threshold value, all the Snapshot pairs in the DP pool in which the threshold value is set are changed to the Failure status. At the same time, the replication data and the management information are released and the usable capacity of the DP pool recovers. Until the usage rate of the DP pool recovers to over -5% of the Replication Data Released threshold value, all pair operations except for pair deletion cannot be performed.

Consistency Groups (CTG)

Application data often spans more than one volume. With Snapshot, it is possible to manage operations spanning multiple volumes as a single group. In a group, all primary logical volumes are treated as a single entity.

Managing Snapshot primary volumes as a group allows multiple operations to be performed on grouped volumes concurrently. Write order is guaranteed across application logical volumes, since snapshots can be taken at the same time, thus ensuring consistency.

By making multiple pairs belong to the same group, the pair operation is possible in units of groups. In the group whose Point-in-time attribute of the group is enabled, the backup data of the S-VOL created in units of groups is the data of the same time.

For setting a group, specify a new group number for a group to be assigned after pair creation when creating a Snapshot pair. The maximum of 1,024 groups can be created in Snapshot.

A group name can be assigned to a group. You can select one pair belonging to the created group and assign a group name arbitrarily by using the pair edit function.



- If CCI is used, a group whose Point-in-Time attribute is disabled can be created. In Navigator 2, only the group whose Point-in-Time attribute is enabled can be created.
 - You cannot change the group specified at the time of the pair creation. To change it, delete the pair once, and specify another group when creating a pair again.
-

Command devices

The command device is a user-selected, dedicated logical volume on the disk array, which functions as the interface to the CCI software. Snapshot commands are issued by CCI (HORCM) to the disk array command device.

A command device must be designated in order to issue Snapshot commands. The command device must be defined in the HORCM_CMD section of the configuration definition file for the CCI instance on the attached host. 128 command devices can be designated for the disk array. You can designate command devices using Navigator 2.



NOTE: Volumes set for command devices must be recognized by the host. The command device volume size must be greater than or equal to 33 MBs.

Differential data management

When a split operation is performed, the V-VOL is assured through management of differential data (locations of Write data from a host contained in the P-VOL and V-VOL) and reference to the bit map performed in a Read operation by the host (see [Figure 7-4](#)). The extent of one bit in the bit map is equivalent to 64 kB. Therefore:

- Even an update of a single kB update requires a data transfer as large as 64 kB to copy from the P-VOL to the DP pool.




Figure 7-4: Differential data

Snapshot pair status

Snapshot displays the pair status of all Snapshot volumes. [Figure 7-5](#) shows the Snapshot pair status transitions and the relationship between the pair status and the Snapshot operations.




Figure 7-5: Snapshot pair status transitions

[Table 7-1 on page 7-17](#) lists and describes the Snapshot pair status conditions.

If a volume is not assigned to a Snapshot pair, its status is Simplex. When the pair is created, the pair status becomes Paired. If the pair is split in this status, the pair status becomes Split and can access the V-VOL. When the Create Pair button is pushed with the **Split the pair immediately after creation is completed** option specified to it, statuses of the P-VOL and the V-VOL change from Simplex to Split.

It is possible to access the P-VOL or V-VOL in the Split state. The pair status changes to Failure (interruption) when the V-VOL cannot be created or updated, or when the V-VOL data cannot be retained due to a disk array failure. Also, if a similar failure occurs when the restoration is instructed and the pair status is the Reverse Synchronizing status, the pair status becomes Failure(R). The P-VOL whose pair status is Failure(R) cannot Read/Write. When the Delete Pair button is pushed, the pair is split and the pair status changes to Simplex.

Table 7-1: Snapshot pair status

Pair status	Description	P-VOL access	V-VOL access
Simplex	If a volume is not assigned to a Snapshot pair, its status is Simplex. If the created pair is deleted, the pair status becomes Simplex. Note that the Simplex volume is not displayed on the list of the Snapshot pair. The P-VOL in the Simplex status accepts I/O operations of Read/Write. The V-VOL does not accept any Read/Write I/O operations.	Read and write.	Read/write is not available.
Paired	The data of the P-VOL and the V-VOL is in the same status. However, since Read to the V-VOL in the Paired status and the Write access cannot be performed, it is actually the same as Simplex.	Read and write.	Read/write is not available.
Split	The P-VOL data at the time of the pair splitting is retained in the V-VOL. When a change of the P-VOL data occurs, the P-VOL data at the time of the split instruction is retained as the V-VOL data. The P-VOL and V-VOL accept Read/Write I/O operations.	Read and write.	Read and write. A Read/Write instruction is not acceptable when the P-VOL is being restored.
Reverse Synchronizing	Data is copied from the V-VOL to the P-VOL for the area where the difference between the P-VOL and the V-VOL exists. When multiple pairs are created for one P-VOL, if a failure occurs when a certain pair is in the Synchronizing status or if a pair in the Reverse Synchronizing status is deleted, different pairs in the same P-VOL all become Failure.	Read and write.	Read/write is not available
Failure	Failure is a status in which the P-VOL data at the time of the split instruction cannot be retained in the V-VOL due to a failure in the disk array. In this status, I/O operations of Read/Write concerning the P-VOL are accepted as before. The V-VOL data is invalidated at this point of time. To resume the split pair, it is required to execute the pair creation again and split the pair once. However, data of the V-VOL created is not the former version that was invalidated, but the P-VOL data at the time of the new pair splitting.	Read and write	Read/write is not available
Failure (R)	This is a state in which the P-VOL data becomes unjustified due to a Failure during restoration (in Reverse Synchronizing status). The P-VOL cannot perform either the Read or Write access. To make it access, it is necessary to delete the pair.	Read/write is not available	Read/write is not available

Pair status	Description	P-VOL access	V-VOL access
Threshold Over	A status in which the replication depletion alert of DP pool reaches the threshold. However, Threshold Over internally operates as Split. To reference the pair status, you are able to recognize as Threshold Over. You can reduce the usage rate of the DP pool by adding the DP pool capacity, deleting the unnecessary Snapshot pair, or deleting the unnecessary DP-VOL.	Read and write.	Read and write. A Read/Write instruction is not acceptable when the P-VOL is being restored

Interfaces for performing Snapshot operations

Snapshot can be operated using the following interfaces:

- Navigator 2 GUI (Hitachi Storage Navigator Modular 2 Graphical User Interface) is a browser-based interface from which Snapshot can be setup, operated, and monitored. The GUI provides the simplest method for performing operations, requiring no previous experience. Scripting is not available.
- CLI (Hitachi Storage Navigator Modular 2 Command Line Interface), from which Snapshot can be setup and all basic pair operations can be performed—create, split, resynchronize, restore, swap, and delete. The GUI also provides these functionalities. CLI also has scripting capability.
- CCI (Hitachi Command Control Interface), used to display volume information and perform all copying and pair-managing operations. CCI provides a full scripting capability which can be used to automate replication operations. CCI requires more experience than the GUI or CLI. CCI is required on Windows 2000 Server for performing mount/unmount operations.

HDS recommends using the GUI to begin operations for new users with no experience with CLI or CCI. Users who are new to replication software but have CLI experience in managing disk arrays may want to continue using CLI, though the GUI is an option. The same recommendation applies to CCI users.



NOTE: Hitachi Replication Manager is used to manage and integrate Copy-on-Write. It provides a GUI topology view of the Snapshot system, with monitoring, scheduling, and alert functions. For more information on purchasing Replication Manager, visit the Hitachi Data Systems website <http://www.hds.com/products/storage-software/hitachi-replication-manager.html>

Installing Snapshot

Snapshot must be installed on the Hitachi Unified Storage using a license key. It can also be disabled or uninstalled. This chapter provides instructions for performing these tasks.

- ❑ [System requirements](#)
- ❑ [Installing or uninstalling Snapshot](#)
- ❑ [Enabling or disabling Snapshot](#)

System requirements

This topic describes minimum system requirements and supported platforms. System requirements

The following table shows the minimum requirements for Snapshot. See [Snapshot specifications on page B-2](#) for additional information.

Table 8-1: Storage system requirements

Item	Requirements
Firmware	Firmware: Version 0916/B or more is required.
Storage Navigator Modular 2	Version 21.50 or more is required for the management PC.
CCI	Version 01-27-03/02 or later is required for the host when CCI is used for the Snapshot operation.
Command devices	Maximum of 128. The command device is required only when CCI is used for Snapshot operation. The command device volume size must be greater than or equal to 33 MB.
DP pool	Maximum of 64 for HUS 150/130 and of 50 for HUS 110. • One per controller required; two per controller highly recommended. • One or more pairs can be assigned to a DP pool.
Volume size	V-VOL size must equal P-VOL size.
Number of controllers	Two

Supported platforms

The following table shows the supported platforms and operating system versions required for Snapshot.

Table 8-2: Supported platforms

Platforms	Operating system version
SUN	Solaris 8 (SPARC)
	Solaris 9 (SPARC)
	Solaris 10 (SPARC)
	Solaris 10 (x86)
	Solaris 10 (x64)
PC Server (Microsoft)	Windows 2000
	Windows Server 2003 (IA32)
	Windows Server 2008 (IA32)
	Windows Server 2003 (x64)
	Windows Server 2008 (x64)
	Windows Server 2003 (IA64)
	Windows Server 2008 (IA64)
HP	HP-UX 11i V1.0 (PA-RISC)
	HP-UX 11i V2.0 (PA-RISC)
	HP-UX 11i V3.0 (PA-RISC)
	HP-UX 11i V2.0 (IPF)
	HP-UX 11i V3.0 (IPF)
	Tru64 UNIX 5.1
IBM®	AIX 5.1
	AIX 5.2
	AIX 5.3
Red Hat	Red Hat Linux AS2.1 (IA32)
	Red Hat Linux AS/ES 3.0 (IA32)
	Red Hat Linux AS/ES 4.0 (IA32)
	Red Hat Linux AS/ES 3.0 (AMD64/EM64T)
	Red Hat Linux AS/ES 4.0 (AMD64/EM64T)
	Red Hat Linux AS/ES 3.0 (IA64)
	Red Hat Linux AS/ES 4.0 (IA64)
SGI	IRIX 6.5.x

Installing or uninstalling Snapshot

A key code or key file is required to install or uninstall. If you do not have the key file or code, you can obtain it from the download page on the HDS Support Portal, <http://support.hds.com>.

- Installation instructions are provided here for Navigator 2 GUI.
- For CLI instructions, see [Operations using CLI on page B-6](#) (advanced users only).

Before installing or uninstalling Snapshot, verify that the Storage system is operating in a normal state. Installation/un-installation cannot be performed if a failure has occurred.

Installing Snapshot

1. In the Navigator 2 GUI, click the disk array in which you will install Snapshot.
2. Click **Show & Configure array**.
3. Select the **Install License** icon in the **Common array Task**.




The **Install License** screen appears.



4. Select the **Key File** or **Key Code** option, and then enter the file name or key code. You may **Browse** for the key file.
5. A screen appears, requesting a confirmation to install Snapshot option. Click **Confirm**.

Install License



6. A completion message appears. Click **Close**.

Install License



Installation of Snapshot is now complete.



NOTE: Snapshot requires the DP pool of Hitachi Dynamic Provisioning (HDP). If HDP is not installed, install HDP.


Uninstalling Snapshot

Snapshot pairs must be released and their status returned to **Simplex** before uninstalling (the status of all volumes are **Simplex**). The key code or key file provided with the optional feature is required. Once uninstalled, Snapshot cannot be used again until it is installed using the key code or key file.

The replication data is deleted after the pair deletion is completed. The replication data deletion may be operated in the background at the time of the pair deletion. Check that the DP pool capacity is recovered after the pair deletion. If it is recovered, the replication data has been deleted.

All Snapshot volumes (V-VOL) must be deleted.

1. In the Navigator 2 GUI, click the check box for the disk array where you want to uninstall Snapshot.
2. Click **Show & Configure disk array**.
3. In the tree view, click **Settings**, then select the **Licenses** icon.



4. On the Licenses screen, click **De-install License**.



5. When you uninstall the option using the key code, click the **Key Code** option, and then set up the key code. When you uninstall the option using the key file, click the **Key File** option, and then set up the path for the key file name. Use Browse to set the path to a key file correctly. Click OK.
6. A message appears. Click **Close**.
Un-installation of Snapshot is now complete.

Enabling or disabling Snapshot

Once Snapshot is installed, it can be enabled or disabled.

In order to disable Snapshot, all Snapshot pairs must be released (the status of all volumes are Simplex).

The replication data is deleted after the pair deletion is completed. The replication data deletion may be operated in the background at the time of the pair deletion. Check that the DP pool capacity is recovered after the pair deletion. If it is recovered, the replication data has been deleted

All Snapshot volumes (V-VOL) must be deleted.

(For instructions using CLI, see [Enabling or disabling Snapshot on page B-10.](#))

1. In the Navigator 2 GUI, select the disk array where you want to enable Snapshot, and click **Show & Configure disk array**.
2. In the tree view, click **Settings**, then click **Licenses**.
3. Select **SNAPSHOT** in the Licenses list, then click **Change Status**.

The Change License screen appears.



4. To enable, check the **Enable: Yes** box.
To disable, clear the **Enable: Yes** box.
5. Click **OK**.
6. A message appears, click **Confirm**.

Reconfigure Memory



7. A message appears, click **Close**.



Snapshot setup

This chapter provides required information for setting up your system for Snapshot. It includes:

Planning and design

- ☐ [Planning and design](#)
- ☐ [Plan and design workflow](#)
- ☐ [Assessing business needs](#)
- ☐ [DP pool capacity](#)
- ☐ [DP pool consumption](#)
- ☐ [Calculating DP pool size](#)
- ☐ [Pair assignment](#)
- ☐ [Operating system host connections](#)
- ☐ [Array functions](#)

Configuration

- ☐ [Configuration workflow](#)
- ☐ [Setting up the DP pool](#)
- ☐ [Setting the replication threshold \(optional\)](#)
- ☐ [Setting up the Virtual Volume \(V-VOL\) \(manual method\) \(optional\)](#)
- ☐ [Deleting V-VOLs](#)
- ☐ [Setting up the command device \(optional\)](#)
- ☐ [Setting the system tuning parameter \(optional\)](#)

Planning and design

A snapshot ensures that volumes with bad or missing data can be restored. With Copy-on-Write Snapshot, you create copies of your production data that can be used for backup and other uses.

Creating a copy system that fully supports business continuity is best done when Snapshot is configured to match your business needs.

This chapter guides you in planning a configuration that meets organization needs and the workload requirements of your host application.

Plan and design workflow

The Snapshot planning effort consists of determining the number of V-VOLs required by your organization, the V-VOLs' lifespan, that is, how long they must be held before being updated again, the frequency that snapshots are taken, and the size of the DP pool. This information is found by analyzing business needs and measuring write workload sent by the host application to the primary volume.

The plan and design workflow consists of the following:

- Assess business needs.
- Determine how often a snapshot should be taken.
- Determine how long the snapshot should be held.
- Determine the number of snapshot copies required per P-VOL.
- Measure production system write workload.

These objectives are addressed in detail in this chapter. Two other tasks are required before your design can be implemented, which are also addressed in this chapter:

- When you have established your Snapshot system design, the system's maximum allowed capacity must be calculated. This has to do with how the disk array manages storage segments.
- Equally important in the planning process are the ways that various operating systems interact with Snapshot.

Assessing business needs

Business needs have to do with how long back-up data needs to be retained and what the business or organization can tolerate when disaster strikes.

The following organizational priorities help determine the following:

- How often a snapshot should be made (frequency)
- How long a snapshot (the V-VOL) should be held (lifespan)
- The number of snapshots (V-VOLs) that will be required for the P-VOL.

Copy frequency

How often copies need to be made is determined by how much data could be lost in a disaster before business is significantly impacted.

To determine how often a snapshot should be taken

- Decide how much data could be lost in a disaster without significant impact to the business.

Ideally, a business desires no data loss. But in the real world, disasters occur and data is lost. You or your organization's decision makers must decide the number of business transactions, the number of hours required to key in lost data, and so on.

If losing 4 hours of business transaction is acceptable, but not more, backups should be planned every 4 hours. If 24 hours of business transaction can be lost, backups may be planned every 24 hours.

Determining how often copies should be made is one of the factors used to determine DP pool size. The more time that elapses between snapshots, the more data accumulates in the DP pool. Copy frequency may need to be modified to reduce the DP pool size.

Selecting a reasonable time between Snapshots

The length of time between snapshots, if too short or too long, can cause problems.

- When short periods are indicated by your company's business needs, consider also that snapshots taken too frequently could make it impossible to recognize logical errors in the storage system. This would result in snapshots of bad data. How long does it take to notice and correct such logical errors? The time span for snapshots should provide ample time to locate and correct logical errors in the storage system.
- When longer periods between snapshots are indicated by business needs, consider that the longer the period, the more data accumulates in the DP pool. Longer periods between backups require more space in the DP pool.

This effect is multiplied if more than one V-VOL is used. If you have two snapshots of the P-VOL, then two V-VOLs are tracking changes to the P-VOL at the same time.

Establishing how long a copy is held (copy lifespan)

Copy lifespan is the length of time a copy (V-VOL) is held, before a new backup is made to the volume. Lifespan is determined by two factors:

- Your organization's data retention policy for holding onto backup copies.
- Secondary business uses of the backup data.

Lifespan based on backup requirements

- If the snapshot is to be used for tape backups, the minimum lifespan must be => the time required to copy the data to tape. For example:

Hours to copy a V-VOL to tape = 3 hours

V-VOL lifespan => 3 hours

- If the snapshot is to be used as a disk-based backup available for online recovery, you can determine the lifespan by multiplying the number of generations of backup you want to keep online by the snapshot frequency. For example:

Generations held = 4

Snapshot frequency = 4 hours

$4 \times 4 = 16$ hours

V-VOL lifespan = 16 hours

Lifespan based on business uses

- If you use snapshot data (the V-VOL) for testing an application, the testing requirements determine the amount of time a snapshot is held.
- If snapshot data is used for development purposes, development requirements may determine the time the snapshot is held.
- If snapshot data is used for business reports, the reporting requirements can determine the backup's lifespan.

Establishing the number of V-VOLs

V-VOL frequency and lifespan determine the number of V-VOLs your system needs per P-VOL.

For example: Suppose your data must be backed up every 12 hours, and business-use of the data in the V-VOL requires holding it for 48 hours. In this case, your Snapshot system would require 4 V-VOLs, since there are four 12-hour intervals during the 48-hour period. This is illustrated in [Figure 9-1](#).




Figure 9-1: V-VOL frequency and lifespan

DP pool capacity

You need to calculate how much capacity must be allocated for the DP pool to have TCE pairs. The capacity required will automatically be taken from the free portion of the DP pool as needed when old data is sent to the DP pool. However, the capacity of the DP pool is not unlimited. So you still need to consider how much capacity is left in the pool for Snapshot.

Using Snapshot consumes DP pool capacity with replication data and management information stored in DP pools, which are differential data between a P-VOL and an S-VOL and information to manage the replication data, respectively. On the other hand, some pair operations, such as pair deletion, recover the usable capacity of the DP pool by removing unnecessary replication data and management information from the DP pool. The following sections show the occasions when replication data and management information increase and decrease, and also how much DP pool capacity they consume.

DP pool consumption

Table 9-1 shows when replication data and management information increase and decrease. An increase in the replication data and management information leads to a decrease in the capacity of the DP pool that Snapshot pairs are using. And a decrease in the replication data and management information recovers the DP pool capacity used by Snapshot pairs.

Table 9-1: Increases and decreases in DP pool consumption

Data type	Data increase	Data decrease
Replication Data	Write on P-VOL/V-VOL.	Execution of pair resync, restore, deletion, and pair status change to PSUE
Management Information	Read/Write on new areas of P-VOL/V-VOL The management information will not increase when Read/Write is executed on the areas where Read/Write has ever been executed	Delete all belonging to the P-VOL The management information will not decrease when pair deletion does not delete all the pairs belonging to a P-VOL.

How much DP pool capacity the replication data and management information need depends on several factors, such as the capacity of a P-VOL and the number of generations.

Determining DP pool capacity

The following indicates how much DP pool capacity the replication data and management information need depending on several factors such as the capacity of a P-VOL and the number of generations.

Replication data

The replication data increases with increasing writes on the P-VOL/V-VOL of a pair in Split status. The formula for the amount of replication data for one V-VOL paired with P-VOL is:

$$\text{P-VOL capacity} \times (100 - \text{rate of coincidence (Note 1)}) \div 100$$

Calculation example for 100 GB of P-VOL and rate of coincidence of 50%:

$$100 \text{ GB} \times (100 - 50) \div 100 = 50 \text{ GB}$$

**NOTES:**

1. The rate of coincidence of data between P-VOL and V-VOL. It indicates 100% once the pair status changes to Split as there is no replication data with the pair, which is the maximum value for the rate of coincidence.
 2. The replication data consumes DP pool capacity by 1 GB. For example, even when the actual amount of replication data stored in a DP pool is less than 1 GB, you will see that 1 GB of the DP pool appears to be consumed.
 3. When one P-VOL is paired with multiple V-VOLs, the amount of replication data can be less than one of replication data for one V-VOL × number of V-VOLs, because replication data can be shared between multiple V-VOLs that belongs to the same P-VOL.
-

Management information

The management information increases with an increase in the P-VOL capacity, the number of generations and the amount of replication data. The following tables show the maximum amount of management information per P-VOL (see Note 1). The management information is shared between Snapshot and TCE.

The generation numbers in the following tables are the maximum number of generations that have been created and split for a P-VOL, not the current number of generations.

**NOTES:**

1. The amount of management information when Read/Write has been executed on the entire area of P-VOL and V-VOLs.
 2. The maximum number of generations that have ever been created and split for a P-VOL, not the current number of generations. For example, even if you reduce the number of generations for a P-VOL from 200 to 100 by pair deletion, the P-VOL will retain the management information for the 200 generations. You can release the management information for all of 200 generations by deleting all of the 200 generations.
 3. If you create a Snapshot-TCE cascade configuration, the management information for Snapshot is only needed as listed in [Table 9-2 on page 9-10](#) through [Table 9-4 on page 9-11](#). See [Figure 9-3 on page 9-9](#) for an example.
-




Figure 9-2: No Snapshot-TCE cascade configuration




Figure 9-3: Snapshot-TCE cascade configuration

Table 9-2: Maximum amount of management information: 100 % of P-VOL capacity

P-VOL capacity	Generation numbers (Note 2)					
	1 to 60	1 to 120	1 to 360	1 to 600	1 to 850	1 to 1024
50 GB	5 GB	5 GB	7 GB	10 GB	11 GB	14 GB
100 GB		6 GB	10 GB	15 GB	19 GB	23 GB
250 GB	7 GB	10 GB	20 GB	31 GB	41 GB	52 GB
500 GB	10 GB	15 GB	36 GB	58 GB	80 GB	99 GB
1 TB	15 GB	25 GB	68 GB	112 GB	155 GB	195 GB
2 TB	27 GB	48 GB	133 GB	220 GB	305 GB	385 GB
4 TB	49 GB	92 GB	263 GB	435 GB	606 GB	766 GB
8 TB	94 GB	180 GB	522 GB	866 GB	1,208 GB	1,528 GB
16 TB	184 GB	356 GB	1,041 GB	1,727 GB	2,413 GB	3,052 GB
32 TB	364 GB	708 GB	2,079 GB	3,450 GB	4,822 GB	6,101 GB
64 TB	725 GB	1,411 GB	4,154 GB	6,898 GB	9,642 GB	12,200 GB
128 TB	1,445 GB	2,818 GB	8,305 GB	13,793 GB	19,281 GB	24,392 GB

Table 9-3: Maximum amount of management information: 50 % of P-VOL capacity

P-VOL capacity	Generation numbers (Note 2)					
	1 to 60	1 to 120	1 to 360	1 to 600	1 to 850	1 to 1024
50 GB	5 GB	5GB	7 GB	9 GB	10 GB	12 GB
100 GB		6 GB	9 GB	14 GB	17 GB	21 GB
250 GB	7 GB	9 GB	18 GB	28 GB	37 GB	47 GB
500 GB	9 GB	14 GB	32 GB	52 GB	71 GB	89 GB
1 TB	14 GB	23 GB	61 GB	100 GB	138 GB	174 GB
2 TB	24 GB	42 GB	118 GB	195 GB	271 GB	344 GB
4 TB	44 GB	82 GB	233 GB	386 GB	537 GB	683 GB
8 TB	84 GB	160 GB	463 GB	767 GB	1,070 GB	1,363 GB
16 TB	163 GB	316 GB	922 GB	1,530 GB	2,137 GB	2,721 GB
32 TB	323 GB	627 GB	1,841 GB	3,056 GB	4,271 GB	5,440 GB
64 TB	642 GB	1,250 GB	3,679 GB	6,109 GB	8,539 GB	10,876 GB
128 TB	1,280 GB	2,495 GB	7,355 GB	12,215 GB	17,075 GB	21,747 GB

Table 9-4: Maximum amount of management information: 25 % of P-VOL capacity

P-VOL capacity	Generation numbers					
	1 to 60	1 to 120	1 to 360	1 to 600	1 to 850	1 to 1024
50 GB	5 GB		7 GB	9 GB	10	12 GB
100 GB	5 GB	6 GB	9 GB	13 GB	16 GB	19 GB
250 GB	7 GB	9 GB	17 GB	26 GB	35 GB	44 GB
500 GB	9 GB	13 GB	30 GB	49 GB	67 GB	84 GB
1 TB	13 GB	22 GB	57 GB	94 GB	129 GB	164 GB
2 TB	23 GB	40 GB	111 GB	183 GB	254 GB	323 GB
4 TB	41 GB	76 GB	218 GB	361 GB	503 GB	642 GB
8 TB	79 GB	150 GB	433 GB	718 GB	1,001 GB	1,280 GB
16 TB	153 GB	296 GB	863 GB	1,431 GB	1,999 GB	2,556 GB
32 TB	302 GB	587 GB	1,722 GB	2,859 GB	3,995 GB	5,109 GB
64 TB	600 GB	1,169 GB	3,441 GB	5,715 GB	7,988 GB	10,215 GB
128 TB	1,197 GB	2,334 GB	6,880 GB	11,426 GB	15,972 GB	20,425 GB

Calculating DP pool size

You need to calculate how much capacity must be allocated for the DP pool to have Snapshot pairs. The capacity required will automatically be taken from the free portion of the DP pool as needed when old data is sent to the DP pool. However, the capacity of the DP pool is not unlimited. So you still need to consider how much capacity is left in the pool for Snapshot.

The factors that determine the DP pool capacity include:

- A total capacity of P-VOLs
- A number of generations (a number of V-VOLs)
- An interval of the pair splitting (a period for holding the V-VOL)
- An amount of data updated during the interval and the spare capacity (safety rate)

The method for calculating the DP pool capacity is:

$$\text{DP pool capacity} = \text{P-VOL capacity} \times (\text{amount of renewed data} \times \text{safety rate}) \times \text{a number of generations}$$

When restoring backup data stored in a tape device, add more than the P-VOL capacity, the recommendation is 1.5 times of the P-VOL capacity or more) to the DP pool capacity computed from the above formula. This will provide sufficient free DP pool capacity, larger than the P-VOL capacity.

Typically, the rate of updated data amount per day is approximately 10%.

If one V-VOL was created for 1 TB P-VOL and pair splitting was given once a day, the recommended DP pool capacity will be approximately 250 GB, with the safety rate as approximately 2.5 times, considering the variance in amount of data renewal due to locality of access and operations.

When five V-VOLs are made per P-VOL of 1 TB and one pair splitting is issued to one of The five V-VOLs a day (each of the V-VOLs holds data for a period of five days), the recommended DP pool capacity is about 1.2 TB (five times the capacity in when only one V-VOL is made).

The recommended value of the DP pool capacity when a capacity of one P-VOL is 1 TB is shown in the following table. Multiply the value in [Table 9-5 on page 9-13](#) by a value of a capacity per V-VOL in TB.

The value shown above is merely a standard because the amount of data actually accumulated in the DP pool varies depending on the application, amount of processed data, and time zone of operation, etc. If a DP pool has a small capacity, it will become full and all V-VOLs will be placed in the Failure status. When introducing Snapshot, provide a DP pool with a sufficient capacity and verify the DP pool capacity beforehand. Monitor the used capacity because it changes, depending on the system operation.

Table 9-5: Recommended value of the DP pool capacity (when the P-VOL capacity is 1 TB)

An interval of Pair Splitting*	V-VOL Number (n)					
	1	2	3	4	5	6 to 14
From one to four hours	0.10 TB	0.20 TB	0.30 TB	0.40 TB	0.50 TB	0.10 x n TB
From four to eight hours	0.15 TB	0.30 TB	0.45 TB	0.60 TB	0.75 TB	0.15 x n TB
From eight to 12 hours	0.20 TB	0.40 TB	0.60 TB	0.80 TB	1.00 TB	0.20 x n TB
From 12 to 24 hours	0.25 TB	0.50 TB	0.75 TB	1.00 TB	1.25 TB	0.25 x n TB

*An interval of pair splitting means a time between pair splitting issued to the designated P-VOL. When there is only one V-VOL, the interval of the pair splitting is as long as the period for retaining the V-VOL. When there are two or more V-VOLs, the interval of the pair splitting multiplied by the number of the V-VOLs is the period for retaining the one V-VOL.

Construct a system in which the interval of pair splitting is less than one day (see [Figure 9-4](#)). It becomes difficult, depending on system environment to estimate the amount of data accumulated in the DP pool when the interval of the pair splitting is long.




Figure 9-4: Interval of pair splitting

When setting two or more pairs (P-VOLs) per DP pool, count the DP pool capacity by calculating a DP pool capacity necessary for each pair and adding the calculated values together.

A capacity of the DP pool can be expanded through addition of the RAID group(s) in the online status (while the Snapshot pair is created).

When returning the backup data from a tape device to the V-VOL, a free DP pool with a capacity larger than the P-VOL capacity is required. More than 1.5 times the P-VOL capacity is recommended. See [Figure 9-5](#).



NOTE: A volume with a SAS drives, a volume with SSD/FMD drives, and a volume with a SAS7.2K drives cannot coexist in the DP pool.




Figure 9-5: DP pool capacity

Requirements and recommendations for Snapshot Volumes

Please review the following key rules and recommendations regarding P-VOLS, V-VOLS, and DP pools. See [Snapshot specifications on page B-2](#), for general specifications required for Snapshot.

- Primary volumes must be set up prior to making Snapshot copies.
- Assign four or more disks to Snapshot volumes for optimal host and copying performance.
- Volumes used for other purposes should not be assigned as a primary volume. If such a volume must be assigned, move as much of the existing write workload to non-Snapshot volumes as possible.

- If multiple P-VOLs are located in the same drive, the status of the pairs should stay the same (Simplex, Paired, and Split). When status differs, performance is difficult to estimate.

Pair assignment

- Do not assign a frequently updated volume to a pair

When the pair status is Split, the old data is copied to a DP pool when writing to a primary volume. Because the load on the processor in the controller is increased, the writing performance becomes limited. When the writing load becomes heavier due to: a large number of write operations, the writing of data with a large block size, frequent write I/O instructions, and continuous writing, the effect becomes the greater. Therefore, be strict in selection of the volume to which Snapshot is applied. When applying Snapshot to a volume bearing a heavy writing load, it is necessary to consider making loads on the other volumes lighter.

- Use a small number of volumes within the same RAID group.

When volumes are assigned to the same RAID group and used as primary volumes, there may be situations where the host I/O for one of the volumes causes restriction on host I/O performance of the other volume(s) due to drive contention. Therefore, it is recommended that you assign few (one or two) primary volumes to the same RAID group. When creating pairs within the same RAID group, standardize the controllers that control volumes in the same RAID group.

- Make an exclusive RAID group for a DP pool.

When another volume is assigned to a RAID group to which a DP pool has been assigned, loads on drives are increased and their performance is restricted because primary volumes correspond to the DP pool in common. Therefore, use a RAID group, to which a DP pool is assigned, for the DP pool only. There can be multiple DP pools in a disk array. Please use different RAID groups for each DP pool.

- Use the SAS drives or the SSD/FMD drives.

When a P-VOL and DP pool are located in a RAID group made up of SAS7.2K drives, the performance of a host I/O is reduced because of the lower performance of the SAS7.2K drive. Therefore, you should assign a primary volume to a RAID group consists of SAS drives or the SSD/FMD drives.

- Assign four or more disks to the data disks.

When there aren't enough data disks in the RAID group, the host performance and copying performance is reduced because read and write operations are restricted. When operating pairs with Snapshot, it is recommended that you use a volume consisting of four or more data disks.

RAID configuration for volumes assigned to Snapshot

Please observe the following regarding RAID levels when setting up Snapshot pair volumes, DP pools, and command devices.

- More than one pair may exist on the same RAID group on the disk array. However, when more than two pair are assigned to the same group, the impact on performance increases. Therefore, it is recommended that when creating pairs within the same RAID group, you should standardize the controllers that control volumes in the RAID group.
- Performance is best when P-VOL and DP pool are assigned to a RAID group consisting of SAS drives or SSD/FMD drives.
- Locate a P-VOL and associated data pool in different ECC groups within a RAID group, as shown in [Figure 9-6](#). When they are in the same ECC group, performance decreases and the chance of failure increase.

Pair resynchronization and releasing

In resynchronization processing and deletion processing, the processing to release the data stored in the DP pool is executed even after the pair status is changed. This processing repeats the short-time release processing and the suspension. In the time zone where the host I/O load is low and the processor use rate of the array is low, the processor spends much time for the release processing. Meanwhile, the processor use rate becomes temporarily high, but the effect on the host I/O is limited.

Therefore, perform the resynchronization and the deletion in the time zone where the host I/O load is low. If the resynchronization and the deletion have to be executed in the time zone where the host I/O load is high by necessity, make the number of pairs to be resynchronized and deleted at one time small and, at the same time, check the load to the array by the processor use rate and the I/O response time and perform the resynchronization and the deletion at intervals as much as possible.

Locating P-VOLS and DP pools

Do not locate the P-VOL and the DP pool within the same ECC group of the same RAID group because:

- A single drive failure causes a degenerated status in the P-VOL and DP pool.
- Performance decreases because processes, such as access to the P-VOL and data copying to the DP pool, are concentrated on a single drive.




Figure 9-6: Locating P-VOL and DP pool in separate ECC Groups

If multiple volumes are set within the same drive and each pair states differs, it is difficult to estimate the performance in order to design the system operational settings. For example, when the VOL0 and VOL2 are both P-VOLs and exist within the same group in the

same drive (V-VOLs are located in different drive group), and when the VOL0 is in Reverse synchronizing status and the VOL2 is in Split status.




Figure 9-7: Locating multiple volumes within the same drive column

If you have set a single volume per drive group, retain the status of pairs (such as Simplex and Split) when setting multiple Snapshot pairs. If each Snapshot pair status differs, it becomes difficult to estimate the performance when designing the system operational settings.


For optimal performance, a P-VOL should be located in a RAID group which contains SAS drives or SSD/FMD drives. When a P-VOL or DP pool is located in a RAID group made up of SAS7.2K drives, the host I/O performance is lessened due to the decreased performance of the SAS7.2K drive. You should assign a primary volume to a RAID group consisting of SAS drives or SSD/FMD drives.

It is recommended to set separate DP pools to the replication data DP pool and the management area DP pool

While using Snapshot, the replication data DP pool and the management area DP pool are frequently accessed. So if the replication data DP pool and the management area DP pool are set to the same DP pool, it negatively can affect the overall performance of Snapshot as the identical DP pool are quite frequently accessed. Therefore, it is highly recommended to set separate DP pools to the replication data DP pool and the management area DP pool. See [Figure 9-8 on page 9-19](#).

However, when using CCI for pair creation, you cannot specify separate DP pools for the replication data DP pool and the management area DP pool. Use HSNM2 for pair creation.

Not recommended configuration: Replication data and management area information are stored in the same DP pool.



Recommended configuration: Replication data and management area information are stored in the separate DP pool.




Figure 9-8: Set separate DP pools

Command devices

When two or more command devices are set within the one disk array, assign them to their respective RAID groups. If they are assigned to the same RAID group both command devices become unavailable due to a system malfunction, such as a drive failure.

Operating system host connections

The following sections provide recommendations and restrictions for Snapshot volumes.

Veritas Volume Manager (VxVM)

A host cannot recognize both a P-VOL and its V-VOL at the same time. Map the P-VOL and V-VOL to separate hosts.

AIX

- A host cannot recognize both a P-VOL and its V-VOL at the same time. Map the P-VOL and V-VOL to separate hosts.
- Multiple V-VOLs per P-VOL cannot be recognized from the same host. Limit host recognition to one V-VOL.

Linux and LVM configuration

A host cannot recognize both a P-VOL and its V-VOL at the same time. Map the P-VOL and V-VOL to separate hosts.

Tru64 UNIX and Snapshot configuration

When rebooting the host, the pair should be split or un-recognized by the host. Otherwise, a system reboot takes a longer amount of time.

Cluster and path switching software

Do not make the V-VOL an object of cluster or path switching software.

Windows Server and Snapshot configuration

- Multiple V-VOLs per P-VOL cannot be recognized from the same host. Limit host recognition to one V-VOL.
- In order to make a consistent backup using a storage-based replication such as Snapshot, you must have a way to flush the data residing on the server memory to the disk array so that the source volume of the replication has the complete data. You can flush the data on the server memory using the umount command of CCI to unmount the volume. When using the umount command of CCI for unmount, use the mount command of CCI for mount. (For more detail about mount/umount command, see the *Hitachi Unified Storage Command Control Interface Installation and Configuration Guide*. If you are using Windows Server 2003, mountvol /P to flush data on the server memory when un-mounting the volume is supported. Understand the specification of the command and run a test before using it for your operation. In Windows Server 2008, use the umount command of CCI to flush the data on the memory of

the server at the time of the unmount. Do not use the mountvol command of Windows standard. Moreover, do not use the directory mount at the time of the mount and only use the mount by the drive letters. Refer to the *Hitachi Unified Storage Command Control Interface Installation and Configuration Guide* for details of the restrictions of Windows Server 2008 when using the mount/umount command.

- If you recognize the P-VOL and S-VOL on Windows Server 2008 at the same time, it may cause an error because the P-VOL and S-VOL have the same disk signature. When the P-VOL and S-VOL have the same data, split the pair and then rewrite the disk signature so that they can retain different disk signatures. You can use the uniqueid command to rewrite a disk signature. See the *Hitachi Unified Storage Command Control Interface Installation and Configuration Guide* for details.
- When a path becomes detached, which can be caused by a controller detachment or interface failure and remains detached for longer than one minute, the command device may not be recognized when path recovery is made. Execute the "re-scan the disks" function of Windows to make the recovery. Restart CCI if Windows cannot access the command device even if CCI is able to recognize it.
- Windows Server may write for the un-mounted volume. If a pair is resynchronized while retaining the data to the V-VOL on the memory of the server, the compatible backup cannot be collected. Therefore, execute the sync command of CCI immediately before re-synchronizing the pair for the un-mounted V-VOL.

Microsoft Cluster Server (MSCS)

- A host cannot recognize both a P-VOL and its V-VOL at the same time. Map the P-VOL and V-VOL to separate hosts.
- When setting the V-VOL to be recognized by the host, use the CCI mount command rather than Disk Administrator.
- Do not place the MSCS Quorum Disk in CCI.
- Assign the exclusive command device to each host.

Windows Server and Dynamic Disk

In an environment of Windows Server, you cannot use Snapshot pair volumes as dynamic disk. The reason is because if you restart Windows or use the Rescan Disks command after creating or re-synchronizing a Snapshot pair, there are cases where the V-VOL is displayed as Foreign in Disk Management and become inaccessible

UNMAP Short Length Mode

Enable UNMAP Short Length Mode when connecting to Windows 2012. If you do not enable it, UNMAP commands may not be completed due to a time-out.

Windows 2000

- A P-VOL and V-VOL *cannot* be made into a dynamic disk on Windows Server 2000.
- Multiple V-VOLs per P-VOL cannot be recognized from the same host. Limit host recognition to one V-VOL.
- When mounting a volume, use the CCI mount command, even if using the Navigator 2 GUI or CLI to operate the pairs. Do not use the Windows mountvol command because the data residing in server memory is not flushed. The CCI mount command does flush data in server memory, which is necessary for Snapshot operations. For more information, see the *Hitachi Unified Storage Command Control Interface Installation and Configuration Guide*.

VMware and Snapshot configuration

When creating a backup of the virtual disk in the vmfs format using Snapshot, shut down the virtual machine that accesses the virtual disk, and then split the pair.

If one volume is shared by multiple virtual machines, shutdown all the virtual machines that share the volume when creating a backup. Therefore, it is not recommended to share one volume by multiple virtual machines in the configuration that creates a backup using Snapshot.

The VMware ESX has a function to clone the virtual machine. Although the ESX clone function and Snapshot can be linked, cautions are required for the performance at the time of execution. For example, when the volume which becomes the ESX clone destination is a Snapshot P-VOL pair whose pair status is Split, since the old data is written to the DP pool for writing to the P-VOL, the time required to clone may become longer and the clone may be terminated abnormally in some cases.

To avoid this, make the Snapshot pair status Simplex, and create and split the pair after executing the ESX clone. Do the same for executing the functions such as migrating the virtual machine, deploying from the template, inflating the virtual disk and Space Reclamation.

- V-VOL does not support the VAAI Unmap command.

With the command "esxcli storage core device vaai status get", the Delete Status shows "unsupported" for a datastore created on a V-VOL. Also, with the command "esxcli storage core device list", the "Thin Provisioning Status" shows "unknown".




Figure 9-9: VMware ESX

- UNMAP Short Length Mode

It is recommended that you enable UNMAP Short Length Mode when connecting to VMware. If you do not enable it, UNMAP commands may not be completed due to a time-out.

Array functions

Identifying P-VOL and V-VOL volumes on Windows

In the Navigator 2 GUI, the P-VOL and V-VOL are identified by their volume number. In Windows, volumes are identified by their HLUN. The following instructions provide procedures for the iSCSI and fibre channel interfaces. To understand the mapping of a volume on Windows, proceed as follows:

1. Identify the HLUN of your Windows disk.
 - a. From the Windows Server Control Panel, select Computer Management/Disk Administrator.
 - b. Right-click the disk whose HLUN you want to know, then select **Properties**. The number displayed to the right of "LUN" in the dialog window is the HLUN.
2. Identify HLUN-to-VOL Mapping for the iSCSI interface as follows. (If using fibre channel, skip to Step 3.)
 - a. In the Navigator 2 GUI, select the desired disk array.
 - b. Select the disk array and click the **iSCSI Targets** icon in the **Groups** tree.



WARNING! Your host group changes will be applied to multiple ports. This change will delete existing host group mappings and corresponding Host Group IDs, corrupting or removing data associated with the host groups. To keep specified host groups you do not want to remove, please cancel this operation and make changes to only one host group at a time.

- c. Click the **iSCSI Target** that the volume is mapped to.
 - d. Click **Edit Target**.
 - e. The list of volumes that are mapped the iSCSI Target is displayed and you can confirm the VOL that is mapped to the H-LUN.
3. For the Fibre Channel interface:
 - a. Start Navigator 2.
 - b. Select the disk array and click the **Host Groups** icon in the **Groups** tree.
 - c. Click the **Host Group** that the volume is mapped to.
 - d. Click **Edit Host Group**.
 - e. The list of volume that is mapped to the Host Group is displayed and you can confirm the VOL that is mapped to the H-LUN.

Volume mapping

When you use CCI, you cannot pair a P-VOL and V-VOL when their mapping information has not been defined in the configuration definition file. To prevent a host from recognizing a P-VOL or V-VOL, use Volume Manager to either map them to a port that is not connected to the host or map them to a host group that does not have a registered host. If you use Storage Navigator instead of Volume Manager, you need only perform this task with either the P-VOL or V-VOL.

Concurrent use of Cache Partition Manager

When Snapshot is used together with Cache Partition Manager, please refer to the *Hitachi Unified Storage Operations Guide*.

Concurrent use of Dynamic Provisioning

The DP-VOLs can be set for a Snapshot P-VOL. SnapShot and Dynamic Provisioning can be used together. Refer to the *Hitachi Unified Storage Dynamic Provisioning Configuration Guide* for the detailed information regarding Dynamic Provisioning. Hereinafter, the volume created in the RAID group is called a normal volume, and the volume created in the DP pool is called a DP-VOL.

Observe the following items:

- Pair status at the time of DP pool capacity depletion
When the DP pool is depleted after operating the Snapshot pair that uses the DP-VOL, the pair status of the pair concerned may be a Failure. [Table 9-6](#) shows the pair statuses before and after the DP pool capacity depletion. When the pair status becomes a Failure caused by the DP pool capacity depletion, add the DP pool capacity whose capacity is depleted, and execute the pair operation again.

Table 9-6: Pair statuses before and after the DP Pool capacity depletion

Pair statuses before the DP pool capacity depletion	Pair Statuses after the DP pool capacity depletion belonging to P-VOL
Simplex	Simplex
Reverse Synchronizing	Reverse Synchronizing Failure R*
Paired	Paired
Split	Split Failure*
Failure	Failure

Table 9-6: Pair statuses before and after the DP Pool capacity depletion

Pair statuses before the DP pool capacity depletion	Pair Statuses after the DP pool capacity depletion belonging to P-VOL
* When write is performed to the P-VOL or V-VOL to which the capacity depletion DP pool belongs, the copy cannot be continued and the pair status becomes a Failure	

- DP pool status and availability of pair operation
When using the DP-VOL for a P-VOL of the Snapshot pair, the pair operation may not be executed depending on the status of the DP pool to which the DP-VOL belongs. [Table 9-7](#) shows the DP pool status and availability of the Snapshot pair operation. When the pair operation fails due to the DP pool status, correct the DP pool status and execute the pair operation again.

Table 9-7: DP pool statuses and availability of Snapshot pair operation

Pair operation	DP pool statuses, DP pool capacity statuses, and DP pool optimization statuses					
	Normal	Capacity in growth	Capacity depletion	Regressed	Blocked	DP in optimization
Create pair	YES	YES	NO	YES	NO	YES
Create pair (split option)	YES	YES	NO	YES	NO	YES
Split pair	YES	YES	NO	YES	NO	YES
Resync pair	YES	YES	NO	YES	NO	YES
Restore pair	YES	YES	NO	YES	NO	YES
Delete pair	YES	YES	YES	YES	YES	YES

- Ensuring usable capacity
When the DP pool was created or the capacity was added, the formatting operates for the DP pool. If pair creation, pair resynchronization, or restoration is performed during the formatting, depletion of the usable capacity may occur. Since the formatting progress is displayed when checking the DP pool status, check if the sufficient usable capacity is secured according to the formatting progress, and then start the operation.
- Operation of the DP-VOL during Snapshot use
When using the DP-VOL for a P-VOL of Snapshot, any of the operations among the capacity growing, capacity shrinking, volume deletion, and Full Capacity Mode changing of the DP-VOL

in use cannot be executed. To execute the operation, delete Snapshot pair of which the DP-VOL to be operated is in use, and then execute it again.

- Operation of the DP pool during Snapshot use

When using the DP-VOL for a P-VOL of Snapshot, the DP pool to which the DP-VOL in use belongs cannot be deleted. To execute the operation, delete the Snapshot pair of which the DP-VOL is in use belonging to the DP pool to be operated, and then execute it again. The attribute edit and capacity addition of the DP pool can be executed usually regardless of the Snapshot pair.

- Caution for DP pool formatting, pair resynchronization, and pair deletion

Continuously performing DP pool formatting, pair resynchronization, or pair deletion to a pair with a lot of replication data or management information can lead to temporary depletion of the DP pool, where used capacity (%) + capacity in formatting (%) = about 100%, and it makes the pair change to Failure. Perform pair resynchronization and pair deletion when sufficient available capacity has been ensured.

- Cascade connection

A cascade can be performed on the same conditions as the normal volume (refer to [Cascading Snapshot with True Copy Extended on page 27-27](#)).

Concurrent use of Dynamic Tiering

These are the considerations for using the DP pool whose tier mode is enabled by using Dynamic Tiering. For the detailed information related to Dynamic Tiering, refer to *Hitachi Unified Storage 100 Dynamic Tiering User's Guide*. Other considerations are common with Dynamic Provisioning.

In the replication data DP pool and the management area DP pool whose tier mode is enabled, the replication data and the management information are not placed in the Tier configured by SSD/FMD. Because of this, use caution with the following:

- The DP pool whose tier mode is enabled and configured only by SSD/FMD cannot be specified as the replication data DP and the management area DP pool.
- The total of the free capacity of the Tier configured by the drive other than SSD/FMD in the DP pool is the total free capacity for the replication data or the management information.
- When the free space of the replication data DP pool and the management area DP pool whose tier mode is enabled decreases, recover the free space of Tier other than SSD/FMD.

When the replication data and the management information are stored in the DP pool whose tier mode is enabled, they are first assigned to 2nd Tier.

The area where the replication data and the management information are assigned is out of the relocation target.

User data area of cache memory

Because Snapshot uses DP pools to work, Dynamic Provisioning/Dynamic Tiering are necessary to operate at the same time. Employing Dynamic Provisioning/Dynamic Tiering acquires some portion of the installed cache memory, which reduces the user area of the cache memory. [Table 9-8 on page 9-29](#) shows the cache memory secured capacity and the user data area when using Dynamic Provisioning/Dynamic Tiering.

The performance effect by reducing the user data area is shown when a large amount of sequential writes were executed at the same time, but it is deteriorated by a few percent when writing 100 volumes at the same time.

Table 9-8: User data area capacity of cache memory

Array type	Cache memory per CTL	DP capacity mode	Management capacity for Dynamic Provisioning	Management capacity for Dynamic Tiering	User data capacity		
					When Dynamic Provisioning is enabled	When Dynamic Provisioning and Dynamic Tiering are Enabled	When Dynamic Provisioning and Dynamic Tiering are disabled
HUS 110	4 GB/CTL	Not supported	420 MB	50 MB	1,000 MB	960 MB	1,420 MB
HUS 130	8 GB/CTL	Regular Capacity	640 MB	200 MB	4,020 MB	3,820 MB	4,660 MB
		Maximum Capacity	1,640 MB	200 MB	3,000 MB	2,800 MB	4,660 MB
	16 GB/CTL	Regular Capacity	640 MB	200 MB	10,640 MB	10,440 MB	11,280 MB
		Maximum Capacity	1,640 MB	200 MB	9,620 MB	9,420 MB	11,280 MB
HUS 150	8 GB/CTL	Regular Capacity	1,640 MB	200 MB	2,900 MB	2,700 MB	4,540 MB
	16 GB/CTL	Regular Capacity	1,640 MB	200 MB	9,520 MB	9,320 MB	11,160 MB
		Maximum Capacity	3,300 MB	200 MB	7,860 MB	7,660 MB	11,160 MB

Limitations of dirty data flush number

This setting determines the number of times processing is executed for flushing the dirty data in the cache to the drive at the same time. This setting is effective when Snapshot is enabled. When all the volumes in the disk array are created in the RAID group of RAID 1 or RAID 1+0 (SAS drives configured and in the DP pool), if this setting is enabled, the dirty data flush number is limited even though Snapshot is enabled. When the dirty data flush number is limited, the response time in I/O, which has a low load and high Read rate, shortens. Note that when TrueCopy or TCE are unlocked at the same time, this setting is not effective.

See [Setting the system tuning parameter \(optional\) on page 9-36](#) or for CLI, see [Setting the system tuning parameter \(optional\) on page B-14](#) for the setting method about the Dirty Data Flush Number Limit.

Load balancing function

The Load balancing function applies to a SnapShot pair.

Enabling Change Response for Replication Mode

If background copy to the pool is timed-out for some reason while write commands are being executed on the P-VOL in Split state, or if restoration from the V-VOL to the P-VOL is timed-out while read commands are being executed on the P-VOL in Reverse Synchronizing state, the array returns Medium Error (03) to the host. Some hosts receiving Medium Error (03) may determine the P-VOL is inaccessible, and stop accessing it. In such cases, enabling the Change Response for Replication Mode makes the array return Aborted Command (0B) to the host. When the host receives Aborted Command (0B), it will retry the command to the P-VOL and the operation will continue.

Configuring Snapshot

This topic describes the steps for configuring Snapshot.

Configuration workflow

Setup for Snapshot consists of assigning volumes for the following:

- DP pool
- V-VOL(s)
- Command device (if using CCI)

The P-VOL should be set up in the disk array prior to Snapshot configuration.

Refer to the following for requirements and recommendations:

- [Requirements and recommendations for Snapshot Volumes on page 9-14](#)
- [Operating system host connections on page 9-20](#)
- [System requirements on page 8-2](#)
- [Snapshot specifications on page B-2](#)

Setting up the DP pool

For directions on how to set up a DP pool, refer to the *Hitachi Unified Storage Dynamic Provisioning Configuration Guide*.

To set the DP pool capacity, see the [Calculating DP pool size on page 9-12](#).

Setting the replication threshold (optional)

To set the Depletion Alert and/or the Replication Data Released of the replication threshold:

1. Select the **Volumes** icon in the **Group** tree view.
2. Select the **DP Pools** tab.
3. Select the DP pool number for the replication threshold that you want to set.

HUS110_91200017 > Groups > Volumes

Volumes	RAID Groups	DP Pools

Rows/Page: 25 Page: 1 of 1										
☐	DP Pool	Tier Mode	RAID Level	Capacity			Current Utilization Percent	Current Over Provisioning Percent	Replication	
				Total	Consumed	Needing Preparation			Available Capacity	Current Utilization Percent
☐		N/A	RAID5(4D+1P)	1.0TB	52.0GB	0.0MB	4%	5%	1.0TB	4%

4. Click **Edit Pool Attribute**.

Edit Pool Attribute

Refresh Information

Help

DP Pool-000

HUS110_91200017 > Groups > Volumes > DP Pool-000

Properties

Notifications

Volumes

Assigned Drives

DP RAID Groups

DP Pool		 000	
Tier Mode		N/A	
RAID Level		RAID5	
Combination		4D+1P	
Page Size		32MB	
Stripe Size		256KB	
Drive Type/RPM		SAS/10K	
Status		Normal	
Reconstruction Progress		N/A	
Capacity	Total		1.0TB
	Replication Available		1.0TB
	Consumed by	Total	52.0GB
		User Data	52.0GB
		Replication Data	0.0MB
		Management Area	0.0MB
Needing Preparation		0.0MB	

5. Enter the **Replication Depletion Alert Threshold** and/or the **Replication Data Released Threshold** in the **Replication** field.

Edit DP Pool Attribute - 000 Help

DP Pool Attribute Properties
Edit the properties of the DP pool.

Basic

DP Pool: 000
Tier Mode: N/A
RAID Level: RAID5(4D+1P)
Stripe Size: 256KB
Drive Type/RDM: SAS/10K
Total Capacity: 1.0TB
Consumed Capacity: 52.0GB

→ DP Pool Consumed Capacity:
Early Alert Threshold: 40 %
Depletion Alert Threshold: 50 %
From 1 to 99
Send notification by E-mail alert and SNMP if the current value exceeds threshold: ☒ Yes

* Over Provisioning:
Warning Threshold: 100 %
Limit Threshold: 130 %
From 50 to 1000
Send notification by E-mail alert and SNMP if the current value exceeds threshold: ☐ Yes
Limit Enforcement: ☒ Enable

→ Replication:
Replication Depletion Alert Threshold: 40 %
Replication Data Released Threshold: 95 %
From 1 to 99

* Required field

OK Cancel

6. Click **OK**.



NOTE: For instructions using CLI see [Setting the replication threshold \(optional\) on page B-11](#)

7. A message appears. Click **Close**.

Edit DP Pool Attribute - 000



DP pool attribute(s) edited successfully.

Close

Setting up the Virtual Volume (V-VOL) (manual method)


Since the Snapshot volume (V-VOL) is automatically created at the time of the pair creation, it is not always necessary to set the V-VOL. It is also possible to create the V-VOL before the pair creation (in advance) and perform the pair creation with the created V-VOL.

Prerequisites

- The V-VOL volume must be the same size as the P-VOL.

To assign volumes as V-VOLs

1. In Navigator 2 GUI, select the desired disk array, then click the **Show & Configure disk array** button.
2. Select the **Snapshot Volumes** icon in the **Setup** tree view of the **Replication** tree view.
3. Click **Create Snapshot VOL**. The Create Snapshot Volume screen appears.



4. Enter the **VOL** to be used for the V-VOL. You can use any unused **VOL** that matches the P-VOL in size. The lowest available volume number is the default.
5. Enter the V-VOL size in the **Capacity** field. The Capacity range is 1 MB - 128 TB.
6. Click **OK**.
7. A message appears, Snapshot volume created successfully. Click **Close**.

Deleting V-VOLs

Prerequisites

- In order to delete the V-VOL, the pair state must be Simplex.
1. Select **Snapshot Volumes** in the **Setup** tree view. The Snapshot **Volumes** list displays.

2. Select a V-VOL you want to delete in the **Snapshot Volumes** list.
3. Click **Delete Snapshot VOL**.
4. A message appears, Snapshot volumes deleted successfully. Click **Close**.

Setting up the command device (optional)

CCI can be used in place of the Navigator 2 GUI and/or CLI to operate Snapshot. CCI interfaces with Hitachi Unified Storage through the command device, which is a dedicated logical volume. A command device must be designated in order to issue Snapshot commands.

Prerequisites

- Setup of the command device is required only if using CCI.
- Volumes assigned to a command device must be recognized by the host.
- The command device must be defined in the HORCM_CMD section of the configuration definition file for the CCI instance on the attached host.
- The command device should be a minimum of 33 MB.
- 128 command devices can be designated per disk array.

References

- [RAID configuration for volumes assigned to Snapshot on page 9-16](#).
- [Snapshot specifications on page B-2](#).
- *Hitachi Unified Storage Command Control Interface Installation and Configuration Guide* for more information about command devices.

To set up a command device

The following procedure employs the Navigator 2 GUI. To use CCI, see [Setting the command device on page B-22](#)

1. In the Settings tree view, select **Command Devices**. The Command Devices screen displays.
2. Select Add **Command Device**. The Add Command Device screen displays.
3. In the Assignable Volumes box, click the check box for the **VOL** you want to add as a command device. A command device must be at least 33 MB.
4. Click the **Add** button. The screen refreshes with the selected volume listed in the Command Device column.
5. Click **OK**.


Setting the system tuning parameter (optional)

This is a setting to determine whether to limit the number of times processing is executed for flushing the dirty data in the cache to the drive at the same time.

To set the Dirty Data Flush Number Limit of a system tuning parameters:


1. Select the **System Tuning** icon in the **Tuning Parameter** of the **Performance** tree view.

The **System Tuning** screen appears.



2. Click **Edit System Tuning Parameters**.

The **Edit System Tuning Parameters** screen appears.



3. Select the **Enable** option of the **Dirty Data Flush Number Limit**.
4. Click **OK**.

5. A message displays. Click **Close**.

Edit System Tuning Parameters



System tuning parameters edited successfully.



Using Snapshot

This chapter describes Snapshot copy operations. The Snapshot workflow includes the following:

- ❑ [Snapshot workflow](#)
- ❑ [Confirming pair status](#)
- ❑ [Create a Snapshot pair to back up your volume](#)
- ❑ [Splitting pairs](#)
- ❑ [Updating the V-VOL](#)
- ❑ [Restoring the P-VOL from the V-VOL](#)
- ❑ [Deleting pairs and V-VOLs](#)
- ❑ [Editing a pair](#)
- ❑ [Use the V-VOL for tape backup, testing, and reports](#)

Snapshot workflow


A typical Snapshot workflow consists of the following:

- Check pair status. Each operation requires a pair to have a specific status.
- Create the pair, in which the P-VOL pairs with the V-VOL but the V-VOL still does not retain a snapshot of the P-VOL.
- Split the pair, which created a snapshot of the P-VOL in the V-VOL and allows use of the data in the S-VOL by secondary applications.
- Update the pair, which take a new snapshot in the V-VOL
- Restore the P-VOL from the S-VOL.
- Delete a pair.
- Edit pair information.


For an illustration of basic Snapshot operations, see [Figure 7-1 on page 7-3](#).

Confirming pair status


1. Select the **Local Replication** icon in the Replication tree view.



2. The Pairs list appears. The pair with the secondary volume without the volume number is not displayed. To display the pair with the secondary volume without the volume number, open the Primary Volumes tab and select the primary volume of the target pair.



3. The list of the primary volumes is displayed in the Primary Volumes tab.



4. When the primary volume is selected, all the pairs of the selected primary volume including the secondary volume without the volume number are displayed.

Primary Volume - 0001
MUS110_91200026 > Replication > Local Replication > Primary Volume - 0001

Pairs								
Rows/Page: 25 Page 1 of 1								
☐	Pair Name	Secondary Volume	Status	DP Pool Replication Data	Management Area	Copy Type	Group Number : Group Name	Point-in-Time
☒	SS_LU0001_LU0901	N/A	Paired(100%)	000	000	Snapshot	---:{Ungrouped}	N/A
☐	SS_LU0001_LU0901_11111410264792	N/A	Paired(100%)	000	000	Snapshot	---:{Ungrouped}	N/A

Create Pair Edit Pair Split Pair Resync Pair Restore Pair Delete Pair Filter Filter Off

- **Pair Name:** The pair name displays.
- **Primary Volume:** The primary volume number displays
- **Secondary Volume:** The secondary volume number displays. The secondary volume without the volume number is displayed as N/A.
- **Status:** The pair status and the identical rate (%) display. See Note.
 - **Simplex:** A pair is not created.
 - **Reverse Synchronizing:** Update copy (reverse) is in progress.
 - **Paired:** pair is created or resynchronization is complete.
 - **Split:** A pair split.
 - **Failure:** A failure occurs.
 - **Failure(R):** A failure occurs in restoration.
 - **---**: Other than above.
- **DP Pool:**
 - **Replication Data:** A DP pool number displays.
 - **Management Area:** A DP pool number displays
- **CopyType:** Snapshot or ShadowImage displays.
- **Group Number:** A group number, group name, or ---:{Ungrouped} displays.
- **GroupName**
- **Point-in-Time:** A point-in-time attribute displays.
- **Backup Time:** Acquired backup time or N/A displays.
- **Split Description:** A character strings appears when you specify **Attach description to identify the pair upon split**. If

this is not specified **Attach description to identify the pair upon split, N/A** displays.

- **MU Number:** MU numbers used in CCI displays.



NOTE: The identical rate displayed with the pair status shows what percentage of data the P-VOL and V-VOL currently share. When write operations from the host are executed on the P-VOL or V-VOL, the differential data is copied to the DP pool in order to maintain the snapshot of the P-VOL, which leads to a decline in the identical ratio for the pair. Note that when a P-VOL is paired with multiple V-VOLs, the only pair which has been split most recently among all the pairs with the P-VOL shows the accurate identical ratio. The identical ratios for the other pairs show an estimated identical ratio that can be used to know roughly how much data is shared between P-VOL and V-VOL.

When an additional pair is created and split to an existing P-VOL, there are cases where the identical rate of the pair, which had been split most recently among the pairs with the P-VOL, can be reduced. The identical rates of pairs can be referred to from the pair information on HSNM2.

There are cases where the identical rates of pairs with the P-VOL can fluctuate for some time when restore starts.

Create a Snapshot pair to back up your volume

For instructions using CLI, see [Creating Snapshot pairs using CLI on page B-14.](#))

When you create a pair, very little time elapses until the pair is established. During this time, the P-VOL remains accessible to the host, but the V-VOL is unavailable until the snapshot is complete and the pair is split.

The create pair procedure allows you to set copy pace, assign the pair to a group (and create a group), and automatically split the pair after it is created.

Prerequisites

- Create a DP pool to be used by Snapshot. It is recommended to create two DP pools.
- Make sure the primary volume is set up on the disk array. See [Snapshot specifications on page B-2](#) for primary volume specifications.



NOTE: If you prefer it, you can also set up the V-VOL when using the Create Pair method. See [Setting up the Virtual Volume \(V-VOL\) \(manual method\) \(optional\) on page 9-34](#)

To create a pair using the create pair procedure

1. In Navigator 2 GUI, select the desired disk array, then click the **Show & Configure disk array** button.
2. From the Replication tree, select the **Local Replication** icon. The **Create Pair** screen appears.

Create Pair Help

Local Pair Property

Enter property information as desired.

* Copy Type : ☐ ShadowImage ☒ Snapshot

Basic **Advanced**

Select the primary volume and the secondary volume to pair and give the pair a unique name. Select a data DP pool that the pair will use.

Pair Name :

31 characters or less (alphanumeric characters, special symbols "%", "*", "+", "-", "/", "=", "@", "_", ".", "!", or "~")
Pair Name is specified automatically in the case of blank.

* Select a Primary Volume :

Primary Volume				
Rows/Page: 25 1 2 3 Page 1 of 1 4 5 6				
VOL	Capacity	RAID Group	DP Pool	RAID Level
☒ 0000	10.0GB	N/A	000	RAID6(8D+2P)
☐ 0001	10.0GB	N/A	000	RAID6(8D+2P)

Filter Filter Off

* Select a Secondary Volume :

☐ Unassign

☒ Assign :

From 0 to max (array model dependent)

DP Pool :

☒ Automatic

☐ Manual : Replication Data (Capacity, %Used) :

Management Area (Capacity, %Used) :

* Required field

OK Cancel

3. Select the **Create Pair** button. The Create Pair screen displays.
4. In the Copy Type area, click the **Snapshot** option. There may be a brief delay while the screen refreshes.
5. In the **Pair Name** box, enter a name for the pair.



NOTE: When not specifying pair name, the following are assigned automatically:

When the S-VOL is assigned: SS_LUXXXX_LUYYYY

When the S-VOL is not assigned:

SS_LUXXXX_LUNONE_nnnnnnnnnnnnnnnnn

XXXX: LU number of P-VOL (four digits with 0)

YYYY: LU number of V-VOL (four digits with 0)

nnnnnnnnnnnnnnnnnn: 14-digit optional numbers

("Year/Month/Day/hour/minute/second/millisecond" at time of pair creation)

6. Select a primary volume. To display all volumes, use

 or  button.

7. Select a secondary volume.
 - Select it using the option **Unassign** or **Assign**. When **Assign** is selected, enter the volume number of the secondary volume in the text box.

- When **Unassign** is selected, the pair creation is performed with the secondary volume without the automatically created volume number. When **Assign** is selected and the volume number of the already existed secondary volume is entered, the pair creation is performed with the secondary volume with the entered volume number. When **Assign** is selected and the unused volume number is entered, the secondary volume with the entered volume number is automatically created and performs the pair creation.



NOTE: The VOL may be different from the volume's H-LUN. Refer to [Cluster and path switching software on page 9-20](#) to map VOL to H-LUN.

8. Select the **DP Pool**, using the **Automatic** or **Manual** option.
 - If **Manual** is selected, select the replication data DP pool and the management area DP pool from the drop-down list.
 - If **Automatic** is selected, the DP pool to be used is automatically selected. When the primary volume is a normal volume, the DP pool with the smallest number in the existing DP pools is selected as the replication data DP pool and the management area DP pool. When the primary volume is the DP-VOL, the DP pool to which the primary volume belongs is selected as the replication data DP pool and the management area DP pool.
9. Click the **Advanced** tab.

Create Pair [Help]

Local Pair Property

Select a copy type of local replication and enter property information as desired.

* Copy Type : ☐ ShadowImage ☒ Snapshot

Basic **Advanced**

Enter any additional properties for the pair.

Copy Pace : Medium

* After pair creation, add the pair to a group.- Group Assignment:

☒ (Ungrouped)

☐ New or existing Group Number : [] From 0 to 1023

☐ Existing Group Name : [] 31 characters or less (alphanumeric characters, special symbols: %, *, ~, ^, &, ", ', /, =, @, ", ", ", ", ", ", or " ")

Automatically split the pair immediately after they are created: ☐ Yes

* NU Number :

☒ Automatic

☐ Manual : [] From 0 to 1023

* Required field

[OK] [Cancel]

10. From the **Copy Pace** drop-down list, select the speed that copies will be made. Copy pace is the speed at which a pair is created or resynchronized. Select one of the following:
- **Slow** — The process takes longer when host I/O activity is heavy. The time of copy or resync completion cannot be guaranteed.
 - **Medium** — (Recommended) The process is performed continuously, but the time of completion cannot be guaranteed. The pace differs depending on host I/O activity.
 - **Fast** — The copy/resync process is performed continuously and takes priority. Host I/O performance is restricted. The time of copy/resync completion is guaranteed.
11. In the **Group Assignment** area, you have the option of assigning the new pair to a consistency group. See [Consistency Groups \(CTG\) on page 7-11](#) for a description. Do one of the following:
- If you do not want to assign the pair to a consistency group, leave the **Ungrouped** button selected.
 - To create a group and assign the new pair to it, click the **New or existing Group Number** button and enter a new number for the group in the box.
 - To assign the pair to an existing group, enter the consistency group number in the **Group Number** box, or enter the group name in the **Existing Group Name** box.



NOTE: Add a Group Name for a consistency group as follows:

- After completing the create pair procedure, on the **Pairs** screen, check the box for the pair belonging to the group.
- Click the **Edit Pair** button.
- On the Edit Pair screen, enter the **Group Name** then click **OK**.


-
12. In the **Split the pair...** field, do one of the following:
- Click the **Yes** box to split the pair immediately. A snapshot will be taken and the V-VOL will become a mirror image of the P-VOL at the time of the split.
 - Leave the **Yes** box unchecked to create the pair. The V-VOL will stay up-to-date with the P-VOL until the pair is split.
13. When specifying a specific MU number, select **Manual** and specify the MU number in the range **0 - 1032**.
14. Click **OK**, then click **Close** on the confirmation screen that appears. The pair has been created.

Splitting pairs

To split the Snapshot pairs:

1. Select the **Local Replication** icon in the **Replication** tree view.
2. Select the pair you want to split the pair in the **Pairs** list.
3. Click **Split Pair**.

The **Split Pair** screen appears.



4. Enter a character strings to the **Attach description to identify the pair upon split** if necessary.
5. Click **OK**.
6. A confirmation message appears. Click **Close**.

Updating the V-VOL

Updating the V-VOL means to take a new snapshot. Two steps are involved when you update the V-VOL: resynchronizing the pair and then splitting the pair.

- Resynchronizing is necessary because after a pair split, no new updates to the P-VOL are copied to the V-VOL. When a pair is resynchronized, the V-VOL becomes a virtual mirror image again of the P-VOL.
- Splitting the pair completes of the snapshot. The V-VOL and the P-VOL are the same at the time the split occurs. After the split the V-VOL does not change. The V-VOL can then be used for tape-backup and for operations by a secondary host.


To update the V-VOL

(For instructions using CLI, see [Splitting Snapshot Pairs on page B-15](#).)

1. In Navigator 2 GUI, select the desired disk array, then click the **Show & Configure disk array** button.
2. From the Replication tree, select the **Local Replication** icon. The **Pairs** screen displays.
3. Select the pair that you want to update and click the **Resync Pair** button at the bottom of the screen. The operation may take several minutes, depending on the amount of data.
4. Check the Yes, I have read the above warning and want to re-synchronize selected pairs. check box, and click Confirm.



5. When the Resync is completed, click the **Split Pair** button.



This operation is completed quickly. When finished, the V-VOL is updated.



NOTE: Differential data is deleted from the DP pool when a V-VOL is updated. Time required for deletion of DP pool data is proportional to P-VOL capacity and P-VOL-to-V-VOL ratio. For a 100 GB P-VOL with a 1:1 ratio, it takes about five minutes. For a ratio of 1 P-VOL to 32 V-VOLs, deletion time is about 15 minutes.

Making the host recognize secondary volumes with no volume number

To make the host recognize secondary volumes which do not have a volume number:

1. When you create the Snapshot pairs, do not specify the Secondary Volume.
2. Confirm the pair status, and search the secondary volumes you want to make the host recognize from the Pair Name, the Backup Time, or the Split Description.
3. Assign the volume number to the secondary volume. See [Editing a pair on page 10-15](#).
4. Assign the H-LUN to the secondary volume.

Restoring the P-VOL from the V-VOL

Snapshot allows you to restore your P-VOL to a previous point in time from any Snapshot image (V-VOL). The amount of time it takes to restore your data depends on the size of the P-VOL, the amount of data that has changed, and your Hitachi Unified Storage model.

When you restore the P-VOL:

- There is a short period when data in the V-VOL is validated to insure that the restoration will complete successfully.
 - During the validation stage, the host cannot access the P-VOL.
 - Even when no differential data exists, restoration is not completed immediately. It takes from 6 to 15-minutes for the disk array to search for differentials. The time depends on the copy pace you have defined.
- Once validation is complete:
 - Copying from V-VOL to P-VOL is performed in the background.
 - Pair status is Reverse Synchronizing.
 - The P-VOL is available for read/write from the host.

The restore command can be issued to 128 P-VOLs at the same time, but actual copying is performed on a maximum of four per controller for HUS 110, and eight per controller for HUS 130 or HUS 150. When background copying can be executed, the copies are completed in the order the command was issued.

To restore the P-VOL from the V-VOL

(For instructions using CLI, see [Restoring V-VOL to P-VOL using CLI on page B-16.](#))

1. Shut down the host application.
2. Un-mount the P-VOL from the production server.
3. In the Storage Navigator 2 GUI, select the **Local Replication** icon in the Replication tree view.
4. In the GUI, select the pair to be restored in the **Pairs** list.
Advanced users using the Navigator 2 CLI, please refer to [Restoring V-VOL to P-VOL using CLI on page B-16.](#)
5. Click **Restore Pair**. View subsequent screen instructions by clicking the Help button.

Deleting pairs and V-VOLs


You can delete a pair or a V-VOL when you no longer need them.

- **Pair:** When a pair is deleted, the primary and virtual volumes return to their **SIMPLEX** state. Both are available for use in another pair.
- **V-VOL:** The pair must be deleted before a V-VOL is deleted. A V-VOL with no volume number is automatically deleted when the pair is deleted.


To delete a pair

(For instructions using the Storage Navigator 2 CLI, see [Deleting Snapshot pairs on page B-17](#).)

1. Select the **Local Replication** icon in the Replication tree view.
2. In the GUI, select the pair you want to delete in the **Pairs** list.
3. Click **Delete Pair**. A confirmation message appears.



4. Check the **Yes, I have read the above warning and agree to delete selected pairs.** check box, and click **Confirm**.
5. A confirmation message appears. Click **Close**.



To delete a V-VOL

1. Make sure that the pair is deleted first. The pair status must be **SIMPLEX** to delete the V-VOL.
2. Select the **Snapshot Volumes** icon in the **Setup** tree view.
3. In the Volumes for Snapshot list, select the V-VOL that you want to delete.
4. Click **Delete Volume for Snapshot**. A message appears.
5. Click **Close**. The **V-VOL is deleted**.

Editing a pair

You can edit certain information concerning a pair.

- For pairs, you can change the name, assignment/deprivation of volume number to secondary volume, group name, and copy pace.

To edit a pair

(For instructions using Navigator 2 CLI, see [Changing pair information on page B-18.](#))

1. In the Navigator 2 GUI, select the **Local Replication** icon in the Replication tree view.
2. In the GUI, select the pair that you want to edit in the **Pairs** list.
3. Click the **Edit Pair** button. Change the Pair Name, Group Name, and/or Copy Pace if necessary. You can view screen instructions for specific information by clicking the Help button.

Edit Pair - SS_LU0001_LU0901 Help

Local Pair Property

This screen shows basic properties for a pair. Click OK after editing, otherwise click Cancel.

* Pair Name :
31 characters or less (alphanumeric characters, special symbols "%", "*", "+", "-", ".", "/", "=", "@", "_", ":", "[", or "]")

Primary Volume : 0001

* Secondary Volume : ☒ Unassign
☐ Assign :
From 0 to max (array model type dependent)

Copy Type : SnapShot

Group Number : ***

Group Name :
31 characters or less (alphanumeric characters, special symbols "%", "*", "+", "-", ".", "/", "=", "@", "_", ":", "[", or "]")

Copy Pace :

OK Cancel

4. When assigning a volume number to the secondary volume without the volume number, select **Assign** and enter the volume number to be assigned in the text box.

When depriving the volume number from the secondary volume with the volume number, select **Unassign**.

When the volume number to be assigned to the secondary volume is already assigned to the secondary volume which is in the pair relationship with the same primary volume, deprive the volume number from the already assigned secondary volume and assign the specified secondary volume.



NOTE: If the volume number is deprived from the secondary volume, the host cannot recognize it. Check that the host does not access and deprive the volume number

5. Click **OK**.
6. A confirmation message appears. Click **Close**.

Edit Pair



Use the V-VOL for tape backup, testing, and reports

Your snapshot image (V-VOL) can be used to fulfill a number of data management tasks performed on a secondary server. These management tasks include backing up production data to tape, using the data to develop or test an application, generating reports, populating a data warehouse, and so on.

Whichever task you are performing, the process for preparing and making your data available is the same. The following process can be performed using the Navigator 2 GUI or CLI, in combination with an operating system scheduler. The process should be performed during non-peak hours for the host application.

To use the V-VOL for secondary functions

1. Un-mount the V-VOL. This is only required if the V-VOL is currently being used by a host server.
2. Resync the pair before stopping or quiescing the host application. This is done to minimize the down time of the production application.
 - GUI users, please see the resync pair instruction in [Updating the V-VOL on page 10-11](#).
 - For instructions using CLI, see the resync pair instruction in [Splitting Snapshot Pairs on page B-15](#).



NOTE: Some applications can continue to run during a backup operation, while others must be shut down. For those that stay running (placed in backup mode or quiesced rather than shut down), there may be a performance slowdown on the P-VOL.

3. When pair status becomes "Paired", shut down or quiece (quiet) the production application, if possible.
4. Split the pair. Doing this insures that the copy will contain the latest mirror image of the P-VOL.

- GUI users, please see the split pair instruction in [Updating the V-VOL on page 10-11](#).
 - For instructions using CLI, please see the split pair instruction in [Splitting Snapshot Pairs on page B-15](#).
5. Un-quiesce or start up the production application so that it is back in normal operation mode.
 6. Mount the (V-VOL on the server if previously unmounted).
 7. Run the backup program using the snapshot image (V-VOL).



NOTE: When performing read operations against the snapshot image (V-VOL), you are effectively reading from the P-VOL. This extra I/O on the P-VOL affects the performance.

Tape backup recommendations

Securing a tape-backup of your V-VOLs is recommended because it allows for restoration of the P-VOL in the event of pair failure (and thus, an invalid V-VOL). This section outlines a general scenario for backing up two V-VOLs:

- A P-VOL is copied to two V-VOLs everyday at 12 midnight and 12 noon.
- Tape backups of the two V-VOLs are made when few I/O instructions are issued by a host. Generally, host I/O should be less than 100 IOPS.
- Backup to tape should use two ports simultaneously.
- The total capacity of each V-VOL must be 1.5 TB or smaller. When the total capacity is 1 TB, time required for backing up (at a speed of 100 MB/sec) is 3 hours.
- DP pool capacity should be increased to 1.5-times the capacity of the P-VOL. This is because all the data that is restored from tape to the V-VOL becomes differential data in relation to the P-VOL. Therefore, the DP pool should be as large or larger than this. It is recommended that the DP pool be sized to 1.5 times the P-VOL capacity as a safety precaution.

[Figure 10-1 on page 10-18](#) illustrates this example scenario.




Figure 10-1: Tape backup

Restoring data from a tape backup

Data can be restored from a tape backup to a V-VOL or P-VOL. Restoring to the V-VOL results in less impact on your Snapshot system and on performance.

- If restoring to the V-VOL:
 - Data must be restored to the V-VOL from which the tape backup was made.
 - Data can then be restored from the V-VOL to the P-VOL. The P-VOL must be unmounted before its restoration.
 - DP pool capacity should be 1.5-times the capacity of the P-VOL.
- If restoring to the P-VOL:
 - Unmount the P-VOL.

- Return all pairs to Simplex (recommended in order to reduce build-up of data in the DP pool and impact to performance). See [Figure 10-2](#).




Figure 10-2: Direct restoration of P-VOL

- Use this method when the V-VOL is in Failure status (V-VOL data is corrupt), or when capacity of the DP pool is less than 1.5-times the capacity of the P-VOL.




Figure 10-3: Restoring backup data from a tape device

Quick recovery backup

The steps to perform quick recovery backup are similar to tape backup.

This restoration method uses backup data within the same disk array backup.

When a software failure (caused by a wrong operation by a user or an application program bug) occurs, perform restoration, selecting backup data you want to return from the V-VOL being retained. See [Figure 10-4](#) and [Figure 10-5](#) on [page 10-21](#).

It is necessary to un-mount the P-VOL once before restoring the P-VOL from the V-VOL.




Figure 10-4: Quick recovery backup




Figure 10-5: Quick recovery backup

Monitoring and troubleshooting Snapshot

It is important that a DP pool's capacity is sufficient to handle the replication data sent to it from the P-VOLs associated with it. If a DP pool should become full, the associated V-VOLs are invalidated, and backup data is lost.

This chapter provides information and instructions for monitoring and maintaining the Snapshot system.

- ❑ [Monitoring Snapshot](#)
- ❑ [Troubleshooting](#)

Monitoring Snapshot

The Snapshot DP pool must have sufficient capacity to handle the write workload demands placed on it. You can check that the DP pool is large enough to handle workload by monitoring pair status and DP pool usage.

Monitoring pair status

To monitor pair status, see [Confirming pair status on page 10-3](#).

Monitoring pair failure

In order to monitor whether SnapShot pairs operate correctly and the data is retained in V-VOLs, you must check pair status regularly. When a hardware failure occurs or the DP pool is shortage, the pair status is changed to **Failure** and the V-VOL data is not retained. Check that the pair status is other than Failure. When the pair status is **Failure**, it is required to restore the status. See [Pair failure on page 11-9](#).

For Snapshot, the following processes are executed when the pair failure occurs.

Table 11-1: Pair failure results

Management software	Results
Navigator 2	A message is displayed in the event log
	The pair status is changed to Failure or Failure(R)
CCI	The pair status is changed to PSUE.
	An error message is output to the system log file. (For UNIX® system and the Windows Server, the syslog file and eventlog file are shown respectively.)

When the pair status is changed to Failure or Failure(R), a trap is reported with SNMP Agent Support Function

When using CCI, the following message is output to the event log. For the details, refer to *Hitachi Unified Storage Command Control Interface Installation and Configuration Guide*.

Table 11-2: CCI system log message

Message ID	Condition	Cause
HORCM_102	The volume is suspended in code 0006	The pair status was suspended due to code 0006.

Monitoring pair failure using a script

When SNMP Agent Support Function not used, it is necessary to monitor the pair failure using a Windows Server script that can be performed using Navigator 2 CLI commands.

The following is a script for monitoring the two pairs (SI_LU0001_LU0002 and SI_LU0003_LU0004) and informing the user when pair failure occurs. The following script is activated every several minutes. The disk array must be registered to beforehand.

```
echo OFF
REM Specify the registered name of the arrays
set UNITNAME=Array1
REM Specify the name of target group (Specify "Ungroup" if the
pair doesn't belong to any group)
set G_NAME=Ungrouped
REM Specify the name of target pair
set P1_NAME=SI_LU0001_LU0002
set P2_NAME=SI_LU0003_LU0004
REM Specify the value to inform "Failure"
set FAILURE=14

REM Checking the first pair
:pair1
aureplicationmon -unit %UNITNAME% -evwait -si -pairname
%P1_NAME% -gname %G_NAME% -nowait
if errorlevel %FAILURE% goto pair1_failure
goto pair2
:pair1_failure
<The procedure for informing a user>*

REM Checking the second pair
aureplicationmon -unit %UNITNAME% -evwait -si -pairname
%P2_NAME% -gname %G_NAME% -nowait
if errorlevel %FAILURE% goto pair2_failure
goto end
:pair2_failure
<The procedure for informing a user>*

:end
%
```

Monitoring DP pool usage

The Snapshot DP pool must have sufficient capacity to handle the write workload demands placed on it. You can monitor DP pool usage by locating the desired DP pool and reviewing the percentage of the DP pool that is being used.

If DP pool usage rate (DP pool usage/DP pool capacity) exceeds the replication Depletion Alert threshold value (default is 40% and can be set from 1 to 99%) in Snapshot, in Navigator 2 a message is displayed in the event log and the status of the split pair using the target DP pool becomes Threshold Over. In CCI, the status of the PFUS pair using the target DP pool becomes PFUS.

If DP pool usage rate (DP pool usage/DP pool capacity) exceeds the Replication Data Released threshold value (default is 95% and can be set from 1 to 99%) in Snapshot, all the Snapshot pairs existing in the relevant DP pool are changed to the Failure status, the replication data and the management information used by the Snapshot pair are cancelled, and the usable capacity of the DP pool recovers.

In order to prevent the DP pool from exceeding the threshold value and the pair statuses from changing to Threshold over or Failure, you must monitor the used DP pool capacity. Even when a hardware maintenance contract is in effect (including the term of guarantee without charge), you must monitor the DP pool capacity so that it does not exceed the threshold value.

When there is a risk of exceeding the threshold value, expand the DP pool capacity (see [Expanding DP pool capacity on page 11-7](#)), or secure a free capacity of the DP pool by releasing V-VOL pairs that are not required to retain the data.

- Processing at the time of replication Depletion Alert threshold value over of the DP pool

If DP pool usage rate (DP pool usage/DP pool capacity) exceeds the replication Depletion Alert threshold value (default is 40% and can be set from 1 to 99%) in Snapshot, the following processes are executed. Also, the E-mail Alert Function and SNMP Agent Support Function will work to notify you of the event happening.

Table 11-3: Processing at the time of replication depletion alert threshold value over

Management software	Results
Navigator 2	A message is displayed in the event log.
	The status of the Split pair using the target DP pool becomes Threshold Over.
CCI	The status of the PSUS pair using the target DP pool becomes PFUS.

- Processing at the time of Replication Data Released threshold value over of the DP pool

If DP pool usage rate (DP pool usage/DP pool capacity) exceeds the Replication Data Released threshold value (default is 95% and can be set from 1 to 99%) in Snapshot, all the Snapshot pairs existing in the relevant DP pool are changed to the Failure status, the replication data and the management information used by the Snapshot pair are cancelled and the usable capacity of the DP pool recovers. Also, any operations except for pair deletion cannot be performed until the Replication Data Released threshold over is lifted.

- Method for informing a user that the threshold value of the replication is exceeded
 - In order to notify of the risk of the DP pool shortage in advance, E-Mail Alert and a trap is reported with SNMP Agent Support Function
 - You can get the pair status as a returned value using the pairvolchk -ss command of CCI. When the status is PFUS, the returned value is 28. (When the volume is specified, the values for the P-VOL and V-VOL are 28 and 38 respectively.) For the details of the pairvolchk command, refer to the *Hitachi Unified Storage Command Control Interface Installation and Configuration Guide*.
 - Monitoring on the used DP pool capacity is necessary for each DP pool
 - The capacity of the DP pool being used (rate of use) can be referred to through CCI or Navigator 2. It is recommended not only to monitor the DP pool threshold value but also to monitor and manage the hourly transition of the used capacity of the DP pool. For details of procedure for referring to the rate of DP pool capacity used, refer to the *Hitachi Unified Storage Command Control Interface Installation and Configuration Guide*.

Using a script to monitor DP pool usage

When SNMP Agent Support Function is not used, it is necessary to monitor the DP pool threshold over by a script that can be performed using Navigator 2 CLI commands.

The following is a script for monitoring the two pairs (SS_LU0001_LU0002 and SS_LU0003_LU0004) on a Windows host and informing the user when DP pool threshold over occurs. The following script is activated every several minutes. The disk array must be registered beforehand.


```

echo OFF
REM Specify the registered name of the arrays
set UNITNAME=Array1
REM Specify the name of target group (Specify "Ungroup" if the
pair doesn't belong to any group)
set G_NAME=Ungrouped
REM Specify the name of target pair
set P1_NAME=SS_LU0001_LU0002
set P2_NAME=SS_LU0003_LU0004
REM Specify the value to inform "Threshold over"
set THRESHOLDOVER=15

REM Checking the first pair
:pair1
aureplicationmon -unit %UNITNAME% -evwait -ss -pairname
%P1_NAME% -gname %G_NAME% -nowait
if errorlevel %THRESHOLDOVER% goto pair1_thresholdover
goto pair2
:pair1_thresholdover
<The procedure for informing a user>*

REM Checking the second pair
aureplicationmon -unit %UNITNAME% -evwait -ss -pairname
%P2_NAME% -gname %G_NAME% -nowait
if errorlevel %THRESHOLDOVER% goto pair2_thresholdover
goto end
:pair2_thresholdover
<The procedure for informing a user>*

:end

```

Expanding DP pool capacity

When the DP pool usage gets close to the replication threshold values, you can expand the DP pool capacity to accommodate more replication data from Snapshot pairs. Refer to the *Hitachi Unified Storage Dynamic Provisioning Configuration Guide* for the details on expanding DP pool.

Other methods for lowering DP pool usage

When a DP pool is in danger of exceeding its Replication Data Released threshold value, the following actions can be taken as alternatives or in addition to expanding the DP pool:

- Delete one or more V-VOLs. With fewer V-VOLs, less data accumulates in the DP pool.
- Reduce V-VOL lifespan. By holding snapshots for a shorter length of time, less data accumulates, which relieves the load on the DP pool.
- A re-evaluation of your Snapshot system's design may show that not enough DP pool space was originally allocated. See [Chapter 9, Snapshot setup](#) for more information.

Troubleshooting

Two types of problems can be experienced with a Snapshot system: pair failure and DP pool capacity exceeded. This topic describes the causes and provides solutions for these problems.

- ❑ [Pair failure](#)
- ❑ [DP Pool capacity exceeds replication threshold value](#)

Pair failure

You can monitor the status of a DP pool whose associated pairs' status is changed to Failure. Pair failure is caused by one of the following reasons:

- A hardware failure occurred that affects pair or DP pool volumes
- A DP pool's capacity usage has exceeded the Replication Data Released threshold value.

To determine the cause of pair failure

1. Check the status of the DP pool whose associated pairs' status is changed to Failure. Using Navigator 2, confirm the message displayed in Event Log tab in Alert & Events window. Check the status of the DP pool used by pairs whose status has been changed to Failure.
 - a. When the message "I6D000 DP pool does not have free space (DP pool-xx)" (xx is the number of the DP pool) is displayed, the pair failure is considered to have occurred due to shortage of the DP pool.
 - b. If the DP pool usage does not exceed the Replication Data Released threshold value, the pair failure is due to hardware failure.

DP pool capacity usage exceeds the Replication Data Released threshold value

If a DP pool's capacity usage exceeds the Replication Data Released threshold value, release all pairs that are using the DP pool among pairs whose status is Failure. The DP pool's exceeded capacity is considered to have occurred because of a problem of the system configuration. Review the configuration, including the DP pool capacity and the number of V-VOLs, after deleting the pairs. Execute the operation to restore the status of the Snapshot pair after reviewing the configuration. You must perform all operations for the restoration when a pair failure has occurred due to the DP pool's exceeded capacity.

Pair failure due to hardware failure

If a pair failure occurs because of a hardware failure, maintain the disk array first. Recover the pair from the failure by a pair operation after the failure of the array has been removed. Also, a pair operation may be necessary for the maintenance work of the array. For example, when formatting of a volume where a failure occurred is required and the volume is a Snapshot P-VOL, the formatting must be done after the pair is released. Even if the maintenance personnel maintain the array, the work by the service personnel is limited to the failure recovery and you must perform the operation to restore the status of a Snapshot pair.

To restore the status of the Snapshot pair, create the pair again after releasing the pair.

The procedure for restoring the pair differs according to the cause, see [Figure 11-1 on page 11-11](#). [Table 11-4 on page 11-12](#) shows the work responsibility schedule for the service personnel and a user.




Figure 11-1: Pair failure recovery procedure

Table 11-4: Operational notes for Snapshot operations

Action	Action taken by
Monitoring pair failure.	User
Confirm the Event Log message using Navigator 2 (Confirming the DP pool).	User
Verify the status of the array.	User
Call maintenance personnel when the array malfunctions.	User
For other reasons, call the Hitachi support center.	User (only for users that are registered in order to receive a support)
Split the pair.	User
Hardware maintenance.	Hitachi Customer Service
Reconfigure and recover the pair.	User

In addition, check the pair status immediately before the occurrence of the pair failure. In the case where the failure occurs when the pair status is Reverse Synchronizing (during restoration from a V-VOL to a P-VOL), the coverage of the data assurance and the detailed procedure for restoring the pair status differ from a case where a failure occurs when the pair status is other than Reverse Synchronizing. [Table 11-5](#) shows the data assurance and the procedure for restoring the pair when a pair failure occurs.

When the pair status is Reverse Synchronizing, data copying for the restoration is being done in the background. Therefore, when the restoration is performed normally, a host recognizes P-VOL data as if it were replaced with V-VOL data from immediately after the start of the restoration, but when a pair failure has occurred, it is impossible to make the host recognize the P-VOL as if it were replaced with the V-VOL and the P-VOL data becomes invalid because copying to the P-VOL is not completed.

Table 11-5: Data assurance and the method for recovering the pair

State of failure	Data assurance	Action taken after failure or failure(R)
Failure (State before Failure other than Reverse Synchronizing)	P-VOL: Assured V-VOL: Not assured	Split the pair, and then create a pair again. Even if the P-VOL data is assured, there may be a case where a pair has already been released because a failure such as a multiple drive blockade has occurred in a volume that configures a DP pool being used by the pair. In such case, confirm that the data exists in the P-VOL, and then create a pair. Incidentally, the V-VOL data generated is not the one invalidated previously but the P-VOL data at the time when the pair was newly created.
Failure(R) (State before Failure Reverse Synchronizing)	P-VOL: Not assured V-VOL: Not assured	Split the pair, restore the backup data to P-VOL, and then create a pair. There may be a case where a pair has already been released because a failure such as a double drive failure has occurred in a volume that configures a P-VOL or a DP pool. In such case, confirm that the backup data restoration has been completed to the P-VOL, and then create a pair. Incidentally, the V-VOL data generated is not the one invalidated previously but the P-VOL data at the time when the pair was newly created.

DP Pool capacity exceeds replication threshold value

When a DP pool capacity that is used exceeds the replication Depletion Alert value, the status of pairs using the DP pool becomes Threshold over. Even when the pair status is changed to Threshold over, the pairs operate as they are in the Split status, but it is necessary to secure the DP pool capacity early because it is highly possible that the DP pool is exhausted. The operation to secure the DP pool capacity is performed by a user. To secure the DP pool capacity, release the pairs that are using the DP pool or expand the DP pool capacity.

When releasing a pair, back up the V-VOL data to a tape device, if necessary, before releasing the pair because the data of the V-VOL of the pair to be released becomes invalid.

To expand the DP pool capacity (see [Setting up the DP pool on page 9-31](#)), add one or more RAID groups to the DP pool.

When the DP pool usage capacity exceeds the Replication Data Released threshold value, all the Snapshot pairs using the relevant DP pool are changed to the Failure status. After securing sufficient DP pool capacity, perform pair creation again for all the Snapshot pairs in the Failure status after pair deletion.

Cases and solutions using DP-VOLs

When configuring a Snapshot pair using the DP-VOL as a pair target volume, the Snapshot pair status may become Failure depending on the combination of the pair status and the DP pool status shown in [Table 11-6](#). Perform the recovery method shown in [Table 11-6](#) for all the DP pools to which the P-VOLs and the S-VOLs where the pair failures have occurred belong.

Table 11-6: Cases and solutions using the DP-VOLs

Pair status	DP pool status	Cases	Solutions
Paired Synchronizing	Formatting	Although the DP pool capacity is being added, the format progress is slow and the required area cannot be allocated.	Wait until the formatting of the DP pool for total capacity of the DP-VOLs created in the DP pool is completed.
	Capacity Depleted	The DP pool capacity is depleted and the required area cannot be allocated.	To make the DP pool status normal, perform the DP pool capacity growing and DP pool optimization, and increase the DP pool free capacity.

Recovering from pair failure due to a hardware failure

You can recover from a pair failure that is due to a hardware failure.

Prerequisite

1. Review the information log to see what the hardware failure is.
2. Restore the Storage system. See Navigator 2 program Help for details.

To recover the Snapshot system after a hardware failure


1. When the systems is restored, delete the pair. See [Deleting pairs and V-VOLs on page 10-14](#) for more information.
2. Re-create the pair. See [Create a Snapshot pair to back up your volume on page 10-6](#).

Confirming the event log

You may need to confirm the event log to analyze the cause of problems.

HSMN2 GUI procedure to display the event log

1. Select the **Alerts & Events** icon. The **Alerts & Events** screen appears.
2. Click the **Event Log** tab in the **Alerts & Events** screen. To search particular message or error detail code, use the browsers.



HSMN2 CLI procedure to display the event log

1. Register the array to confirm the event log on the command prompt
2. Execute the `auinfomsg` command and confirm the event log.

```
% auinfomsg -unit array-name
Controller 0/1 Common
08/21/2012 17:00:37 00 I1H600 PSUE occurred[Snapshot]
08/21/2012 17:00:37 00 IAIQ00 The change of pair status
failed[Snapshot](CTG-20,code-0009)
08/21/2012 16:58:29 00 I1H600 PSUE occurred[Snapshot]
08/21/2012 16:58:29 00 IAIQ00 The change of pair status
failed[Snapshot](CTG-00,code-0009)
08/21/2012 16:01:21 00 I6HQ00 LU has been created(LU-0052)
08/21/2012 16:01:19 00 I6HQ00 LU has been created(LU-0051)
%
```

3. When searching the specified messages or error detail codes, store the output result in the file and use the search function of the text editor as shown below.

```
% auinfomsg -unit array-name>infomsg.txt  
%
```

Snapshot fails in a TCE S-VOL - Snapshot P-VOL cascade configuration

The method of issuing a pair split to Snapshot pairs on the remote array using HSNM2 or CCI is described in [Snapshot cascade configuration local and remote backup operations on page 27-85](#). As described, you can issue and reserve a pair split to a Snapshot group that is cascaded with a TCE group while the TCE group is in Paired state. However, there are cases where Snapshot pairs for which a pair split has been reserved can change to Failure state when a pair operation is performed to the Snapshot/ TCE pair or the Snapshot/ TCE pair state change after the pair split has been reserved for Snapshot. See [Pair failure on page 11-9](#) for the factors that can make Snapshot pairs change to Failure state. You can check error codes shown in the Event Log on HSNM2 to find out the factor that made the Snapshot pairs change to Failure state.

When Snapshot pairs for which a pair split has been reserved for change to Failure state due to some factor, you need to delete the Snapshot pairs that changed to Failure state and create the Snapshot pairs again. You must remove the factor (that you find by checking error codes) before performing a pair split to the Snapshot pairs.

See [Confirming the event log on page 11-16](#) for the procedure to check the Event Log.

Message contents of the Event Log

In the message, a time when an error occurred, a controller number, a message, and an error code are displayed. The error message that appears when Snapshot pairs for which a pair split has been reserved change to **Failure** state is "The change of pair status failed[SnapShot]".




Figure 11-2: Error message example

Table 11-7 shows the error codes and the corresponding factor to each error code, so you can understand what makes Snapshot pairs change to Failure state.

Table 11-7: Error codes and corresponding factors

Error codes	Corresponding factors
0001	A pair creation has been issued to a TCE group that is cascaded with a Snapshot group for which a pair split has been reserved.
0002	A pair resync has been issued to a TCE group or a TCE pair in a TCE group that is cascaded with a Snapshot group for which a pair split has been reserved.
0003	A pair split has been issued to a TCE pair in a TCE group that is cascaded with a Snapshot group for which a pair split has been reserved.
0004	On the local array, a pair deletion has been issued to a TCE pair in a TCE group that is cascaded with a Snapshot group for which a pair split has been reserved.
0005	On the remote array, a pair deletion has been issued to a TCE group or a TCE pair in a TCE group that is cascaded with a Snapshot group for which a pair split has been reserved.
0006	A pair operation that causes an S-VOL to change to takeover state has been issued to a TCE group or a TCE pair in a TCE group that is cascaded with a Snapshot group for which a pair split has been reserved.
0007	Using CCI, a pairsplit -mScas has been issued to a TCE group that is cascaded with a Snapshot group for which a pair split has been reserved.

Table 11-7: Error codes and corresponding factors

Error codes	Corresponding factors
0008	A planned shutdown or power loss occurred before a pair split that has been reserved for a Snapshot group is actually executed.
0009	A pair split that has been reserved for a Snapshot group has been timed out.
000A	The local DP pool for a TCE group that is cascaded with a Snapshot group for which a pair split has been reserved has been depleted.
000B	The pair state of Snapshot pairs is not Paired when a pair split that has been reserved for the Snapshot group is actually executed.
000C	The pair state of a TCE group that is cascaded with a Snapshot group for which a pair split has been reserved is not Paired when the reserved pair split is actually executed.
000D	The max number of Snapshot generations has already been created when the reserved pair split is actually executed.
000E	The status of the Replication Data DP Pool or Management Area DP Pool for a Snapshot group for which a pair split has been reserved for is other than Normal/Regression. Or Replication Data Released Threshold for the DP pool is exceeded. Or the DP pool is depleted.
000F	The firmware on the local array does not support this feature when a TCE group has been deleted which is cascaded with a Snapshot group for which a pair split has been reserved for.
0010	The pair state of a TCE group that is cascaded with a Snapshot group for which a pair split has been reserved is not Paired when the TCE group is deleted.

TrueCopy Remote Replication theory of operation

A broken link, an accidentally erased file, the force of nature: negative occurrences cause problems to storage systems. When access to critical data is interrupted, a business can suffer irreparable harm.

Hitachi TrueCopy Remote Replication helps you keep critical data backed up in a remote location, so that negative incidents do not have a lasting impact.

The key topics in this chapter are:

- ❑ [TrueCopy Remote Replication](#)
- ❑ [How TrueCopy works](#)
- ❑ [Typical environment](#)
- ❑ [TrueCopy interfaces](#)
- ❑ [Typical workflow](#)
- ❑ [Operations overview](#)

TrueCopy Remote Replication

TrueCopy Remote Replication creates a duplicate of a production volume to a secondary volume located at a remote site. Data in a TrueCopy backup stays synchronized with the data in the local disk array. This happens when data is written from the host to the local disk array then to the remote system, via fibre channel or iSCSI link. The host holds subsequent output until acknowledgement is received from the remote disk array for the previous output.

When a synchronized pair is split, writes to the primary volume are no longer copied to the secondary side. Doing this means that the pair is no longer synchronous. Output to the local disk array is cached until the primary and secondary volumes are re-synchronized. When resynchronization takes place, only the changed data is transferred, rather than the entire primary volume. This reduces copy time.

TrueCopy can be teamed with ShadowImage or Snapshot, on either or both local and remote sites. These in-system copy tools allow restoration from one or more additional copies of critical data.

Besides disaster recovery, TrueCopy backup copies can be used for development, data warehousing and mining, or migration applications.



NOTE: TrueCopy Remote cannot be used together with TCE. TrueCopy Remote volumes can be cascaded with ShadowImage or Snapshot volumes.

How TrueCopy works

A TrueCopy “pair” is created when you:

- Select a volume on the local disk array that you want to copy.
- Create or identify the volume on the remote disk array that will contain the copy.
- Connect the local and remote disk arrays with a fibre channel or iSCSI link
- Copy all primary volume data to the secondary volume.

Under normal TrueCopy operations, all data written to the primary volume is copied to the secondary volume, ensuring that the secondary copy is a complete and consistent backup.

If the pair is split, the primary volume continues being updated, but data in the secondary volume remains as it was at the time of the split. At this time:

- The secondary volume becomes available for read/write access by secondary host applications.
- Changes to primary and secondary volumes are tracked by differential bitmaps.
- The pair can be made identical again by re-synchronizing changes from primary-to-secondary or secondary-to-primary.

To plan a TrueCopy system an understanding is needed of its components.

Typical environment

A typical configuration consists of the following elements. Many but not all require user set up.

- Two disk arrays—one on the local side connected to a host, and one on the remote side connected to the local disk array. Connections are made via fibre channel or iSCSI.
- A primary volume (P-VOL) on the local disk array that is to be copied to the secondary volume (S-VOL) on the remote side. Primary and secondary volumes may be composed of several volumes.
- A DMLU on local and remote disk arrays, which hold TrueCopy information.
- Interface and command software, used to perform TrueCopy operations. Command software uses a command device (volume) to communicate with the disk arrays.

Figure 12-1 shows a typical TrueCopy environment.




Figure 12-1: TrueCopy components

Volume pairs

As described above, original data is stored in the P-VOL and the remote copy is stored in the S-VOL. The pair can be paired, split, re-synchronized, and returned to the simplex state. When synchronized, the volumes are paired;

when split, new data sent is to the P-VOL but held from the S-VOL. When re-synchronized, changed data is copied to the S-VOL. When necessary, data in the S-VOL can be copied to the P-VOL (P-VOL restoration). Volumes on the local and remote disk arrays must be defined and formatted prior to pairing.

Remote Path

TrueCopy operations are carried out between local and remote disk arrays connected by a fibre channel or iSCSI interface. A data path, referred to as the remote path, connects the port from the local disk array that executes the volume replication to the port on the remote disk array. User setup is required on the local disk array.

Differential Management LU (DMLU)

A DMLU (Differential Management Logical Unit) is an exclusive volume for storing the differential data at the time when the volume is copied. The DMLU in the disk array is treated in the same way as the other volumes. To create a TrueCopy pair, it is necessary to prepare one DMLU in the local and remote array. The differential information of all TrueCopy pairs is managed by this DMLU.

However, a volume that is set as the DMLU is not recognized by a host (it is hidden).

As shown in [Figure 12-2 on page 12-5](#), the array accesses the differential information stored in the DMLU and refers to/updates it in the copy processing to synchronize the P-VOL and the S-VOL and to manage the difference of the P-VOL and the S-VOL.

The creatable pair capacity is dependent on the DMLU capacity. If the DMLU does not have enough capacity to store the pair differential information, the pair cannot be created. In this case, a pair can be added by expanding the DMLU. The DMLU capacity 10 GB minimum and 128 GB maximum. See [Setting up the DMLU on page 14-20](#) for the number of creatable pairs according to the capacity and the total capacity of the volumes to be paired.

DMLU precautions:

- The volume belonging to RAID 0 cannot be set as a DMLU.
- When a failure occurs in the DMLU, all the pairs of ShadowImage, TrueCopy, and/or Volume Migration are changed to Failure. Therefore, secure sufficient redundancy for the RAID group to which the DMLU is located.
- In the status where the pair status is Split, Split Pending, or Reverse Synchronizing, the I/O performance of the DMLU may affect the host I/O performance to the volume which configures the pair. Using RAID 1+0 can decrease the effect to the host I/O performance.
- When setting the unified volume as a DMLU, it cannot be set if the capacity of each unified volume becomes less than 1 GB on an average. For example, when setting a volume of 10 GB as a DMLU, if the volume consists of 11 sub-volumes, it cannot be set as a DMLU.

- The volume assigned to the host cannot be set as a DMLU.
- In the DMLU expansion not using Dynamic Provisioning, select a RAID group which meets the following conditions:
 - The drive type and the combination are the same as the DMLU
 - A new volume can be created
 - A sequential free area for the capacity to be expanded exists
- When either pair of ShadowImage, TrueCopy, or Volume Migration exist, the DMLU cannot be removed.




Figure 12-2: DMLU

Command devices

The command device is a user-selected, dedicated logical volume on the disk array, which functions as the interface to the CCI software. TrueCopy commands are issued by CCI (HORCM) to the disk array command device.

A command device must be designated in order to issue TrueCopy commands. The command device must be defined in the HORCM_CMD section of the configuration definition file for the CCI instance on the attached host. 128 command devices can be designated for the array. You can designate command devices using Navigator 2.



NOTE: Volumes set for command devices must be recognized by the host. The command device volume size must be greater than or equal to 33 MBs.

Consistency group (CTG)

Application data often spans more than one volume. With TrueCopy, it is possible to manage operations spanning multiple volumes as a single group. In a group, all primary logical volumes are treated as a single entity.

Managing primary volumes as a group allows TrueCopy operations to be performed on all volumes in the group concurrently. Write order in secondary volumes is guaranteed across application logical volumes. User setup is required.

Since multiple pairs can belong to the same group, pair operation is possible in units of groups. For example, in the group in which the Point-in-time attribute is enabled, the backup data of the S-VOL is created at the same time.

For setting a group, specify a new group number for a group to be assigned after pair creation when creating a TrueCopy pair. The maximum of 1,024 groups can be created in TrueCopy.

A group name can be assigned to a group. You can select one pair belonging to the created group and assign a group name arbitrarily by using the pair edit function.

TrueCopy interfaces

TrueCopy can be operated using of the following interfaces:

- The GUI (Hitachi Storage Navigator Modular 2 Graphical User Interface) is a browser-based interface from which TrueCopy can be setup, operated, and monitored. The GUI provides the simplest method for performing operations, requiring no previous experience. Scripting is not available.
- CLI (Hitachi Storage Navigator Modular 2 Command Line Interface), from which TrueCopy can be setup and all basic pair operations can be performed—create, split, resynchronize, restore, swap, and delete. The GUI also provides these functions. CLI also has scripting capability.
- CCI (Hitachi Command Control Interface (CCI), used to display volume information and perform all copying and pair-managing operations. CCI provides a full scripting capability which can be used to automate replication operations. CCI requires more experience than the GUI or CLI. CCI is required for performing failover and fall back operations. It is also required on Windows 2000 Server and Windows Server 2008 for mount/unmount operations.

HDS recommends using the GUI to begin operations for new users with no experience with CLI or CCI. Users who are new to replication software but have CLI experience in managing disk arrays may want to continue using CLI, though the GUI is an option. The same recommendation applies to CCI users.



NOTE: Hitachi Replication Manager can be used to manage and integrate TrueCopy. It provides a GUI representation of the TrueCopy system, with monitoring, scheduling, and alert functions. For more information, visit the Hitachi Data Systems website.

Typical workflow

Designing, creating, and using a TrueCopy system consists of the following tasks:

- **Planning:** you assemble the necessary components of a TrueCopy system. This includes establishing path connections between the local and remote disk arrays, volume sizing and RAID configurations, understanding how to use TrueCopy concurrently with ShadowImage and/or Copy-on-Write, and other necessary pre-requisite information and tasks.
- **Design:** you gather business requirements and write workload data to size TrueCopy remote path bandwidth to fit your organization's requirements.
- **Configuration:** you implement the system and create an initial pair.
- **Operations:** you perform copy and maintenance operations.
- **Monitoring the system**
- **Troubleshooting**

Operations overview

The basic TrueCopy operations are shown in [Figure 12-3](#). They consist of creating, splitting, resynchronizing, swapping, deleting a pair.

- **Create Pair.** This establishes the initial copy using two volumes that you specify. Data is copied from the P-VOL to the S-VOL. The P-VOL remains available to the host for read and write throughout the operation. Writes to the P-VOL are duplicated to the S-VOL. The pair status changes to Paired when the initial copy is complete.
- **Split.** The S-VOL is made identical to the P-VOL and then copying from the P-VOL stops. Read/write access becomes available to and from the S-VOL. While the pair is split, the disk array keeps track of changes to the P-VOL and S-VOL in track maps. The P-VOL remains fully accessible in Split status.
- **Resynchronize pair.** When a pair is re-synchronized, changes in the P-VOL since the split is copied to the S-VOL, making the S-VOL identical to the P-VOL again. During a resync operation, the S-VOL is inaccessible to hosts for write operations; the P-VOL remains accessible for read/write. If a pair was suspended by the system because of a pair failure, the entire P-VOL is copied to the S-VOL during a resync.
- **Swap pair.** The pair roles are reversed.
- **Delete pair.** The pair is deleted and the volumes return to Simplex status.




Figure 12-3: TrueCopy pair operations

See the individual procedures and more detailed information in [Chapter 15, Using TrueCopy Remote](#).

Installing TrueCopy Remote

This chapter provides procedures for installing and setting up TrueCopy using the Navigator 2 GUI. CLI and CCI instructions are included in this manual in the appendixes.

- ❑ [System requirements](#)
- ❑ [Installation procedures](#)

System requirements

The minimum requirements for TrueCopy are listed below.

Table 13-1: Environment and requirements of TrueCopy

Item	Contents
Environment	• Firmware: Version 0916/A or higher is required. • Navigator 2: version 22.0 or higher is required for management PC. • CCI: Version 01-27-03/02 or higher is required for host only when CCI is used for the operation of True Copy.
Requirements	• Number of controllers: 2 (dual configuration) • DMLU is required for 1 of each array. The DMLU size must be greater than or equal to 10 GB to less than 128 GB. • Number of arrays: 2 • Two license keys for TrueCopy • Size of volume: The P-VOL size must equal the S-VOL volume size. • The command device is required only when CCI is used for the operation of TrueCopy. The command device volume size must be greater than or equal to 33 MB.

See [TrueCopy specifications on page C-2](#) for additional information.

Installation procedures

TrueCopy is an extra-cost option; it must be installed and enabled on the local and remote disk arrays.

Before proceeding, verify that the disk array is operating in a normal state. Installation/un-installation cannot be performed if a failure has occurred.

The following sections provide instructions for installing, enabling/disabling, and uninstalling TrueCopy.

Installing TrueCopy Remote

Prerequisites

- A key code or key file is required to install or uninstall TrueCopy. If you do not have the key file or code, you can obtain it from the download page on the HDS Support Portal, <http://support.hds.com>.
- TrueCopy cannot be installed if more than 239 hosts are connected to a port on the disk array.

To install TrueCopy

1. In the Navigator 2 GUI, click the check box for the disk array where you want to install TrueCopy, then click the **Show & Configure disk array** button.
2. Under Common disk array Tasks, click **Install License**.



3. The Install License screen displays.



4. Select the **Key File** or **Key Code** option, then enter the file name or key code. You may browse for the Key File.
5. Click **OK**.

6. Click **Confirm** on the subsequent screen to proceed, then click **Close** on the installation complete message.

Enabling or disabling TrueCopy Remote

TrueCopy is automatically enabled when it is installed. You can disable and re-enable it.

Prerequisites

Before disabling TrueCopy:

- TrueCopy pairs must be deleted and the status of the volumes must be Simplex.
- The remote path must be deleted.
- TrueCopy cannot be enabled if more than 239 hosts are connected to a port on the disk array.

To enable / disable TrueCopy

1. In the Navigator 2 GUI, click the check box for the disk array, then click the **Show & Configure disk array** button.
2. In the tree view, click **Settings**, then click **Licenses**.
3. Select **TrueCopy** in the Licenses list.
4. Click **Change Status**. The Change License screen displays.
5. To disable, clear the **Enable: Yes** check box.
To enable, check the **Enable: Yes** check box.
6. Click **OK**.
7. A message appears confirming that TrueCopy is disabled. Click **Close**.


Uninstalling TrueCopy Remote

Prerequisites

- TrueCopy pairs must be deleted. Volume status must be Simplex.
- To uninstall TrueCopy, the key code or key file provided with the optional feature is required. Once uninstalled, TrueCopy cannot be used (locked) until it is again installed using the key code or key file. If you do not have the key code, you can obtain it from the download page on the HDS Support Portal, <http://support.hds.com>.

To uninstall TrueCopy

1. In the Navigator 2 GUI, click the check box for the disk array, then click the **Show & Configure disk array** button.
2. In the navigation tree, click **Settings**, then click **Licenses**.



3. On the Licenses screen, select **TrueCopy** in the Licenses list and click the **De-install License** button.

4. On the De-Install License screen, enter the code in **Key Code** box, and then click **OK**.



The image shows a 'De-Install License' dialog box. At the top, there is a title bar with 'De-Install License' and a 'Help' button. Below the title bar is a green header with the text 'License Property'. The main area contains the instruction 'Enter the information for the license to be de-installed.' followed by two radio buttons under the label '*De-install with:'. The first radio button is for 'Key File:' and is currently unselected; it has a text box and a 'Browse...' button below it with the instruction 'Input the Key File Name.' The second radio button is for 'Key Code:' and is selected; it has a text box below it with the instruction 'Input the Key Code.' At the bottom left, there is a note '* Required field'. At the bottom right, there are 'OK' and 'Cancel' buttons. A mouse cursor is pointing at the 'OK' button.

5. On the confirmation screen, click **Close**.

TrueCopy Remote setup

This chapter provides required information for setting up your system for TrueCopy Remote. It includes:

Planning and design

- ☐ [Planning for TrueCopy](#)
- ☐ [The planning workflow](#)
- ☐ [Planning disk arrays](#)
- ☐ [Planning volumes](#)
- ☐ [Operating system recommendations and restrictions](#)
- ☐ [Calculating supported capacity](#)

Setup procedures

- ☐ [Setup procedures](#)
- ☐ [Changing the port setting](#)
- ☐ [Determining remote path bandwidth](#)
- ☐ [Remote path requirements, supported configurations](#)
- ☐ [Remote path configurations for Fibre Channel](#)
- ☐ [Remote path configurations for iSCSI](#)
- ☐ [Connecting the WAN Optimization Controller](#)
- ☐ [Supported connections between various models of arrays](#)

Planning for TrueCopy

Planning a TrueCopy system requires an understanding of the components that are used in the remote backup environment and awareness of their requirements, restrictions, and recommendations.

- ❑ [The planning workflow](#)
- ❑ [Planning disk arrays](#)
- ❑ [Planning volumes](#)
- ❑ [Operating system recommendations and restrictions](#)
- ❑ [Calculating supported capacity](#)
- ❑ [Setup procedures](#)
- ❑ [Concurrent use of Dynamic Provisioning](#)

The planning workflow

Implementing a TrueCopy system requires setting up local and remote disk arrays, TrueCopy volumes, the remote path that connects the volumes, and interface(s).

A planning workflow can be organized in the following manner:

- Planning disk arrays for TrueCopy.
- Planning volume setup, which consists of:
 - Understanding TrueCopy primary and secondary volume specifications, recommendations, and restrictions,
 - Understanding how to use unified volumes to create P-VOLs and S-VOLs (optional)
 - Cascading TrueCopy volumes with ShadowImage or Snapshot volumes (optional).
 - Specifications, recommendations, and restrictions for DMLUs
 - Specifications, recommendations, and restrictions for command devices (only required if CCI is used)
- Planning remote path connections, which includes:
 - Reviewing supported path configurations (covered in [Changing the port setting on page 14-26](#))
 - Measuring write workload to determine the bandwidth that is required (covered in [Changing the port setting on page 14-26](#))

Planning disk arrays

Hitachi Unified Storage can be connected with Hitachi Unified Storage, AMS2100, AMS2300, AMS2500, WMS100, AMS200, AMS500, or AMS1000. Any combination of disk array may be used on the local and remote sides. When using the earlier model disk arrays, please observe the following:

- The maximum number of pairs between different model disk arrays is limited to the maximum number of pairs supported by the smallest disk array.
- The firmware version for WMS 100, AMS 200, AMS 500, or AMS 1000 must be 0780/A or later when connecting with HUS.
- The firmware version for AMS 2100, AMS2300, AMS2500 must be 08B7/A or later when connecting with HUS.
- The bandwidth of the remote path to a AMS2100, AMS2300, AMS2500, WMS 100, AMS 200, AMS 500, or AMS 1000 must be 20 Mbps or more.
- The pair operations for WMS 100, AMS 200, AMS 500, or AMS 1000 cannot be performed using the Navigator 2 GUI.
- AMS2100, AMS2300, AMS2500, WMS 100, AMS 200, AMS 500, or AMS 1000 cannot use functions that are newly supported by Hitachi Unified Storage.

Planning volumes

Please review the recommendations in the following subsections before setting up TrueCopy volumes.

Also, review:

- [System requirements on page 13-2](#)
- [TrueCopy specifications on page C-2](#)

Volume pair recommendations

- Because data written to a primary volume is also written to a secondary volume at a remote site synchronously, performance is impacted according to the distance to the remote site. Assigning volumes to primary or secondary volumes should be limited to those not required to return a quick response to a host.
- The number of volumes within the same RAID group should be limited. Pair creation or resynchronization for one of the volumes may impact I/O performance for the others because of contention between drives.
When creating two or more pairs within the same RAID group, standardize the controllers for the volumes in the RAID group. Also, perform pair creation and resynchronization when I/O to other volumes in the RAID group is low.
- For a P-VOL, use SAS drives. When a P-VOL is located in a RAID group containing SAS7.2K drives, I/O host performance, pair formation, and pair resynchronization decrease due to the lower performance of the SAS7.2K drive. Therefore, it is best to assign a P-VOL to a RAID group that consists of SAS drives.
- Assign an volume consisting of four or more data disks, otherwise host and/or copying performance may be lowered.
- Volumes used for pair volumes should have stripe size of 64 kB and segment size of 16 kB.
- When using the SAS7.2K drives, make the data disks between 4D and 6D.
- When TrueCopy and Snapshot are cascaded, the Snapshot DP pool activity influences host performance and TrueCopy copying. Therefore, assign a volume of SAS drives and assign four or more disks (which have higher performance than SAS7.2K drives), to a DP pool.
- Limit the I/O load on both local and remote disk arrays to maximize performance. Performance on the remote disk array affects performance on both the local system and the synchronization of volumes.
- Synchronize Cache Execution Mode must be turned off on the remote disk array to prevent possible data path failure.

Volume expansion

A unified volume can be used as a TrueCopy P-VOL or S-VOL. When using TrueCopy with Volume Expansion, please observe the following:

- P-VOL and S-VOL capacities must be equal, though the number of volumes composing them (unified volumes may differ, as shown in [Figure 14-1](#)).




Figure 14-1: Capacity same, number of VOLS different

- A unified volume composed of 128 or more volumes cannot be assigned to the P-VOL or S-VOL, as shown in [Figure 14-2](#).




Figure 14-2: Number of volumes restricted

- P-VOLs and S-VOLs made of unified volumes can be assigned to different RAID levels and have a different number of data disks, as shown in [Figure 14-3 on page 14-6](#).




Figure 14-3: Combination of RAID levels

- A TrueCopy P-VOL or S-VOL cannot be used to compose a unified volume.
- The volumes created in the RAID group with different drive types cannot be unified, as shown in [Figure 14-4 on page 14-7](#). Unify volumes consisting of drives of the same type.




Figure 14-4: Unifying volumes of different drive type not supported

Operating system recommendations and restrictions

The following sections provide operating system recommendations and restrictions.

Host time-out

I/O time-out from the host to the disk array should be more than 60 seconds. You can figure I/O time-out by increasing the remote path time limit times 6. For example, if the remote path time-out value is 27 seconds, set host I/O time-out to 162 seconds (27 x 6) or more.

P-VOL, S-VOL recognition by same host on VxVM, AIX®, LVM


VxVM, AIX®, and LVM do not operate properly when both the P-VOL and S-VOL are set up to be recognized by the same host. The P-VOL should be recognized one host on these platforms, and the S-VOL recognized by another.

Setting the Host Group options

When MC/Service Guard is used on HP server, connect the host group (fibre channel) or iSCSI Target to HP server as follows:


For fibre channel interfaces

1. In the Navigator 2 GUI, access the disk array and click **Host Groups** in the Groups tree view.



WARNING! Your host group changes will be applied to multiple ports. This change will delete existing host group mappings and corresponding Host Group IDs, corrupting or removing data associated with the host groups. To keep specified host groups you do not want to remove, please cancel this operation and make changes to only one host group at a time.

2. Click the check box for the **Host Group** that you want to connect to the HP server.
3. Click **Edit Host Group**.



The Edit Host Group screen appears.

WWNs | Volumes | **Options**

Select or check options for host group.
When Platform or Middleware is set, Mode Setting follow automatically.

Platform: **hp-ux** Middleware: **not specify**

Mode Settings:

Common Setting: **Standard**

Additional Settings:

HP-UX Mode:	☒ Yes
PSUE Read Reject Mode:	☒ Yes
Mode Parameters Changed Notification Mode:	☐ Yes
NACA Mode:	☐ Yes
Task Management Isolation Mode:	☐ Yes
Unique Reserve Mode 1:	☐ Yes
Port-ID Conversion Mode:	☐ Yes
Tru Cluster Mode:	☐ Yes
Product Serial Response Mode:	☐ Yes
Same Node Name Mode:	☐ Yes
CCHS Mode:	☐ Yes
NOP-In Suppress Mode:	☐ Yes
S-VOL Disable Advanced Mode:	☐ Yes
Discovery CHAP Mode:	☐ Yes
Unique Extended COPY Mode:	☐ Yes
Unique Write Same Mode:	☐ Yes
DP Depletion Detail Reply Mode:	☐ Yes
Unit Attention Change Mode:	☐ Yes
Unique Compare and Write Mode:	☐ Yes

* Required field


4. Select the **Options** tab.
5. From the Platform drop-down list, select **HP-UX**. Doing this causes "Enable HP-UX Mode" and "Enable PSUE Read Reject Mode" to be selected in the Additional Setting box.
6. Click OK. A message appears, click Close.

Edit Host Group - Port 0A:000




For iSCSI interfaces

1. In the Navigator 2 GUI, access the disk array and click **iSCSI Targets** in the Groups tree view.



2. The iSCSI Targets screen appears.
3. Click the check box for the iSCSI Targets that you want to connect to the HP server.
4. Click **Edit Target**.



The Edit iSCSI Target screen appears.

The screenshot shows the 'Options' tab for editing an iSCSI target. It includes a 'Platform' dropdown set to 'hp-ux' and a 'Middleware' dropdown set to 'not specify'. The 'Mode Settings' section has a 'Common Setting' dropdown set to 'Standard'. The 'Additional Settings' section contains a list of modes with checkboxes, including 'HP-UX Mode' and 'PSUE Read Reject Mode', both of which are checked. A footer note states '* Required field'.

5. Select the **Options** tab.
6. From the Platform drop-down list, select **HP-UX**. Doing this causes "Enable HP-UX Mode" and "Enable PSUE Read Reject Mode" to be selected in the Additional Setting box.
7. Click **OK**. A message appears, click **Close**.

Windows 2000 Servers

- A P-VOL and S-VOL *cannot* be made into a dynamic disk on Windows 2000 Server.
- Native OS mount/dismount commands can be used for all platforms, **except Windows 2000 Server and Windows Server 2008**. The native commands on this environment do not guarantee that all data buffers are completely flushed to the volume when dismounting. In these instances, you must use Hitachi's Command Control Interface (CCI) to perform volume mount/unmount operations. See the *Hitachi Unified Storage Command Control Interface Installation and Configuration Guide* for more information.

Windows Server

Volume mount:

- In order to make a consistent backup using a storage-based replication such as Truecopy Remote, you must have a way to flush the data residing on the server memory to the array, so that the source volume of the replication has the complete data. You can flush the data on the server memory using the umount command of CCI to unmount the volume. When using the umount command of CCI for unmount, use the mount command of CCI for mount. (For more detail about mount/umount command, see the *Hitachi Unified Storage Command Control Interface Installation and Configuration Guide*.)
- If you are using Windows Server 2003, mountvol /P to flush data on the server memory when un-mounting the volume is supported. Please understand the specification of the command and run sufficient test before you use it for your operation.
- In Windows Server 2008, refer to the *Hitachi Unified Storage Command Control Interface Installation and Configuration Guide* for the detail of the restrictions of Windows Server 2008 when using the mount/umount command.
- Windows Server may write for the un-mounted volume. If a pair is resynchronized while remaining the data to the S-VOL on the memory of the server, the compatible backup cannot be collected. Therefore, execute the sync command of CCI immediately before re-synchronizing the pair for the un-mounted S-VOL.
- Refer to the *Hitachi Unified Storage Command Control Interface Installation and Configuration Guide* for more detail about the CCI commands.

Volumes recognized by the host:

- If you recognize the P-VOL and S-VOL on Windows Server 2008 at the same time, it may cause an error because the P-VOL and S-VOL have the same disk signature. When the P-VOL and S-VOL have the same data, split the pair and then rewrite the disk signature so that they can retain different disk signatures. You can use the uniqueid command to rewrite a disk signature. See the *Hitachi Unified Storage Command Control Interface Installation and Configuration Guide* for the detail.

Command devices

- -When a path detachment, which is caused by a controller detachment or interface failure, continues for longer than one minute, the command device may be unable to be recognized at the time when recovery from the path detachment is made. To make the recovery, execute the "re-scanning of the disks" of Windows. When Windows cannot access the command device although CCI is able to recognize the command device, restart CCI.

Dynamic Disk in Windows Server

In an environment of Windows Server, you cannot use TrueCopy pair volumes as dynamic disk. This is because if you restart Windows or use the Rescan Disks command after creating or re-synchronizing a TrueCopy pair, there are cases where the S-VOL is displayed as **Foreign** in Disk Management and become inaccessible.

UNMAP Short Length Mode

Enable UNMAP Short Length Mode when connecting to Windows 2012. If you do not enable it, UNMAP commands may not be completed due to a time-out.

Identifying P-VOL and S-VOL in Windows

In Navigator 2, the P-VOL and S-VOL are identified by their volume number. In Windows, volumes are identified by HLUN. These instructions provide procedures for the fibre channel and iSCSI interfaces. To confirm the H-LUN:

1. From the Windows Server 2003 Control Panel, select Computer Management/Disk Administrator.
2. Right-click the disk whose HLUN you want to know, then select **Properties**. The number displayed to the right of "LUN" in the dialog window is the HLUN.

For Fibre Channel interface:

Identify HLUN-to-VOL Mapping for the Fibre Channel interface, as follows:

1. In the Navigator 2 GUI, select the desired disk array.
2. In the array tree, click the **Group** icon, then click **Host Groups**.



WARNING! Your host group changes will be applied to multiple ports. This change will delete existing host group mappings and corresponding Host Group IDs, corrupting or removing data associated with the host groups. To keep specified host groups you do not want to remove, please cancel this operation and make changes to only one host group at a time.

3. Click the **Host Group** to which the volume is mapped.

4. On the screen for the host group, click the **Volumes** tab. The volumes mapped to the Host Group display. You can confirm the VOL that is mapped to the H-LUN.

For iSCSI interface:

Identify HLUN-to-VOL mapping for the iSCSI interface as follows.

1. In the Navigator 2 GUI, select the desired array.
2. In the array tree that displays, click the **Group** icon, then click the **iSCSI Targets** icon in the Groups tree.
3. On the iSCSI Target screen, select an iSCSI target.
4. On the target screen, select the **Volumes** tab. Find the identified HLUN. The VOL displays in the next column.
5. If the HLUN is not present on a target screen, on the iSCSI Target screen, select another iSCSI target and repeat Step 4.

VMware and TrueCopy configuration


When creating a backup of the virtual disk in the vmfs format using TrueCopy, shutdown the virtual machine which accesses the virtual disk, and split the pair.

If one volume is shared by multiple virtual machines, shut down all the virtual machines that share the volume when creating a backup. It is not recommended to share one volume by multiple virtual machines in the configuration that creates a backup using TrueCopy.

The VMware ESX has a function to clone the virtual machine. Although the ESX clone function and TrueCopy can be linked, cautions are required for the performance at the time of execution. For example, when the volume which becomes the ESX clone destination is a TrueCopy P-VOL pair whose pair status is Paired, since the data is written to the S-VOL for writing to the P-VOL, the time required for a clone may become longer and the clone may be terminated abnormally in some cases. To avoid this, we recommend the operation to make the TrueCopy pair status Split or Simplex and to resynchronize or create the pair after executing the ESX clone. Also, it is the same for executing the functions such as migration the virtual machine, deploying from the template and inflating the virtual disk.

UNMAP Short Length Mode

It is recommended that you enable UNMAP Short Length Mode when connecting to VMware. If you do not enable it, UNMAP commands may not be completed due to a time-out.



Volumes recognized by the same host restrictions

Windows Server

The target volume for TrueCopy has to recognize and release a drive with **mount** and **unmount** command of CCI, instead of your specifying the drive letter.

Cannot use **mountvol** command of Windows Server because the data is not flush when releasing. For more detail, see the *Command Control Interface (CCI) Reference Guide*.

Cannot combine with a path switching software.

AIX

Not available. If you have set the P-VOL and the S-VOL to be recognized by the same host, the VxVM, AIX, and LVM will not operate properly. Set only the P-VOL of TrueCopy to be recognized by the host and let another host recognize the S-VOL.

Concurrent use of Dynamic Provisioning

The DP-VOLs can be set for a P-VOL or an S-VOL in TrueCopy.

This section describes the points to remember when using TrueCopy and Dynamic Provisioning together. Refer to the *Hitachi Unified Storage Dynamic Provisioning Configuration Guide* for detailed information about Dynamic Provisioning. Hereinafter, the volume created in the RAID group is called a normal volume, and the volume created in the DP pool is called a DP-VOL.

- When using a DP-VOL as a DMLU

Check that the free capacity (formatted) of the DP pool to which the DP-VOL belongs is more than or equal to the capacity of the DP-VOL which can be used as the DMLU, and then set the DP-VOL as a DMLU. If the free capacity of the DP pool is less than the capacity of the DP-VOL which can be used as the DMLU, the DP-VOL cannot be set as a DMLU.

- Volume type that can be set for a P-VOL or an S-VOL of TrueCopy

The DP-VOL can be used for a P-VOL or an S-VOL of TrueCopy. [Table 14-1 on page 14-15](#) shows a combination of a DP-VOL and a normal volume that can be used for a P-VOL or an S-VOL of TrueCopy.

Table 14-1: Combination of a DP-VOL and a normal volume

TrueCopy P-VOL	TrueCopy S-VOL	Contents
DP-VOL	DP-VOL	Available. The P-VOL and S-VOL capacity can be reduced compared to the normal volume. (Note 1) When both the P-VOL and the S-VOL use DP-VOLs, a pair cannot be created by combining the DP-VOLs that have different setting of Enabled/Disabled for Full Capacity Mode.
DP-VOL	Normal volume	Available. In this combination, copying after pair creation takes about the same time it takes when the normal volume is a P-VOL. Moreover, when executing the swap, the DP pool of the same capacity as the normal volume (original S-VOL) is used. After the pair is split and reclaim zero page, the S-VOL capacity can be reduced.
Normal volume	DP-VOL	Available. When the pair status is Split, the S-VOL capacity can be reduced compared to the normal volume by reclaim zero page.



NOTES:

1. When both the P-VOL and the S-VOL use DP-VOLs, a pair cannot be created by combining the DP-VOLs which have different settings of Enabled/Disabled for Full Capacity Mode.
2. Depending on the volume usage, the consumed capacity of the P-VOL and the S-VOL may differ even in the Paired status. Execute the DP Optimization and zero page reclaim as needed.
3. The consumed capacity of the S-VOL may be reduced due to the resynchronization.

- Pair status at the time of DP pool capacity depletion

When the DP pool is depleted after operating the TrueCopy pair that uses the DP-VOL, the pair status of the pair concerned may be a Failure.

[Table 14-2 on page 14-16](#) shows the pair statuses before and after the

DP pool capacity depletion. When the pair status becomes a Failure caused by the DP pool capacity depletion, add the DP pool capacity whose capacity is depleted, and execute the pair operation again.

Table 14-2: Pair Statuses before the DP pool capacity depletion and pair statuses after the DP pool capacity depletion

Pair statuses before the DP pool capacity depletion	Pair statuses after the DP pool capacity depletion belonging to P-VOL	Pair statuses after the DP pool capacity depletion belonging to data pool
Simplex	Simplex	Simplex
Synchronizing	Synchronizing Failure*	Failure
Paired	Paired Failure *	Failure
Split	Split	Split
Failure	Failure	Failure
* When write is performed to the P-VOL to which the capacity depletion DP pool belongs, the copy cannot be continued and the pair status becomes a Failure.		

- DP pool status and availability of pair operation

When using the DP-VOL for a P-VOL or an S-VOL or a data pool of the TrueCopy pair, the pair operation may not be executed depending on the status of the DP pool to which the DP-VOL belongs. [Table 14-3](#) shows the DP pool status and availability of the TrueCopy pair operation. When the pair operation fails due to the DP pool status, correct the DP pool status and execute the pair operation again.

Table 14-3: DP pool statuses and availability of pair operation

Pair operation	DP pool statuses, DP pool capacity statuses, and DP pool optimization statuses					
	Normal	Capacity in growth	Capacity depletion	Regressed	Blocked	DP in optimization
Create pair	YES ¹	YES	YES ¹ YES ²	YES	NO	YES
Split pair	YES	YES	YES	YES	YES	YES
Resync pair	YES ¹	YES	YES ¹ YES ²	YES	NO	YES
Swap pair	YES ²	YES	YES ²	YES	NO	YES
Delete pair	YES	YES	YES	YES	YES	YES

Table 14-3: DP pool statuses and availability of pair operation

Pair operation	DP pool statuses, DP pool capacity statuses, and DP pool optimization statuses					
	Normal	Capacity in growth	Capacity depletion	Regressed	Blocked	DP in optimization
1. Refer to the status of the DP pool to which the DP-VOL of the S-VOL belongs. If the pair operation causes the DP pool belonging to the S-VOL to be depleted, the pair operation cannot be performed.						
2. Refer to the status of the DP pool to which the DP-VOL of the P-VOL belongs. If the pair operation causes the DP pool belonging to the S-VOL to be depleted, the pair operation cannot be performed.						

- Formatting in the DP pool

When the DP pool was created or the capacity was added, the formatting operates for the DP pool. If pair creation, pair resynchronization, or swapping is performed during the formatting, depletion of the usable capacity may occur. Since the formatting progress is displayed when checking the DP pool status, check if the sufficient usable capacity is secured according to the formatting progress, and then start the operation.

- Operation of the DP-VOL during TrueCopy use
When using the DP-VOL for a P-VOL or an S-VOL of TrueCopy, the DP pool to which the DP-VOL in use belongs cannot be deleted. To execute the operation, delete the TrueCopy pair of which the DP-VOL is in use belonging to the DP pool to be operated, and then execute it again. The attribute edit and capacity addition of the DP pool can be executed usually regardless of the TrueCopy pair.
- Operation of the DP pool during TrueCopy use
When using the DP-VOL for a P-VOL or an S-VOL of TrueCopy, the DP pool to which the DP-VOL in use belongs cannot be deleted. To execute the operation, delete the TrueCopy pair of which the DP-VOL is in use belonging to the DP pool to be operated, and then execute it again. The attribute edit and capacity addition of the DP pool can be executed regardless of the TrueCopy pair.
- Cascade connection

A cascade can be performed with the same conditions as the normal volume.

Concurrent use of Dynamic Tiering

The considerations for using the DP pool or the DP-VOL whose tier mode is enabled by using Dynamic Tiering are described. For the detailed information related to Dynamic Tiering, refer to *Hitachi Unified Storage 100 Dynamic Tiering User's Guide*. Other considerations are common with Dynamic Provisioning.

- When using a DP-VOL whose tier mode is enabled as a DMLU

When using the DP-VOL whose tier mode is enabled as DMLU, check that the free capacity (formatted) of the Tier other than SSD/FMD of the DP pool to which the DP-VOL belongs is more than or equal to the DP-VOL used as DMLU, and then set it. At the time of the setting, the entire capacity of DMLU is assigned from 1st Tier. However, the Tier configured by SSD/FMD is not assigned to DMLU. Furthermore, the area assigned to DMLU is out of the relocation target.

Load balancing function

The Load balancing function applies to a TrueCopy pair.

Enabling Change Response for Replication Mode

When write commands are being executed on the P-VOL in Paired state, if background synchronization copy is timed-out for some reason, the array returns Hardware Error (04) to the host. Some hosts receiving Hardware Error (04) may determine the P-VOL inaccessible and stop accessing it. In such cases, enabling the Change Response for Replication Mode makes the array return Aborted Command (0B) to the host. When the host receives Aborted Command (0B), it will retry the command to the P-VOL and the operation will continue.

Calculating supported capacity

Table 14-4 shows the maximum capacity of the S-VOL by the DMLU capacity in TB. The maximum capacity of the S-VOL is the total value of the S-VOL capacity of ShadowImage, TrueCopy, and Volume Migration.

The maximum capacity shown in Table 14-4 is the value smaller than the pair creatable capacity displayed in Navigator 2. That is because the pair creatable capacity in Navigator 2 is treated not as the real capacity but as the value rounded up by 1.5 TB unit, not as the actual capacity when calculating the S-VOL capacity. The maximum capacity (the capacity of which the pair can be surely created) reduced by the capacity capable of rounding up by the number of S-VOLs becomes the capacity shown in Table 14-4 .

Table 14-4: Maximum S-VOL capacity / DMLU capacity in TB

S-VOL number	DMLU Capacity				
	10 GB	32GB	64GB	96 GB	128 GB
2	256				
32	1,031	3,411	4,096		
64	983	3,363	6,327	7,200	
128	887	3,267	6,731		
512	311	2,691	6,1552		
1,024	N/A	1,923	5,387		
4,096	N/A	N/A	779	4,241	7,200

Setup procedures

The following sections provide instructions for setting up the DMLU and Remote Path.

(For CCI users, TrueCopy/CCI setup includes configuring the command device, the configuration definition file, the environment variable, and volume mapping. See [Operations using CLI on page C-5](#) for instructions.)

Setting up the DMLU


When the DMLU (differential management logical unit) is not set up prior to using TrueCopy, you must set it up. The DMLU is an exclusive volume for storing the differential data at the time when the volume is copied. The DMLU in the array is treated in the same way as the other volumes. However, a volume that is set as the DMLU is not recognized by a host (it is hidden).

Prerequisites

- Capacity from a minimum of 10 GB to a maximum of 128 GB in units of GB. Recommended size is 64 GB. The pair creatable capacity of ShadowImage, TrueCopy, and Volume Migration differs depending on the capacity.
- DMLUs must be set up on both the local and remote the disk arrays.
- One DMLU is required on each disk array; two are recommended, the second used as backup. No more than two DMLUs can be installed.
- Please also review specifications for DMLUs in [TrueCopy specifications on page C-2](#).
- Must be other than RAID 0.
- When a failure occurs in the DMLU, all the pairs of ShadowImage, TrueCopy, and/or Volume Migration are changed to Failure. Therefore, secure sufficient redundancy for the RAID group to which the DMLU is located
- In the status where the pair status is Split, Split Pending, or Reverse Synchronizing, the I/O performance of the DMLU may effect the host I/O performance to the volume which configures the pair. Using RAID 1+0 or SSD/FMD drives can decrease the effect to the host I/O performance.
- Stripe size 64 KB, 256 KB. However, when the stripe size is 256 KB, the volume in the configuration more than or equal to 17D+2P cannot be set in the DMLU
- When the volume is unified, the capacity of each unified volume becomes the average of 1 GB or more
- It is not assigned to the host
- It is not specified as the command device
- A pair is not configured
- It is not specified as a reserve volume

To define the DMLU

1. In the Navigator 2 GUI, select the disk array where you want to set up the DMLU.
2. Select the **DMLU** icon in the Setup tree view of the Replication tree view.
3. The **Differential Management Logical Units** list appears.
4. Click **Add DMLU**. The Add DMLU screen appears.



5. Select the VOL that you want to assign as DMLUs, and then click **OK**. A confirmation message appears.
6. Select the **Yes, I have read...** check box, then click **Confirm**. When a success message appears, click **Close**.

To add the DMLU

1. In the Navigator 2 GUI, select the disk array where you want to set up the DMLU.
2. Select the **DMLU** icon in the Setup tree view of the Replication tree view.
3. The **Differential Management Logical Units** list appears.

4. Select the VOL you want to add, and click **Add DMLU Capacity**. The **Add DMLU Capacity** screen appears.

Add DMLU Capacity - 1003 Help

Differential Management Logical Unit Capacity Property

Enter the information for logical unit added to a DMLU.

LUN : 1003
 Total Capacity: 64.0GB
 RAID Group: 000
 RAID Level: RAID5(4D+1P)
 Drive Type: SAS

*New Capacity: 96 GB
 To 128GB

Use RAID Group:

Assignable RAID Groups			
RAID Group	RAID Level	Capacity	
		Total	Free
000	RAID5(4D+1P)	4.0TB	4.0TB

* Required field

OK Cancel

5. Enter a capacity to the **New Capacity** and click **OK**.
6. A confirmation message appears. Click **Close**.

To remove the designated DMLU

1. In the Navigator 2 GUI, select the disk array where you want to set up the DMLU.
2. Select the **DMLU** icon in the Setup tree view of the Replication tree view.
3. The **Differential Management Logical Units** list appears.
4. Select the VOL you want to remove, and click **Remove DMLU**.
5. A confirmation message appears. Click **Close**.

Adding or changing the remote port CHAP secret (iSCSI only)

Challenge-Handshake Authentication Protocol (CHAP) provides a level of security at the time that a link is established between the local and remote disk arrays. Authentication is based on a shared secret that validates the identity of the remote path. The CHAP secret is shared between the local and remote disk arrays.

Prerequisites

- The disk array IDs for local and remote disk arrays are required.

To add a CHAP secret

This procedure is used to add CHAP authentication manually on the remote disk array.

1. On the remote disk array, navigate down the GUI tree view to **Replication / Setup / Remote Path**. The Remote Path screen appears. (Though you may have a remote path set, it does not show up on the remote disk array. Remote paths are set from the local disk array.)
2. Click the **Remote Port CHAP** tab. The Remote Port CHAP screen appears.
3. Click the **Add Remote Port CHAP** button. The Add Remote Port CHAP screen appears.
4. Enter the **Local disk array ID**.
5. Enter CHAP **Secrets** for Remote Path 0 and Remote Path 1, following on-screen instructions.
6. Click **OK** when finished.

To change a CHAP secret

1. Split the TrueCopy pairs, after confirming first that the status of all pairs is Paired.
 - To confirm pair status, see [Monitoring pair status on page 16-3](#).
 - To split pairs, see [Splitting pairs on page 15-9](#).
2. On the local disk array, delete the remote path. Be sure to confirm that the pair status is Split before deleting the remote path. See [Deleting the remote path on page 15-13](#).
3. Add the remote port CHAP secret on the remote disk array. See the instructions above.
4. Re-create the remote path on the local disk array. See [Setting the remote path](#).

For the CHAP secret field, select **manually** to enable the CHAP Secret boxes so that the CHAP secrets can be entered. Use the CHAP secret added on the remote disk array.
5. Resynchronize the pairs after confirming that the remote path is set. See [Resynchronizing pairs on page 15-10](#).

Setting the remote path

Data is transferred between the P-VOL and S-VOL on the remote path. The remote path is set up on the local disk array.

Set one path for one controller, two paths total

Figure 14-5 on page 14-24 shows the combinations of a controller and a port. As illustrated in the following figure, the combination of two CTL0s, or two CTL1s on the local and the remote side, i.e. Combination 1 is available, and the combination of CTL0 and CTL1 on the both sides, i.e. Combination 2 is available for FC only.




Figure 14-5: Combinations of a controller and a port

Prerequisites

- Two paths are recommended; one from controller 0 and one from controller 1.
- Some remote path information cannot be edited after the path is set up. To make changes, it is necessary to delete the remote path then set up a new remote path with the changed information.
- Both local and remote disk arrays must be connected to the network for the remote path.
- The remote disk array ID will be required on the GUI screen. The remote disk array ID is shown on the main disk array screen.
- Network bandwidth will be required.
- For iSCSI, the following additional information is required:
 - Remote IP address, listed in the remote disk array's GUI Settings/IP Settings. You can identify the IP address for the remote path in the IPv4 or IPv6 format. Be sure to use the same format when specifying the port IP addresses for the remote path for the local array and the remote array.
 - TCP port number. You can see this by navigating to the remote disk array's GUI's Settings/IP Settings/selected port screen.
 - CHAP secret (if specified on the remote disk array—see [Adding or changing the remote port CHAP secret \(iSCSI only\)](#) on page 14-23 for more information).

To set up the remote path

1. On the local disk array, from the navigation tree, click **Replication**, then click **Setup**. The Setup screen appears.
2. Click **Remote Path**; on the Remote Path screen click the **Create Remote Path** button. The Create Remote Path screen appears.
3. For Interface Type, select **Fibre** or **iSCSI**.
4. Enter the **Remote disk array ID**.
5. Enter the **Remote path name**.
6. Enter **Bandwidth**. Select **Over 1000,0Mbps** in the **Bandwidth** for over 1000 Mbps network bandwidth. When connecting the array directly to host's HBA, set the bandwidth according to the transfer rate.
7. (iSCSI only) In the CHAP secret field, select **Automatically** to allow TCE to create a default CHAP secret, or select **manually** to enter previously defined CHAP secrets. The CHAP secret must be set up on the remote disk array.
8. In the two remote path boxes, Remote Path 0 and Remote Path 1, select local ports. Select the port number (0E and 1E) that connected to the remote path. For iSCSI, enter the **Remote Port IP Address** and **TCP Port No.** for the remote disk array's controller 0 and 1 ports. The IPv4 or IPv6 format can be used to specify the IP address.
9. Click **OK**.

Changing the port setting

If the port setting is changed during the firmware update, the remote path may be blocked or the remote pair may be changed to Failure. Change the port setting after completing the firmware update.

If the port setting is changed in the local array and the remote array at the same time, the remote path may be blocked or the remote pair may be changed to Failure. Change the port setting by taking an interval of 30 seconds or more for every port change.

Remote path design

A remote path must be designed to adequately manage your organization's data throughput. This topic provides instructions for analyzing business requirements, measuring write-workload, and calculating your system's changed data over a given period. Remote path configurations are also provided.

- ❑ [Determining remote path bandwidth](#)
- ❑ [Remote path requirements, supported configurations](#)

Determining remote path bandwidth

Bandwidth for the TrueCopy remote path is based on the amount of production output to the primary volume. Sufficient bandwidth must be present to handle the transfer of all workload levels, from average MB/sec to peak MB/sec. Planning bandwidth also accounts for growth and a safety factor.

Measuring write-workload

To determine the bandwidth necessary to support a TrueCopy system, the peak workload must be identified and understood. Workload data is collected using performance monitoring software on your operating system. This data is best collected over a normal monthly cycle. It should include end-of month, quarter, or year, or times when workload is heaviest.

To collect workload data

1. Using your operating system's performance monitoring software, collect the following:
 - I/O per second (IOPS).
 - Disk-write bytes-per-second for every physical volume that will be replicated.

The data should be collected at 5-10 minute intervals, over a 4-6 week period. The period should include periods when demand on the system is greatest.

2. At the end of the period, convert the data to MB-per-second, if it is not already so. Import the data into a spreadsheet tool. [Figure 14-6 on page 14-29](#) shows graphed data throughput in MB per second.




Figure 14-6: Data throughput in MB/sec

Figure 14-7 shows IOPS throughput over the same period.




Figure 14-7: IOPS throughput

3. Locate the highest peak to determine the greatest MB-per-second workload.
4. Be aware of extremely high peaks. In some cases, a batch job, defragmentation, or other process could be driving workload to abnormally high levels. It is sometimes worthwhile to review the processes that are running. After careful analysis, it may be possible to lower or even eliminate some spikes by optimizing or streamlining high-workload processes. Changing the timing of a process may lower workload; another option may be to schedule a suspension of the TrueCopy pair (split the pair) when a spiking process is active.
5. With peak workload established, take into consideration the following:
 - Channel extension. Extending fibre channel over IP telecommunication links changes workloads.
 - The addition of the IP headers and conversion from fibre channel's 2112-byte frames to the 1500-byte Maximum Transfer Unit of Ethernet add approximately 10% to the amount of data transferred.
 - Compression is also a factor. The exact compression ratio is dependent on the compressibility of the data and speed of the telecommunications link. Hitachi Data Systems uses 1.8:1 as a compression rule-of-thumb, though real-life ratios are typically higher.
 - Projected growth rate accounts for the increase expected in write workload over a 1, 2, or 3 year period.
 - Safety factor adds extra bandwidth for unusually high spikes that might occur.
6. The bandwidth must be at least as large as the peak MB/sec, including channel extension overhead and compression ratios.

Optimal I/O performance versus data recovery

An organization's business requirements affect I/O performance and data recovery. Business demands indicate whether a pair must be maintained in the synchronized state or can be split. Understanding the following comparisons is useful in refining bandwidth requirements.

When a pair is synchronizing, the host holds any new writes until confirmation that the prior write is copied to both the P-VOL and the S-VOL. This assures a synchronous backup, but the resulting latency impacts I/O performance.

On the other hand, a pair in Split status has no impact on host I/O. This advantage is offset if a failure occurs. In a pair failure situation, data saved to the P-VOL while the pair is split is not copied to the S-VOL, and may therefore be lost. When performance is the priority, data recovery is impacted.

To deal with these competing demands, an organization must determine its recovery point objective (RPO), which is how much data loss is tolerable before the business suffers significant impact. RPO shows the point back in time from the disaster point that a business must recover to. A simple illustration is shown in [Figure 14-8](#). The RPO value shows you whether a pair can be split, and how long it can stay split.




Figure 14-8: Determining recovery point

- If no data loss can be tolerated, the pair is never split and remains in the synchronous, paired state.
- If one hour of data loss can be tolerated, a pair may be split for one hour.
- If 8 hours of data loss can be tolerated, a pair may be split for 8 hours.

Finding the recovery point means exploring the number of lost business transactions business can survive, determining the number of hours that may be required to key-in or otherwise recover lost data.

Performance and data recovery are also affected when the TrueCopy system is cascaded with ShadowImage. The following sections describe the impact of cascading.

Remote path requirements, supported configurations

The remote path is the connection used to transfer data between the local array and remote array. TrueCopy supports fibre channel and iSCSI port connectors and connections.

The following kinds of networks are used with TrueCopy:

- Local Area Network (LAN), for system management. Fast Ethernet is required for the LAN.
- Wide Area Network (WAN) for the remote path. For best performance:
 - A fibre channel extender is required.
 - iSCSI connections may require a WAN Optimization Controller (WOC).

Figure 14-9 shows the basic TrueCopy configuration with a LAN and WAN. More specific configurations are shown in [Remote path configurations on page 14-34](#).




Figure 14-9: Remote path configuration

Requirements are provided in the following:

- [Management LAN requirements on page 14-33](#)
- [Remote path requirements on page 14-33](#)
- [Remote path configurations on page 14-34](#)
- [Fibre Channel extender on page 14-40](#).

Management LAN requirements

Fast Ethernet is required for an IP LAN.

Remote path requirements

This section discusses the TrueCopy remote path requirements for a WAN connection. This includes the following:

- Types of lines
- Bandwidth
- Distance between local and remote sites
- WAN Optimization Controllers (WOC) (optional)

For instructions on assessing your system's I/O and bandwidth requirements, see:

- [Measuring write-workload on page 14-28](#)
- [Determining remote path bandwidth on page 14-28](#)

[Table 14-5](#) provides remote path requirements for TrueCopy. A WOC may also be required, depending on the distance between the local and remote sites and other factors listed in [Table 14-11 on page 14-46](#).

Table 14-5: Remote path requirements

Item	Requirements
Bandwidth	• Bandwidth must be guaranteed. • Bandwidth must be 1.5 Mb/s or more for each remote path. 100 Mb/s recommended. • Requirements for bandwidth depend on an average inflow from the host into the array.
Remote Path Sharing	• The remote path must be dedicated for TrueCopy pairs. • When two or more pairs share the same path, a WOC is recommended for each pair.

[Table 14-6](#) shows types of WAN cabling and protocols supported by TrueCopy and those not supported.

Table 14-6: Supported and un-supported WAN types

	WAN Types
Supported	Dedicated Line (T1, T2, T3 etc)
Not-supported	ADSL, CATV, FTTH, ISDN

Remote path configurations

One remote path must be set up per controller, two paths per array. With two paths, an alternate path is available in the event of link failure during copy operations.

Paths can be constructed from:

- Local controller 0 to remote controller 0 or 1
- Local controller 1 to remote controller 0 or 1

Paths can connect a port A with a port B, and so on.

The following sections describe Fibre channel and iSCSI path configurations. Recommendations and restrictions are included.

Remote path configurations for Fibre Channel

The Hitachi Unified Storage array supports direct and switch Fibre Channel connections only. Hub connections are not supported.

- Direct connection (loop only) is a direct link between the local and remote arrays.
- Switch connections push data from the local array through a fibre channel link across a WAN to the remote switch and fibre channel to the remote array. Switch connections increase throughput between the arrays. F-Port (Point-to-Point) and FL-Port (Loop) switch connections are supported.

Direct connection

A direct connection is a standard point-to-point fibre channel connection between ports, as shown in [Figure 14-10](#). Direct connections are typically used for systems 500 meters to 10 km apart.




Figure 14-10: Direct connection, two hosts

Recommendations

- Optimal performance occurs when the paths are connected to parallel controllers, that is, local controller 0 is connected with remote controller 0, and so on.
- Between a host and array, only one path is required. However, two are recommended, with one available for a backup.
- When connecting the local array and the remote array directly, set the transfer rate of fibre channel to the fixed rate (the same setting of 2 G bps, 4 G bps, and 8 G bps) for each array, following the table below.

Table 14-7: Transfer rates

Transfer rate of the port of the directly connected local array	Transfer rate of the port of the directly connected remote array
2 Gbps	2Gbps
4 Gbps	4 Gbps
8 Gbps	8 Gbps

- When connecting the local array and the remote array directly and setting the transfer rate to Auto, the remote path may be blocked. If the remote path is blocked, change the transfer rate of the fixed rate.

Fibre Channel switch connection 1

Switch connections increase throughput between the disk arrays.

When a pair is created on TrueCopy, it is necessary to connect two hosts with the LAN for communicating between CCIs on the host associated with the local array and on the host associated with the remote array. If there are two hosts, one host activates both CCIs on the local side and on the remote side; it is not necessary to connect two hosts with the LAN.

Figure 14-11 shows path configurations using a single switch per array.




Figure 14-11: Switch connection, 1 host, 2 hosts

Fibre Channel switch connection 2

Between a host and an array only one remote path is acceptable. If a configuration has two remote paths as illustrated in [Figure 14-12](#), a remote path can be switched when a failure in a remote path or the controller blockage occurs.

[Figure 14-12](#) shows multiple switches per array.




Figure 14-12: Multiple switches with two hosts

Recommendations

When two hosts exist, a LAN is required to provide communication between local and remote CCIs, when used. In this case, the local host activates the CCIs on the local and remote side.

The array must be connected with a switch as follows ([Table 14-8 on page 14-38](#)).

Table 14-8: Connections between Array and a Switch

Mode of array	Switch		
	For 8 G bps	For 4 G bps	For 2 G bps
Auto Mode	From the viewpoint of the performance, one path/controller between the array and a switch is acceptable, as illustrated above. The same port is available for the host I/O and for copying data of TCE.	See the left column	See the left column
8 G bps Mode			Not available
4 G bps Mode			Not available
2 G bps Mode			See the left column

One-Path-Connection between Arrays

When a pair is created on TrueCopy, it is necessary to connect two hosts with the LAN for communicating between CCIs on the host associated with the local array and on the host associated with the remote array. If there are two hosts, one host activates both CCIs on the local side and on the remote side; it is not necessary to connect two hosts with the LAN. See [Table 14-13](#).

If a failure occurs in a switch or a remote path, a remote path cannot be switched. Therefore, this configuration is not recommended.




Figure 14-13: Fibre Channel one path connection

Fibre Channel extender

Distance limitations can be reduced when channel extenders are used with fibre channel. This section provides configurations and recommendations for using channel extenders, and provides information on WDM and dark fibre.

Figure 14-14 shows two remote paths using two FC switches, Wavelength Division Multiplexor (WDM) extender, and dark fibre to make the connection to the remote site.




Figure 14-14: Fibre Channel switches, WDM, Dark Fibre Connection

Recommendations

- Two remote paths are recommended between local and remote arrays. In the event of path failure, data copying is automatically shifted to the alternate path.
- WDM has the same speed as fibre channel; however, response time increases when distance between sites increases.

For more information on WDM, see [Appendix D, Wavelength Division Multiplexing \(WDM\) and dark fibre](#).

Path and switch performance

Performance guidelines for two or more paths between hosts, switches, and arrays are shown in [Table 14-9](#).

Table 14-9: Performance guidelines for paths and switches

Array mode	Switch		
	8 G bps	4 G bps	2 G bps
Auto Mode	One path per controller between the array and a switch is sufficient for both host I/O and TrueCopy operations. (Shown in Figure 14-12 .)	Same as 8 G bps/Auto Mode.	Same as 8 G bps/Auto Mode.
8 G bps Mode			Not available
4 G bps Mode			Not available
2 G bps Mode			Same as 8 G bps/Auto Mode

Port transfer rate for Fibre Channel

The communication speed of the fibre channel port on the array must match the speed specified on the host port. These two ports—fibre channel port on the array and host port—are connected via fibre channel cables. Each port on the array must be set separately.

Table 14-10: Setting port transfer rates

	If the host port is set to	Set the array port to
Manual mode		
	2 G bps	2 G bps
	4 G bps	4 G bps
	8 G bps	8 G bps
Auto mode	2 G bps	Auto, with max of 2 G bps
	4 G bps	Auto, with max of 4 G bps
	8 G bps	Auto, with max of 8 G bps

- When using a direct connection, Auto mode may cause blockage of the data path. In this case, change the transfer rate of Manual mode. Maximum speed is ensured using the manual settings.

Specify port transfer rate in Navigator 2 GUI, on the **Edit FC Port** screen (**Settings/FC Settings/port/Edit Port** button).



NOTE: If your remote path is a direct connection, do not modify the transfer rate until after the remote pair is split. This causes remote path failure.

Find details on communication settings in the *Hitachi Unified Storage Hardware Installation and Configuration Guide*.

Remote path configurations for iSCSI

When using the iSCSI interface at the connection between the arrays, the types of cables and switches used for the Gigabit Ethernet and the 10 Gigabit Ethernet differ. For the Gigabit Ethernet, use the LAN cable and the LAN switch. For the 10 Gigabit Ethernet, use the Fibre cable and the switch usable for 10 Gigabit Ethernet.

The iSCSI remote path can be set up in the following configurations:

- Direct connection
- Local Area Network (LAN) switch connections
- Wide Area Network (WAN) connections
- WAN Optimization Controller (WOC) connections

Direct iSCSI connection

When setting the local and remote arrays in the same site at the time of TrueCopy configuration or data restoration, you can connect both arrays directly with the LAN cable. [Figure 14-15](#) shows the configuration where the arrays are directly connected with the LAN cables. One path is allowed between the host and the array. If there are two paths as illustrated in [Figure 14-15](#), a failure occurs in a path, the other path can take over

Recommendations

- Two paths should be configured from the host to the disk array. This provides a backup path in the event of path failure.
- When a large amount of data is to be copied to the remote site, the initial copy between local side and remote systems may be performed at the same location. In this case, category 5e9 or 6 copper LAN cable is recommended.




Figure 14-15: Direct iSCSI connection

Single LAN switch, WAN connection

Figure 14-16 on page 14-44 shows two remote paths using one LAN switch and network to the remote array.




Figure 14-16: Single-Switch connection

Recommendations

- This configuration is not recommended because a failure in a LAN switch or WAN would halt operations.
- Separate LAN switches and paths should be used for host-to-array and array-to-array, for improved performance.

Multiple LAN switch, WAN connection

Figure 14-17 shows two remote paths using multiple LAN switches and WANs to make the connection to the remote site.




Figure 14-17: Multiple-Switch and WAN connection

Recommendations

- Separate the switches, using one for the host I/O and another for the remote copy. If you use one switch for both host I/O and remote copy, the performance may deteriorate.
- Two remote paths should be set. When a failure occurs in one path, the data copy can continue with the other path.

Connecting the WAN Optimization Controller

The WAN Optimization Controller (WOC) is an appliance which can accelerate long-distance TCP/IP communication. WOC prevents performance degradation of TrueCopy in case that there is a long distance between the local site and the remote site. In addition, in case there are two or more pairs of the local and remote arrays sharing the same WAN, the WOC guarantees available bandwidth for each pair.

- Use [Table 14-11](#) to determine whether the TrueCopy system requires the addition of a WOC.
- [Table 14-12](#) shows the requirements for WOC.

Table 14-11: Conditions when WOC is required

Item	Condition
Latency, Distance	If round trip time is 5 ms or more, or distance between the local site and the remote site is 100 miles (160 km) or further, WOC is highly recommended.
WAN Sharing	If there are two or more pairs of the local and remote arrays sharing the same WAN, WOC is recommended for each pair.

Table 14-12: WOC requirements

Item	Requirements
LAN Interface	Gigabit Ethernet, 10 Gigabit Ethernet, or fast Ethernet must be supported.
Performance	Data transfer capability must be equal to or more than bandwidth of WAN.
Functions	• Function which drops data transfer rate to a value input by a user must be supported. The function is called shaping, throttling or rate limiting. • Data compression must be supported • TCP acceleration must be supported.

Switches and WOCs connection (1)

Figure 14-18 shows the configuration when local and remote arrays are connected via the switches and WOCs.




Figure 14-18: WOC system configuration (1)

Recommendations

- Two remote paths should be set. Using another path (a switch, WOC or WAN) for every remote path can automatically continue the data copy with another remote path when a failure occurs in one path.
- When WOC provides a port of Gigabit Ethernet or 10 Gigabit Ethernet, the switch connected directly to (for example, Port 0B and Port 1B) each array is not required. Connect port of each array to WOC directly.

Switches and WOCs connection (2)

Figure 14-19 shows a configuration example in which two remote paths use the common network when connecting the local and remote arrays via the switch and WOC.




Figure 14-19: WOC system configuration (2)

Recommendations

- Two remote paths should be set. However, if a failure occurs in the path (a switch, WOC or WAN) used commonly by two remote paths (path 0 and path 1), the paths of path 0 and path 1 are blocked. As a result, path switching is impossible and the data copy cannot be continued.
- When WOC provides two or more ports of Gigabit Ethernet or 10 Gigabit Ethernet, the switch connected directly to (for example, Port 0B and Port 1B) each array is not required. Connect port of each array to WOC directly.

Two sets of a pair connected via the switch and WOC (1)

Figure 14-20 shows a configuration when two sets of a pair of the local array and remote array exist and are connected via the switch and WOC.




Figure 14-20: Two sets of a pair connected via the switch and WOC (1)

Recommendations

- Two remote paths should be set for each array. Using another path (a switch, WOC or WAN) for every remote path can automatically continue the data copy with another remote path when a failure occurs in one path.
- When WOC provides two or more ports of Gigabit Ethernet or 10 Gigabit Ethernet, the switch connected directly to (for example, Port 0B and Port 1B) each array is not required. Connect port of each array to WOC directly.
- When the switch supports VLAN, you can make the switch connected directly to Port 0B of the local array 1 and that of the local array 2 the same switch. In this case, you add the port where Port 0B of the local array 1 is connected directly and the port where WOC1 is connected

directly to the same VLAN (hereinafter called VLAN 1). Furthermore, add the power where Port 0B of the local array 2 is connected directly and the port where WOC3 is connected directly to the same VLAN (hereinafter called VLAN2). It is necessary to make VLAN1 and VLAN2 another VLAN. Do the same for Port 1B of the local array and remote array.

Two sets of a pair connected via the switch and WOC (2)

Figure 14-21 shows a configuration example in which two sets of a pair of the local array and remote array exist and they are connected via the switch and WOC.




Figure 14-21: Two sets of a pair connected via the switch and WOC (2)

Recommendations

- When WOC provides two or more ports of Gigabit Ethernet or 10 Gigabit Ethernet, the switch connected directly to (for example, Port 0B and Port 1B) each array is not required. Connect port of each array to WOC directly.
- When the switch supports VLAN, you can make the switch connected directly to Port 0B of the local array 1 and that of the local array 2 the same switch. In this case, you add the port where Port 0B of the local array 1 is connected directly and the port where WOC1 is connected directly to the same VLAN (hereinafter called VLAN 1). Furthermore, add the port where Port 0B of the local array 2 is connected directly and the port where WOC3 is connected directly to the same VLAN (hereinafter called VLAN2). It is necessary to make VLAN1 and VLAN2 another VLAN. Do the same for Port 1B of the local array and remote array.

Using the remote path — best practices

The following best practices are provided to reduce and eliminate path failure.

- If both arrays are powered off, power-on the remote array first.
- When powering down both arrays, turn off the local array first.
- Before powering off the remote array, change pair status to Split. In Paired or Synchronizing status, a power-off results in Failure status on the remote array.
- If the remote array is not available during normal operations, a blockage error results with a notice regarding SNMP Agent Support Function and TRAP. In this case, follow instructions in the notice.

Path blockage automatically recovers after restarting. If the path blockage is not recovered when the array is READY, contact Hitachi Customer Support.

- Power off the arrays before setting or changing the fibre transfer rate

Remote processing

- Consideration for remote processing:

When a write I/O instruction received at a local site is executed at a remote site synchronously, performance attained at the remote site directly affects performance that is attained at the local site. The performance attained at the local site or the system is lowered when the remote site is overloaded, due to a large number of updates, etc.

Therefore, carefully monitor the load on the remote site as well as the local site.

- When using DP-VOL for the P-VOL and S-VOL of TrueCopy and executing I/O to the P-VOL when the pair status is Synchronizing or Paired, check that there is enough free capacity in the DP pool to which the S-VOL belongs (entire capacity of DP pool x progress of formatting - consumed capacity), and then execute it.

Although the format status of the belonging DP pool is formatting, if the DP-VOL formatting is completed, the DP-VOL can create a pair. If the pair status is Synchronizing or Paired and dual writing is executed in the S-VOL, a new area may be required for the S-VOL. However, if the required area cannot be secured in the DP pool, it must wait until the DP pool formatting is progressed, and the I/O performance may be extremely deteriorated due to the waiting time.

- In bidirectional TrueCopy Remote, the operation from HSNM2 is inhibited in both arrays where both directions are Paired and Synchronizing statuses.

In the configuration where both sites can be local and remote, when written by the host for each pair of both sites whose pair status is Paired or Synchronizing, the load of the array increases because the dual writing processing operates on the remote side of the array of both sites. In such status, when the operation is executed repeatedly at the same time from

HSNM2 for both arrays, the load of the array further increases and the I/O performance by the host may deteriorate. Therefore, when written by the host for each pair of both sites whose pair status of both sites is Paired or Synchronizing, do not execute the operation from HSNM2 at the same time for both arrays and execute the operation for each array.

Supported connections between various models of arrays

Hitachi Unified Storage can be connected with Hitachi Unified Storage, AMS2100, AMS2300, AMS2500, WMS100, AMS200, AMS500, or AMS1000.

Table 14-13 shows the supported connections between various models of arrays.

Table 14-13: Supported connections between various models of arrays

Local array	Remote Array					
	WMS100	AMS200	AMS500	AMS1000	AMS2000	Hitachi Unified Storage
WMS100	Supported	Supported	Supported	Supported	Supported	Supported
AMS200	Supported	Supported	Supported	Supported	Supported	Supported
AMS500	Supported	Supported	Supported	Supported	Supported	Supported
AMS1000	Supported	Supported	Supported	Supported	Supported	Supported
AMS2000	Supported	Supported	Supported	Supported	Supported	Supported
Hitachi Unified Storage	Supported	Supported	Supported	Supported	Supported	Supported

Restrictions on supported connections

- The maximum number of pairs that can be created is limited to the maximum number of pairs supported by the arrays, whichever is fewer.
- The firmware version of WMS100, AMS200, AMS500, or AMS1000 must be 0787/A or later when connecting with Hitachi Unified Storage.
- If a Hitachi Unified Storage as the local array connects to an AMS2010, AMS2100, AMS2300, or AMS2500 with under 08B7/A as the remote array, the remote path will be blocked along with the following message:
 - For Fibre Channel connection:
The target of remote path cannot be connected(Port-xy)
Path alarm(Remote-X,Path-Y)
 - For iSCSI connection:
Path Login failed
- The firmware version of AMS2010, AMS2100, AMS2300, or AMS2500 must be 08B7/A or later when connecting with Hitachi Unified Storage.
- The bandwidth of the remote path to WMS100, AMS200, AMS500, or AMS1000 must be 20 Mbps or more.
- The pair operation of WMS100, AMS200, AMS500, or AMS1000 cannot be done from Navigator 2.
- WMS100, AMS200, AMS500, or AMS1000 cannot use the functions that are newly supported by Hitachi Unified Storage.

Using TrueCopy Remote

This chapter provides procedures for performing basic TrueCopy operations using the Navigator 2 GUI. Appendixes with CLI and CCI instructions for the same operations are included in this manual.

- ❑ [TrueCopy operations](#)
- ❑ [Pair assignment](#)
- ❑ [Checking pair status](#)
- ❑ [Creating pairs](#)
- ❑ [Splitting pairs](#)
- ❑ [Resynchronizing pairs](#)
- ❑ [Swapping pairs](#)
- ❑ [Editing pairs](#)
- ❑ [Deleting pairs](#)
- ❑ [Deleting the remote path](#)
- ❑ [Operations work flow](#)
- ❑ [TrueCopy ordinary split operation](#)
- ❑ [TrueCopy ordinary pair operation](#)
- ❑ [Data migration use](#)
- ❑ [TrueCopy disaster recovery](#)
- ❑ [Resynchronizing the pair](#)
- ❑ [Data path failure and recovery](#)

TrueCopy operations

Basic TrueCopy operations consists of the following. See [TrueCopy disaster recovery on page 15-19](#) for disaster recovery procedures.

- Always check pair status. Each operation requires the pair to be in a specific status.
- Create the pair, in which the S-VOL becomes a duplicate of the P-VOL.
- Split the pair, which separates the P-VOL and S-VOL and allows read/write access to the S-VOL.
- Re-synchronize the pair, in which the S-VOL again mirrors the on-going, current data in the P-VOL.
- Swap pairs, which reverses pair roles.
- Delete a pair.
- Edit pair information.

Pair assignment

- Do not assign a volume (required for a quick response to a host) to a pair.

For a TrueCopy pair, data written to a P-VOL is also written to an S-VOL at a remote site synchronously. Therefore, performance of a write operation instructed by a host is lowered according to the distance to the remote site. Select the TrueCopy pair carefully. Observe the matter described above, particularly when a volume required for a high-performance response is required

- Assign a small number of volumes within the same RAID group.

When volumes are assigned to the same RAID group and used as pair volumes, pair creation or resynchronization of one volume affects the performance of a host I/O, pair creation, and/or resynchronization of the other pair, so that the performance may be restricted due to drive contention. Therefore, it is best to assign a small number (one or two) of volumes to be paired to the same RAID group.

- For a P-VOL, use the SAS drives or SSD/FMD drives.

When a P-VOL is located in a RAID group consisting of the SAS7.2K drives, performance of a host I/O, pair creation, and pair resynchronization, etc., is lowered because of the lower performance of the SAS7.2K drives. Therefore, it is recommended to assign a P-VOL to a RAID group consisting of the SAS drives or SSD/FMD drives.

- Assign four or more disks to the data disks.

When the data disks that compose a RAID group are not sufficient, it affects the host performance and/or copying performance adversely because reading/writing from/to the drives is restricted. Therefore, when operating pairs with ShadowImage, it is recommended that you use a volume consisting of four or more data disks.

- When using the SAS7.2K drives, make the data disks between 4D and 6D

When the number of data disks, which configures a RAID group, is large in the case where the SAS7.2K drives are used, the copying performance is affected. Therefore, it is recommended to use a volume with the number of data disks between 4D and 6D for the TrueCopy volume in the case where the SAS7.2K drives are used.

- When cascading TrueCopy and Snapshot pairs, assign a volume of the SAS drives or the SSD/FMD drives and assign four or more disks to a DP pool.

When TrueCopy and Snapshot are cascaded, performance of the drives composing a DP pool influences the performance of the host operation and copying. Therefore, it is best to assign a volume of SAS drives or SSD/FMD drives and assign four or more disks (which have higher performance) than SAS7.2K drives, to a DP pool.

Checking pair status

Each TrueCopy operation requires a specific pair status. Before performing any operation, check pair status.

- Find status requirements for each operation under the Prerequisites sections.
- To view a pair's current status in the GUI, please refer to [Monitoring pair status on page 16-3](#).

Pairs Operations

TrueCopy operations can be performed from the UNIX/PC host using CCI software and/or Navigator 2.

Confirm that the state of the remote path is **Normal** before doing a pair operation. If you do a pair operation when a remote path is not set up or its state is **Diagnosing** or **Blocked**, the pair operation may not be completed correctly.

Creating pairs

A TrueCopy pair consists of primary and secondary volume whose data stays synchronized until the pair is split. During the create pair operation, the following takes place:

- All data in the local P-VOL is copied to the remote S-VOL.
- The P-VOL remains available to the host for read/write throughout the copy operation.
- Pair status is Synchronizing while the initial copy operation is in progress.
- Status changes to Paired when the initial copy is complete.

- New writes to the P-VOL continue to be copied to the S-VOL in the Paired status.

Prerequisite information and best practices

- In the remote array, create a volume with the same capacity as that of a volume to be backed up in the local array
- Logical units for volumes to be paired must be in Simplex status.
- DMLUs must be set up.
- The create pair and resynchronize operations affect performance on the host. Therefore:
 - Perform the operation when I/O load is light.
 - Limit the number of pairs that you create simultaneously within the same RAID group to two.
 - If a TrueCopy pair is cascaded with ShadowImage, and the pair of one or the other is in Paired or Synchronizing status, place the other in Split status to lower the impact on performance.
 - If you have two TrueCopy pairs on the same two disk arrays and the pairs are bi-directional, perform copy operations at different times to lower the impact on performance.
 - Monitor write-workload on the remote disk array as well as on the local disk array. Performance on the remote disk array affects performance on the local disk array, since TrueCopy operations are slowed down by unrelated remote operations. Performance backup reverberates across the two systems.
 - Use a copy pace that matches your priority for either performance or copying speed.

The following sections discuss options in the Create Pair procedure.

Copy pace

Copy pace is the speed at which data is copied during pair creation or re-synchronization. You select the copy pace on the GUI procedure when you create or resync a pair (if using CLI, you enter a copy pace parameter).

Copy pace impacts host I/O performance. A slow copy pace has less impact than a medium or fast pace. The pace is divided on a scale of 1 to 15 (in CCI only), as follows:

- Slow — between 1-5. The process takes longer when host I/O activity is heavy. The amount of time to complete an initial copy or resync cannot be guaranteed.
- Medium — between 6-10. (Recommended) The process is performed continuously, but the amount of time to complete the initial copy or resync cannot be guaranteed. Actual pace varies according on host I/O activity.
- Fast — between 11-15. The copy/resync process is performed continuously and takes priority. Host I/O performance is restricted. The amount of time to complete an initial copy or resync is guaranteed.

You can change the copy pace which was once set later by using the edit function. You may change it when you feel that the creation time takes a long time in the pace specified at the time of the creation or the effect on the host I/O is significant because the copy processing is given priority.

Fence level

The Fence Level setting determines whether the host is denied access to the P-VOL if a TrueCopy pair is suspended due to an error.

You must decide whether you want to bring the production application(s) to a halt if the remote site is down or inaccessible.

There are two synchronous fence-level settings:

- **Never** – The P-VOL will never be fenced. “Never” ensures that a host never loses access to the P-VOL, even if all TrueCopy copy operations are stopped. Once the failure is corrected, a full re-copy may be needed to ensure that the S-VOL is current. This setting should be used when I/O performance out-weighs data recovery.
- **Data** – The P-VOL will be fenced if an update copy operation fails. “Data” insures that the S-VOL remains identical to the P-VOL. This is done by preventing the host from writing to the P-VOL during a failure. This setting should be used for critical data.

Operation when the fence level is “never”

The file systems for UNIX and Windows Server do not have the writing log (or the journal file). Even though "data" is set for the fence level, the file sometimes does not correspond to the directory. Therefore, "never" is set for the fence level.

In this case, the data of S-VOL is used after **fsck** and **chkdsk** are executed. The data, however, is not guaranteed completely. Therefore, we recommend a configuration which saves the complete data in the P-VOL or the S-VOL that is cascaded by using ShadowImage on the remote side.

Creating pairs procedure

To create a pair

1. In the Navigator 2 GUI, click the check box for the local disk array, then click the **Show & Configure disk array** button.
2. Select the **Remote Replication** icon in the **Replication** tree view. The Remote Replication screen appears.
3. Click the **Create Pair** button. The Create Pair screen appears.

Create Pair Help

Remote Pair Properties

Enter property information as desired.

CopyType : TrueCopy

Basic **Advanced**

Specify the primary volume and the secondary volume to pair and input a unique pair name. And specify a group to assign after pair creation.

Pair Name :

31 characters or less (alphanumeric characters, special symbols "%", "*", "+", "-", ".", ":", "/", "=", "@", "_", "!", "[", or "]")

Remote Array : **Assignable Remote Arrays**

Remote Array ID	Remote Path Name
91200026	Array_91200026

* Select a Primary Volume : **Primary Volume**

VOL	Capacity	RAID Group	DP Pool
0001	50.0GB	N/A	000
0002	10.0GB	000	N/A

* Secondary Volume :

After pair creation, add the pair to a group. Group Assignment:

☒ (Ungrouped)

☐ New or existing Group Number :

from 0 to 255

☐ Existing Group Name :

31 characters or less (alphanumeric characters, special symbols "%", "*", "+", "-", ".", ":", "/", "=", "@", "_", "!", "[", or "]")

Available Pair Capacity : 0.0MB

* Required field

4. Enter a **Pair Name**, if desired. Omitting a Pair Name, the default Pair Name (TC_LUxxxxx_LUyyyyy: xxxx is Primary Volume, yyyy is Secondary Volume) is created (which can be changed via the Edit Pair screen).
5. Select a Primary Volume. To display all volumes, use the scroll buttons. VOL may be different from H-LUN, which is recognized by the host. Confirm the mapping of VOL and H-LUN.
6. In the **Group Assignment** area, you have the option of assigning the new pair to a group. (For a description, see [Consistency group \(CTG\) on page 12-6](#).) Do one of the following:
 - If you do not want to assign the pair to a group, leave the **Ungrouped** button selected.

- To create a group and assign the new pair to it, click the **New or existing Group Number** button and enter a new number for the group in the box.
- To assign the pair to an existing group, enter its number in the **Group Number** box, or enter the group name in the **Existing Group Name** box.



NOTE: When a group is created, future pairs can be added to it. You can also name the group on the Edit Pair screen. See [Editing pairs on page 15-11](#) for details.

7. Click the **Advanced** tab.

8. From the **Copy Pace** dropdown list, select the speed at which copies will be made. Select Slow, Medium, or Fast. See [Copy pace on page 15-4](#), for more information.
9. In the **Do initial copy from the primary volume...** field, leave **Yes** checked to copy the primary to the secondary volume.
Clear the check box to create a pair without copying the P-VOL at this time. Do this when the S-VOL is already a copy of the P-VOL.
10. Select a **Fence Level** of Never or Data. See [Fence level on page 15-6](#) for more information.
11. Click **OK**.
12. Check the **Yes, I have read...** message then click Confirm.
13. When the success message appears, click **Close**.

Splitting pairs

All data written to the P-VOL is copied to the S-VOL when the pair is in Paired status. This continues until the pair is split. Then, updates continue to be written to the P-VOL but not the S-VOL. Data in the S-VOL is frozen at the time of the split, and the pair is no longer synchronous.

When a pair is split:

- Data copying to the S-VOL is completed so that the data is identical with P-VOL data. The time it takes to perform the split depends on the amount of differential data copied to the S-VOL.
- If the pair is included in a group, all pairs in the group are split.

After the Split Pair operation:

- The secondary volume becomes available for read/write access by secondary host applications.
- Separate track tables record updates to the P-VOL and to the S-VOL.
- The pair can be made identical again by re-synchronizing the pair.

Prerequisites

The pair must be in Paired status.

To split the pair

1. In Navigator 2 GUI, select the desired disk array, then click the **Show & Configure disk array** button.
2. From the Replication tree, select the **Remote Replication** icon. The **Pairs** screen appears.
3. Click the check box for the pair you want to split, then click **Split Pair**. The Split Pair screen appears.



4. Select the **Option** you want for the S-VOL, Read/Write, which makes the S-VOL available to be written to by a secondary application, or Read Only, which prevents it from being written to by a secondary application.
5. Click **OK** and **Close**.

Resynchronizing pairs

Re-synchronizing a pair that has been split updates the S-VOL so that it is again identical with the P-VOL. Differential data accumulated on the local disk array since the last pairing is updated to the S-VOL.

- Pair status during a re-synchronizing is Synchronizing.
- Status changes to Paired when the resync is complete.
- If P-VOL status is Failure and S-VOL status is Takeover or Simplex, the pair cannot be recovered by resynchronizing. The pair must be deleted and created again.

Prerequisites

- The pair must be in Split status.
- The prerequisites for creating a pair apply to resynchronizing. See [Creating pairs on page 15-3](#).

To resync the pair

1. In Navigator 2 GUI, select the desired disk array, then click the **Show & Configure disk array** button.
2. From the Replication tree, select the **Remote Replication** icon. The **Pairs** screen appears.
3. Select the pair you want to resync.
4. Click the **Resync Pair** button. View further instructions by clicking the Help button, as needed.

Swapping pairs

In a pair swap, primary and secondary-volume roles are reversed. Data flows from remote to local or new disk array.

A pair swap is performed when data in the S-VOL must be used to restore the local disk array, or possibly to a new disk array/volume following disaster.

The swap operation can swap the paired pairs (Paired), the split pairs (Split), the suspended pairs (Failure), or the takeover pairs (Takeover).

Prerequisites and Notes

- To swap the pairs, the remote path must be set for the local array from the remote array.
- You can swap the pairs whose statuses are Paired, Split, or Takeover
- The pair swap is executed by the command to remote array. Confirm that the target of the command is remote array.
- As long as swap is performed from Navigator 2 on the remote array, no matter how many times swap is performed, the copy direction will not return to the original direction (P-VOL on the local array and S-VOL on the remote array).
- When the pair is swapped, P-VOL pair status changes to **Failure**.

To swap TrueCopy pairs

1. In Navigator 2 GUI, connect to the remote disk array, then click the **Show & Configure disk array** button.
2. From the Replication tree, select the **Remote Replication** icon. The **Pairs** screen appears.
3. Select the pair you want to swap.
4. Click the **Swap Pair** button.
5. On the message screen, check the **Yes, I have read...** box, then click **Confirm**.
6. Click **Close** on the confirmation screen.

Editing pairs

You can edit the name, group name, and copy pace for a pair. A group created with no name can be named from the Edit Pair screen.

To edit pairs

1. In Navigator 2 GUI, select the desired disk array, then click the **Show & Configure disk array** button.
2. From the Replication tree, select the **Remote Replication** icon. The **Pairs** screen appears.
3. Select the pair that you want to edit.
4. Click the **Edit Pair** button.
5. Make any changes, then click **OK**.

6. On the confirmation message, click **Close**.



NOTE: Edits made on the local disk array are not reflected on the remote disk array. To have the same information reflected on both disk arrays, it is necessary to edit the pair on the remote disk array also.

Deleting pairs

When a pair is deleted, transfer of differential data from P-VOL to S-VOL is completed, then the volumes become Simplex. The pair is no longer displayed in the Remote Replication pair list on Navigator 2 GUI.

- A pair can be deleted regardless of its status. However, data consistency is not guaranteed unless status prior to deletion is Paired.
- If the operation fails, the P-VOL nevertheless becomes Simplex. Transfer of differential data from P-VOL to S-VOL is terminated.
- Normally, a Delete Pair operation is performed on the local disk array where the P-VOL resides. However, it is possible to perform the operation from the remote disk array, though with the following results:
 - Only the S-VOL becomes Simplex.
 - Data consistency in the S-VOL is not guaranteed.
 - If during the pair deletion from the local array, if only the P-VOL becomes Simplex and only the S-VOL remains in the remote array, perform the pair deletion from the remote array.
 - The P-VOL does not recognize that the S-VOL is in Simplex status. When the P-VOL tries to send differential data to the S-VOL, it recognizes that the S-VOL is absent and the pair becomes Failure.
 - When the pair status changes to Failure, the status of the other pairs in the group also becomes Failure.
 - From the remote disk array, this Failure status is not seen and pair status remains Paired.
 - When executing the pair deletion in the batch file or the script, insert a five-second wait before executing the next processing step.
 - Pair creation of TrueCopy which specified the volume specified as the S-VOL of the deleted pair
 - Pair creation of Volume Migration which specified the volume specified as the S-VOL of the deleted pair
 - Deletion of the volume specified as the S-VOL of the deleted pair
 - Shrinking of the volume specified as the S-VOL of the deleted pair
 - Removing of the DMLU
 - Expanding capacity of the DMLU

An example batch file with a five-second wait is:

```
ping 127.0.0.1 -n 5 > nul
```

To delete a pair

1. In Navigator 2 GUI, select the desired disk array, then click the **Show & Configure disk array** button.
2. From the Replication tree, select the **Remote Replication** icon. The **Pairs** screen appears.
3. Select the pair you want to delete in the Pairs list, and click **Delete Pair**.
4. On the message screen, check the **Yes, I have read...** box, then click **Confirm**.
5. Click **Close** on the confirmation screen.

Deleting the remote path

When the remote path becomes unnecessary, delete the remote path.

Prerequisites

- The pair status of the volumes using the remote path to be deleted must be Simplex or Split
- Do not do a pair operation for a TrueCopy pair when the remote path for the pair is not set up, because the pair operation may not complete correctly.



NOTE: When performing the planned shutdown of the remote array, the remote path should not necessarily be deleted. Change all the TrueCopy pairs in the array to the Split status, and then perform the planned shutdown of the remote array. After restarting the array, perform the pair resynchronization. However, when the Warning notice to the Failure Monitoring Department at the time of the remote path blockade or the notice by the SNMP Agent Support Function or the E-mail Alert Function is not desired, delete the remote path, and then turn off the power of the remote array.

To delete the remote path

1. Connect the local array, and select the **Remote Path** icon in the **Setup** tree view in the **Replication** tree. The **Remote Path** list appears.
2. Select the remote path you want to delete in the **Remote Path** list and click **Delete Path**.
3. A message appears. Click **Close**.

Operations work flow

TrueCopy is a function for the synchronous remote copy between the volumes in the arrays connected via the remote path. Using TrueCopy, the data received from the host is written into arrays on local and remote sides simultaneously. The data of arrays on local and remote sides are always synchronized. When the resynchronization is executed, you can save time by transferring the differential data only to the array on the remote side.



NOTE: When turning the power of the array on or off, restarting the array, replacing the firmware, changing the transfer rate setting, and changing the remote array, be careful of the following items.

- When you turn on the array where a path has already been set, turn on the remote array first. Turn on the local array after the remote array is READY.
When you turn off the array where a path has already been set, turn off the local array first, and then turn off the remote array.
- When you restart the array, verify that the array is on the remote side of TrueCopy. When the array on the remote side is restarted, both paths are blocked.
- When the array on the remote side is powered off or restarted when the TrueCopy pair status is Paired or Synchronizing, it is changed to Failure. If you power off or restart the array, do so after changing the TrueCopy pair status to Split. When the power of the remote array is turned off or restarted, the remote path is blocked. However, by changing the pair status to Split, you can prevent the pair status changing to Failure. When the Warning notice to the Failure Monitoring Department at the time of the remote path blockade or the notice by the SNMP Agent Support Function or the E-mail Alert Function is not desired, delete the remote path, and then turn off the power of the remote array.
- You will receive an error/blockage if the remote array is not available. A notice regarding SNMP Agent Support Function and TRAP occurs in path blockade mode.
- Perform the functions in the notice and the check the Failure Monitoring Department in advance. Path blockade automatically recovers after restarting. If the path blockage is not recovered when the array is READY, contact Hitachi Customer Support.
- When there is a fibre channel interface, the local array is directly connected with the remote array paired with the local array, and the setting of the fibre transfer rate must not be modified while the array power is on. If the setting of the fibre transfer rate is modified, a path blockage will occur.
- When you replace the remote array with a different array, ensure that you have deleted all the TrueCopy pairs and the remote path before changing the remote array.

TrueCopy ordinary split operation

When the host I/O is lightly executed (for example, at night) resync is executed. After resync, I/O operation of the database on the local side is stopped and the pair is split (see [Figure 15-1](#)).




Figure 15-1: TrueCopy ordinary Split operation

If the volumes are cascaded via ShadowImage on the remote side and resync of TrueCopy is executed after volumes on ShadowImage split, the data is saved in the secondary side of ShadowImage even though a failure occurs during the resync of TrueCopy (see [Figure 15-2](#).)




Figure 15-2: TrueCopy ordinary split operation

During the Split status of TrueCopy, the host I/O performance does not deteriorate because the host does not write to the remote side. Since the data when TrueCopy is split after resync is not saved, the data from when a pair of TrueCopy is split to when a failure occurs is not saved. Therefore, we recommended this operation to the user who attaches importance to the host I/O performance.

TrueCopy ordinary pair operation

If you want to back up the data and a pair of TrueCopy pairs are split, the backup data at the time remains on the remote side. By cascading volumes with ShadowImage on the remote side, the data when a pair of Shadow Images are split is saved on the secondary side of ShadowImage (see [Figure 15-3](#)).



NOTE: When a copy operation on ShadowImage and TrueCopy is performed, the copy prior mode is recommended.




Figure 15-3: TrueCopy ordinary pair operation

The performance of the write operation from the host during the pair status deteriorates because the array gives a “finish” response to the host after writing to the array on the remote side. Therefore, we recommend this operation to the user who attaches importance to data recovery when failure occurs.

Data migration use

If you want to use the data on the local side from the remote side and the TrueCopy pairs are split directly after the resync operation is performed, the data of the local side at this point in time remains on the remote side (see [Figure 15-4](#)).



NOTE: The copy operations are performed in the copy prior mode.




Figure 15-4: Data migration use

TrueCopy disaster recovery

TrueCopy is designed for you to create and maintain a viable copy of the production volume at a remote location. If the data path or local host or disk array go down, data in the remote copy is available for restoring the primary volume, or for host operations to be switched to the remote site to continue operations.

This topic provides procedures for disaster recovery under varying circumstances.

This chapter describes restoration procedures without and with a ShadowImage or Snapshot backup. Four scenarios are presented with recovery procedures.



NOTE: In the following procedures, ShadowImage or Snapshot pairs are cascaded with TrueCopy and are referred to as the “backup pair”.

Also, CCI is located on a host management server, and the production applications are located on a host production server.

- ☐ [Resynchronizing the pair](#)
- ☐ [Data path failure and recovery](#)
- ☐ [Host server failure and recovery](#)
- ☐ [Production site failure and recovery](#)
- ☐ [Automatic switching using High Availability \(HA\) software](#)
- ☐ [Manual switching](#)
- ☐ [Special problems and recommendations](#)

Resynchronizing the pair

In this scenario, an error has occurred in the replication environment, which causes the TrueCopy pair to suspend with an error. To recover TrueCopy pair operations, proceed as follows:

1. If ShadowImage or Snapshot backups exist on the remote array, re-synchronize the backup pair.
2. Split the backup pair.
3. Confirm that the remote path is operational.
4. Resynchronize the TrueCopy pair.

Data path failure and recovery

In this scenario, a power outage has disabled the local site and the remote path. A decision is made to move host applications to the remote site, where the TrueCopy S-VOL will become the primary volume. When the data path is unavailable, CCI cannot communicate with the opposite server, and the horctakeover command cannot perform the pair resync/swap operation.

1. Shutdown applications on the local production site.
2. Execute the CCI horctakeover command from the remote site. This transfers access to the remote array. The S-VOL status becomes SSWS, which is read/write enabled.
3. At this time, horctakeover attempts a "swaps resync". This fails because the data path is unavailable.
4. Execute the CCI pairdisplay -fc command to confirm the SSWS status.
5. At the remote site, mount the volumes and bring up the applications.
6. When the local site and/or data path is again operational, prepare to re-establish TrueCopy operations, as follows:
 - a. If ShadowImage or Snapshot are used, re-synchronize the backup pair on the remote array. Use the Navigator 2 GUI or CLI to do this if the production server is not available to run the CCI command.
 - b. Confirm that the backup pair status is Paired, then split the pair.
 - c. Create a data path, from the remote array to the local array. See [Defining the remote path on page 5-5](#) for instructions. Confirm that the data path is operational.
7. Perform the CCI pairresync -swaps command at the remote site. This completes the reversal of the PVOL-SVOL relationship. TrueCopy operations now flow from the remote array to the local array.
8. Execute the CCI pairdisplay command to confirm that the P-VOL pair is on the remote array.

The remote site is now the production site.

Next, the production environment is returned to its normal state at the local site. Applications should be moved back to the local production site in a staged manner. The following procedure may be performed hours, days, or weeks after the completion of the previous steps.

9. Shut down the application(s) on the remote server then unmount the volumes.
10. Boot the local production servers. DO NOT MOUNT the volumes or start the applications.
11. Execute the CCI `horctakeover` command on the local management server. Because the data path is operational, this command includes the pair resync/swap operation, which reverses TrueCopy pair roles. The S-VOL becomes read/write disabled.
12. Execute the CCI `pairdisplay` command to confirm that the P-VOL pair is now on the local array.
13. Mount the volumes and start the applications on the local production server.

Host server failure and recovery

In this scenario, the host server fails and no local standby server is available. A decision is made to move the application to the remote site. The data path and local array are still operational.

1. Confirm that the remote server is available and that all CCI `horcm` instances are started on it.
2. Ensure that a live replication path (remote path) from the remote to the local array is in place.
3. Execute the CCI command `horctakeover` from the remote server. Because the data path link is operational, this command includes pair resync/swap, which reverses TrueCopy pair roles.
4. Execute the CCI command `pairdisplay` to confirm that the P-VOL pair is on the remote site.
5. Mount the volumes and start the applications.

The remote site is now the production site.

At a later date, the host server is restored, and the applications are moved back to the production site in a staged manner. The following procedure may be performed hours, days, or weeks after the completion of the previous steps.

6. Shutdown applications and unmount volumes on remote server
7. Boot the local production server. DO NOT mount the volumes or start the applications.
8. Execute the CCI `horctakeover` command from the local management server. Because the data path is operational, this command includes the pair resync/swap operation, which reverses TrueCopy pair roles. The S-VOL becomes read/write disabled.
9. Execute the CCI `pairdisplay` command to confirm that the P-VOL pair is on the local site.
10. Mount the volumes and start the applications.

Host timeout

It is recommended that you set more than 60 seconds for the I/O timeout from the host to the array.

Production site failure and recovery

In this scenario, the production site is unavailable and a decision is made to move all operations to the remote site.

1. Confirm that the necessary remote servers are available and that all CCI `horcm` instances are started.
2. Execute the CCI `horctakeover` command from the remote server.
3. Execute the CCI `pairdisplay -fc` command to confirm that the P-VOL pair is on the remote site, or, if the data path is unavailable, that the S-VOL is in SSWS status, which enables read/write.
4. Mount the volumes and start the applications on the remote site.

The remote site is now the production site.

When the local site is again operational, TrueCopy operations can be restarted.

5. Make sure the data path is deleted on the local array.
6. If ShadowImage is cascaded on the original remote site, re-synchronize the pair.
7. Split the ShadowImage pair.
8. Set up the remote path, on the remote array.
9. Execute the CCI `pairresync -swaps` command from the remote management server.
10. If Step 9 fails, you may need to recreate the paths and delete and create new TrueCopy pairs, which will require full synchronization on all volumes. You may need to contact the support center.
11. When TrueCopy operations are normal and data is flowing from the remote P-VOL to the local S-VOL, production operations can be switched back to the local side again. To do this, start at Step 6 in the [Host server failure and recovery on page 15-21](#).

Automatic switching using High Availability (HA) software

If both the host and the disks (disk array) on the local side are destroyed by a disaster condition (for example, an earthquake), data recovery and continuing operation are executed by the stand-by host and the disks on the remote side.

By installing the High Availability (HA) Software on both the local side and the remote side, when a failure occurs in the host or in the array, although the input and output operation from the host on the local side cannot be executed, operation will continue by automatically switching over to the stand-by host on the remote side. A configuration is shown in [Figure 15-5 on page 15-23](#).




Figure 15-5: Configuration for failover

When a failure occurs on the local side, a host is switched into the stand-by host on the remote side (failover) by HA software. Automatically executing the script of the recovery process in the stand-by host on the remote side enables the host on the remote side to continue the operation.



NOTE: Several minutes are needed for the switching process.

The recovery processes for the database is:

1. Issuing the **takeover** command of CCI from the stand-by host makes the stand-by host possible to access the disk on the remote side
2. By using **REDO** log, the data of the database is recovered.

The recovery processes for the file system is:

1. Issuing the **takeover** command of CCI from the stand-by host on the remote side makes the stand-by host possible to access the disk on the remote side.

2. For UNIX, the file of *fsck* and Windows Server is recovered by executing *chkdsk*.

Manual switching

As mentioned in this section, if the host on the local side can access disk arrays on both the local side and the remote side via the fibre channel, the stand-by host on the remote side can be OFF. If the host on the local side can access the disks on both the local side and the remote side, it is not necessary to connect with the host on the local side and the stand-by host on the remote side with LAN (see Figure 5.5)




Table 15-1: Disaster recovery

When a failure occurs on the local side, executing the script for the recovery process enables continuous operation by the stand-by host.

The recovery processes for the database is:

1. Issuing the **takeover** command of CCI from the stand-by host makes the stand-by host possible to access the disk on the remote side
2. By using **REDO** log, the data of the database is recovered.

The recovery processes for the file system is:

1. Issuing the **takeover** command of CCI from the stand-by host on the remote side makes the stand-by host possible to access the disk on the remote side.
2. For UNIX, the file of fsck and Windows is recovered by executing chkdisk.

Special problems and recommendations

- Before path management software (such as Hitachi Dynamic Link Manager) handles a recovery, the remote path must be manually recovered and the pair resynchronized. Otherwise the management software freezes when attempting to write to the P-VOL.
- The file systems for UNIX and Windows 2000 Server /Windows Server 2003/2008 do not have writing logs or journal files. Though “data” may be specified as the fence level, the file sometimes does not correspond to the directory. In this case, “fsck” and “chkdisk” must be executed before the S-VOL can be used, though data consistency cannot be completely guaranteed. The work-around is to cascade TrueCopy with ShadowImage, which saves the complete P-VOL or S-VOL to the ShadowImage S-VOL.

For more information on performing system recovery using CCI, see the *Hitachi Unified Storage Command Control Interface Installation and Configuration Guide*.

For information on cascading with ShadowImage or Snapshot, see [Cascading ShadowImage on page 27-2](#) and [Cascading Snapshot on page 27-19](#).

Monitoring and troubleshooting TrueCopy Remote

This chapter provides information and instructions for monitoring and troubleshooting the TrueCopy system.

- ❑ [Monitoring and maintenance](#)
- ❑ [Monitoring pair failure](#)
- ❑ [Troubleshooting](#)

Monitoring and maintenance

Monitoring the TrueCopy system is an ongoing operation that should be performed to maintain your pairs as intended.

- When you want to perform a pair command, first check the pair's status. Each operation requires a specific status or statuses.
- Pair-status changes when an operation is performed. Check status to see that TrueCopy pairs are operating correctly and that data is updated from P-VOLs to S-VOLs in the Paired status, or that differential data management is performed in the Split status.

When a hardware failure occurs, a pair failure may occur as a result. When a pair failure occurs, the processes in [Table 16-1](#) are executed:

Table 16-1: Processes when pair failure occurs

Management Software	Results
Navigator 2	A message is displayed in the event log.
	The pair status is changed to Failure.
CCI	An error message is output to the system log file, as shown in Table 16-2 .
	The pair status is changed to PSUE. (On UNIX the syslog file appears; on Windows 2000, the eventlog file display.)
SNMP Agent Support Function	A trap is reported.

Table 16-2: CCI system log message when PSUE

Message ID	Condition	Cause
HORCM_102	The volume is suspended in code 0006.	The pair status was suspended due to code 0006.

See the maintenance log section in the *Hitachi Unified Storage Command Control Interface Installation and Configuration Guide* for more information.

A sample script is provided for CLI users in [Operations using CLI on page C-5](#).

Monitoring pair status

The pair status changes as a result of operations on the TrueCopy pair. You can find out how an array is controlling the TrueCopy pair from the pair status. You can also detect failures by monitoring the pair status. Figure 16-1 shows the pair status transitions.

The pair status of a pair with the reverse pair direction (S-VOL to P-VOL) changes in the same way that the pair with the original pair direction (P-VOL to S-VOL) changes. Look at the figure with the reverse pair direction in mind. Once the resync copy completes, the pair status changes to Paired.





Figure 16-1: Pair status transitions and operations

Monitoring using the GUI is done at the user's discretion. Monitoring should be repeated frequently. Email notifications of problems can be set up using the GUI.

To monitor pair status using the GUI

1. In Navigator 2 GUI, select the desired disk array, then click the **Show & Configure disk array** button.
2. From the Replication tree, select the **Remote Replication** icon in the Replication tree view.



The **Pairs** screen appears.

Name	Local VOL	Attribute	Remote Array ID	Remote Path Name	Remote VOL	Status	DP Pool	Replication Data	Management Area
TC_LU0000_LU0000	0000	Primary	91200027	Array_91200027	0000	Paired(0%)	N/A	N/A	N/A

- Name: The pair name is displayed.
- Local VOL: The local side VOL is displayed.
- Attribute: The volume type (Primary or Secondary) is displayed.
- Remote Array ID: The remote array ID is displayed.
- Remote Path Name: The remote path name is displayed.
- Remote VOL: The remote side VOL is displayed.
- Status: The pair status is displayed. For each pair status meaning, see [Table 16-3 on page 16-5](#). The percentage denotes the progress rate (%) when the pair status is Synchronizing. When the pair status is Paired, it denotes the coincidence rate (%) of a P-VOL and an S-VOL.

When the pair status is Split, it denotes the coincidence rate (%) of current data and the data at the time of pair splitting.

- DP Pool:
 - Replication Data: A Replication Data DP pool number displays.
 - Management Area: A Management Area DP pool number displays.
 - Copy Type: TrueCopy Extended Distance is displayed.
 - Group Number: A group number and group name is displayed.
 - Group Name
3. Locate the pair whose status you want to review in the **Pair** list. Status descriptions are provided in [Table 16-3](#). You can click the **Refresh Information** button (not in view) to make sure data is current.
- The percentage that appears with each status shows how close the S-VOL is to being completely paired with the P-VOL.
 - The Attribute column shows the pair volume for which status is shown.

Table 16-3: TrueCopy pair status

Pair status	Description	P-VOL access	S-VOL access
Simplex	The volume is not assigned to the pair. If the created pair is deleted, the pair status becomes Simplex. Note that the Simplex volume is not displayed on the Remote Replication pair list. The disk array accepts read and write operations for Simplex volumes.	Read: Yes Write: Yes	Read: Yes Write: Yes
Synchronizing	Copying from the P-VOL to the S-VOL is in process. If the split pair is resynchronized, only the differential data of the P-VOL is copied to the S-VOL. If the pair at the time of the pair creation is resynchronized, entire P-VOL is copied to the S-VOL.	Read: Yes Write: Yes	Read: Yes, mount operation disabled. Write: No
Paired	The copy operation is complete.	Read: Yes Write: Yes	Read: Yes, mount operation disabled. Write: No
Split	The copy operation is suspended. The disk array starts accepting write operations for P-VOL and S-VOL. When the pair is resynchronized, the disk array executes the differential data copying from P-VOL to S-VOL.	Read: Yes Write: Yes	Read: Yes Write: Yes

Table 16-3: TrueCopy pair status (Continued)

Pair status	Description	P-VOL access	S-VOL access
Takeover	Takeover is a transitional status after Swap Pair is executed. Immediately after the pair is changed to Takeover status, the pair relationship is swapped and copy from the new P-VOL to the new S-VOL is started. Only the S-VOL has this status. The S-VOL in the Takeover status can perform the Read/Write access from the host.	–	Read: Available Write: Available
Failure	A failure occurs and copy operations are suspended forcibly. • If Data is specified as the fence level, the disk array rejects all write I/O. Read I/O is also rejected if PSUE Read Reject is specified. • If Never is specified, read/write I/O continues as long as the volume is unblocked • S-VOL read operations are accepted but not write operations. • To recover, resynchronize the pair (might require copying entire P-VOL).' • See Fence level on page 15-6 for more information.	Read: Yes/ No Write: Yes/No	Read: Yes Write: No

For CCI status see [Operations using CLI on page C-5](#).

Status Narrative : If a volume is not assigned to a TrueCopy pair, its status is Simplex. When a TrueCopy pair is being created, the status of the P-VOL and S-VOL is Synchronizing. When the copy operation is complete, status becomes Paired. If the system cannot maintain Paired status for any reason, the pair status changes to Failure. When the Split Pair operation is complete, the pair status changes to Split and the S-VOL can be written to. When you start a Resync Pair operation, the pair status changes to Synchronizing. When the operation is completed, the pair status changes to Paired. When you delete a pair, pair status changes to Simplex.



NOTE: Pair status for the P-VOL can differ from status for the S-VOL. If the remote path breaks down when pair status is Paired, the local disk array becomes Failure because it cannot send data to the remote disk array. The remote disk array remains Paired though there is no write I/O from the P-VOL.

Monitoring pair failure

It is necessary to check the pair statuses regularly to ensure that TrueCopy pairs are operating correctly and data is updated from P-VOLs to S-VOLs in the Paired status, or that differential data management is performed in the Split status. When a hardware failure occurs, the failure may cause a pair failure and may change the pair status to Failure. Check that the pair status is other than Failure. When the pair status is Failure, you must restore the status. See [Pair failure on page 16-10](#).

For TrueCopy, the following processes are executed when the pair failure occurs:

Table 16-4: Pair failure results

Management software	Results
Navigator 2	A message is displayed in the event log
	The pair status is changed to Failure.
CCI	The pair status is changed to PSUE.
	An error message is output to the system log file. (For UNIX® system and the Windows Server, the syslog file and eventlog file are shown respectively.)

When the pair status is changed to Failure, a trap is reported with SNMP Agent Support Function

When using CCI, the following message is output to the event log. For the details, refer to *Hitachi Unified Storage Command Control Interface Installation and Configuration Guide*.

Table 16-5: CCI system log message

Message ID	Condition	Cause
HORCM_102	The volume is suspended in code 0006	The pair status was suspended due to code 0006.

Monitoring of pair failure using a script

When SNMP Agent Support Function not used, it is necessary to monitor the pair failure using a Windows Server script that can be performed using Navigator 2 CLI commands.

The following is a script for monitoring the two pairs (SI_LU0001_LU0002 and SI_LU0003_LU0004) and informing the user when pair failure occurs. The following script is activated every several minutes. The disk array must be registered to beforehand.

```
echo OFF
REM Specify the registered name of the arrays
set UNITNAME=Array1
REM Specify the name of target group (Specify "Ungroup" if the
pair doesn't belong to any group)
set G_NAME=Ungrouped
REM Specify the name of target pair
set P1_NAME=SI_LU0001_LU0002
set P2_NAME=SI_LU0003_LU0004
REM Specify the value to inform "Failure"
set FAILURE=14

REM Checking the first pair
:pair1
aureplicationmon -unit %UNITNAME% -evwait -si -pairname
%P1_NAME% -gname %G_NAME% -nowait
if errorlevel %FAILURE% goto pair1_failure
goto pair2
:pair1_failure
<The procedure for informing a user>*

REM Checking the second pair
aureplicationmon -unit %UNITNAME% -evwait -si -pairname
%P2_NAME% -gname %G_NAME% -nowait
if errorlevel %FAILURE% goto pair2_failure
goto end
:pair2_failure
<The procedure for informing a user>*

:end
%
```

Monitoring the remote path

Monitor the remote path to ensure that data copying is unimpeded. If a path is blocked, the status is Detached, and data cannot be copied. You can adjust remote path bandwidth to improve data transfer rate.

To monitor the remote path

1. In the Replication tree, click **Setup**, then **Remote Path**. The Remote Path screen appears.
2. Review statuses and bandwidth. Path statuses can be Normal, Blocked, or Diagnosing. When Blocked or Diagnosing is displayed, data cannot be copied.
3. Take corrective steps as needed.

Troubleshooting

Pair failure

A pair failure occurs when one of the following takes place:

- A hardware failure occurs.
- Forcible release is performed by the user. This occurs when you halt a Pair Split operation. The disk array places the pair in Failure status.

If the pair was not forcibly suspended, the cause is hardware failure.

To restore pairs after a hardware failure

1. If volumes for the P-VOL and S-VOL were re-created after the failure, the pairs must be re-created.
2. If the volumes were recovered and it is possible to resync the pair, then do so. If resync is not possible, delete then re-create the pairs.
3. If a P-VOL restore was in progress during a hardware failure, delete the pair, restore the P-VOL if possible, and create a new pair.

Restoring pairs after forcible release operation

Create or re-synchronize the pair. When an existing pair is re-synchronized, the entire P-VOL is re-copied to the S-VOL.



NOTE: A TrueCopy operation may stop when the format command is run on Navigator 2 or Windows 2000 Server and Windows Server 2003/2008. This operation causes the Format, Synchronize Cache, or Verify operations to be performed, which in turn cause TrueCopy path blockage. The copy operation restarts when the Verify, Format, or Synchronize Cache operation is completed.

Recovering from a pair failure

Figure 16-2 shows a workflow to be done when a pair failure occurs from determining the factor to restore of the pair status by pair operation. Table 16-6 on page 16-12 shows the work responsibility schedule for the service personnel and a user.




Figure 16-2: Recovery from a pair failure

Table 16-6: Operational notes for TrueCopy operations

Action	Action taken by
Monitoring pair failure.	User
Verify that operation of suspends the pair by the user.	User
Verify the status of the array.	User
Call maintenance personnel when the array malfunctions.	User
For other reasons, call the Hitachi support center.	User (only for users that are registered in order to receive a support)
Hardware maintenance.	Hitachi Customer Service
Reconfigure and recover the pair.	User

Cases and solutions Using the DP-VOLs

When configuring a TrueCopy pair using the DP-VOL as a pair target volume, the TrueCopy pair status may become Failure depending on the combination of the pair status and the DP pool status shown in [Table 16-7](#). Perform the recovery method shown in [Table 16-7](#) for all the DP pools to which the P-VOLs and the S-VOLs where the pair failures have occurred belong.

Table 16-7: Cases and solutions using the DP-VOLs

Pair status	DP pool status	Cases	Solutions
Paired Synchronizing	Formatting	Although the DP pool capacity is being added, the format progress is slow and the required area cannot be allocated.	Wait until the formatting of the DP pool for total capacity of the DP-VOLs created in the DP pool is completed.
	Capacity Depleted	The DP pool capacity is depleted and the required area cannot be allocated.	To make the DP pool status normal, perform the DP pool capacity growing and DP pool optimization, and increase the DP pool free capacity.

TrueCopy Extended Distance theory of operation

When both fast performance and geographical distance capabilities are vital, Hitachi TrueCopy Extended Distance (TCE) software for Hitachi Unified Storage provides bi-directional, long-distance, remote data protection. TrueCopy Extended Distance supports data copy, failover and multi generational recovery without affecting your applications.

The key topics in this chapter are:

- ❑ [How TrueCopy Extended Distance works](#)
- ❑ [Configuration overview](#)
- ❑ [Operational overview](#)
- ❑ [Typical environment](#)
- ❑ [TCE Components](#)
- ❑ [TCE interfaces](#)

How TrueCopy Extended Distance works

With TrueCopy Extended Distance (TCE), you create a copy of your data at a remote location. After the initial copy is created, only changed data transfers to the remote location.

You create a TCE copy when you:

- Select a volume on the production disk array that you want to replicate
- Create a volume on the remote disk array that will contain the copy
- Establish a Fibre Channel or iSCSI link between the local and remote disk arrays
- Make the initial copy across the link on the remote disk array.

During and after the initial copy, the primary volume on the local side continues to be updated with data from the host application. When the host writes data to the P-VOL, the local disk array immediately returns a response to the host. This completes the I/O processing. The disk array performs the subsequent processing independently from I/O processing.

Updates are periodically sent to the secondary volume on the remote side at the end of the “update cycle”. This is a time period established by the user. The cycle time is based on the recovery point objective (RPO), which is the amount of data in time (2-hours’ worth, 4 hour’s worth) that can be lost after a disaster, until the operation is irreparably damaged. If the RPO is two hours, the business must be able to recover all data up to two hours before the disaster occurred.

When a disaster occurs, storage operations are transferred to the remote site and the secondary volume becomes the production volume. All the original data is available in the S-VOL, from the last completed update. The update cycle is determined by your RPO and by measuring write-workload during the TCE planning and design process.

For a detailed discussion of the disaster recovery process using TCE, please refer to [Process for disaster recovery on page 20-21](#).

Configuration overview

The local array and remote array are connected with remote lines such as DWDM (Dense Wavelength Division Multiplexing) lines. The local array contains the P-VOL, which stores the data of applications that run on the host. The remote array contains the S-VOL, which is a remote copy of the P-VOL.

Operational overview

If the host writes data to the P-VOL when a TCE pair has been created for the P-VOL and S-VOL (See [Figure 17-1](#) (1)) the local array immediately returns a response to the host (2). This completes the I/O processing. The array performs the subsequent processing independently from the I/O processing.

If new data is written to update data that has not been transferred to the S-VOL, the local array copies the un-transferred data to the DP pool (3). When the data was already transferred or the transfer was unnecessary, it is over-written.

The local array transmits the data written by the host to the S-VOL as update data (4). The remote array returns a response to the local array when it has received the update data (5).

If the update data from the local array updates internal pre-determined data of the S-VOL, the remote array copies that data to the DP pool (6).

The local array and remote array accomplish asynchronous remote copy by repeating the above processing.




Figure 17-1: TCE operational overview

Typical environment

A typical configuration consists of the following elements. Many but not all require user set up.

- Two disk arrays—one on the local side connected to a host, and one on the remote side connected to the local disk array. Connections are made via Fibre Channel or iSCSI.
- A primary volume on the local disk array that is to be copied to the secondary volume on the remote side.
- Interface and command software, used to perform TCE operations. Command software uses a command device (volume) to communicate with the disk arrays.

Figure 17-2 shows a typical TCE environment.




Figure 17-2: Typical TCE environment

TCE Components

To operate TCE, software including TCE license, Navigator 2, and CCI is required in addition to hardware including the two arrays, the PC/WSs (for hosts and servers), and the cables.

Navigator 2 is mainly used to set up the TCE configuration, operate pairs, and do maintenance. CCI is used mainly for the operation of volume pairs of TCE.

Volume pairs

When the initial TCE copy is completed, the production and backup volumes are said to be "Paired". The two paired volumes are referred to as the primary volume (P-VOL) and secondary volume (S-VOL). Each TCE pair consists of one P-VOL and one S-VOL. When the pair relationship is established, data flows from the P-VOL to the S-VOL.

While in the Paired status, new data is written to the P-VOL and then periodically transferred to the S-VOL, according to the user-defined update cycle.

When a pair is "split", the data flow between the volumes stops. At this time, all the differential data that has accumulated in the local disk array since the last update is copied to the S-VOL. This insures that its data is the same as the P-VOL's and is consistent and usable data.

TCE performs remote copy operations for logical volume pairs established by the user. Each TCE pair consists of one primary volume (P-VOL) and one secondary volume (S-VOL), which are located in the arrays that are connected by a Fibre Channel interface or iSCSI. The TCE P-VOLs are the primary volumes that contain original data. The TCE S-VOLs are the secondary or mirrored volumes that contain backup data. Because the data transfer to the S-VOL is done regularly, some differences between the P-VOL and S-VOL data are made in a pair that is receiving the host I/O instruction.

During TCE operations, the P-VOLs remain available to all hosts for read and write I/O operations. An exception to this includes when the volume is impossible to access (for example, a volume blockage). The S-VOLs become available for write operations from the hosts only after the pair has been split.

Depending on how the pair is split, the S-VOL is available for both read and write I/O.

The pairsplit operation takes some time until it is completed, because it is required to reflect the P-VOL data at the time of the reception of the instruction on the S-VOL.

When a TCE volume pair is created, the data on the P-VOL is copied to the S-VOL and the initial copy is completed. After the initial copy is completed, differential data is copied regularly in a cycle specified by a user. If you need to access an S-VOL, you can "split" the pair to make the S-VOL accessible

While a TCE pair is split, the array keeps track of all changes to the P-VOL and S-VOL. When a pair is resynchronized, the differential data of the P-VOL is copied to the S-VOL (in order to update the P-VOL and the S-VOL) and the regular copy of the differential data from the P-VOL to the S-VOL is started again.

Remote path

The remote path is a path between the local array and the remote array, which is used for transferring data between P-VOLs and S-VOLs. The remote path has two paths, path 0 and path 1. The interface type of two remote paths between the arrays needs to be the same. A minimum of two paths are required and a maximum of two paths is supported, 1 per controller. See [Figure 17-3](#).




Figure 17-3: Switching the paths

Alternative path

Two paths must be set to avoid stopping (suspending) the copy operation due to a single point malfunction in the path. A single path for each controller on the local and remote disk array must be set and a duplex path for each pair is allocated.

To avoid malfunction, the path can be automatically switched from the main path to the alternative path from the local disk array.

Confirming the path condition

TCE supports a function that periodically issues commands between the disk arrays and monitors the path status. When a path status is blocked due to path malfunction, its status will be reported from the LED (no status is reported for temporary command error).

Port connection and topology for Fibre Channel Interface

The disk array supports direct or switch connection only. Hub connection is not supported. A connection via Hub is not supported even if the connection is FL-Port of Fabric (switch).

It is necessary to designate command device(s) before a remote path setting.

If a failure such as a controller blockage has occurred, a path setting cannot be performed.

The remote path on TCE is specified using Navigator 2.

You cannot change the remote path information once it is set. When changing it, delete the set remote path information and reset it again as new remote path information.

It is necessary to change all the pairs of TCE to Split before remote path deleting.

For topology, see [Table 17-1](#).

Table 17-1: Supported topology

#	Port connection	Topology	Local	Remote
1	Direct	Point to Point	Not available	Not available
2		Loop	Available	Available
3	HUB	Loop	Not available	Not available
4	Switch	Point to Point F-Port	Available	Available
5		Loop FL-Port	Available	Available

Port transfer rate for Fibre Channel

Set the transfer rate of the Fibre Channel of the array to a value corresponding to the transfer rate of the equipment connected directly to the array according to [Table 17-2](#) for each port.

Table 17-2: Transfer rate

Transfer rate of the opposed connection equipment of the array for each port	Transfer rate of the array
Fixed rate: 2 G bps	2 G bps
Fixed rate: 4 G bps	4 G bps
Fixed rate: 8 G bps	8 G bps
Auto (Maximum rate: 2 G bps)	2 G bps
Auto (Maximum rate: 4 G bps)	4 G bps
Auto (Maximum rate: 8 G bps)	8 G bps



NOTE: When making the transfer rate of the array “Auto”, it may not link up at the maximum rate depending on the connection equipment. Check the transfer rate with Navigator 2 when starting the array, Switch and HBA, and so forth. If it differs from the maximum rate, change it to the fixed rate or pull out/insert the cable

DP pools

TCE retains the differential data to be transferred to the S-VOL by saving it in the DP pool in the local array. The data, transferred to the S-VOL in order to provide for the case where the S-VOL data is demanded to be used because of a failure on the P-VOL side, etc., is guaranteed as saved in a DP pool in the remote array. The differential data is called replication data and the area to store the differential data is called a replication data DP pool.

The replication data DP pool is necessary for each of the local array and the remote array. Furthermore, the area for managing the pairs of which replication data is in which P-VOL is called a management area DP pool, and it is necessary for each of the local array and the remote array.

Up to 64 DP pools (HUS 130/HUS 150) or up to 50 DP pools (HUS 110) can be created per disk array and a DP pool to be used by a certain P-VOL is specified when a pair is created. A DP pool to be used can be specified for each P-VOL. Two or more TCE pairs can share a single DP pool.

There are the replication depletion alert threshold value and the replication data release threshold value for the DP pools.

You must specify the capacity of the DP pool. You need to specify the capacity that is enough to support the practical use, taking the amount of the differential data and the cycle in consideration.

When the DP pool overflows:

- The DP pool in the local array becomes full. When the status of the TCE pair is Paired, the P-VOL is changed to Pool Full. When the pair status is Synchronizing, the P-VOL is changed to Failure. An overflow of the DP pool in the local array has no effect on the S-VOL status.
- The DP pool in the remote array becomes full. When the status of the TCE pair is Paired, the P-VOL is changed to Failure and the S-VOL is changed to Pool Full. When the pair status is Synchronizing, the P-VOL is changed to Failure and the S-VOL is changed to Inconsistent.



NOTE: When even one RAID group assigned to a DP pool is damaged, all the pairs using the DP pool are placed in the **Failure** status.

The Replication threshold value can be set in the DP pool. In the Replication threshold value, the replication Depletion Alert threshold value and the Replication Data Released threshold value can be set, although TCE does not refer to the Replication Depletion Alert threshold value. The threshold value to be set is the ratio of the usage of the DP pool for the entire capacity

of the DP pool. Setting the Replication threshold value helps prevent the DP pool from becoming depleted by TCE. Although TCE does not refer to the Replication Depletion Alert threshold value, always set a larger Replication Data Released threshold value than the replication Depletion Alert threshold value.

The replication Data Released threshold value cannot be set within the range of -5% of the Replication Depletion Alert threshold value.

When the usage rate of the replication data DP pool or management area DP pool for the P-VOL reaches the Replication Data Released threshold value, the pair status of the P-VOL changes to Pool Full. The replication data for the P-VOL is released at the same time and the usable capacity of the DP pool recovers. When the usage rate of the replication data DP pool or management area DP pool for the S-VOL reaches the Replication Data Released threshold value, the pair status of the S-VOL changes to Pool Full.

Until the usage rate of the DP pool recovers to over -5% of the Replication Data Released threshold value, pair creation, pair resynchronization and pair swap cannot be performed.

Guaranteed write order and the update cycle

S-VOL data must have the same order in which the host updates the P-VOL. When write order is guaranteed, the S-VOL has data consistency with the P-VOL.

As explained in the previous section, data is copied from the P-VOL and local DP pool to the S-VOL following the update cycle. When the update is complete, S-VOL data is identical to P-VOL data *at the end of the cycle*. Since the P-VOL continues to be updated while and after the S-VOL is being updated, S-VOL data and P-VOL data are not identical.

However, the S-VOL and P-VOL can be made identical when the pair is split. During this operation, all differential data in the local P-VOL and DP pool is transferred to the S-VOL, as well as all cached data in host memory. This cached data is flushed to the P-VOL, then transferred to the S-VOL as part of the split operation, thus ensuring that the two are identical.

If a failure occurs during an update cycle, the data in the update is inconsistent. Write order in the S-VOL is nevertheless guaranteed — at the point-in-time of the previous update cycle, which is stored in the remote DP pool.

Figure 17-4 shows how S-VOL data is maintained at one update cycle back of P-VOL data.

Managing primary volumes as a group allows TCE operations to be performed on all volumes in the group concurrently. Write order in secondary volumes is guaranteed across application logical volumes. [Figure 17-5](#) shows TCE operations with a group.

By making multiple pairs belong to the same group, the pair operation is possible in units of groups. In the group whose Point-in-time attribute of the group is enabled, the backup data of the S-VOL created in units of groups is the data of the same time.

For setting a group, specify a new group number for a group to be assigned after pair creation when creating a TCE pair. The maximum of 16 groups can be created in TCE. While a group number from 0 to 255 can be used for TCE, the max number of groups that can be actually created is 16. Note that the group number that is being used by TrueCopy cannot be used for TCE as group numbers are shared between TrueCopy and TCE.

A group name can be assigned to a group. You can select one pair belonging to the created group and assign a group name arbitrarily by using the pair edit function.




Figure 17-5: TCE operations with group

In this illustration, observe the following:

- The P-VOLs belong to the same group. The host updates the P-VOLs as required (1).

- The local disk array identifies the differential data in the P-VOLs when the cycle is started (2) in an atomic manner. The differential data of the group of the P-VOLs are determined at time T2.
- The local disk array transfers the differential data to the corresponding S-VOLs (3). When all differential data is transferred, each S-VOL is identical to its P-VOL at time T2 (4).
- If pairs are split or deleted, the local disk array stops the cycle update for the group. Differential data between P-VOLs and S-VOLs is determined at that time. All differential data is sent to the S-VOLs, and the split or delete operations on the pairs completes. S-VOLs maintain data consistency across pairs in the group.

Command Devices

The command device is a user-selected, dedicated logical volume on the disk array, which functions as the interface to the CCI software. TCE commands are issued by CCI (HORCM) to the disk array command device.

A command device must be designated in order to issue TCE commands. The command device must be defined in the HORCM_CMD section of the configuration definition file for the CCI instance on the attached host. 128 command devices can be designated for the disk array. You can designate command devices using Navigator 2.



NOTE: Volumes set for command devices must be recognized by the host. The command device volume size must be greater than or equal to 33 MBs.

TCE interfaces

TCE can be setup, used and monitored using of the following interfaces:

- The GUI (Hitachi Storage Navigator Modular 2 Graphical User Interface), which is a browser-based interface from which TCE can be setup, operated, and monitored. The GUI provides the simplest method for performing operations, requiring no previous experience. Scripting is not available.
- CLI (Hitachi Storage Navigator Modular 2 Command Line Interface), from which TCE can be setup and all basic pair operations can be performed—create, split, resynchronize, restore, swap, and delete. The GUI also provides these functions. CLI also has scripting capability.
- CCI (Hitachi Command Control Interface (CCI)), which is used to display volume information and perform all copying and pair-managing operations. CCI provides a full scripting capability which can be used to automate replication operations. CCI requires more experience than the GUI or CLI. CCI is required for performing failover and fall back operations, and, on Windows 2000 Server, mount/unmount operations.

HDS recommends using the GUI to begin operations for new users with no experience with CLI or CCI. Users who are new to replication software but have CLI experience in managing disk arrays may want to continue using CLI, though the GUI is an option. The same recommendation applies to CCI users.

Installing TrueCopy Extended

This chapter provides TCE installation and setup procedures using the Navigator 2 GUI. Instructions for CLI and CCI can be found in the appendixes.

- ❑ [TCE system requirements](#)
- ❑ [Installation procedures](#)

TCE system requirements

This section describes the minimum TCE requirements.

Table 18-1: TCE requirements

Item	Minimum requirements
Firmware version	Firmware: Version 0916/A or higher is required
Navigator 2 version	Navigator 2: Version 21.60 or higher is required for management PC.
CCI version	Version 01-27-03/02 or later is required host only when CCI is used for the operation of TCE.
Number of arrays	2
Supported array models	HUS 150, HUS 130, HUS 110
TCE and Dynamic Provisioning license keys	Two license keys for TCE and Dynamic Provisioning.
Number of controllers:	2 (dual configuration)
DP pool	DP pool (local and remote)

Installation procedures

The following sections provide instructions for installing, enabling/disabling, and uninstalling TCE. Please note the following:

- TCE must be installed on the local and remote disk arrays.
- Before proceeding, verify that the disk array is operating in a normal state. Installation/un-installation cannot be performed if a failure has occurred.
- TCE and TrueCopy cannot be used together because their licenses are independent from each other.
- When the interface is iSCSI, you cannot install TCE if 240 or more hosts are connected to a port. Reduce the number of hosts connecting to one port to 239 or less and install TCE.

Installing TCE

Prerequisites

- A key code or key file is required to install or uninstall TCE. If you do not have the key file or code, you can obtain it from the download page on the HDS Support Portal, <http://support.hds.com>.
- When the interface is iSCSI and you install TCE, the maximum number of connectable hosts per port becomes 239.

To install TCE

1. In the Navigator 2 GUI, click the array in which you will install TCE.
2. Click **Show & Configure array**.
3. Select the **Install License** icon in the **Common array Task**.




The **Install License** screen appears.



4. Select the **Key File** or **Key Code** option, and then enter the file name or key code. You may **Browse** for the key file.
5. A screen appears, requesting a confirmation to install TCE option. Click **Confirm**.

Install License



6. A completion message appears. Click **Close**.

Install License



7. Installation of TCE is now complete.



NOTE: TCE needs the DP pool of Dynamic Provisioning. If Dynamic Provisioning is not installed, install Dynamic Provisioning.

Enabling or disabling TCE


TCE is automatically enabled when it is installed. You can disable or re-enable it.

Prerequisites

- To enable when using TCE with iSCSI, there must be fewer than 240 hosts connected to a port on the disk array.
- When disabling TCE:
 - pairs must be deleted and the status of the volumes must be Simplex.
 - The path settings must be deleted.

To enable or disable TCE

1. In the Navigator 2 GUI, click the check box for the disk array, then click the **Show & Configure disk array** button.
2. In the tree view, click **Settings**, then click **Licenses**.
3. Select **TC-Extended** in the licenses list.
4. Click **Change Status**. The Change License screen appears.



5. To **disable**, clear the **Enable: Yes** check box.
To **enable**, check the **Enable: Yes** check box.
6. Click **OK**.
7. A message appears. Click **Close**.
Enabling or disabling of TCE is now complete.


Uninstalling TCE

Prerequisite

- TCE pairs must be deleted. Volume status must be Simplex.
- The path settings must be deleted.
- A key code or key file is required. If you do not have the key file or code, you can obtain it from the download page on the HDS Support Portal, <http://support.hds.com>.

To uninstall TCE

1. In the Navigator 2 GUI, click the check box for the disk array where you will uninstall TCE, then click the **Show & Configure disk array** button.
2. Select the **Licenses** icon in the **Settings** tree view.



The **Licenses** list appears.

3. Click **De-install License**.

The **De-Install License** screen appears.



4. When you uninstall the option using the key code, click the **Key Code** option, and then set up the key code. When you uninstall the option using the key file, click the **Key File** option, and then set up the path for the key file name. Use Browse to set the path to a key file correctly. Click **OK**.
5. A message appears, click **Close**.
The Licenses list appears.
Un-installation of TCE is now complete.

TrueCopy Extended Distance setup

This chapter provides required information for setting up your system for TrueCopy Extended Distance. It includes:

Planning and design

- ❑ [Plan and design — sizing DP pools and bandwidth](#)
- ❑ [Plan and design — remote path](#)
- ❑ [Plan and design—disk arrays, volumes and operating systems](#)

Setup procedures

- ❑ [Setup procedures](#)
- ❑ [Setting up DP pools](#)
- ❑ [Setting the replication threshold \(optional\)](#)
- ❑ [Setting the cycle time](#)
- ❑ [Adding or changing the remote port CHAP secret](#)
- ❑ [Setting the remote path](#)
- ❑ [Deleting the remote path](#)
- ❑ [Operations work flow](#)

Plan and design — sizing DP pools and bandwidth

This topic provides instructions for measuring write-workload and sizing DP pools and bandwidth.

- ❑ [Plan and design workflow](#)
- ❑ [Assessing business needs — RPO and the update cycle](#)
- ❑ [Measuring write-workload](#)
- ❑ [DP pool size](#)
- ❑ [Determining bandwidth](#)
- ❑ [Performance design](#)

Plan and design workflow

You design your TCE system around the write-workload generated by your host application. DP pools and bandwidth must be sized to accommodate write-workload. This topic helps you perform these tasks as follows:

- Assess business requirements regarding how much data your operation must recover in the event of a disaster.
- Measure write-workload. This metric is used to ensure that DP pool size and bandwidth are sufficient to hold and pass all levels of I/O.
- Calculate DP pool size. Instructions are included for matching DP pool capacity to the production environment.
- Calculate remote path bandwidth: This will make certain that you can copy your data to the remote site within your update cycle.

Assessing business needs — RPO and the update cycle

In a TCE system, the S-VOL will contain nearly all of the data that is in the P-VOL. The difference between them at any time will be the differential data that accumulates during the TCE update cycle.

This differential data accumulates in the local DP pool until the update cycle starts, then it is transferred over the remote data path.

Update cycle time is a uniform interval of time during which differential data copies to the S-VOL. You will define the update cycle time when creating the TCE pair.

The update cycle time is based on:

- the amount of data written to your P-VOL
- the maximum amount of data loss your operation could survive during a disaster.

The data loss that your operation can survive and remain viable determines to what point in the past you must recover.

An hour's worth of data loss means that your recovery point is one hour ago. If disaster occurs at 10:00 am, upon recovery your restart will resume operations with data from 9:00 am.

Fifteen minutes worth of data loss means that your recovery point is 15 minutes prior to the disaster.

You must determine your recovery point objective (RPO). You can do this by measuring your host application's write-workload. This shows the amount of data written to the P-VOL over time. You or your organization's decision-makers can use this information to decide the number of business transactions that can be lost, the number of hours required to key in lost data and so on. The result is the RPO.

Measuring write-workload

Bandwidth and DP pool size are determined by understanding the write-workload placed on the primary volume from the host application.

- After the initial copy, TCE only copies changed data to the S-VOL.
- Data is changed when the host application writes to storage.
- Write-workload is a measure of changed data over a period of time.

When you know how much data is changing, you can plan the size of your DP pools and bandwidth to support your environment.

Collecting write-workload data

Workload data is collected using your operating system's performance monitoring feature. Collection should be performed during the busiest time of month, quarter, and year so you can be sure your TCE implementation will support your environment when demand is greatest. The following procedure is provided to help you collect write-workload data.

To collect workload data

1. Using your operating system's performance monitoring software, collect the following:
 - Disk-write bytes-per-second for every physical volume that will be replicated.
 - Collect this data at 10 minute intervals and over as long a period as possible. Hitachi recommends a 4-6 week period in order to accumulate data over all workload conditions including times when the demands on the system are greatest.
2. At the end of the collection period, convert the data to MB/second and import into a spreadsheet tool. In [Figure 19-1 on page 19-5](#), column C shows an example of collected raw data over 10-minute segments.

	A	B	C	D	E	F	G	H
	Sample #	Time - 10 min. segments	Raw Data - MB/sec collected per 10 min. segment	30 Min Rolling Ave	60 Min Rolling Ave	24 Hour Rolling Avg		Raw Data Project Growth
1								
2								15.00%
3	1	0:00	0.863					0.992
4	2	0:10	2.81					3.232
5	3	0:20	0.858		1.24	2.17		0.987
6	4	0:30	0.45		1.56	2.17		0.518
7	5	0:40	0.813		1.36	2.17		0.935
8	6	0:50	1.63		1.56	2.17		1.875
9	7	1:00	2.81		2.24	2.17		3.232
10	8	1:10	1.58		2.30	2.18		1.817
11	9	1:20	2.05		2.36	2.18		2.358
12	10	1:30	4.57		2.48	2.17		5.256
13	11	1:40	1.13		2.66	2.17		1.300
14	12	1:50	2		2.72	2.16		2.300
15	13	2:00	3.56		2.30	2.16		4.094
16	14	2:10	2.64		2.26	2.19		3.036
17	15	2:20	2.39		2.34	2.19		2.749

Figure 19-1: Write-Workload spreadsheet

Fluctuations in write-workload can be seen from interval to interval. To calculate DP pool size, the interval data will first be averaged, then used in an equation. (Your spreadsheet at this point would have only rows B and C populated.)

DP pool size

You need to calculate how much capacity must be allocated for the DP pool to have TCE pairs. The capacity required will automatically be taken from the free portion of the DP pool as needed when old data is sent to the DP pool. However, the capacity of the DP pool is not unlimited. So you still need to consider how much capacity is left in the pool for TCE.

Using TCE consumes DP pool capacity with replication data and management information stored in DP pools, which are differential data between a P-VOL and an S-VOL and information to manage the replication data, respectively. On the other hand, some pair operations such as pair deletion recover the usable capacity of the DP pool by removing unnecessary replication data and management information from the DP pool. The following sections show when replication data and management information increase and decrease as well as how much DP pool capacity they consume.

DP pool consumption

Table 19-1 shows when the replication data and management information increases and decreases. An increase in the replication data and management information leads to a decrease in the capacity of the DP pool that TCE pairs are using. And a decrease in the replication data and management information recovers the DP pool capacity used by TCE pairs.

Table 19-1: DP pool consumption

	Occasions for increase	Occasions for decrease
Replication Data	Cycle copying, execution of pair resync	Cycle copying, after cycle copying completed
Management Information	Creating pair, cycle copying	Deleting pair

How much capacity TCE consumes

The replication data increases as Write operations are being executed to the P-VOL on the local array during the cycle copy. On the remote array, it increases with the amount of the cycle copy. . At a maximum, the same amount of replication data is needed for the P-VOL or S-VOL.

The management information increases with the P-VOL capacity. Table 3.2 indicates the amount of management information depending on the P-VOL capacity. Table 19-2 indicates management information per P-VOL or S-VOL. Refer to [Management information on page 9-8](#) for the amount of management information for a SnapShot-TCE cascade configuration.

Table 19-2: Capacity of the management information

P-VOL capacity	Management information
50 GB	5 GB
100 GB	
250 GB	7 GB
500 GB	9 GB
1 TB	13 GB
2 TB	23 GB
4 TB	41 GB
8 TB	79 GB
16 TB	153 GB
32 TB	302GB
64 TB	600 GB
128 TB	1,197 GB

Determining bandwidth

The purpose of this section is to ensure that you have sufficient bandwidth between the local and remote disk arrays to copy all your write data in the time-frame you prescribe. The goal is to size the network so that it is capable of transferring estimated future write workloads.

TCE requires two remote paths, each with a minimum bandwidth of 1.5 Mbs.

To determine the bandwidth

1. Graph the data in column "C" in the [Write-Workload spreadsheet on page 19-5](#).
2. Locate the highest peak. Based on your write-workload measurements, this is the greatest amount of data that will need to be transferred to the remote disk array. Bandwidth must accommodate maximum possible workload to insure that the system does not become subject to its capacity being exceeded. This would cause further problems, such as the new write data backing up in the DP pool, update cycles becoming extended, and so on.
3. Though the highest peak in your workload data should be used for determining bandwidth, you should also take notice of extremely high peaks. In some cases a batch job, defragmentation, or other process could be driving workload to abnormally high levels. It is sometimes worthwhile to review the processes that are running. After careful analysis, it may be possible to lower or even eliminate some spikes by optimizing or streamlining high-workload processes. Changing the timing of a process may lower workload.
4. Although bandwidth can be increased, Hitachi recommends that projected growth rate be factored over a 1, 2, or 3 year period.

[Table 19-3](#) shows TCE bandwidth requirements.

Table 19-3: Bandwidth requirements

Average inflow	Bandwidth requirements	WAN types
.08 -.149 MB/s	1.5 Mb/s or more	T1
.15 -.299 MB/s	3 Mb/s or more	T1 x two lines
.3 -.599 MB/s	6 Mb/s or more	T2
.6 - 1.199 MB/s	12 Mb/s or more	T2 x two lines
1.2 - 4.499 MB/s	45 Mb/s or more	T3
4.500 - 9.999 MB/s	100 Mb/s or more	Fast Ethernet

Performance design

A system using TCE is made up of many types of components, such as local and remote arrays, P-VOLs, S-VOLs, DP pools, and lines. If a performance bottleneck occurs in just one of these components, the entire system breaks down. If the balance between inflow to the P-VOL and outflow from the P-VOL to the S-VOL is not good, differential data accumulates on the local array, making it impossible for the S-VOL to be used for recovery purposes.

Accordingly, when a system using TCE is built, performance design that takes into account the performance balance of the entire system is necessary. The purpose of performance design using TCE is to find a system configuration in which the average inflow to the P-VOL and the average outflow to the S-VOL match.

Figure 19-2 shows the locations of the major performance bottlenecks in a system using TCE. In addition to these, performance bottlenecks can occur on a front-end path, but these are not problems specific to TCE and are therefore not discussed.




Figure 19-2: TCE performance bottlenecks

Table 19-4 shows the effects of performance bottlenecks on the inflow speed and outflow speed. If the processor of the local array is a bottleneck, not only does the host I/O processing performance drop, the performance of processing that transfers data to the remote array also deteriorates. If the inflow and outflow speeds are not up to the target values due to a processor bottleneck of the local array, corrective action such as replacing the array controller with a higher-end model is required. Locations of performance bottlenecks and effects on inflow/outflow speeds

Table 19-4: Locations of performance bottlenecks and effects on inflow / outflow speeds

No.	Bottleneck location	Inflow speed	Outflow speed
1	Processor of local array	Yes	Yes
2	P-VOL drive	Yes	Yes
3	P-VOL DP pool drive	Yes	Yes
4-1	Line (bandwidth)	No	Yes
4-2	Line (delay time)	No	Yes
5	Processor of remote array	No	Yes
6	S-VOL drive	No	Yes
7	S-VOL DP pool drive	No	Yes

Yes: Adverse effects

No: No adverse effects

The effects on the inflow speed and outflow speed of bottlenecks at each location are explained in Table 19-5.

Table 19-5: Bottleneck description

No.	Type of bottleneck	Description
1	Local array processor	The local array processor handles host I/O processing, processing to copy data to a DP pool, and processing to transfer data to the remote array. If processor of the local array is overloaded, the inflow speed and/or outflow speed drops.
2	P-VOL drive	There are many I/Os issued on P-VOL such as reading or writing data in response to a host I/O request, reading data when it is to be copied to a DP pool, and reading data when it is transferred to the remote array. If the P-VOL load increases, the inflow speed and/or outflow speed drops.

No.	Type of bottleneck	Description
3	P-VOL DP pool drive	There are many I/Os issued on a DP pool at local array such as writing data when it is copied from P-VOL or reading data when it is transferred to remote. Because data is copied only when data that has not been transferred is updated during each cycle, the amount of data to be saved per cycle is small compared with the S-VOL DP pool. When the local side DP pool load increases, the inflow speed and/or outflow speed drops.
4-1	Line (bandwidth)	The bandwidth of the line limits the maximum data transfer rate from the local array to the remote array.
4-2	Line (delay time)	Because there are only 32 ongoing data transfers at a time per a controller of local array, the longer the delay time, the greater the drop of the outflow speeds.
5	Processor of remote array	The remote array processor handles processing incoming data from local array and coping of pre-determined data to a DP pool. The higher the load of the processor of the remote array, the greater the drop in outflow speed.
6	S-VOL drive	There are many I/Os issued on S-VOL such as writing data in response to a data transfer from a local array and reading data when it is to be copied to a DP pool. If the S-VOL load increases, the outflow speed drops.
7	S-VOL DP pool drive	There are many I/Os issued on a pool at remote array such as writing data when it is copied from P-VOL. When the remote side DP pool load increases, the outflow speed drops.

Plan and design — remote path

A remote path is required for transferring data from the local disk array to the remote disk array. This topic provides network and bandwidth requirements, and supported remote path configurations.

- ❑ [Remote path requirements](#)
- ❑ [Remote path configurations](#)
- ❑ [Using the remote path — best practices](#)

Remote path requirements

The remote path is the connection used to transfer data between the local disk array and remote disk array. TCE supports Fibre Channel and iSCSI port connectors and connections. The connections you use must be either one or the other; they cannot be mixed.

The following kinds of networks are used with TCE:

- Local Area Network (LAN), for system management. Fast Ethernet is required for the LAN.
- Wide Area Network (WAN) for the remote path. For best performance:
 - A Fibre Channel extender is required.
 - iSCSI connections may require a WAN Optimization Controller (WOC).

Figure 19-3 shows the basic TCE configuration with a LAN and WAN.




Figure 19-3: Remote path configuration

Requirements are provided for the following:

- [Management LAN requirements on page 19-13](#)
- [Remote data path requirements on page 19-13](#)
- [Remote path configurations on page 19-14](#)

- [Fibre Channel extender connection on page 19-20](#)

Management LAN requirements

Fast Ethernet is required for an IP LAN.

Remote data path requirements

This section discusses the TCE remote path requirements for a WAN connection. This includes the following:

- Types of lines
- Bandwidth
- Distance between local and remote sites
- WAN Optimization Controllers (WOC) (optional)

For instructions on assessing your system's I/O and bandwidth requirements, see:

- [Measuring write-workload on page 19-4](#)
- [Determining bandwidth on page 19-7](#)

[Table 19-6](#) shows remote path requirements for TCE. A WOC may also be required, depending on the distance between the local and remote sites and other factors listed in [Table 19-11](#).

Table 19-6: Remote data path requirements

Item	Requirements
Bandwidth	• Bandwidth must be guaranteed. • Bandwidth must be 1.5 Mb/s or more for each pair. 100 Mb/s recommended. • Requirements for bandwidth depend on an average inflow from the host into the disk array. • See Table 19-3 on page 19-7 for bandwidth requirements.
Remote path sharing	• The remote path must be dedicated for TCE pairs. • When two or more pairs share the same path, a WOC is recommended for each pair.

[Table 19-7](#) shows types of WAN cabling and protocols supported by TCE and those not supported.

Table 19-7: Supported, not Supported WAN types

WAN Types	
Supported	Dedicated Line (T1, T2, T3 etc)
Not-supported	ADSL, CATV, FTTH, ISDN

Remote path configurations

TCE supports both Fibre Channel and iSCSI connections for the remote path.

- Two remote paths must be set up, one per controller. This ensures that an alternate path is available in the event of link failure during copy operations.
- Paths can be configured from:
 - Local controller 0 to remote controller 0 or 1
 - Local controller 1 to remote controller 0 or 1
- Paths can connect a port A with a port B, and so on. Hitachi recommends making connections between the same controller/port, such as port 0B to 0B, and 1 B to 1 B, for simplicity. Ports can be used for both host I/O and replication data.

The following sections describe supported Fibre Channel and iSCSI path configurations. Recommendations and restrictions are included.

Fibre Channel

The Fibre Channel remote data path can be set up in the following configurations:

- Direct connection
- Single Fibre Channel switch and network connection
- Double Fibre Channel switch and network connection
- Wavelength Division Multiplexing (WDM) and dark fibre extender

The disk array supports direct or switch connection only. Hub connections are not supported.

The connection via a switch supports both F-Port (Point-to-Point) and FL-Port (Loop).

General recommendations

The following is recommended for all supported configurations:

- TCE requires one path between the host and local disk array. However, two paths are recommended; the second path can be used in the event of a path failure.

Direct connection

Figure 19-4 illustrates two remote paths directly connecting the local and remote disk arrays. This configuration can be used when distance is very short, as when creating the initial copy or performing data recovery while both disk arrays are installed at the local site.




Figure 19-4: Direct Fibre Channel connection

Recommendations

- When connecting the local array and the remote array directly, set the transfer rate of fibre channel to the fixed rate (the same setting of 2 G bps, 4 G bps, and 8 G bps) for each array, following the table below.

Table 19-8: Transfer rates

Transfer rate of the port of the directly connected local array	Transfer rate of the port of the directly connected remote array
2 Gbps	2Gbps
4 Gbps	4 Gbps
8 Gbps	8 Gbps

- When connecting the local array and the remote array directly and setting the transfer rate to Auto, the remote path may be blocked. If the remote path is blocked, change the transfer rate of the fixed rate.

Fibre Channel switch connection 1

Switch connections increase throughput between the disk arrays.

When a pair is created on TCE, it is necessary to connect two hosts with the LAN for communicating between CCIs on the host associated with the local array and on the host associated with the remote array. If there are two hosts, one host activates both CCIs on the local side and on the remote side; it is not necessary to connect two hosts with the LAN. See [Figure 19-5](#).




Figure 19-5: Fibre Channel switch connection 1

Fibre Channel switch connection 2

Between a host and an array only one remote path is acceptable. If a configuration has two remote paths as illustrated in [Figure 19-6](#), a remote path can be switched when a failure in a remote path or the controller blockage occurs.




Figure 19-6: Fibre Channel switch connection 2

The array must be connected with a switch as follows ([Table 19-9](#)).

Table 19-9: Connections between Array and a Switch

Mode of array	Switch		
	For 8 G bps	For 4 G bps	For 2 G bps
Auto Mode	From the viewpoint of the performance, one path/controller between the array and a switch is acceptable, as illustrated above. The same port is available for the host I/O and for copying data of TCE.	See the left column	See the left column
8 G bps Mode			Not available
4 G bps Mode			Not available
2 G bps Mode			See the left column

One-Path-Connection between Arrays

When a pair is created on TCE, it is necessary to connect two hosts with the LAN for communicating between CCIs on the host associated with the local array and on the host associated with the remote array. If there are two hosts, one host activates both CCIs on the local side and on the remote side; it is not necessary to connect two hosts with the LAN. See [Figure 19-7](#).

If a failure occurs in a switch or a remote path, a remote path cannot be switched. Therefore, this configuration is not recommended.




Figure 19-7: Fibre Channel one path connection

Fibre Channel extender connection

Channel extenders convert Fibre Channel to FCIP or iFCP, which allows you to use IP networks and significantly improve performance over longer distances.

If distance extends, response time will increase. Nevertheless, the distance does not affect the response time from a host. However, the longer distance deteriorates the performance of data transfer for copying and may cause a pair failure, depending on the delay time, bandwidth, and line quality.

Figure 19-8 illustrates two remote paths using two Fibre Channel switches, Wavelength Division Multiplexor (WDM) extender, and dark fibre to make the connection to the remote site.




Figure 19-8: Fibre Channel switches, WDM, Dark Fibre connection

Recommendations

- Only qualified components are supported.

For more information about WDM, see [Wavelength Division Multiplexing \(WDM\) and dark fibre on page D-45](#).

Port transfer rate for Fibre Channel

The communication speed of the Fibre Channel port on the disk array must match the speed specified on the host port. These two ports—Fibre Channel port on the disk array and host port—are connected via the Fibre Channel cable. Each port on the disk array must be set separately.

Table 19-10: Setting port transfer rates

Mode	If the host port is set to	Set the remote disk array port to
Manual mode	1 Gbps	1 Gbps
	2 Gbps	2 Gbps
	4 Gbps	4 Gbps
	8 Gbps	8 Gbps
Auto mode	2 Gbps	Auto, with max of 2 Gbps
	4 Gbps	Auto, with max of 4 Gbps
	8 Gbps	Auto, with max of 8 Gbps

Maximum speed is ensured using the manual settings.

You can specify the port transfer rate using the Navigator 2 GUI, on the **Edit FC Port** screen (Settings/FC Settings/port/Edit Port button).



NOTE: If your remote path is a direct connection, make sure that the disk array power is off when modifying the transfer rate to prevent remote path blockage.

Find details on communication settings in the *Hitachi Unified Storage Hardware Installation and Configuration Guide*.

iSCSI

When using the iSCSI interface at the connection between the arrays, the types of cables and switches used for the Gigabit Ethernet and the 10 Gigabit Ethernet differ. For the Gigabit Ethernet, use the LAN cable and the LAN switch. For the 10 Gigabit Ethernet, use the Fibre cable and the switch usable for 10 Gigabit Ethernet.

The iSCSI remote data path can be set up in the following configurations:

- Direct connection
- Local Area Network (LAN) switch connections
- Wide Area Network (WAN) connections
- WAN Optimization Controller (WOC) connections

Recommendations

The following is recommended for all supported configurations:

- Two paths should be configured from the host to the disk array. This provides a backup path in the event of path failure.

Direct connection

Figure 19-9 illustrates two remote paths directly connecting the local and remote disk arrays with the LAN cable. Direct connections are used when the local and remote disk arrays are set up at the same site. Figure 19-9 shows the configuration when the arrays are directly connected with the LAN cables. One path is allowed between the host and the array. If there are two paths and a failure occurs in a path, the other path can take over.




Figure 19-9: Direct iSCSI connection

Recommendations

- When a large amount of data is to be copied to the remote site, the initial copy between local side and remote systems may be performed at the same location.

Single LAN switch, WAN connection

Figure 19-10 shows two remote paths using one LAN switch and network to the remote disk array.




Figure 19-10: Single-Switch connection

Recommendations

- This configuration is not recommended because a failure in a LAN switch or WAN would halt operations.
- Separate LAN switches and paths should be used for host-to-disk array and disk array-to-disk array, for improved performance.

Connecting Arrays via Switches

Figure 19-11 shows the configuration in which the local and remote arrays are connected via the switches.




Figure 19-11: Connecting arrays with switches

Recommendations

- We recommend you separate the switches, using one for the host I/O and another for the remote copy. If you use one switch for both host I/O and remote copy, the performance may deteriorate.
- Two remote paths should be set. When a failure occurs in one path, the data copy can continue with the other path.

WAN optimization controller (WOC) requirements

WAN Optimization Controller (WOC) is a network appliance that enhances WAN performance by accelerating long-distance TCP/IP communications. TCE copy performance over longer distances is significantly increased when WOC is used. A WOC guarantees bandwidth for each line.

- Use [Table 19-11](#) to determine whether your TCE system requires the addition of a WOC.
- [Table 19-12](#) shows the requirements for WOCs.

Table 19-11: Conditions requiring a WOC

Item	Condition
Latency, Distance	If round trip time is 5 ms or more, or distance between the local site and the remote site is 100 miles (160 km) or further, WOC is highly recommended.
WAN Sharing	If two or more pairs share the same WAN, A WOC is recommended for each pair.

Table 19-12: WOC requirements

Item	Requirements
LAN Interface	Gigabit Ethernet 10 Gigabit Ethernet, or fast Ethernet must be supported.
Performance	Data transfer capability must be equal to or more than bandwidth of WAN.
Functions	• Traffic shaping, bandwidth throttling, or rate limiting must be supported. These functions reduce data transfer rates to a value input by the user. • Data compression must be supported. • TCP acceleration must be supported.

Combining the Network between Arrays

We recommend you separate the switches, using one for the host I/O and another for the remote copy. If you use one switch for both host I/O and remote copy, the performance may deteriorate.

Two remote paths should be set. However, if a failure occurs in the path (a switch or WAN) used commonly by two remote paths (path 0 and path 1), both path 0 and path 1 are blocked. As a result, the path switching becomes impossible and the data copy cannot be continued.

When WOC provides a port of Gigabit Ethernet or 10 Gigabit Ethernet, the switch connected directly to Port 0B and Port 1B in each array is not required. Connect port of each array to WOC directly.

Figure 19-12 shows the configuration that two remote paths use the common network when connecting the local and remote arrays via the switch.




Figure 19-12: Connection via the IP network

Connections with multiple switches, WOCs, and WANs

Figure 19-13 illustrates two remote connections using multiple switches, WOCs, and WANs to make the connection to the remote site.




Figure 19-13: Connections using multiple switches, WOCs, and WANs

Recommendations

- When WOC provides a port of Gigabit Ethernet or 10 Gigabit Ethernet, the switch connected directly to Port 0B and Port 1B in each array is not required. Connect port of each array to WOC directly.
- Using separate LAN switch, WOC and WAN for each remote path ensures that data copy automatically continues on the second path in the event of a path failure.

Multiple array connections with LAN switch, WOC, and single WAN

Figure 19-14 shows two local arrays connected to two remote disk arrays, each via a LAN switch and WOC.




Figure 19-14: Multiple array connection using single WAN

Recommendations

- When WOC provides two or more ports of Gigabit Ethernet or 10 Gigabit Ethernet, the switch connected directly to (for example in Figure 5.14, Port 0B and Port 1B) each array is not required. Connect port of each array to WOC directly.
- Two remote paths should be set. However, if a failure occurs in the path (a switch, WOC or WAN) used commonly by two remote paths (path 0 and path 1), the paths of path 0 and path 1 are blocked. As a result, the path switching is not possible and the data copy cannot be continued.

Multiple array connections with LAN switch, WOC, and two WANs

Figure 19-15 shows two local arrays connected to two remote disk arrays, each via switches and WOCs.




Figure 19-15: Multiple array connection using two WANs

Recommendations

- Two remote paths should be set for each array. Using another path (a switch, WOC or WAN) for every remote path can automatically continue the data copy with another remote path when a failure occurs in one path.
- When WOC provides two or more ports of Gigabit Ethernet or 10 Gigabit Ethernet, the switch connected directly to (for example in Port 0B and Port 1B) each array is not required. Connect port of each array to WOC directly.
- You can reduce the number of switches by using a switch with VLAN capability. If a VLAN switch is used, port 0B of local disk array 1 and WOC1 should be in one LAN (VLAN1); port 0B of local disk array 2 and WOC3 should be in another LAN (VLAN2). Connect the VLAN2 port directly to Port 0B of the local disk array 2 and WOC3.

Local and remote array connection by the switches and WOC

Figure 19-16, shows two sets of a pair of the local array and remote array connected via the switch and WOC.




Figure 19-16: Local and remote array connection by the switches and WOC

Recommendations

- When WOC provides two or more ports of Gigabit Ethernet or 10 Gigabit Ethernet, the switch connected directly to (for example in Port 0B and Port 1B) each array is not required. Connect port of each array to WOC directly.
- You can reduce the number of switches by using a switch with VLAN capability. If a VLAN switch is used, port 0B of local disk array 1 and WOC1 should be in one LAN (VLAN1); port 0B of local disk array 2 and WOC3 should be in another LAN (VLAN2). Connect the VLAN2 port directly to Port 0B of the local disk array 2 and WOC3.

Using the remote path — best practices

The following best practices are provided to reduce and eliminate path failure.

- If both disk arrays are powered off, power-on the remote disk array first.
- When powering down both disk arrays, turn off the local disk array first.
- Before powering off the remote disk array, change pair status to Split. In Paired or Synchronizing status, a power-off results in Failure status on the remote disk array.
- If the remote disk array is not available during normal operations, a blockage error results with a notice regarding SNMP Agent Support Function and TRAP. In this case, follow instructions in the notice.

Path blockage automatically recovers after restarting. If the path blockage is not recovered when the disk array is READY, contact Hitachi Customer Support.

- Power off the disk arrays before performing the following operations:
 - Changing the microcode program (firmware)
 - Setting or changing the fibre transfer rate

Plan and design—disk arrays, volumes and operating systems

This topic provides the information you need to prepare your disk arrays and volumes for TCE operations.

- ❑ [Planning workflow](#)
- ❑ [Supported connections between various models of arrays](#)
- ❑ [Planning volumes](#)
- ❑ [Operating system recommendations and restrictions](#)

Planning workflow

Planning a TCE system consists of determining business requirements for recovering data, measuring production write-workload and sizing DP pools and bandwidth, designing the remote path, and planning your disk arrays and volumes. This topic discusses disk arrays and volumes as follows:

- Requirements and recommendations for using previous versions of AMS with Hitachi Unified Storage.
- Volume set up: volumes must be set up on the disk arrays before TCE is implemented. Volume requirements and specifications are provided.
- Operating system considerations: Operating systems have specific restrictions for replication volumes pairs. These restrictions plus recommendations are provided.
- Maximum Capacity Calculations: Required to make certain that your disk array has enough capacity to support TCE. Instructions are provided for calculating your volumes' maximum capacity.

Supported connections between various models of arrays

Hitachi Unified Storage can be connected with Hitachi Unified Storage, AMS2000, AMS500, or AMS1000. Table 19-13 shows whether connections between various models of arrays are supported or not.

Table 19-13: Supported connections between various models of arrays

Local array	Remote array					
	WMS100	AMS200	AMS500	AMS1000	AMS2000	Hitachi Unified Storage
WMS100	NO	NO	NO	NO	NO	NO
AMS200	NO	NO	NO	NO	NO	NO
AMS500	NO	NO	YES	YES	YES	YES
AMS1000	NO	NO	YES	YES	YES	YES
AMS2000	NO	NO	YES	YES	YES	YES
Hitachi Unified Storage	NO	NO	YES	YES	YES	YES

Connecting HUS with AMS500, AMS1000, or AMS2000

- The maximum number of pairs that can be created is limited to the maximum number of pairs supported by the disk arrays, whichever is fewer.
- When connecting the HUS to each other, if the firmware version of the remote array is under 0916/A, the remote path will be blocked along with the following message:
The firmware version of AMS500/1000 must be 0787/B or later when connecting with HUS100.
- If a HUS as the local array connects to a WMS100, AMS200, AMS500, or AMS1000 with under 0787/B as the remote array, the remote path will be blocked along with the following message:
The firmware version of AMS2000 must be 08B7/B or later when connecting with HUS100
- If a Hitachi Unified Storage as the local array connects to an AMS2010, AMS2100, AMS2300, or AMS2500 with under 08B7/B as the remote array, the remote path will be blocked along with the following message:
 - For Fibre Channel connection:
The target of remote path cannot be connected(Port-xy)
Path alarm(Remote-X,Path-Y)
 - For iSCSI connection:
Path Login failed
- The bandwidth of the remote path to AMS500/1000 must be 20 Mbps or more.
- The pair operation of AMS500/1000 cannot be done from Navigator 2.

- Because AMS500 or AMS1000 has only one data pool per controller, the user cannot specify which data pool to use. For that reason, when connecting AMS500 or AMS1000 with HUS, the specifications about the data pools are:
 - When AMS500 or AMS1000 is the local array, the DP pool 0 is selected if the VOL of the S-VOL is even, and the DP pool 1 is selected if it is odd. In the configuration that the volume numbers of the S-VOL include odd pairs and even pairs, both DP pool 0 and DP pool 1 are required.
 - When HUS is the local array, the data pool number is ignored even if specified. The data pool 0 is selected the owner controller of the S-VOL is 0, and data pool 1 is selected if it is 1.
- AMS500, AMS1000, or AMS2000 cannot use the functions that are newly supported by Hitachi Unified Storage.

Planning volumes

Please review the recommendations in the following sections before setting up TCE volumes. Also, review [TCE system specifications on page D-2](#).

Prerequisites and best practices for pair creation

- Both arrays must be able to communicate with each other via their respective controller 0 and controller 1 ports.
- Bandwidth for the remote path must be known.
- Local and remote arrays must be able to communicate with the Hitachi Storage Navigator 2 server, which manages the arrays.
- The remote disk array ID is required during both initial copy procedures. This is listed on the highest-level GUI screen for the disk array.
- The create pair and resynchronize operations affect performance on the host. Best practice is to perform the operation when I/O load is light.
- For bi-directional pairs (host applications at the local and remote sites write to P-VOLs on the respective disk arrays), creating or resynchronizing pairs may be performed at the same time. However, best practice is to perform the operations one at a time to lower performance impact.
- Use SAS drives or SSD/FMD drives for the primary volume.
- Use SAS drives or SSD/FMD drives for the secondary volume and DP pool.
- Assign a DP pool volume to a distinct RAID group. When another volume is assigned to the same RAID group to which a DP pool volume belongs, the load on drives increases and their performance is reduced. Therefore, it is recommended to assign a DP pool volume to an exclusive RAID group. When there are multiple DP pool volumes in an array, different RAID groups should be used for each DP pool volume.

Volume pair and DP pool recommendations

- The P-VOL and S-VOL must be identical in size, with matching block count. To check block size, in the Navigator 2 GUI, navigate to Groups/RAID Groups/Volumes tab. Click the desired volume. On the popup window that appears, review the Capacity field. This shows block size.
- The number of volumes within the same RAID group should be limited. Pair creation or resynchronization for one of the volumes may impact I/O performance for the others because of contention between drives.

When creating two or more pairs within the same RAID group, standardize the controllers for the volumes in the RAID group. Also, perform pair creation and resynchronization when I/O to other volumes in the RAID group is low.

- Assign a volume consisting of four or more data disks, otherwise host and copying performance may be lowered.
- Limit the I/O load on both local and remote disk arrays to maximize performance. Performance on each disk array also affects performance on the other disk array, as well as DP pool capacity and the synchronization of volumes. Therefore, it is best to assign a volume of SAS drives, SAS7.2K drives, or SSD/FMD drives, and assign four or more disks to a DP pool.

Operating system recommendations and restrictions

The following sections provide operating system recommendations and restrictions.

Host time-out

I/O time-out from the host to the disk array should be more than 60 seconds. You can figure I/O time-out by increasing the remote path time limit times 6. For example, if the remote path time-out value is 27 seconds, set host I/O time-out to 162 seconds (27 x 6) or more.

P-VOL, S-VOL recognition by same host on VxVM, AIX®, LVM

VxVM, AIX®, and LVM do not operate properly when both the P-VOL and S-VOL are set up to be recognized by the same host. The P-VOL should be recognized one host on these platforms, and the S-VOL recognized by a different host.

HP server

When MC/Service Guard is used on a HP server, connect the host group (Fibre Channel) or the iSCSI Target to HP server as follows:


For Fibre Channel interfaces

1. In the Navigator 2 GUI, access the disk array and click **Host Groups** in the Groups tree view. The Host Groups screen appears.
2. Click the check box for the **Host Group** that you want to connect to the HP server.



WARNING! Your host group changes will be applied to multiple ports. This change will delete existing host group mappings and corresponding Host Group IDs, corrupting or removing data associated with the host groups. To keep specified host groups you do not want to remove, please cancel this operation and make changes to only one host group at a time.

- Click **Edit Host Group**.



The **Edit Host Group** screen appears.

WWNs | Volumes | Options

Select or check options for host group.
When Platform or Middleware is set, Mode Setting follow automatically.

Platform: hp-ux | Middleware: not specify

Mode Settings:

Common Settings: Standard

Additional Setting:

HP-UX Mode:	☒ Yes
PSUE Read Reject Mode:	☒ Yes
Mode Parameters Changed Notification Mode:	☐ Yes
NACA Mode:	☐ Yes
Task Management Isolation Mode:	☐ Yes
Unique Reserve Mode 1:	☐ Yes
Port-ID Conversion Mode:	☐ Yes
Tru Cluster Mode:	☐ Yes
Product Serial Response Mode:	☐ Yes
Same Node Name Mode:	☐ Yes
CCHS Mode:	☐ Yes
NOP-In Suppress Mode:	☐ Yes
S-VOL Disable Advanced Mode:	☐ Yes
Discovery CHAP Mode:	☐ Yes
Unique Extended COPY Mode:	☐ Yes
Unique Write Same Mode:	☐ Yes
DP Depletion Detail Reply Mode:	☐ Yes
Unit Attention Change Mode:	☐ Yes
Unique Compare and Write Mode:	☐ Yes

* Required field

- Select the **Options** tab.
- From the Platform drop-down list, select **HP-UX**. Doing this causes **Enable HP-UX Mode**, **Enable PSUE Read Reject Mode**, and **Enable PSUE Read Reject Mode** to be selected in the Additional Setting box.
- Click **OK**. A message appears, click **Close**.

For iSCSI interfaces

1. In the Navigator 2 GUI, access the disk array and click **iSCSI Targets** in the Groups tree view. The iSCSI Targets screen appears.
2. Click the check box for the iSCSI Targets that you want to connect to the HP server.
3. Click **Edit Target**. The Edit iSCSI Target screen appears.
4. Select the **Options** tab.
5. From the Platform drop-down list, select **HP-UX**. Doing this causes "Enable HP-UX Mode" and "Enable PSUE Read Reject Mode" to be selected in the Additional Setting box.
6. Click **OK**. A message appears, click **Close**.

Windows Server 2000

- A P-VOL and S-VOL *cannot* be made into a dynamic disk on Windows Server 2000 and Windows Server 2008.
- Native OS mount/dismount commands can be used for all platforms, except Windows Server 2000. The native commands on this environment do not guarantee that all data buffers are completely flushed to the volume when dismounting. In these instances, you must use CCI to perform volume mount/unmount operations. For more information on the CCI mount/unmount commands, see the *Hitachi Unified Storage Command Control Interface Installation and Configuration Guide*.

Windows Server 2003 or 2008

- A P-VOL and S-VOL *can* be made into a dynamic disk on Windows Server 2003.
- When mounting a volume, use Volume{GUID} as an argument of the CCI mount command (if used for the operation). The Volume{GUID} can be used in CCI versions 01-13-03/00 and later.
- In Windows Server 2008, refer to the *Hitachi Unified Storage Command Control Interface Installation and Configuration Guide* for the restrictions when the mount/unmount command is used.
- Windows may write for the un-mounted volume. If a pair is resynchronized while retaining data on the S-VOL on the server memory, the compatible backup cannot be collected. Therefore, execute the CCI sync command immediately before re-synchronizing the pair for the un-mounted S-VOL.
- In Windows Server 2008, set only the P-VOL of TCE to be recognized by the host and let another host recognize the S-VOL.
- (CCI only) When describing a command device in the configuration definition file, specify it as Volume{GUID}.
- (CCI only) If a path detachment is caused by controller detachment or Fibre Channel failure, and the detachment continues for longer than one minute, the command device may not be recognized when recovery occurs. In this case, execute the "re-scanning of the disks" in

Windows. If Windows cannot access the command device, though CCI recognizes the command device, restart CCI.

Windows Server and TCE configuration volume mount

In order to make a consistent backup using a storage-based replication such as TCE, you must have a way to flush the data residing on the server memory to the array, so that the source volume of the replication has the complete data.

You can flush the data on the server memory using the **umount** command of CCI to unmount the volume. When using the **umount** command of CCI for unmount, use the **mount** command of CCI for mount.

If you are using Windows Server™ 2003, **mountvol /P** to flush data on the server memory when un-mounting the volume is supported. Please understand the specification of the command and run sufficient test before you use it for your operation.

In Windows Server™ 2008, refer to the *Command Control Interface (CCI) Reference Guide* for the restrictions when mount/unmount command is used.

Windows Server may write for the un-mounted volume. If a pair is resynchronized while remaining the data to the S-VOL on the memory of the server, the compatible backup cannot be collected. Therefore, execute the **sync** command of CCI immediately before re-synchronizing the pair for the un-mounted S-VOL.

For more detail about the CCI commands, see the *Command Control Interface (CCI) Reference Guide*.

Volumes to be recognized by the same host

If you recognize the P-VOL and S-VOL on Windows Server 2008 at the same time, it may cause an error because the P-VOL and S-VOL have the same disk signature. When the P-VOL and S-VOL have the same data, split the pair and then rewrite the disk signature so that they can retain different disk signatures. You can use the **uniqueid** command to rewrite a disk signature. See the *Hitachi Unified Storage Command Control Interface Installation and Configuration Guide* for details.

Identifying P-VOL and S-VOL in Windows

In Navigator 2, the P-VOL and S-VOL are identified by their volume number. In Windows, volumes are identified by HLUN. These instructions provide procedures for the fibre channel and iSCSI interfaces. To confirm the HLUN:

1. From the Windows Server 2003 Control Panel, select Computer Management/Disk Administrator.

2. Right-click the disk whose HLUN you want to know, then select **Properties**. The number displayed to the right of "VOL" in the dialog window is the HLUN.

For Fibre Channel interface:

Identify HLUN-to-VOL Mapping for the Fibre Channel interface, as follows:

1. In the Navigator 2 GUI, select the desired disk array.
2. In the array tree, click the **Group** icon, then click **Host Groups**.



WARNING! Your host group changes will be applied to multiple ports. This change will delete existing host group mappings and corresponding Host Group IDs, corrupting or removing data associated with the host groups. To keep specified host groups you do not want to remove, please cancel this operation and make changes to only one host group at a time.

3. Click the **Host Group** to which the volume is mapped.
4. On the screen for the host group, click the **Volumes** tab. The volumes mapped to the Host Group display. You can confirm the VOL that is mapped to the H-LUN.

For iSCSI interface:

Identify HLUN-to-VOL Mapping for the iSCSI interface as follows.

1. In the Navigator 2 GUI, select the desired array.
2. In the array tree that displays, click the **Group** icon, then click the **iSCSI Targets** icon in the Groups tree.
3. On the iSCSI Target screen, select an iSCSI target.
4. On the target screen, select the **Volumes** tab. Find the identified HLUN. The VOL displays in the next column.
5. If the HLUN is not present on a target screen, on the iSCSI Target screen, select another iSCSI target and repeat Step 4.

Dynamic Disk in Windows Server

- In an environment of Windows Server, you cannot use TCE pair volumes as dynamic disk. The reason for this restriction is because in this case if you restart Windows or use the Rescan Disks command after creating or re-synchronizing a TCE pair, there are cases where the S-VOL is displayed as foreign in disk management and becomes inaccessible.

UNMAP Short Length Mode

Enable UNMAP Short Length Mode when connecting to Windows 2012. If you do not enable it, UNMAP commands may not be completed due to a time-out.

VMware and TCE configuration

When creating a backup of the virtual disk in the vmfs format using TCE, shutdown the virtual machine that accesses the virtual disk, and then split the pair.

If one volume is shared by multiple virtual machines, shutdown all the virtual machines that share the volume when creating a backup. It is not recommended to share one volume by multiple virtual machines in the configuration that creates a backup using TCE.

The VMware ESX has a function to clone the virtual machine. Although the ESX clone function and TCE can be linked, cautions are required for the performance at the time of execution. For example, when the volume which becomes the ESX clone destination is a TCE P-VOL pair whose pair status is Paired, since the data is written to the S-VOL for writing to the P-VOL, the time required for a clone may become longer and the clone may be terminated abnormally in some cases. To avoid this, we recommend the operation to make the TCE pair status Split or Simplex and to resynchronize or create the pair after executing the ESX clone. Also, it is the same for executing the functions such as migration the virtual machine, deploying from the template and inflating the virtual disk.




Figure 19-17: VMware ESX

UNMAP Short Length Mode

It is recommended that you enable UNMAP Short Length Mode when connecting to VMware. If you do not enable it, UNMAP commands may not be completed due to a time-out.

Changing the port setting

If the port setting is changed during the firmware update, the remote path may be blocked or the remote pair may be changed to Failure. Change the port setting after completing the firmware update.

If the port setting is changed in the local array and the remote array at the same time, the remote path may be blocked or the remote pair may be changed to Failure. Change the port setting by taking an interval of 30 seconds or more for every port change.

Concurrent use of Dynamic Provisioning

The DP-VOLs can be set for a P-VOL or an S-VOL of TCE.

The points to keep in mind when using TCE and Dynamic Provisioning together are described here. Refer to the *Hitachi Unified Storage Hitachi Unified Storage Dynamic Provisioning Configuration Guide* for detailed information about Dynamic Provisioning. Hereinafter, the volume created in the RAID group is called a normal volume and the volume created in the DP pool is called a DP-VOL.

- Volume type that can be set for a P-VOL or an S-VOL of TCE.

The DP-VOL can be used for a P-VOL or an S-VOL of TCE. [Table 19-14](#) shows a combination of a DP-VOL and a normal volume that can be used for a P_VOL or an S-VOL of TCE.

Table 19-14: Combination of a DP-VOL and a normal volume

TCE P-VOL	TCE S-VOL	Contents
DP-VOL	DP-VOL	Available. The P-VOL and S-VOL capacity can be reduced compared to the normal volume. (Note 1)
DP-VOL	Normal volume	Available. In this combination, copying after pair creation takes about the same time it takes when the normal volume is a P-VOL. Moreover, when executing the swap, the DP pool of the same capacity as the normal volume (original S-VOL) is used. After the pair is split and reclaim zero page, the S-VOL capacity can be reduced.
Normal volume	DP-VOL	Available. When the pair status is Split, the S-VOL capacity can be reduced compared to the normal volume by reclaim zero page.



NOTES:

1. When creating a TCE pair using the DP-VOLs, in the P-VOL or the S-VOL specified at the time of the TCE pair creation, the DP-VOLs whose Full Capacity Mode is enabled and disabled cannot be mixed.
2. Depending on the volume usage, the consumed capacity of the P-VOL and the S-VOL may differ even in the Paired status. Execute the DP Optimization and zero page reclaim as needed.
3. The consumed capacity of the S-VOL may be reduced due to the resynchronization.

- Pair status at the time of DP pool capacity depletion

When the DP pool is depleted after operating the TCE pair that uses the DP-VOL, the pair status of the pair concerned may be a Failure.

[Table 19-15](#) shows the pair statuses before and after the DP pool

capacity depletion. When the pair status becomes a Failure caused by the DP pool capacity depletion, add the DP pool capacity whose capacity is depleted, and execute the pair operation again.

Table 19-15: Pair statuses before and after the DP pool capacity depletion

Pair statuses before the DP pool capacity depletion belonging to P-VOL or S-VOL	Pair statuses after the DP pool capacity depletion belonging to P-VOL		Pair statuses after the DP pool capacity depletion belonging to S-VOL	
	P-VOL pair status	S-VOL pair status	P-VOL pair status	S-VOL pair status
Simplex	Simplex	Simplex	Simplex	Simplex
Synchronizing	Synchronizing	Synchronizing	Failure ²	Synchronizing
Reverse Synchronizing	Reverse Synchronizing	Reverse Synchronizing	Failure ²	Reverse Synchronizing
Paired	Paired Failure ¹	Paired Failure	Failure ²	Paired
Split	Split	Split	Split	Split
Failure	Failure	Failure	Failure	Failure
Notes: 1. When write is performed to the P-VOL to which the capacity depletion DP pool belongs, the copy cannot be continued and the pair status becomes a Failure. 2. The remote path on the local array will fail.				

- DP pool status and availability of pair operation

When using the DP-VOL for a P-VOL or an S-VOL of the TCE pair, the pair operation may not be executed depending on the status of the DP pool to which the DP-VOL belongs. [Table 19-16](#) and [Table 19-17](#) show the DP pool status and availability of the TCE pair operation. When the pair operation fails due to the DP pool status, correct the DP pool status and execute the pair operation again YES.

Table 19-16: DP pool for P-VOL statuses and availability of pair operation

Pair Operation	DP pool statuses, DP pool capacity statuses, and DP pool optimization statuses					
	Normal	Capacity in growth	Capacity depletion	Regressed	Blocked	DP in optimization
Create pair	YES*	YES	NO *	YES	NO	YES
Split pair	YES	YES	YES	YES	YES	YES
Resync pair	YES *	YES	NO *	YES	NO	YES
Swap pair	YES *	YES	NO *	YES	YES	YES
Delete pair	YES	YES	YES	YES	YES	YES
* Refer to the status of the DP pool to which the DP-VOL of the S-VOL belongs. If the pair operation causes the DP pool belonging to the S-VOL to be fully depleted, the pair operation cannot be executed.						

YES indicates a possible case

NO indicates an unsupported case

Table 19-17: DP pool for S-VOL statuses and availability of pair operation

Pair operation						
	Normal	Capacity in growth	Capacity depletion	Regressed	Blocked	DP in optimization
Create pair	YES *	YES	NO *	YES	NO	YES
Split pair	YES	YES	YES	YES	YES	YES
Resync pair	YES *	YES	NO *	YES	YES	YES
Swap pair	YES*	YES	YES *	YES	NO	YES
Delete pair	YES	YES	YES	YES	YES	YES
* Refer to the status of the DP pool to which the DP-VOL of the P-VOL belongs. If the pair operation causes the DP pool belonging to the S-VOL to be fully depleted, the pair operation cannot be executed.						

YES indicates a possible case

NO indicates an unsupported case

- Formatting the DP pool

When the DP pool was created or the capacity was added, the formatting operates for the DP pool. If pair creation, pair resynchronization, or swapping is performed during the formatting, depletion of the usable capacity may occur. Since the formatting progress is displayed when checking the DP pool status, check if the sufficient usable capacity is secured according to the formatting progress, and then start the operation.

- Operation of the DP-VOL during TCE use

When using the DP-VOL for a P-VOL or an S-VOL of TCE, any of the operations among the capacity growing, capacity shrinking, volume deletion, and Full Capacity Mode changing volume deletion of the DP-VOL in use cannot be executed. To execute the operation, delete the TCE pair of which the DP-VOL to be operated is in use, and then perform it again.

- Operation of the DP pool during TCE use

When using the DP-VOL for a P-VOL or an S-VOL of TCE, the DP pool to which the DP-VOL in use belongs cannot be deleted. To execute the operation, delete the TCE pair of which the DP-VOL is in use belonging to the DP pool to be operated, and then execute it again. The attribute edit and capacity addition of the DP pool can be executed usually regardless of the TCE pair.

- Caution for DP pool formatting, pair resynchronization, and pair deletion

Continuously performing DP pool formatting, pair resynchronization, or pair deletion to a pair with a lot of replication data or management information can lead to temporary depletion of the DP pool, where $\text{used capacity (\%)} + \text{capacity in formatting (\%)} = \text{about 100\%}$, and it makes the pair change to Failure. Perform pair resynchronization and pair deletion when sufficient available capacity has been ensured.

- Cascade connection

A cascade can be performed on the same conditions as the normal volume. See [Cascading TCE on page 27-78](#).

- Pool shrink

Pool shrink is not possible for the replication data DP pool and management area DP pool. If you need to shrink the pool, delete all the pairs that use the DP pool.

Concurrent use of Dynamic Tiering

The considerations for using the DP pool whose tier mode is enabled by using Dynamic Tiering are described. For the detailed information related to Dynamic Tiering, refer to *Hitachi Unified Storage 100 Dynamic Tiering User's Guide*. Other considerations are common with Dynamic Provisioning.

In the replication data DP pool and the management area DP pool whose tier mode is enabled, the replication data and the management information are not placed in the Tier configured by SSD/FMD. Therefore, note the following points:

- The DP pool whose tier mode is enabled and configured only by SSD/FMD cannot be specified as the replication data DP and the management area DP pool.
- The total of the free capacity of the Tier configured by the drive other than SSD/FMD in the DP pool is the total free capacity for the replication data or the management information.
- When the free space of the replication data DP pool and the management area DP pool whose tier mode is enabled decreases, recover the free space of Tier other than SSD/FMD.
- When the replication data and the management information are stored in the DP pool whose tier mode is enabled, they are first assigned to 2nd Tier.
- The area where the replication data and the management information are assigned is out of the relocation target.

Load balancing function

The Load balancing function applies to a TCE pair. When the pair state is Paired and cycle copy is being performed, the load balancing function does not work.

Enabling Change Response for Replication Mode

If background copy to the pool times-out for some reason while write commands are being executed on the P-VOL in Paired state, or if restoration to the S-VOL times-out while read commands are being executed on the S-VOL in Paired Internally Busy state (including Busy state), the array returns Medium Error (03) to the host. Some hosts receiving Medium Error (03) may determine the P-VOL or S-VOL inaccessible, and stop accessing it. In such cases, enabling the Change Response for Replication Mode makes the array return Aborted Command (0B) to the host. When the host receives Aborted Command (0B), it will retry the command to the P-VOL or S-VOL, and the operation will continue.

User data area of cache memory

Because TCE uses DP pools to work, Dynamic Provisioning/Dynamic Tiering are necessary to operate at the same time. Employing Dynamic Provisioning/Dynamic Tiering acquires some portion of the installed cache memory, which reduces the user area of the cache memory. [Table 19-18](#) shows the cache memory secured capacity and the user data area at the time of using the program product.

The performance effect by reducing the user data area is shown when a large amount of sequential write was executed at the same time, but it is deteriorated by a few percent at the time of writing 100 volumes at the same time.

Table 19-18: User data area capacity of cache memory

ray pe	Cache memory per CTL	DP capacity mode	Managem ent capacity for Dynamic Provisioni ng	Managem ent capacity for Dynamic Tiering	User data capacity		
					When Dynamic Provisioning is enabled	When Dynamic Provisioning and Dynamic Tiering are Enabled	When Dynamic Provisioning and Dynamic Tiering are disabled
110	4 GB/CTL	Not supported	420 MB	50 MB	1,000 MB	960 MB	1,420 MB
130	8 GB/CTL	Regular Capacity	640 MB	200 MB	4,020 MB	3,820 MB	4,660 MB
		Maximum Capacity	1,640 MB	200 MB	3,000 MB	2,800 MB	4,660 MB
	16 GB/CTL	Regular Capacity	640 MB	200 MB	10,640 MB	10,440 MB	11,280 MB
		Maximum Capacity	1,640 MB	200 MB	9,620 MB	9,420 MB	11,280 MB
150	8 GB/CTL	Regular Capacity	1,640 MB	200 MB	2,900 MB	2,700 MB	4,540 MB
	16 GB/CTL	Regular Capacity	1,640 MB	200 MB	9,520 MB	9,320 MB	11,160 MB
		Maximum Capacity	3,300 MB	200 MB	7,860 MB	7,660 MB	11,160 MB

Setup procedures

The following sections provide instructions for pair assignment , setting up the DP pools, replication threshold, CHAP secret (iSCSI only), and remote path.

Pair Assignment

- Use a small number of volumes within the same RAID group.

When volumes are assigned to the same RAID group and used as pair volumes, there may be a case where a pair creation or resynchronization for one of the volumes causes restrictions to be placed on performance of host I/O, pair creation, resynchronization, etc. for the other volumes because of contention between drives. It is recommended to assign one or two volumes to be paired to the same RAID group.

- Use SAS drives or SSD/FMD drives for the primary volume.

When a P-VOL is located in a RAID group consisting of the SAS7.2K drives, performance of a host I/O, pair creation, and pair resynchronization, etc. is lowered because of the lower performance of the SAS7.2K drive. Therefore, it is recommended to assign a primary volume to a RAID group consisting of the SAS drives or the SSD/FMD drives. **Ordinary Paired** status with a configuration of the SAS7.2K drives is not recommended.

- Use SAS drives or SSD/FMD drives for the secondary volume and DP pool.

When the SAS7.2K drive is used for the S-VOL or the primary or secondary DP pool, there is a higher possibility that a suspension failure could occur in the TCE pair. Data transfer is also slower than when a SAS drive or an SSD/FMD drive is used. Therefore, it is recommended to use the SAS drive or the SSD/FMD drive. Ordinary **Paired** status with a configuration of the SAS7.2K drives is not recommended.

- Perform pair creation and resynchronization when I/O load is minimal.

When a pair is newly created, or resynchronization of a pair is done when a TCE pair is in the **Paired** status, the volume of differential data may become larger because the transfer of the differential data between the TCE pair is delayed. When the delay continues, the data in the DP pool could overflow and the TCE pair is split and placed in the **Pool Full** status. It is recommended to perform the system operation such as pair creation or pair resynchronization when I/O load is minimal (in the night or on a holiday). It is also recommended to limit the number of the pairs in the **Synchronizing** status to one or two at a time.

- Assign a DP pool to a distinct RAID group.

When another volume is assigned to the same RAID group to which a DP pool belongs, the load on drives increases and their performance is reduced. Therefore, it is recommended to assign a DP pool to an exclusive RAID group. When there are multiple DP pools in an array, different RAID groups should be used for each DP pool volume.

- It is recommended to set separate DP pools to the replication data DP pool and the management area DP pool

While using TCE, the replication data DP pool and the management area DP pool are frequently accessed. So if the replication data DP pool and the management area DP pool are set to the same DP pool, it can negatively affect the overall performance of Snapshot, as the identical DP pool is quite frequently accessed. Therefore, it is highly recommended to set separate DP pools to the replication data DP pool and the management area DP pool. See [Figure 19-18 on page 19-52](#).

However, when using CCI for pair creation, you cannot specify separate DP pools for the replication data DP pool and the management area DP pool. Use HSNM2 for pair creation.




Figure 19-18: Set separate DP pools

- Assign four or more disks to the data disks.

When the data disks that comprise a RAID group are not sufficient, it adversely affects the host performance and/or copying performance by restricting read/write operations to the drives. Therefore, when operating pairs with TCE, it is recommended to use a volume consisting of four or more data disks.

- For bi-directional TCE pairs, perform a pair creation/resynchronization in each of the directions one after the other.

In a configuration where each site hosts as both local and remote site (as shown in [Figure 19-19 on page 19-53](#)), creation and resynchronization of a pair may be performed from the both sites at the same time. In such a case,

the time interval between pair creation until completion of the resynchronization becomes longer and the influence on performance of the other operation becomes greater. This is due to reading and writing of data in parallel at the both sites. Therefore, it is recommended to perform pair creation and resynchronization in each of the two directions one after the other.

When constructing a configuration, consider that there may be a case where the cycles in the both directions overlap each other because the differential data of the TCE pair is regularly copied in the cycle even when the pair is in the **Paired** status.




Figure 19-19: Bidirectional TCE Operation

Setting up DP pools


For directions on how to set up a DP pool, refer to the *Hitachi Unified Storage Dynamic Provisioning Configuration Guide*.

To set the DP pool capacity, see the [DP pool size on page 19-5](#).

Setting the replication threshold (optional)

To set the Depletion Alert and/or the Replication Data Released of the replication threshold:

- 1. Select the **Volumes** icon in the **Group** tree view.
- 2. Select the **DP Pools** tab.
- 3. Select the DP pool number for the replication threshold that you want to set.



- Click **Edit Pool Attribute**.

DP Pool-000
HUS150_92012345 > Groups > Volumes > DP Pool-000

Properties		Notifications	Volumes	Assigned Drives	DP RAID Groups
DP Pool	000				
Tier Mode	Disabled				
RAID Level	RAID5				
Combination	8D+2P				
Page Size	32MB				
Stripe Size	256KB				
Drive Type/RPM	SAS/10K				
Status	Normal				
Reconstruction Progress	N/A				
Capacity	Total	8.0TB			
	Replication Available	8.0TB			
	Consumed by	Total	2.0TB		
		User Data	683.0GB		
		Replication Data	410.0GB		
	Management Area	512.0GB			
Needing Preparation	10.6GB				

- Enter the **Replication Depletion Alert Threshold** and/or the **Replication Data Released Threshold** in the **Replication** field.

DP Pool Attribute Properties
Edit the properties of the DP pool.

Basic

DP Pool:	000
Tier Mode:	Disabled
RAID Level:	RAID5(8D+2P)
Stripe Size:	256KB
Drive Type/RPM:	SAS/10K
Total Capacity:	8.0TB
Consumed Capacity:	2.0TB

* DP Pool Consumed Capacity:

Early Alert Threshold: %
 Depletion Alert Threshold: %
From 1 to 99

Send notification by E-mail alert and SNMP if the current value exceeds threshold: ☒ Yes

* Over Provisioning:

Warning Threshold: %
 Limit Threshold: %
From 50 to 1000

Send notification by E-mail alert and SNMP if the current value exceeds threshold: ☐ Yes

Limit Enforcement: ☐ Enable

* Replication:

Replication Depletion Alert Threshold: %
 Replication Data Released Threshold: %
From 1 to 99

* Required field

- Click **OK**.

7. A message appears. Click **Close**.



Setting the cycle time

Set the cycle time in which the remote copy of the differential data of the pair in the Paired status is made using Navigator 2. The cycle time is to be set for each array. The shortest limit value that can be set is calculated as number of CTGs of local array or remote array × 30 seconds.



NOTE: The copy may take a time longer than the cycle time specified, depending on a scale of the amount of the differential data or because of a low line speed.

To set the cycle time :

1. Select **Options** icon in the **Setup** tree view of the **Replication** tree view. The **Options** screen appears.



2. Click **Edit Options**. The **Edit Options** screen appears.



3. Enter the value to set for the cycle time in the text box of the cycle time. The lower limit is 30 seconds.
4. Click **OK**.
5. The confirmation message is displayed. Click Close Click **Close**.



Adding or changing the remote port CHAP secret

(For disk arrays with iSCSI connectors only)

Challenge-Handshake Authentication Protocol (CHAP) provides a level of security at the time that a link is established between the local and remote disk arrays. Authentication is based on a shared secret that validates the identity of the remote path. The CHAP secret is shared between the local and remote disk arrays.

- CHAP authentication is automatically configured with a default CHAP secret when the TCE Setup Wizard is used. You can change the default secret if desired.
- CHAP authentication is not configured when the Create Pair procedure is used, but it can be added.

Prerequisites

- Disk array IDs for local and remote disk arrays are required.

To add a CHAP secret

This procedure is used to add CHAP authentication manually on the remote disk array.

1. On the remote disk array, navigate down the GUI tree view to **Replication / Setup / Remote Path**. The Remote Path screen appears. (Though you may have a remote path set, it does not show up on the remote disk array. Remote paths are set from the local disk array.)
2. Click the **Remote Port CHAP** tab. The Remote Port CHAP screen appears.
3. Click the **Add Remote Port CHAP** button. The Add Remote Port CHAP screen appears.
4. Enter the **Local disk array ID**.
5. Enter CHAP **Secrets** for Remote Path 0 and Remote Path 1, following on-screen instructions.
6. Click **OK** when finished.
7. The confirmation message appears. Click **Close**.

Add Remote Port CHAP



Remote port CHAP added successfully.

Close

Close

To change a CHAP secret

1. Split the TCE pairs, after confirming first that the status of all pairs is Paired.
 - To confirm pair status, see [Monitoring pair status on page 21-3](#).
 - To split pairs, see [Splitting a pair on page 20-6](#).
2. On the local disk array, delete the remote path. Be sure to confirm that the pair status is Split before deleting the remote path. See [Deleting the remote path on page 21-20](#).
3. Add the remote port CHAP secret on the remote disk array. See the instructions above.
4. Re-create the remote path on the local disk array. See [Setting the remote path on page 19-58](#).

For the CHAP secret field, select **manually** to enable the CHAP Secret boxes so that the CHAP secrets can be entered. Use the CHAP secret added on the remote disk array.
5. Resynchronize the pairs after confirming that the remote path is set. See [Resynchronizing a pair on page 20-8](#).

Setting the remote path

A remote path is the data transfer connection between the local and remote disk arrays.

- Two paths are recommended; one from controller 0 and one from controller 1.
- Remote path information cannot be edited after the path is set up. To make changes, it is necessary to delete the remote path then set up a new remote path with the changed information.
- Use the Create Remote Path procedure, described below.

Prerequisites

- Both local and remote disk arrays must be connected to the network for the remote path.
- The remote disk array ID will be required. This is shown on the main disk array screen.
- Network bandwidth will be required.
- For iSCSI, the following additional information is required:

- For iSCSI array model, you can identify the IP address for the remote path in the IPv4 or IPv6 format. Be sure to use the same format when specifying the port IP addresses for the remote path for the local array and the remote array.
- Set the remote paths from the controller 0 to the other controller 0 and from the controller 1 to the other controller 1.
- Remote IP address, listed in the remote disk array's GUI Settings/IP Settings
- TCP port number. You can see this by navigating to the remote disk array's GUI Settings/IP Settings/selected port screen.
- CHAP secret (if specified on the remote disk array—see [Setting the cycle time on page 19-56](#) for more information).

To set up the remote path

1. On the local disk array, from the navigation tree, select the **Remote Path** icon in the **Setup** tree view in the **Replication** tree.
2. Click **Create Path**. The **Create Remote Path** screen appears.

3. For **Interface Type**, select **Fibre** or **iSCSI**.
4. Enter the **Remote disk array ID**.
Use default value for Remote Path Name: The Remote Path Name named to Array_Remote Array ID.
Enter Remote Path Name Manually: Enter the characters strings the displaying characters.
5. Enter the bandwidth number into the **Bandwidth** field. Select **Over 1000,0Mbps** in the **Bandwidth** for over 1000,0 Mbps network bandwidth. When connecting the array directly to the other array, set the bandwidth according to the transfer rate.

6. (iSCSI only) In the CHAP secret field, select **Automatically** to allow TCE to create a default CHAP secret, or select **manually** to enter previously defined CHAP secrets. The CHAP secret must be set up on the remote disk array.
7. In the two remote path boxes, Remote Path 0 and Remote Path 1, select local ports. For iSCSI, Specify the following items for the Remote Path 0 and Remote Path 1'
 - **Local Port:** Select the port number that connected to the remote path. The IPv4 or IPv6 format can be used to specify the IP address
 - **Remote Port IP Address:** Specify the remote port IP address that connected to the remote path.
8. (iSCSI only) When the CHAP secret specifies to the remote port, enter the specified characters to the CHAP Secret.
9. Click **OK**.
10. A message appears. Click **Close**.

Deleting the remote path

When the remote path becomes unnecessary, delete the remote path.

Prerequisites

- The pair status of the volumes using the remote path to be deleted must be Simplex or Split
- Do not do a pair operation for a TCE pair when the remote path for the pair is not set up, because the pair operation may not complete correctly.



NOTE: When performing the planned shutdown of the remote array, the remote path should not necessarily be deleted. Change all the TCE pairs in the array to the Split status, and then perform the planned shutdown of the remote array. After restarting the array, perform the pair resynchronization. However, when the Warning notice to the Failure Monitoring Department at the time of the remote path blockade or the notice by the SNMP Agent Support Function or the E-mail Alert Function is not desired, delete the remote path, and then turn off the power of the remote array.

To delete the remote path

1. Connect the local array, and select the **Remote Path** icon in the **Setup** tree view in the **Replication** tree. The **Remote Path** list appears.
2. Select the remote path you want to delete in the **Remote Path** list and click **Delete Path**.
3. A message appears. Click **Close**.

Operations work flow

TCE is a function for the asynchronous remote copy between the volumes in the arrays connected via the remote path. Data written in the local array by the host is written in the remote array by TCE regularly. At the time of the pairing or resynchronization, data on the local side can be transferred to the remote side in a short time because only the differential data is transferred to the remote array. Furthermore, it realizes a function to get the snapshot from the remote volume according to an instruction from the host connected to the local array.



NOTE: When turning the power of the array on or off, restarting the array, replacing the firmware, and changing the transfer rate setting, be careful of the following items.

- When you turn on the array where a path has already been set, turn on the remote array first. Turn on the local array after the remote array is READY.
When you turn off the array where a path has already been set, turn off the local array first, and then turn off the remote array.
- When you restart the array, verify that the array is on the remote side of TCE. When the array on the remote side is restarted, both paths are blocked.
- When the array on the remote side is powered off or restarted when the TrueCopy pair status is Paired or Synchronizing, it is changed to Failure. If you power off or restart the array, do so after changing the TrueCopy pair status to Split. When the power of the remote array is turned off or restarted, the remote path is blocked. However, by changing the pair status to Split, it can be prevented that the pair status is changed to Failure. When the Warning notice to the Failure Monitoring Department at the time of the remote path blockade or the notice by the SNMP Agent Support Function or the E-mail Alert Function is not desired, delete the remote path, and then turn off the power of the remote array.
- You will receive an error/blockage if the remote array is not available. When a remote path is blocked, a TRAP occurs, that is, a notification to the SNMP Agent Support Function. The remote path of TCE is recovered from the blockade automatically after the array is restarted. If the remote path blockage is not recovered when the array is READY, contact Hitachi Support. The time until the array status changes to READY after turning on the power is about six minutes or shorter even when the array has the maximum configuration. The time required varies depending on the array configuration.
- Do not change the firmware when the pair status is **Synchronizing** or **Paired**. If you change the firmware, be sure to perform it after splitting the pair.
- When a fibre channel interface, the local array is directly connected with the remote array paired with the local array, the setting of the fibre transfer rate must not be modified while the array power is on. If the setting of the fibre transfer rate is modified, a remote path blockage will occur.

- When you replace the remote array with a different array, ensure that you have deleted all the TCE pairs and the remote path before changing the remote array.

Using TrueCopy Extended

This chapter provides procedures for performing basic TCE operations using the Navigator 2 GUI. Appendixes with CLI and CCI instructions are included in this manual.

- ❑ [TCE operations](#)
- ❑ [Checking pair status](#)
- ❑ [Creating the initial copy](#)
- ❑ [Splitting a pair](#)
- ❑ [Resynchronizing a pair](#)
- ❑ [Swapping pairs](#)
- ❑ [Editing pairs](#)
- ❑ [Deleting pairs](#)
- ❑ [Example scenarios and procedures](#)

TCE operations

Basic TCE operations consist of the following:

- Checking pair status. Each operation requires the pair to be in a specific status.
- Creating the pair, in which the S-VOL becomes a duplicate of the P-VOL.
- Splitting the pair, which stops updates from the P-VOL to the S-VOL and allows read/write of the S-VOL.
- Re-synchronizing the pair, in which the S-VOL again mirrors the on-going, current data in the P-VOL.
- Swapping pairs, which reverses pair roles.
- Deleting a pair.
- Editing pair information.

These operations are described in the following sections. TCE operations can be performed from the UNIX/PC host using CCI software and/or Navigator 2.

Confirm that the state of the remote path is **Normal** before doing a pair operation. If you do a pair operation when a remote path is not set up or its state is **Diagnosing** or **Blocked**, the pair operation may not be completed correctly.

Checking pair status

Each TCE operation requires a specific pair status. Before performing any operation, check pair status.

- Find an operation's status requirement in the Prerequisites sections below.
- To monitor pair status, refer to [Monitoring pair status on page 21-3](#).

Creating the initial copy

Two methods are used for creating the initial TCE copy:

- The GUI setup wizard, which is the simplest and quickest method. Includes remote path setup.
- The GUI Create Pair procedure, which requires more setup but allows for more customizing.

Both procedures are described in this section.

During pair creation:

- All data in the P-VOL is copied to the S-VOL.
- The P-VOL remains available to the host for read/write.
- Pair status is Synchronizing while the initial copy operation is in progress.
- Status changes to Paired when the initial copy is complete.

Create pair procedure

With the Create Pair procedure, you create a TCE pair and specify copy pace, consistency groups, and other options. Please review the prerequisites on page 20-3 before starting. You will be Creating a volume whose size are the same as those of the backup target volume to the remote array.

To create a pair using the Create Pair procedure

1. In Navigator 2 GUI, select the desired disk array, then click the **Show & Configure disk array** button.
2. From the Replication tree, select the **Remote Replication** icon. The **Pairs** screen appears.
3. Select the **Create Pair** button at the bottom of the screen The **Create Pair** screen appears.

Create Pair [Help]

Remote Pair Properties

Enter property information as desired.

CopyType : TrueCopy

Basic **Advanced**

Specify the primary volume and the secondary volume to pair and input a unique pair name. And specify a group to assign after pair creation.

Pair Name :

31 characters or less (alphanumeric characters, special symbols *%, *+, *4*, *- , *., */*, * =*, *@*, *_ , *_, *_, *_, or *_)

Remote Array : **Assignable Remote Arrays**

Remote Array ID	Remote Path Name
91200026	Array_91200026

* Select a Primary Volume : **Primary Volume**

VOL	Capacity	RAID Group	DP Pool
0001	50.0GB	N/A	000
0002	10.0GB	000	N/A

* Secondary Volume :

After pair creation, add the pair to a group. - Group Assignment:

☒ (Ungrouped)

☐ New or existing Group Number :

☐ Existing Group Name :

31 characters or less (alphanumeric characters, special symbols *%, *+, *4*, *- , *., */*, * =*, *@*, *_ , *_, *_, *_, or *_)

Available Pair Capacity : 0.0MB

* Required field

OK Cancel

4. On the Create Pair screen that appears, confirm that the **Copy Type** is TCE and enter a name in the **Pair Name** box following on-screen guidelines. If omitted, the pair is assigned a default name (TCE_LUxxxx_LUyyyy: xxxx is Primary Volume, yyyy is Secondary

Volume). In either case, the pair is named in the local disk array, but not in the remote disk array. On the remote disk array, the pair appears with no name. Add a name using Edit Pair.

5. Select a Primary Volume, and enter a Secondary Volume



NOTE: In Windows 2003 Server, volumes are identified by HLUN. The VOL and H-LUN may be different. See [Identifying P-VOL and S-VOL in Windows on page 19-41](#) to map VOL to HLUN.


6. Select **Automatic** or **Manual** for the DP Pool. When you select the Manual, select a DP Pool Number of local array from the drop-down list.
7. When you select **Manual**, enter a DP Pool Number of remote array.
8. For **Group Assignment**, you assign the new pair to a consistency group.
 - To create a group and assign the new pair to it, click the **New or existing Group Number** button and enter a new number for the group in the box.
 - To assign the pair to an existing group, enter its number in the **Group Number** box, or enter the group name in the **Existing Group Name** box.
 - If you do not want to assign the pair to a consistency group, they will be assigned automatically. Leave the **New or existing Group Number** button selected with no number entered in the box.



NOTE: You can also add a Group Name for a consistency group as follows:

- a. After completing the create pair procedure, on the **Pairs** screen, check the box for the pair belonging to the group.
- b. Click the **Edit Pair** button.
- c. On the Edit Pair screen, enter the **Group Name** and click **OK**.

9. Select the **Advanced** tab.



10. From the **Copy Pace** drop-down list, select a pace. Copy pace is the rate at which a pair is created or resynchronized. The time required to complete this task depends on the I/O load, the amount of data to be copied, cycle time, and bandwidth. Select one of the following:
 - Slow — The option takes longer when host I/O activity is high. The time to copy may be quite lengthy.
 - Medium — (Recommended - default) The process is performed continuously, but copying does not have priority and the time to completion is not guaranteed.
 - Fast — The copy/resync process is performed continuously and has priority. Host I/O performance will be degraded. The time to copy can be guaranteed because it has priority.

You can change the set copy pace later by using the pair edit function. You may change it when you feel that the creation time takes a long time in the pace specified at the time of the creation or the effect on the host I/O is significant because the copy processing is given priority.
11. In the **Do initial copy from the primary volume...** field, leave **Yes** checked to copy the primary to the secondary volume. All the data of the P-VOL is copied to the corresponding S-VOL in the initial copy. Furthermore, the P-VOL data updated during the initial copy is also reflected in the S-VOL. Therefore, when the pair status becomes Paired, it is guaranteed that the data of the P-VOL and the S-VOL is the same.

Clear the check box to create a pair without copying the P-VOL at this time, and thus reduce the time it takes to set up the configuration for the pair. Use this option also when data in the primary and secondary volumes already match. The system treats the two volumes as paired even though no data is presently transferred. Resync can be selected manually at a later time when it is appropriate.
12. Click **OK**, then click **Close** on the confirmation screen that appears. The pair has been created.
13. A confirmation message appears. Click **Close**.

Splitting a pair

Data is copied to the S-VOL at every update cycle until the pair is split.

- When the split is executed, all differential data accumulated in the local disk array is updated to the S-VOL.
- After the split operation, write updates continue to the P-VOL but not to the S-VOL.

After the Split Pair operation:

- S-VOL data is consistent to P-VOL data at the time of the split. The S-VOL can receive read/write instructions.
- The TCE pair can be made identical again by re-synchronizing from primary-to-secondary or secondary-to-primary.

The pair must be in Paired status. The time required to split the pair depends on the amount of data that must be copied to the S-VOL so that the data is current with the P-VOL's data.

The following can be specified as an option at the time of the pair splitting

- S-VOL accessibility. Set the access to the S-VOL after split. You can select either Read/Write possible or Read only possible. The default is Read/Write possible.
- Instruction of the status transition to the S-VOL. If the forcible Takeover is specified, the S-VOL is changed to the Takeover status and Read/Write becomes possible. You can use it when testing if the operation can restart when switching to the S-VOL while the I/O to the P-VOL continues.


When specifying the recovery from Takeover, the S-VOL in the Takeover status is changed to the Split status. When the S-VOL is changed to Takeover by forcible Takeover, for re-synchronizing the S-VOL and the P-VOL, perform the resynchronization after recovering the S-VOL from Takeover to Split.

To split the pair



NOTE: When the pair status is **Paired**, if the local array receives the command to split the pair, it transfers all the differential data remained in the local array to the remote array and then changes the pair status to **Split**. Therefore, even if the array receives the command to split the pair, the pair status might not change to **Split** immediately.

1. In Navigator 2 GUI, select the desired disk array, then click the **Show & Configure disk array** button.
2. From the Replication tree, select the **Remote Replication** icon. The **Pairs** screen displays.
3. Select the pair you want to split.
4. Click the **Split Pair** button at the bottom of the screen. The **Split Pair** screen appears.



NOTE: When split from to the remote array, you can specify the Status change to the secondary volume. In this case select the Forced Takeover or Recover from Takeover

5. If splitting default Options is **Read / Write** for the secondary volume, if you want to protect the write operation, specify **Read Only**.
6. Click **OK**.
7. A confirmation message appears. Click Close.

Resynchronizing a pair

When discarding the backup data retained in the S-VOL by split or recovering the suspended pair (Failure status), perform the pair resynchronization to resynchronize the S-VOL and the P-VOL.

Re-synchronizing a pair updates the S-VOL so that it is again identical with the P-VOL. Differential data accumulated on the local disk array since the last pairing is updated to the S-VOL.

- Pair status during a re-synchronizing is Synchronizing.
- Status changes to Paired when the resync is complete.
- If P-VOL status is Failure and S-VOL status is Takeover or Simplex, the pair cannot be recovered by resynchronizing. It must be deleted and created again.
- Best practice is to perform a resynchronization when I/O load is low, to reduce impact on host activities.

Prerequisites

- The pair must be in Split, Failure, or Pool Full status.

To resync the pair

1. In the Navigator 2 GUI, select the desired disk array, then click the **Show & Configure disk array** button.
2. From the Replication tree, select the **Remote Replication** icon. The **Pairs** screen displays.
3. Select the pair you want to resync.
4. Click the **Resync Pair** button. View further instructions by clicking the Help button, as needed.

Swapping pairs

When the P-VOL data cannot be used and the data retained in the S-VOL as the remote backup is returned to the P-VOL, swap the pair. If it is swapped, the volume which was first a P-VOL is switched to an S-VOL and the volume which was an S-VOL is switched to a P-VOL, and synchronizes the S-VOL after switching and the P-VOL.

In a pair swap, primary and secondary-volume roles are reversed. The direction of data flow is also reversed.

This is done when host operations are switched to the S-VOL, and when host-storage operations are again functional on the local disk array.

Prerequisites and Notes

- To swap the pairs, the remote path must be set for the local array from the remote array.
- The pair swap is executed on the remote disk array.
- As long as swap is performed from Navigator 2 on the remote array, no matter how many times swap is performed, the copy direction will not return to the original direction (P-VOL on the local array and S-VOL on the remote array).
- The pair swap is performed in units of groups. Therefore, even if you select a pair and performed it, the pairs in the group are all swapped.
- When swap the pair, P-VOL pair status changes to **Failure**.

To swap TCE pairs

1. In Navigator 2 GUI, connect to the remote disk array, then click the **Show & Configure disk array** button.
2. From the Replication tree, select the **Remote Replication** icon. The **Pairs** screen displays.
3. Select the pair you want to swap.
4. Click the **Swap Pair** button.
5. On the message screen, check the **Yes, I have read...** box, then click **Confirm**.
6. Click **Close** on the confirmation screen.
7. When the pairs are swapped, the processing to store the S-VOL is executed in the background with the backup data (previous definite data) saved in the DP pool. If this processing takes time, the following error occurs. If the message, "DMER090094: The LU whose pair status is Busy exists in the target group" displays, proceed as follows:
 - a. Check the pair status for each LU in the target group. Pair status will change to Takeover. Confirm this before proceeding. Click the **Refresh Information** button to see the latest status.
 - b. When the pairs have changed to Takeover status. execute the Swap command again.

Editing pairs

You can edit the name, group name, and copy pace for a pair. A group created with no name can be named from the Edit Pair screen.

To edit pairs

1. In Navigator 2 GUI, select the desired disk array, then click the **Show & Configure disk array** button.
2. From the Replication tree, select the **Remote Replication** icon. The **Pairs** screen displays.
3. Select the pair that you want to edit.
4. Click the **Edit Pair** button.
5. Make any changes, then click **OK**.
6. On the confirmation message, click **Close**.



NOTE: Edits made on the local disk array are not reflected on the remote disk array. To have the same information reflected on both disk arrays, it is necessary to edit the pair on the remote disk array also.

Deleting pairs

When a pair is deleted, transfer of differential data from P-VOL to S-VOL is completed, then the volumes become Simplex. The pair is no longer displayed in the Remote Replication pair list on Navigator 2 GUI.

- A pair can be deleted regardless of its status. However, data consistency is not guaranteed unless status prior to deletion is Paired.
- If the operation fails, the P-VOL nevertheless becomes Simplex. Transfer of differential data from P-VOL to S-VOL is terminated.
- Normally, a Delete Pair operation is performed on the local disk array where the P-VOL resides. However, it is possible to perform the operation from the remote disk array, though with the following results:
 - Only the S-VOL becomes Simplex.
 - Data consistency in the S-VOL is not guaranteed.
 - The P-VOL does not recognize that the S-VOL is in Simplex status. When the P-VOL tries to send differential data to the S-VOL, it recognizes that the S-VOL is absent and the pair becomes Failure.
 - When the pair status changes to Failure, the status of the other pairs in the group also becomes Failure.
 - From the remote disk array, this Failure status is not seen and pair status remains Paired.
 - When executing the pair deletion in the batch file or the script, insert a five-second wait before executing the next processing step. An example batch file with a five-second wait is:
 - Pair creation of TrueCopy which specified the volume specified as the S-VOL of the deleted pair
 - Pair creation of Volume Migration which specified the volume specified as the S-VOL of the deleted pair
 - Deletion of the volume specified as the S-VOL of the deleted pair
 - Shrinking of the volume specified as the S-VOL of the deleted pair

```
ping 127.0.0.1 -n 5 > nul
```

To delete a pair

1. In Navigator 2 GUI, select the desired disk array, then click the **Show & Configure disk array** button.
2. From the Replication tree, select the **Remote Replication** icon. The **Pairs** screen displays.
3. Select the pair you want to delete in the Pairs list, and click **Delete Pair**.
4. On the message screen, check the **Yes, I have read...** box, then click **Confirm**.

Click **Close** on the confirmation screen

Example scenarios and procedures

This topic describes four use-cases and the processes for handling them.

- ❑ [CLI scripting procedure for S-VOL backup](#)
- ❑ [Procedure for swapping I/O to S-VOL when maintaining local disk array](#)
- ❑ [Procedure for moving data to a remote disk array](#)
- ❑ [Process for disaster recovery](#)

CLI scripting procedure for S-VOL backup

Snapshot can be used with TCE to maintain timed backups of S-VOL data. The following illustrates and explains how to perform TCE and Snapshot operations.

An example scenario is used in which three hosts, A, B, and C, write data to volumes on the local disk array, as shown in [Figure 20-1](#).

- Although the database is started in Host A, the files that the database deals are stored in a D drive and E drive. The D drive and E drive are actually VOL1 and VOL2. The VOL1 and VOL2 are backed up in the remote array every day at 11 o'clock at night using TCE.
- The remote array stores the backup data for one week.
- In another Host B, VOL3 in the array is used as a D drive. The D drive of host B can be backed up at time (for example, 2 o'clock) different from the backup of the database file of Host A.
- Navigator 2 CLI is required for each host to link up with the application in the host. Also, each host should be connected via LAN for the array management not only in the local array but also in the remote array to operate the Snapshot pair in the remote array.
- A file system application on host B, and a mail server application on host C, store their data as indicated in the graphic.
- The TCE S-VOLs are backed up daily, using Snapshot. Each Snapshot backup is held for seven days. There are seven Snapshot backups.
- The volumes for the other applications are also backed up with Snapshot on the remote disk array, as indicated. These snapshots are made at different times than the database snapshots to avoid performance problems.
- Each host is connected by a LAN to the disk arrays.
- CLI scripts are used for TCE and Snapshot operations.




Figure 20-1: Configuration example for a remote backup system

Scripted TCE, Snapshot procedure

The TCE/Snapshot system shown in [Figure 20-1](#) is set up using the Navigator 2 GUI. Day-to-day operations are handled using CCI or CLI scripts. In this example, CLI scripts are used.

[Table 20-1](#) shows the application operated by the host connected to the local array, the volume used, and the volume of the backup destination indicated in the construction example in [Figure 20-1](#). Each host collects the backup of the volume used to the remote array once a day. It leaves the backup in every other day until one week before by deciding the V-VOL for creating the backup for every day of the week.

Table 20-1: TCE volumes by host application

Host name	Application	Volume to use (backup target volume)	Backup in remote array Snapshot volume			
			For Monday	For Tuesday	...	For Sunday
Host A	Database	VOL1 (D drive)	VOL101	VOL111	...	VOL161
		VOL2 (E drive)	VOL102	VOL112	...	VOL162
Host B	File system	VOL3 (D drive)	VOL103	VOL113	...	VOL163
Host C	Mail server	VOL4 (M drive)	VOL104	VOL114	...	VOL164
		VOL5 (N drive)	VOL105	VOL115	...	VOL165
		VOL6 (O drive)	VOL106	VOL116	...	VOL166

In the procedure example that follows, scripts are executed for host A on Monday at 11 p.m. The following assumptions are made:

- The system is completed.
- The TCE pairs are in Paired status.
- The Snapshot pairs are in Split status.
- Host A uses a Windows operating system.

The variables used in the script are shown in [Table 20-2](#). The procedure and scripts follow.

Table 20-2: CLI script variables and descriptions

#	Variable name	Content	Remarks
1	STONAVM_HOME	Specify the directory in which SNM2 CLI was installed.	When the script is in the directory in which SNM2 CLI was installed, specify ".".
2	STONAVM_RSP_PASS	Be sure to specify "on" when executing SNM2 CLI in the script.	This is the environment variable to enter "Yes" automatically for the inquiry of SNM2 CLI command.
3	LOCAL	Name of the local disk array registered in SNM2 CLI	
4	REMOTE	Name of the remote disk array registered in SNM2 CLI	
5	TCE_PAIR_DB1, TCE_PAIR_DB2	Name of the TCE pair generated at the setup	The default names are as follows. TCE_LUxxxx_LUyyyy xxxx: LUN of P-VOL yyyy: LUN of S-VOL
6	SS_PAIR_DB1_MON SS_PAIR_DB2_MON	Name of the Snapshot pair when creating the backup in the remote disk array on Monday	The default names are as follows. TCE_LUxxxx_LUyyyy xxxx: LUN of P-VOL yyyy: LUN of S-VOL
7	DB1_DIR DB2_DIR	Directory on the host where the volume is mounted	
8	LU1_GUID, LU2_GUID	GUID of the backup target volume recognized by the host	You can search it by the mountvol command of Windows.
9	TIME	Time-out value of the aureplicationmon command	Make it longer than the time taken for the resynchronization of TCE.

1. Specify the variables to be used in the script, as shown below.

```

set STONAVM_HOME=.
set STONAVM_RSP_PASS=on
set LOCAL=Localdisk array
set REMOTE=Remotedisk array
set TCE_PAIR_DB1=TCE_LU0001_LU0001
set TCE_PAIR_DB2=TCE_LU0002_LU0002
set SS_PAIR_DB1_MON=SS_LU0001_LU0101
set SS_PAIR_DB2_MON=SS_LU0002_LU0102

set DB1_DIR=D:\
set DB2_DIR=E:\
set LU1_GUID=xxxxxxxx-xxxx-xxxx-xxxx-xxxxxxxxxxxxxx
set LU2_GUID=yyyyyyyy-yyyy-yyy-yyy-yyy-yyyyyyyyyyyy
set TIME=18000
(To be continued)

```

2. 2.Stop the database and un-mount it to make the data of the backup target volume stationary. raidqry is a CCI command.

```
(Continued from the previous section)
<Stop the access to C:\hus100\DB1 and C:\hus100\DB2>

REM Unmount of P-VOL
raidqry -x umount %DB1_DIR%
raidqry -x umount %DB2_DIR%
(To be continued)
```

3. Split the TCE pair, then check that the pair status becomes Split, as shown below. This updates data in the S-VOL and makes it available for secondary uses, including Snapshot operations.

```
(Continued from the previous section)
REM pair split
aureplicationremote -unit %LOCAL% -split -tce -pairname
%TCE_PAIR_DB1% -gno 0
aureplicationremote -unit %LOCAL% -split -tce -pairname
%TCE_PAIR_DB2% -gno 0
REM Wait until the TCE pair status becomes Split.
aureplicationmon -unit %LOCAL% -evwait -tce -pairname
%TCE_PAIR_DB1% -gno 0 -st split -pvol -timeout %TIME%
aureplicationmon -unit %LOCAL% -evwait -tce -pairname
%TCE_PAIR_DB1% -gno 0 -nowait
IF NOT %ERRORLEVEL% == 13 GOTO ERROR_TCE_Split
aureplicationmon -unit %LOCAL% -evwait -tce -pairname
%TCE_PAIR_DB2% -gno 0 -st split -pvol -timeout %TIME%
aureplicationmon -unit %LOCAL% -evwait -tce -pairname
%TCE_PAIR_DB2% -gno 0 -nowait
IF NOT %ERRORLEVEL% == 13 GOTO ERROR_TCE_Split
(To be continued)
```

4. Mount the P-VOL, and restart the database application, as shown below.

```
(Continued from the previous section)
REM Mount of P-VOL
raidqry -x mount %DB1_DIR% Volume{%LU1_GUID%}
raidqry -x mount %DB2_DIR% Volume{%LU2_GUID%}

<Restart access to C:\hus100\DB1 and C:\hus100\DB2>
(To be continued)
```

5. Resynchronize the Snapshot backup. Then split the Snapshot backup. These operations are shown in the example below.

```
(Continued from the previous section)
REM Resynchronization of the Snapshot pair which is cascaded
aureplicationlocal -unit %REMOTE% -resync -ss -pairname
%SS_PAIR_DB1_MON% -gno 0
aureplicationlocal -unit %REMOTE% -resync -ss -pairname
%SS_PAIR_DB2_MON% -gno 0
REM Wait until the Snapshot pair status becomes Paired.
aureplicationmon -unit %REMOTE% -evwait -ss -pairname
%SS_PAIR_DB1_MON% -gno 0 -st paired -pvol -timeout %TIME%
aureplicationmon -unit %REMOTE% -evwait -ss -pairname
%SS_PAIR_DB1_MON% -gno 0 -nowait
IF NOT %ERRORLEVEL% == 12 GOTO ERROR_SS_Resync
aureplicationmon -unit %REMOTE% -evwait -ss -pairname
%SS_PAIR_DB2_MON% -gno 0 -st paired -pvol -timeout %TIME%
aureplicationmon -unit %REMOTE% -evwait -ss -pairname
%SS_PAIR_DB2_MON% -gno 0 -nowait
IF NOT %ERRORLEVEL% == 12 GOTO ERROR_SS_Resync

REM Pair split of the Snapshot pair which is cascaded
aureplicationlocal -unit %REMOTE% -split -ss -pairname
%SS_PAIR_DB1_MON% -gno 0
aureplicationlocal -unit %REMOTE% -split -ss -pairname
%SS_PAIR_DB2_MON% -gno 0
REM Wait until the Snapshot pair status becomes Split.
aureplicationmon -unit %REMOTE% -evwait -ss -pairname
%SS_PAIR_DB1_MON% -gno 0 -st split -pvol -timeout %TIME%
aureplicationmon -unit %REMOTE% -evwait -ss -pairname
%SS_PAIR_DB1_MON% -gno 0 -nowait
IF NOT %ERRORLEVEL% == 13 GOTO ERROR_SS_Split
aureplicationmon -unit %REMOTE% -evwait -ss -pairname
%SS_PAIR_DB2_MON% -gno 0 -st split -pvol -timeout %TIME%
aureplicationmon -unit %REMOTE% -evwait -ss -pairname
%SS_PAIR_DB2_MON% -gno 0 -nowait
IF NOT %ERRORLEVEL% == 13 GOTO ERROR_SS_Split
(To be continued)
```

6. When the Snapshot backup operations are completed, re-synchronize the TCE pair, as shown below. When the TCE pair status becomes Paired, the backup procedure is completed.

```
(Continued from the previous section)
REM Return the pair status to Paired (Pair resynchronization)
aureplicationremote -unit %LOCAL% -resync -tce -pairname
%TCE_PAIR_DB1% -gno 0
aureplicationremote -unit %LOCAL% -resync -tce -pairname
%TCE_PAIR_DB2% -gno 0
REM Wait until the TCE pair status becomes Paired.
aureplicationmon -unit %LOCAL% -evwait -tce -pairname
%TCE_PAIR_DB1% -gno 0 -st paired -pvol -timeout %TIME%
aureplicationmon -unit %LOCAL% -evwait -tce -pairname
%TCE_PAIR_DB1% -gno 0 -nowait
IF NOT %ERRORLEVEL% == 12 GOTO ERROR_TCE_Resync
aureplicationmon -unit %LOCAL% -evwait -tce -pairname
%TCE_PAIR_DB2% -gno 0 -st paired -pvol -timeout %TIME%
aureplicationmon -unit %LOCAL% -evwait -tce -pairname
%TCE_PAIR_DB2% -gno 0 -nowait
IF NOT %ERRORLEVEL% == 12 GOTO ERROR_TCE_Resync
echo The backup is completed.
GOTO END
(To be continued)
```

7. If pair status does not become Paired within the aureplicationmon command time-out period, perform error processing, as shown below.

```
(Continued from the previous section)
REM Error processing
:ERROR_TCE_Split
< Processing when the S-VOL data of TCE is not determined within
the specified time>
GOTO END
:ERROR_SS_Resync
< Processing when Snapshot pair resynchronization fails and the
Snapshot pair status does not become Paired>
GOTO END
:ERROR_SS_Split
< Processing when Snapshot pair split fails and the Snapshot pair
status does not become Split>
GOTO END
:ERROR_TCE_Resync
< Processing when TCE pair resynchronization does not terminate
within the specified time>
GOTO END

:END
```

Procedure for swapping I/O to S-VOL when maintaining local disk array

The following shows a procedure for temporarily shifting I/O to the S-VOL in order to perform maintenance on the local disk array. In the procedure, host server duties are switched to a standby server.

1. On the local disk array, stop the I/O to the P-VOL.

2. Split the pair, which makes P-VOL and S-VOL data identical.
3. On the remote site, execute the swap pair command. Since no data is transferred, the status is changed to Paired after one cycle time.
4. Split the pair.
5. Restart I/O, using the S-VOL on the remote disk array.
6. On the local site, perform maintenance on the local disk array.
7. When maintenance on the local disk array is completed, resynchronize the pair from the remote disk array. This copies the data that has been updated on the S-VOL during the maintenance period.
8. On the remote disk array, when pair status is Paired, stop I/O to the remote disk array, and un-mount the S-VOL.
9. Split the pair, which makes data on the P-VOL and S-VOL identical.
10. On the local site, issue the pair swap command. When this is completed, the S-VOL in the local disk array becomes the P-VOL again.
11. Business can restart at the local site. Mount the new P-VOL on the local disk array to local host server and restart I/O.

Procedure for moving data to a remote disk array

This section provides a procedure in which application data is copied to a remote disk array, and the copied data is analyzed. An example scenario is used in which:

- A, P-VOL VOL1 and VOL2 of the local array are put in the same CTG and then are copied to S-VOL VOL1 and VOL2 of the remote array as shown in [Figure 20-2](#).
- The P-VOL volumes are in the same consistency group (CTG).
- The P-VOL volumes are paired with the S-VOL volumes, LU1 and LU2 on the remote disk array.
- A data-analyzing application on host D analyzes the data in the S-VOL. Analysis processing is performed once every hour.




Figure 20-2: Configuration example for moving data

Example procedure for moving data

1. Stop the applications that are writing to the P-VOL, then un-mount the P-VOL.
2. Split the TCE pair. Updated differential data on the P-VOL is transferred to the S-VOL. Data in the S-VOL is stabilized and usable after the split is completed.
3. Mount the P-VOL and then resume writing to the P-VOL.
4. Mount the S-VOL.
5. Read and analyze the S-VOL data. The S-VOL data can be updated, but updated data will be lost when the TCE pair is resynchronized. If updated data is necessary, be sure to save the data to a volume other than the S-VOL.
6. Un-mount the S-VOL.
7. Re-synchronize the TCE pair.

Process for disaster recovery

This section explains behaviors and the general process for continuing operations on the S-VOL and then failing back to the P-VOL, when the primary site has been disabled.

In the event of a disaster at the primary site, the cycle update process is suspended and updating of the S-VOL stops. If the host requests an S-VOL takeover (CCI horctakeover), the remote disk array restores the S-VOL using data in the DP pool from the previous cycle.

The Hitachi Unified Storage version of TCE does not support mirroring consistency of S-VOL data, even if the local disk array and remote path are functional. P-VOL and S-VOL data are therefore not identical when takeover is executed. Any P-VOL data updates made during the time the takeover command was issued cannot be salvaged.

Takeover processing

S-VOL takeover is performed when the horctakeover operation is issued by the secondary disk array. The TCE pair is split and system operation can be continued with the S-VOL only. In order to settle the S-VOL data being copied cyclically, it is restored using the data that was pre-determined in the preceding cycle and saved to the DP pool, as mentioned above. The S-VOL is immediately enabled to receive the I/O instruction.

When the SVOL_Takeover is executed, data restoration processing from the DP pool of the secondary site to the S-VOL is performed in the background. During the period from the execution of the SVOL_Takeover until the completion of the data restoration processing, performance of the host I/O for the S-VOL is deteriorated. P-VOL and S-VOL data are not the same after this operation is performed.

For details on the `horctakeover` command, see the *Hitachi Unified Storage Command Control Interface Installation and Configuration Guide*.

Swapping P-VOL and S-VOL

SWAP Takeover ensures that system operation continues by reversing the characteristics of the P-VOL and the S-VOL and swapping the relationship between the P-VOL and S-VOL. After S-VOL takeover, host operations continue on the S-VOL, and S-VOL data becomes updated as a result of I/O operations. When continuing application processing using the S-VOL or when restoring application processing to the P-VOL, the swap function makes the P-VOL up-to-date, by reflecting updated data on the S-VOL to the P-VOL.

Failback to the local disk array

The fallback process involves restarting business operations at the local site. The following shows the procedure after the pair swap is performed.

1. On the remote disk array, after S-VOL takeover and the TCE pair swap command are executed, the S-VOL is mounted, and data restoration is executed (`fsck` for UNIX and `chkdsk` for Windows).
2. I/O is restarted using the S-VOL.
3. When the local site/disk array is restored, the TCE pair is created from the remote disk array. At this time, the S-VOL is located on the local disk array.
4. After the initial copy is completed and status is Paired, I/O to the remote TCE volume is stopped and it is unmounted.
5. The TCE pair is split. This completes transfer of data from the remote volume to the local volume.
6. At the local site, the pair swap command is issued. When this is completed, the S-VOL in the local disk array becomes the P-VOL.
7. Mount the new P-VOL on the local disk array is mounted and I/O is restarted.

Monitoring and troubleshooting TrueCopy Extended

This chapter provides information and instructions for troubleshooting and monitoring the TCE system.

- ❑ [Monitoring and maintenance](#)
- ❑ [Troubleshooting](#)
- ❑ [Correcting DP pool shortage](#)
- ❑ [Cycle copy does not progress](#)
- ❑ [Correcting disk array problems](#)
- ❑ [Correcting resynchronization errors](#)
- ❑ [Using the event log](#)
- ❑ [Miscellaneous troubleshooting](#)

Monitoring and maintenance

This section provides information and instructions for monitoring and maintaining the TCE system.

- ☐ [Monitoring pair status](#)
- ☐ [Monitoring DP pool capacity](#)
- ☐ [Monitoring the remote path](#)
- ☐ [Monitoring cycle time](#)
- ☐ [Changing copy pace](#)
- ☐ [Monitoring synchronization](#)
- ☐ [Routine maintenance](#)

Monitoring pair status

Pair status should be checked periodically to insure that TCE pairs are operating correctly. If the pair status becomes Failure or Pool Full, data cannot be copied from the local disk array to the remote disk array.

Also, status should be checked before performing a TCE operation. Specific operations require specific pair statuses.

The pair status changes as a result of operations on the TCE pair. You can find out how an array is controlling the TCE pair from the pair status. You can also detect failures by monitoring the pair status. [Figure 21-1](#) shows the pair status transitions.

The pair status of a pair with the reverse pair direction (S-VOL to P-VOL) changes in the same way that the pair with the original pair direction (P-VOL to S-VOL) changes. Look at the figure with the reverse pair direction in mind. Once the resync copy completes, the pair status changes to Paired.




Figure 21-1: Pair status transitions

Monitoring using the GUI is done at the user's discretion. Monitoring should be performed frequently. Email notifications can be set up to inform you when failure and other events occur.

To monitor pair status using the GUI

1. In Navigator 2 GUI, select the desired disk array, then click the **Show & Configure disk array** button.
2. From the Replication tree, select the **Remote Replication** icon. The **Pairs** screen displays.

☐	Name	Local VOL	Attribute	Remote Array ID	Remote Path Name	Remote VOL	Status
☒	TCE_LU0000_LU0000	0000	Primary	93000026	Array_93000026	0000	Paired(0%)
☐	TCE_LU0001_LU0001	0001	Primary	93000026	Array_93000026	0001	Paired(20%)
☐	TCE_LU0002_LU0002	0002	Primary	93000026	Array_93000026	0002	Paired(40%)

- Name: The pair name is displayed.
 - Local VOL: The local side VOL is displayed.
 - Attribute: The volume type (Primary or Secondary) is displayed.
 - Remote Array ID: The remote array ID is displayed.
 - Remote Path Name: The remote path name is displayed.
 - Remote VOL: The remote side VOL is displayed.
 - Status: The pair status is displayed. The each pair status meaning, see the section 2.6. The percentage denotes the progress rate (%) when the pair status is Synchronizing. When the pair status is Paired, it denotes the coincidence rate (%) of a P-VOL and an S-VOL. When the pair status is Split, it denotes the coincidence rate (%) of current data and the data at the time of pair splitting.
 - DP Pool:
 - Replication Data: A Replication Data DP pool number displays.
 - Management Area: A Management Area DP pool number displays.
 - Copy Type: TrueCopy Extended Distance is displayed.
 - Group Number: A group number and group name is displayed.
 - Group Name
3. Locate the pair whose status you want to review in the **Pair** list. Status descriptions are provided in [Table 21-2 on page 21-5](#). You can click the **Refresh Information** button (not in view) to make sure data is current.

- The percentage that displays with each status shows how close the S-VOL is to being completely paired with the P-VOL.

The pair status changes as a result of operations on the TCE pair. You can find out how an array is controlling the TCE pair from the pair status. You can also detect failures by monitoring the pair status. [Table 21-1](#) shows the pair accessibility.

The Attribute column shows the pair volume for which status is shown.

Table 21-1: Pair accessibility

Pair Status	Access to P-VOL		Access to S-VOL	
	Read	Write	Read	Write
Simplex	YES	YES	YES	YES
Synchronizing	YES	YES	YES	NO
Paired	YES	YES	YES	NO
Paired:split	YES	YES	YES	NO
Paired:delete	YES	YES	YES	NO
Split	YES	YES	YES	YES or NO
Pool Full	YES		YES	NO
Takeover			YES	YES
Busy			YES	YES
Paired Internally Busy			YES	NO
Inconsistent			NO	NO
Failure	YES	YES		

Table 21-2: Pair status definitions

Pair status	Description	Access to P-VOL	Access to S-VOL
Simplex	If a volume is not assigned to a TCE pair, its status is Simplex. If the created pair is deleted, the pair status becomes Simplex. Note that the Simplex volume is not displayed on the list of the TCE pair	Read/Write	Read/Write
Synchronizing	Copying is in progress, initiated by Create Pair or Resynchronize Pair operations. Upon completion, pair status changes to Paired. Data written to the P-VOL during copying is transferred as differential data after the copying operation is completed. Copy progress is shown on the Pairs screen in the Navigator 2 GUI. If the split pair is resynchronized, only the differential data of the P-VOL is copied to the S-VOL. If the pair at the time of the pair creation is resynchronized, entire P-VOL is copied to the S-VOL.	Read/Write	Read Only

Table 21-2: Pair status definitions (Continued)

Pair status	Description	Access to P-VOL	Access to S-VOL
Paired	The copy is completed and the data of the P-VOL and the S-VOL is the same. In case of the Paired status, the update for the P-VOL is periodically reflected in the S-VOL and the synchronized status is retained in the P-VOL and the S-VOL. If you check the identical rate in the pair information, it is 100%.	Read/Write	Read Only
Paired:split	When a pair-split operation is initiated, the differential data accumulated in the local disk array is updated to the S-VOL before the status changes to Split. Paired:split is a transitional status between Paired and Split.	Read/Write	Read Only
Paired:delete	When a pair-delete operation is initiated, the differential data accumulated in the local disk array is updated to the S-VOL before the status changes to Simplex. Paired:delete is a transitional status between Paired and Simplex.	Read/Write	Read Only
Split	The data of the P-VOL and the S-VOL is not synchronized. All the positions of the update to the P-VOL and the S-VOL are stored in the DP pool as the differential information. You can check the differential amount of the P-VOL and the S-VOL by checking how far the identical rate of the pair information falls below 100%	Read/Write	Read/Write (mountable) or Read only

Table 21-2: Pair status definitions (Continued)

Pair status	Description	Access to P-VOL	Access to S-VOL
Pool Full	Pool Full indicates that the usage rate of the DP pool reaches the Replication Data Released threshold and usable capacity of the DP pool be reducing. When the consumed capacity of the DP pool depleted, updating copy from the P-VOL to the S-VOL cannot continue. If the usage rate of the DP pool for the P-VOL reaches the Replication Data Released threshold while the pair status is Paired, the pair status at the local array where the P-VOL resides changes to this status. In this case the pair status at the remote array remains Paired. While the pair status at the local array is Pool Full, the data written to the P-VOL are managed as differential data. If the usage rate of the DP pool for the S-VOL reaches the Replication Data Released threshold while the pair status is Paired, the pair status at the remote array where the S-VOL resides changes to this status. In this case the pair status at the local array becomes Failure. In order to recover the pair from Pool Full, add the DP pool capacity or reduce the use of the DP pool, and then resynchronize the pair. If a pair in a group has met the condition to become Pool Full, not only the status of this pair but also the statuses all the other pairs in the group become Pool Full. Pool Full is changed in units of CTG. For example, when the DP pool depletion occurs in pool #0, all the pairs which use the DP pool are changed to Pool Full. In addition, all the pairs (using pool #1) in CTG to which the pairs changed to Pool Full belong are changed to Pool Full.	Read/Write	Read Only
Takeover	Takeover is a transitional status after Swap Pair is initiated. The data in the remote DP pool, which is in a consistent state established at the end of the previous cycle, is restored to the S-VOL. Immediately after the pair becomes Takeover, the pair relationship is swapped and copy from the new P-VOL to the new S-VOL is started. Only the S-VOL has this status. The S-VOL in the Takeover status can perform the Read/Write access from the host.		Read/Write

Table 21-2: Pair status definitions (Continued)

Pair status	Description	Access to P-VOL	Access to S-VOL
Paired Internally Busy	Paired Internally Busy is a transitional status after Swap Pair is tried. When Swap Pair is performed and the remote array can communicate with the local array through the remote path, the pair status of the S-VOL becomes Paired Internally Busy. The determined data at the end of the previous cycle is being restored to the S-VOL. Takeover will come after Paired Internally Busy. The time for completing the restoration processing can be estimated by the difference amount of the pair status display items of Navigator 2. This is shown PAIR for CCI.	Read/Write	Read only
Busy	Busy is a transitional status after Swap Pair is tried. When Swap Pair is performed and the remote array can not communicate with the local array through the remote path, the pair status of the S-VOL becomes Busy. It indicates that the determined data at the end of the previous cycle is being restored to the S-VOL. Takeover will come after Busy. This is shown SSWS(R) for CCI.		No Read/Write
Inconsistent	This status on the remote disk array occurs when copying from P-VOL to S-VOL stops due to failure in the S-VOL. The failure includes failure of the HDD that constitutes the S-VOL, or the DP pool for the S-VOL becomes depleted. To recover, resynchronize the pair, which leads to a full volume copy of the P-VOL to the S-VOL.		No Read/Write

Table 21-2: Pair status definitions (Continued)

Pair status	Description	Access to P-VOL	Access to S-VOL
Failure	A failure occurred and the copy operation is suspended forcibly. P-VOL pair status changes to Failure if copying from the P-VOL to the S-VOL can no longer continue. The failure includes HDD failure and remote path failure that disconnects the local disk array and the remote disk array. • Data consistency is guaranteed in the group if the pair status at the local disk array changes from Paired to Failure. • Data consistency is not guaranteed if pair status changes from Synchronizing to Failure. • Data written to the P-VOL is managed as differential data. To recover, remove the cause then resynchronize the pair. When a pair in the group has a factor to be Failure, all the pairs in the group become Failure.	Read/ Write	

Monitoring DP pool capacity

Monitoring DP pool capacity is critical for the following reasons:

- Data copying from the local to remote disk array will halt when:
 - The local DP pool's use rate reaches 90 percent
 - The remote DP pool's capacity is full

Also, the local disk array could be damaged if data copying is stopped for these reasons.

This section provides instructions for

- Monitoring DP pool usage
- Specifying the threshold value
- Adding capacity to the DP pool

Monitoring DP pool usage

When the usage rate of the DP pool in local array reaches the Replication Data Released threshold or the usage rate of the DP pool in remote array reaches the Replication Data Released threshold, the data copy from the local array to the remote array stops. If the local array is damaged while the data copy is stopping, the amount of the data loss increases. Therefore, it is necessary to operate so that the DP pool does not run short.

The monitoring of the DP pool capacity is performed to detect the risk of a DP pool shortage in advance by checking the use rate of the DP pool periodically. Also, the threshold value to warn the decrease of the remaining capacity is set for the DP pool. If the use rate of the DP pool exceeds the threshold value, you will be notified. Add the RAID group to the DP pool to expand the capacity or decrease the number of pairs using the DP pool to store the free capacity of the DP pool.

Checking DP pool status or changing threshold value of the DP pool

Refer to the *Hitachi Unified Storage Dynamic Provisioning Configuration Guide*. For changing the threshold value, see [Setting the replication threshold \(optional\) on page 19-54](#).

Adding DP pool capacity

Refer to the *Hitachi Unified Storage Dynamic Provisioning Configuration Guide*.

Processing when DP pool is Exceeded

TCE provides a DP pool for both the local array and remote array shown in [Figure 21-2](#). How the DP pool is used differs between the local array and remote array.

Local array: When the data is updated by the host before transferring the differential data (replication data) to the S-VOL, the data that has not been transferred is copied to the DP pool.

Remote array: When the update copy is performed to the S-VOL, internal pre-determined S-VOL data is copied to the DP pool. During takeover, the internal pre-determined data is restored from the DP pool onto the S-VOL.

The local array copies data on the P-VOL that has not been transferred to the DP pool if there is write data that updates that data (). When data that has been transferred is updated () by new data, the data that has been transferred is not copied to the DP pool. After completion of cycle, the data copied to the DP pool is deleted.

The remote array copies the internal pre-determined data to the DP pool () when the data is updated in update copy processing (). The copied data is used during takeover processing. The copied data is deleted when cycle update has been completed.

Both the local array and remote array do not copy the data at the time of the data update in the second time or later if it is in the same cycle.




Figure 21-2: How TCE uses the DP pool

Because the DP pool size has an upper limit, the unused capacity of the DP pools of both the local array and remote array is used up if the amount of the data to be copied increases. In addition, data copied by Snapshot also affects DP pool capacity because Snapshot uses the same DP pool.

The remote array deletes the copied data used by Snapshot and changes the Snapshot pair status to Failure () if the capacity of the DP pool used by the remote array reaches the limit (in [Figure 21-3 on page 21-13](#)). However, since the internal pre-determined S-VOL data is not subject to deletion (), the S-VOL can be used for takeover even if the DP pool capacity is exceeded.



NOTE:

1. TCE pair and Snapshot pair share the same DP pool, but data consistency policies after pool over are different. See [Figure 21-3 on page 21-13](#) for more details.
2. In a case of Snapshot, V-VOL data becomes invalid if a pool over occurred. V-VOL data cascading from TCE pair S-VOL also becomes invalid.

I




Figure 21-3: Effects of exceeding the DP pool capacity in the remote array

Monitoring the remote path

Monitor the remote path to ensure that data copying is unimpeded. If a path is blocked, the status is Detached, and data cannot be copied.

You can adjust remote path bandwidth and cycle time to improve data transfer rate.

To monitor the remote path

1. In the Replication tree, click **Setup**, then **Remote Path**. The Remote Path screen displays.
2. Review statuses and bandwidth. Path statuses can be Normal, Blocked, or Diagnosing. When Blocked or Diagnosing is displayed, data cannot be copied.
3. Take corrective steps as needed, using the buttons at the bottom of the screen.

Changing remote path bandwidth

Increase the amount of bandwidth allocated to the remote path when data copying is slower than write-workload. Bandwidth that is slow results in untransferred data accumulating in the DP pool. This in turn can result in a full DP pool, which causes pair failure.

To change bandwidth

1. In the Replication tree, click **Setup**, then click **Remote Path**. The Remote Path screen displays.
2. Select the remote path to change the bandwidth from the **Remote Path** list.
3. Click **Edit Path**. The **Edit Remote Path** screen appears.
4. Enter the bandwidth of the network that the remote path is able to use in the text box. Select **Over 1000,0Mbps** in the **Bandwidth** for over 1000,0 Bbps network bandwidth.
5. Click **OK**.
6. When the confirmation screen appears, click **Close**.

Monitoring cycle time

Cycle time is the interval between updates from the P-VOL to the S-VOL. Cycle time is set to the default of 300 seconds during pair creation.

Cycle time can range between 30 seconds to 3600 seconds. The cycle time is 300 seconds by default and you may set it up to 3600 seconds. The shortest limit value that can be set is calculated as number of CTGs of local array or remote array × 30 seconds. When consistency groups are used, the minimum cycle time increases. For one group the minimum cycle time is 30 seconds, for two groups minimum cycle time is 60 seconds, and so on, up to 64 groups with a minimum of 32 minutes. See [Table 21-3 on page 21-15](#)

Table 21-3: Number of CTGs and minimum cycle time

CTGs number	Minimum cycle time
1	30 seconds
2	1 minute
3	1.5 minutes
16	8 minutes
64	32 minutes

Updated data is copied to the S-VOL at the cycle time intervals. Be aware that this does not guarantee that all differential data can be sent within the cycle time. If the inflow to the P-VOL increases and the differential data to be copied is larger than bandwidth and the update cycle allow, then the cycle expands until all the data is copied.

When the inflow to the P-VOL decreases, the cycle time normalizes again. If you suspect that the cycle time should be modified to improve efficiency, you can reset it.

You learn of cycle time problems through monitoring. Monitoring cycle time can be done by checking group status, using CLI. See [Confirming consistency group \(CTG\) status on page D-23](#) for details.



NOTES:

1. Because the drive spin-up or system copy (operation for ensuring the system configuration) is performed preferentially to the updated copy of TCE, the cycle of TCE is temporarily interrupted if either of these operations is performed. As a result, the corresponding cycle time is lengthened.
 2. If an unpaired CTG occurs due to pair deletion, the number of CTGs may differ between the local array and the remote array. In that case, you can match the number of CTGs in the local array and the remote array by deleting the unpaired CTG.
-

Changing cycle time

To change cycle time

1. In the Replication tree, click **Setup**, then click **Options**. The Options screen appears.
2. Click **Edit Options**. The Edit Options screen appears.
3. Enter the new **Cycle Time** in seconds. The limits are 30 seconds to 3600 seconds.

4. Click **OK**.
5. When the confirmation screen appears, click **Close**.

Changing copy pace

Copy pace is the rate that data is copied during pair creation, resynchronization, and updating. The pace can be slow, medium, and fast. The time that it takes to complete copying depends on pace, the amount of data to be copied, and bandwidth.

To change copy pace

1. Connect to the local disk array and select the **Remote Replication** icon in the Replication tree view.
2. Select a pair from the pair list.
3. Click **Edit Pair**. The Edit Pair screen appears.
4. Select a **copy pace** from the dropdown list.
 - Slow — Copying takes longer when host I/O activity is high. The time to complete copying may be lengthy.
 - Medium — (Recommended) Copying is performed continuously, though it does not have priority; the time to completion is not guaranteed.
 - Fast — Copying is performed continuously and has priority. Host I/O performance is degraded. Copying time is guaranteed.
5. Click **OK**.
6. When the confirmation message appears, click **Close**.

Monitoring synchronization

Monitoring synchronization is monitoring the time difference between the P-VOL data and S-VOL data. If the time difference becomes larger, it means that the RPO performance has decreased. In this case, it is likely that a failure has occurred somewhere in the system or a performance bottleneck has occurred. By detecting the abnormality immediately and by taking appropriate corrective action, you can reduce the risk (such as mounting data loss) in the event of a disaster.

Monitoring synchronization using CCI

To monitor synchronization, use the `pairsyncwait` command of CCI from the local host. Using this command, you can find out the timing of reflecting the update data written to the P-VOL on the S-VOL.

Figure 21-4 shows an example of synchronization monitoring using `pairsyncwait`. In this example, the current Q-Marker is obtained. By measuring the time until that Q-Marker is reflected on the S-VOL, the time difference between the P-VOL data and S-VOL data is estimated.

In TCE, array at local manages two different Q-Markers for P-VOLs and its associated.

S-VOLs in a CTG. When a P-VOL in the CTG is updated by a host, Q-Marker is incremented by one. When a cycle completes, Q-Marker of S-VOLs in a CTG is updated by Q-Marker of P-VOLs recorded at which the cycle started. So if S-VOLs' Q-Marker is larger than or equal to Q-Marker got at Time_0, it means that all differential data updated before Time_0 have been copied to S-VOLs and P-VOL data newer than Time_0 can be read from S-VOL by `pairsplit` or `horctakeover`.

It takes about 2 minutes this means that it took two minutes until reflecting the data written to the P-VOL on the S-VOL.

The abnormal occurrence detection can be automated by making a script, executing it periodically, monitoring the time difference between the P-VOL and the S-VOL, and notifying system administrators that the time difference becomes larger than expected.

```
# date                /* Obtain current time
Fri Mar 22 11:18:58 2008/03/07

# pairsyncwait -g vg01 -nowait/* Obtain current sequence number
UnitID  CTGID  Q-Marker  Status  Q-Num
  0      3    01003408ef  NOWAIT  2

# pairsyncwait -g vg01 -t 100 -m 01003408ef/* Wait with obtained sequence number
UnitID  CTGID  Q-Marker  Status  Q-Num
  0      3    01003408ef  DONE    0

# date
Fri Mar 22 11:21:10 2008/03/07
```

Figure 21-4: Monitoring synchronization

Monitoring synchronization using Navigator 2

The time difference between the P-VOL data and S-VOL data can also be checked by using Navigator 2.

By checking the TCE pair information with Navigator 2, you can see the determination time of the S-VOL. The time difference between the P-VOL and S-VOL can be determined by comparing the current time of the local array and determination time of the S-VOL.

Figure 21-5 shows an example of displaying the current time of array, and Figure 21-6 shows an example of displaying the determination time of the S-VOL. The time difference between the P-VOL and S-VOL can be calculated by subtracting the determination time from the current time.




Figure 21-5: Checking the current time of the local array with Navigator 2 GUI

How asynchronous copies are performed and when each cycle complete can be monitored from Navigator 2. Navigator 2 shows how much data needs to be copied from P-VOL to S-VOL and a prediction of time when it will complete.

```
% aureplicationremote -unit array-name -refer -groupinfo
Group                               CTL Lapsed Time  Difference Size[MB]  Tran
sfer Rate[KB/s]  Transfer Completion
0:TCE_Group1      0  00:00:25      0
200  00:00:30
%
```

Figure 21-6: Prediction of cycle completion time from Navigator 2 CLI

Routine maintenance

You may want to delete a volume pair or remote path. The following sections provide prerequisites and procedures.

Deleting a volume pair

When a pair is deleted, the P-VOL and S-VOL change to Simplex status and the pair is no longer displayed in the GUI Remote Replication pair list.

Please review the following before deleting a pair:

- When a pair is deleted, the primary and secondary volumes return to the Simplex state after the differential data accumulated in the local disk array is updated to the S-VOL. Both are available for use in another pair. Pair status is *Paired:delete* while differential data is transferred.
- If failure occurs when the pair is *Paired:delete*, the data transfer is terminated and the pair becomes *Failure*. While pair status changes to *Failure*, it cannot be resynchronized.
- Deleting a pair whose status is *Synchronizing* causes the status to become *Simplex* immediately. In this case, data consistency is not guaranteed.
- A Delete Pair operation can result in the pair deleted in the local disk array but not in the remote disk array. This can occur when there is a remote path failure or the pair status on the remote disk array is *Busy*. In this instance, wait for the pair status on the remote disk array to become *Takeover*, then delete it.
- Normally, a Delete Pair operation is performed on the local disk array where the P-VOL resides. However, it is possible to perform the operation from the remote disk array, though with the following results:
 - Only the S-VOL becomes *Simplex*.
 - Data consistency in the S-VOL is not guaranteed.
 - The P-VOL does not recognize that the S-VOL is in *Simplex* status. Therefore, when the P-VOL tries to send differential data to the S-VOL, it sees that the S-VOL is absent, and P-VOL pair status changes to *Failure*.
 - When a pair's status changes to *Failure*, the status of the other pairs in the group also becomes *Failure*.
- After an *SVOL_Takeover* command is issued, the pair cannot be deleted until S-VOL data is restored from the remote DP pool.

To delete a TCE pair

1. In the Navigator 2 GUI, select the desired disk array, then click the **Show & Configure disk array** button.

2. From the Replication tree, select the **Remote Replication** icon in the Replication tree view.
3. Select the pair you want to delete in the **Pairs** list.
4. Click **Delete Pair**.

Deleting the remote path

Delete the remote path from the local disk array.

When the planned shutdown is necessary such as maintenance the remote array, delete the remote path.



NOTE: The status of all volumes must not be synchronizing or paired.

Prerequisites

- Pairs must be in Split or Simplex status.

To delete the remote path

1. In the Storage Navigator 2 GUI, select the **Remote Path** icon in the **Setup** tree view in the **Replication** tree.
2. On the Remote Path screen, click the box for the path that is to be deleted.
3. Click the **Delete Path** button.
4. Click **Close** on the Delete Remote Path screen.

TCE tasks before a planned remote disk array shutdown

Before shutting down the remote disk array, do the following:

- Split all TCE pairs. If you perform the shutdown without splitting the pairs, the P-VOL status changes to Failure. In this case, re-synchronize the pair after restarting the remote disk array.
- Delete the remote path (from local disk array).

TCE tasks before updating firmware

Before and after updating an disk array's firmware, perform the following TCE operations:

- TCE pairs must be split before updating the disk array firmware.
- After the firmware is updated, resynchronize TCE pairs.

Troubleshooting

TCE stops operating when any of the following occur:

- Pair status changes to **Failure**
- Pair status changes to **Pool Full**
- Remote path status changes to **Detached**

To track down the cause of the problem and take corrective action.

1. Check the Event Log, which may indicate the cause of the failure. See [Using the event log on page 21-32](#).
2. Check pair status.
 - a. If pair status is Pool Full, please continue with instructions in [TCE troubleshooting on page 21-22](#).
 - b. If pair status is Failure, check the following:
 - Check the status of the local and remote disk arrays. If there is a Warning, please continue with instructions in [Correcting disk array problems on page 21-26](#).
 - Check pair operation procedures. Resynchronize the pairs. If a problem occurs during resynchronization, please continue with instructions in [Correcting resynchronization errors on page 21-30](#).
3. Check remote path status. If status is Detached, please continue with instructions in [Correcting disk array problems on page 21-26](#).

For troubleshooting flow diagrams see [Figure 21-7 on page 21-22](#) and [Figure 21-8 on page 21-23](#).




Figure 21-7: TCE troubleshooting




Figure 21-8: TCE troubleshooting

Correcting DP pool shortage

When the usage rate of the DP pool in local array reaches the Replication Data Released threshold or the usage rate of the DP pool in remote array reaches the Replication Data Released threshold, the data copy from local array to remote array is stopped and the pair status change into **Pool Full**. In that case the pair status must be recovered according to the following directions. The status of Snapshot pairs that uses the DP pool that is shortage change into Failure and the V-VOL data is lost.

When the usage rate of a DP pool reaches the Replication Data Released threshold in the local array, the replication data that has been saved in the DP pool in the local array by TCE is deleted. Also, when the DP pool is depleted, the replication data that has been saved in the DP pool in each array by Snapshot is deleted regardless of whether the array is local or remote

For DP pool troubleshooting flow diagrams see [Figure 21-9 on page 21-24](#) and [Figure 21-10 on page 21-25](#).




Figure 21-9: DP pool troubleshooting




Figure 21-10: DP pool troubleshooting

Cycle copy does not progress

The cycle copy which is executed in Paired status does not start until the system copy completes. See [Confirming consistency group \(CTG\) status on page D-23](#) and if Difference Size does not decrease for a while, it may be possible that the system copy is running. And see [Displaying the event log on page D-24](#). If Message A listed below is displayed, you can confirm that the system copy is actually running and preventing the cycle copy from working. Please wait until the system copy completes without executing any commands on HSNM2. Message B will be displayed when the system copy completes.

Message A (Displayed when the system copy starts):

```
00 I14000 System copy started(Unit-xx,HDU-yy)
```

Message B (Displayed when the system copy completes):

```
00 I14100 System copy completed(Unit-xx,HDU-yy)
xx:Unit number
yy:Drive number
```

Message contents of event log

In the message, a time when an error occurred, a message, and an error detail code are displayed. The error message that indicates the DP pool is shortage is "I6GJ00 DP Pool Consumed Capacity Over (Pool-xx)" (xx is the DP pool number).




Figure 21-11: Error message example

Correcting disk array problems

A problem or failure in an disk array or remote network path can cause pairs to stop copying. Take the following action to correct disk array problems.

1. Review the information log to see what the hardware failure is.
2. Restore the disk array. Drive failures must be corrected by Hitachi maintenance personnel.
3. When the system is restored, recover TCE pairs.

For a detached remote path, parts should be replaced, then the remote path setup again.

For drive multiple failures (shown in [Table 21-4 on page 21-27](#)) the pairs most likely need to be deleted and recreated.

Table 21-4: Failure type and recovery

Failure type	Location	Situation	Recovery procedure	Action taken by
Drive multiple failure	P-VOL	Data not reflected on the S-VOL may have been lost.	Recover the pair after the drive failure is removed.	Drive replacement: Hitachi maintenance personnel Recover the pair. User
	S-VOL	Remote copy cannot be continued because S-VOL cannot be updated.		
	Local DP pool	Remote copy cannot be continued because differential data is not available.		
	Remote DP pool	Takeover to the S-VOL cannot be done because internally pre-determined data of the S-VOL is lost.		
Path detached		Failures have occurred in the secondary array or remote path, you cannot communicate with the secondary array and cannot continue the remote copying	Replace the parts Reconstruct the remote path Recover the remote array.	Hitachi maintenance personnel*
*Contact third party if the extender failed.				

Deleting replication data on the remote array

When the usage rate of the replication data DP pool on the remote array exceeds the Replication Depletion Alert threshold value, replication data stored in the pool will not be deleted. See "Restriction when the replication data DP pool usage exceeds the Replication Depletion Alert threshold" on [Table 21-7 on page 21-33](#) for more details. To fix this non-deletion problem, you can use the following procedure to delete replication data from the pool.

Deleting replication data takes about two hours. You can repeat these steps once every two hours to delete all replication data.

To delete replication data using the HSNM2 GUI

1. Connect to the remote array.
2. Open the **Edit Pool Attribute** screen for the replication data DP pool where the usage rate exceeds the Replication Depletion Alert threshold value. Refer to [Table 21-7 on page 21-33](#) for the information on how to open the **Edit Pool Attribute** screen.
3. Click **OK**.

Deletion of replication data will be performed by clicking **OK** regardless of whether any values in the **Edit Pool Attribute** screen is changed or not.

To delete replication data using the HSNM2 CLI

1. Connect to the remote array.
2. Perform the `audppool -chg` command with any option for the replication data DP pool where the usage rate exceeds the Replication Depletion Alert threshold value. You do not have to change a value with any option to start the deletion of replication data. Refer to [Setting the replication threshold on page D-10](#) for the information on how to use the `audppool` command.

Delays in settling of S-VOL Data

When the amount of data that flows into the primary disk array from the host is larger than outflow from the secondary disk array, more time is required to complete the settling of the S-VOL data, because the amount of data to be transferred increases.

When the settlement of the S-VOL data is delayed, the amount of the data loss increases if a failure in the primary disk array occurs.

Differential data in the primary disk array increases a when:

- The load on the controller is heavy
- An initial or resynchronization copy is made
- The path or controller is switched

DP-VOLs troubleshooting

When configuring a TCE pair using the DP-VOL as a pair target volume, the TCE pair status may become Failure depending on the combination of the pair status and the DP pool status shown in [Table 21-5 on page 21-29](#). Check the pair status and the DP pool status, and perform the countermeasure according to the conditions. For checking the DP pool status, check all the DP pools to which the P-VOLs, the S-VOLs, the local site DP pool, and the remote site DP pool of the pairs where pair failures have occurred belong. Refer to the *Dynamic Provisioning User's Guide* for how to check the DP pool status.

Table 21-5: Cases and solutions using the DP-VOLs

Pair status	DP pool status	Cases	Solutions
Paired Synchronizing	Formatting	Although the DP pool capacity is being added, the format progress is slow and the required area cannot be allocated.	Wait until the formatting of the DP pool for total capacity of the DP-VOLs created in the DP pool is completed.
	Capacity Depleted	The DP pool capacity is depleted and the required area cannot be allocated.	For making the DP pool status normal, perform the DP pool capacity growing and DP pool optimization, and increase the DP pool free capacity.

Correcting resynchronization errors

When a failure occurs after a resynchronization has started, an error message cannot be displayed. In this case, you can check for the error detail code in the Event Log. [Figure 21-12](#) shows an example of the detail code.

The error message for pair resynchronizing is “The change of the remote pair status failed”.




Figure 21-12: Detail code example for failure during resync

[Table 21-6](#) lists error codes that can occur during a pair resync and the actions you can take to make corrections.

Table 21-6: Error codes for failure during resync

Error code	Error contents	Actions to be taken
0307	The disk array ID of the remote disk array cannot be specified.	Check the serial number of the remote disk array.
0308	The volume assigned to a TCE pair cannot be specified.	The resynchronization cannot be performed. Create a pair again after deleting the pair.
0309	Restoration from the DP pool is in progress.	Retry after waiting for a while.
030A	The target S-VOL of TCE is a P-VOL of Snapshot. Besides, the Snapshot pair is being restored or reading/writing is not allowed.	When the Snapshot pair is being restored, execute it after the restoration is completed. When reading/writing is not allowed, execute it after enabling the reading/writing.
030C	The TCE pair cannot be specified in the CTG.	The resynchronization cannot be performed. Create a pair again after deleting the pair.
0310	The status of the TCE pair is Takeover.	
0311	The status of the TCE pair is Simplex.	
031F	The volume of the S-VOL of the TCE is S-VOL Disable.	Check the volume status of in the remote disk array, release the S-VOL Disable, and execute it again.
0320	The target volume in the remote disk array is undergoing the parity correction.	Retry after waiting for a while.

Table 21-6: Error codes for failure during resync (Continued)

Error code	Error contents	Actions to be taken
0321	The status of the target volume in the remote disk array is other than Normal or Regression.	Execute it again after restoring the target volume status.
0322	The number of unused bits is insufficient.	Retry after waiting for a while.
0323	The volume status of the DP pool is other than Normal or Regression.	Retry after the pool volume has recovered.
0324	The S-VOL is undergoing the forced restoration by means of parity.	Retry after making the restoration by means of parity.
0325	The expiration date of the temporary key is expired.	The resynchronization cannot be performed because the trial time limit is expired. Purchase the permanent key.
0326	The disk drives that configure a RAID group, to which a target volume in the remote disk array belongs have been spun down.	Perform the operation again after spinning up the disk drives that configure the RAID group.
032D	The status of the RAID group that includes the S-VOL is not Normal.	Perform the same operation after the status becomes Normal.
032E	The copy operation cannot be performed because write operations to the specified S-VOL on the remote array are not allowed due to DP pool capacity depletion for the S-VOL.	Resolve the DP pool capacity depletion and retry.
032F	The process of reconfigure memory is in progress on the remote array	Retry after the process of reconfigure memory is completed.
0332	The status of the specified Replication Data DP pool for the remote array is other than Normal or Regression.	Check the status of the Relication Data DP pool for the remote array.
0333	The status of the specified Management Area DP pool for the remote array is other than Normal or Regression.	Check the status of the Management Area DP pool for the remote array.
0337	The TCE pair deletion process is running on the Management Area DP pool of the remote array.	Retry after waiting for a while.
0339	The cycle time of the local array is less than the minimum value (number of CTGs of local array or remote array × 30 seconds).	Set the cycle time of the local array to the minimum value or more. Or, delete the unused pairs and re-execute.
033A	The cycle time of the remote array is less than the minimum value (number of CTGs of local array or remote array × 30 seconds).	Set the cycle time of the remote array to the minimum value or more. Or, delete the unused pairs and re-execute.
033B	The replication data DP pool or management area DP pool on the remote array consists of SSD/FMDs only and the Tier mode for the DP pool is enabled.	Add another Tier to the DP pool or specify another DP pool.

Using the event log

Using the event log helps in locating the reasons for a problem. The event log can be displayed using Navigator 2 GUI or CLI.

To display the Event Log using the GUI

- 1. Select the **Alerts & Events** icon. The Alerts & Events screen appears.



- 2. Click the **Event Log** tab. The Event Log displays

Event Log messages show the time when an error occurred, the message, and an error detail code, as shown in [Figure 21-13](#). If the DP pool is full, the error message is "I6D000 data pool does not have free space (Data pool-xx)", where xx is the data pool number.




Figure 21-13: Detail code example for data pool error

Miscellaneous troubleshooting

Table 21-7 contains details on pair and takeover operations that may help when troubleshooting. Review these restrictions to see if they apply to your problem.

Table 21-7: Miscellaneous troubleshooting

Restriction	Description
Restrictions for pair splitting	When a pair split operation is begun, data is first copied from the P-VOL to the S-VOL. This causes a time delay before the status of the pair becomes Split.
	The splitting of the TCE pair cannot be done when the pairsplit -msecs processing is being executed for the CTG.
	When a command to split pairs in each CTG is issued while the pairsplit -msecs processing is being executed for the cascaded Snapshot pair, the splitting cannot be executed for all the pairs in the CTG.
	When a command to split each pair is issued and the target pair is under the completion processing, it cannot be accepted if the Paired to be split is undergoing the end operation.
	When a command to split each pair is issued and the target pair is under the completion processing, it cannot be accepted if the Paired to be split is undergoing the splitting operation.
	When a command to split pairs in each group is issued, it cannot be executed if even a single pair that is being split exists in the CTG concerned.
	When a command to terminate pairs in each group is issued, it cannot be executed if even a single pair that is being split exists in the CTG concerned.
	The pairsplit -P command is not supported.

Table 21-7: Miscellaneous troubleshooting (Continued)

Restriction	Description
Restrictions on execution of the horctakeover (SVOL_Takeover) command	When the SVOL_Takeover operation is performed for a pair by the horctakeover command, the S-VOL is first restored from the DP pool. This causes a time delay before the status of the pair changes.
	The restoration of up to four volumes can be done in parallel for each controller. When restoration of four or more volumes is required, the first four volumes are selected according to an order given in the requirement, but the following volumes are selected in ascending order of the volume numbers.
	Because the SVOL_Takeover operation is performed on the secondary side only, the differential data of the P-VOL that has not been transferred is not reflected on the S-VOL data even when the TCE pair is operating normally.
	When the S-VOL of the pair, to which the instruction to perform the SVOL_Takeover operation is issued, is in the Inconsistent status that does not allow Read/Write operation, the SVOL_Takeover operation cannot be executed. Whether the Split is Inconsistent or not can be referred to using Navigator 2.
	When the command specifies the target as a group, it cannot be executed for all the pairs in the CTG if even a single pair in the Inconsistent status exists in the CTG.
	When the command specifies the target as a pair, it cannot be executed if the target pair is in the Simplex or Synchronizing status.
Restrictions on execution of the pairsplit -msecs command	The pair splitting instruction cannot be issued to the Snapshot pair cascaded with the TCE S-VOL pair in the Synchronizing or Paired status from the host on the secondary side.
	When even a single pair in the CTG is being split or deleted, the command cannot be executed.
	Pairsplit -msecs processing is continued unless it becomes Failure or Pool Full.

Table 21-7: Miscellaneous troubleshooting (Continued)

Restriction	Description
Restrictions on the performance of pair delete operation	When a delete pair operation is begun, data is first copied from the P-VOL to the S-VOL. This causes a time delay before the status of the pair changes. The end processing is continued unless it becomes Failure or Pool Full.
	A pair cannot be deleted it is being split. When a delete pair command is issued to a group, it will not be executed if any of the pairs in the group is being split.
	A pair cannot be deleted when the pairsplit -mscas command is being executed. This applies singly or by the CTG.
	When a delete pair command is issued to a group, it will not be executed if any of the pairs in the group is undergoing the pair split -mscas operation.
	Also in the execution of the pairsplit -R command that requires the secondary disk array to delete a pair, the differential data of the P-VOL that has not been transferred is not reflected on the S-VOL data in the same way as the case of the SVOL_Takeover operation.
	The pairsplit -R command cannot be executed during the restoration of the S-VOL data through the SVOL_Takeover operation.
	The pairsplit -R command cannot be issued to each group when a pair, whose S-VOL data is being restored through the SVOL_Takeover operation, exists in the CTG.
Restrictions while using load balancing	The load balancing function is not applied to the volumes specified as a TCE pair. Since the ownership of the volumes specified as a TCE pair is the same as the ownership of the volumes specified as a DP pool, perform the setting so that the ownership of volumes specified as a DP pool is balanced in advance.

Table 21-7: Miscellaneous troubleshooting (Continued)

Restriction	Description
Restriction when the replication data DP pool usage exceeds the Replication Depletion Alert threshold	When the usage rate of the replication data DP pool on the remote array exceeds the Replication Depletion Alert threshold value, replication data stored in the pool will not be deleted. The replication data transferred to the remote array during a cycle copy is temporarily stored in the replication data DP pool on the remote array. Normally, the replication data stored in the pool is automatically deleted when the cycle copy completes. However, when with an increase in I/O workloads on the P-VOL, the transfer of the replication data is not complete within the cycle copy, then the stored replication data increases. This increase causes the usage rate of the replication data DP pool to exceed the Replication Depletion Alert threshold value, and deletion of the replication data at the end of a cycle copy is not done. As a result, the usage rate of the replication data DP pool is not reduced. To avoid this situation, you need to adjust the cycle time and the amount of I/O workloads so that a cycle copy will complete within the cycle time. Also, a large amount of replication data being transferred during a single cycle copy causes a sudden increase in replication data in the replication data DP pool, which makes it more likely the usage rate could exceed the Replication Depletion Alert threshold value.

TrueCopy Modular Distributed theory of operation

TrueCopy Modular Distributed (TCMD) software expands the capabilities of TrueCopy Extended Distance (TCE) software and it allows up to 8 local arrays connect to a remote array, along with bi-directional, long-distance, remote data protection originating from TCE.

The key topics in this chapter are:

- ❑ [TrueCopy Modular Distributed overview](#)
- ❑ [Distributed mode](#)

TrueCopy Modular Distributed overview

TrueCopy Modular Distributed (hereinafter called "TCMD") is software exclusive for HUS100 to expand the TrueCopy and the TCE function. By expanding the TCE function, you can back up or copy the data on one array to the maximum of eight arrays. You can also backup or copy the data on the maximum of eight arrays to one array. However, since TCMD is software to expand the Copy function, you cannot operate TCMD alone. Each TrueCopy and TCE pair consists of one copy source volume (a primary volume: hereinafter called "P-VOL") and one copy destination volume (a secondary volume: hereinafter called "S-VOL"), and the TCMD pair also consists of one P-VOL and one S-VOL.

Figure 22-1 shows an example of the centralized data backup.

TCMD can collect the backups of the master data at a maximum of eight locations to be stored at the head office. Only one array is prepared for the backup of each remote location for disaster recovery.

The HA (High Availability) configuration is unsupported in TCMD.

TCMD cannot use the data delivery and the centralized backup together.




Figure 22-1: TCMD overview

Figure 22-2 on page 22-3 shows an operation example of TCMD data delivery use. By using TCMD, ShadowImage and SnapShot together, it is possible to distribute data of a volume in the head office to each remote branch.




Figure 22-2: TCMD data delivery

Distributed mode

You can set the Distributed mode on the array by installing TCMD in the array. The Distributed mode can be set to Hub or Edge. Set the Distributed mode on all arrays that configure TCMD. The array on which the Distributed mode is set to Hub is the Hub array. The array on which the Distributed mode is set to Edge is the Edge array.

When TCMD is uninstalled, N/A is displayed for the Distributed mode. The array on which Distributed mode is displayed as N/A is called the Normal array.

Table 22-1 shows the Distributed mode type.

Table 22-1: Distributed mode type

Distributed mode	Meaning	Contents
Hub	The array is the Hub array.	You can set remote paths to two or more Edge arrays.
Edge	The array is the Edge array.	You can set remote path to one Hub array, Edge array, or Normal array.
Normal (N/A)	The array is the Normal array.	You can set remote path to one Edge array or Normal array.

Figure 22-3 on page 22-4 shows the Distributed mode setting example. Before setting the Distributed mode, TCMD must be installed in all arrays shown in Figure 22-3 and the license status must be enabled. The array in which TCMD is installed becomes the Edge array (Array A to Array H). Set the Distributed mode to Hub in only the Array X to be the Hub array.




Figure 22-3: TCMD setting example

Installing TrueCopy Modular Distributed

This chapter provides TCMD installation and setup procedures using the Navigator 2 GUI. Instructions for CLI can be found in the appendix.

- ❑ [TCMD system requirements](#)
- ❑ [Installation procedures](#)

TCMD system requirements

This section describes the minimum TCMD requirements.

Table 23-1: TCMD requirements

Item	Minimum requirements
Environment	When using TCMD by TCE: - Firmware: Version 0917/A or higher is required - Navigator 2: Version 21.70 or higher is required for management PC
	When using TCMD by TrueCopy: : - Firmware: Version 0935/A or higher is required - Navigator 2: Version 23.50 or higher is required for management PC
	When using the iSCSI for the remote path interface: firmware version 0920/B or more and HSNM2 version 21.75 or more for the management PC are required.
	CCI: Version 01-27-03/02 or higher is required for host only when CCI is used for the operation of TrueCopy or TCE.
Requirements	Array Model: HUS 150, HUS 130, HUS 110.
	Number of controllers: 2 (dual configuration).
	The TrueCopy or TCE license key is installed and its status valid on all the arrays.
	Two or more license keys for TCMD license keys.
	Command devices: Minimum 1 (The command device is required only when CCI is used for the copy operation.)

Installation procedures

Since TCMD is an extra-cost option, TCMD cannot usually be selected (locked) when first using the array. To make TCMD available, you must install TCMD and make its function selectable (unlocked).

TCMD can be installed from Navigator 2. This section describes the installation/un-installation procedures performed by using Navigator 2 via the GUI.

For procedures performed by using the Command Line Interface (CLI) of Navigator 2, see [Appendix E, TrueCopy Modular Distributed reference information](#).



NOTE: Before installing or uninstalling TCMD, verify that array is operating in a normal state. If a failure such as a controller blockade has occurred, installation or un-installation cannot be performed.

Installing TCMD

Prerequisites

- Before installing TCMD, TCE or TrueCopy must be installed and the status must be enabled.

To install TCMD

1. In the Navigator 2 GUI, click the array in which you will install TCMD.
2. Click **Show & Configure array**.
3. Select the **Install License** icon in the **Common array Task**.



The **Install License** screen appears.

4. Select the **Key File** or **Key Code** option, and then enter the file name or key code. You may **Browse** for the key file.
5. A screen appears, requesting a confirmation to install TCMD option. Click **Confirm**.

Install License

6. A completion message appears. Click **Close**.

7. The **Licenses** list screen appears. Confirm the **TC-DISTRIBUTED** character strings on the **Licenses** list and ensure its status is **Enabled**.

Installation of TCMD is now complete.

Uninstalling TCMD


To uninstall TCMD, the key code or key file provided with the optional feature is required. Once uninstalled, TCMD cannot be used (locked) until it is again installed using the key code or key file.

Prerequisite

- All of TCE or TrueCopy pairs must be deleted. Volume status must be Simplex.
- All the remote path settings must be deleted.
- All the remote port CHAP secret settings must be deleted.
- A key code or key file is required. If you do not have the key file or code, you can obtain it from the download page on the HDS Support Portal, <https://portal.hds.com>.

To uninstall TCMD


1. In the Navigator 2 GUI, click the check box for the disk array where you will uninstall TCMD, then click the **Show & Configure disk array** button.
2. Select the **Licenses** icon in the **Settings** tree view.



The **Licenses** list appears.

3. Click **De-install License**.

The **De-Install License** screen appears.



4. When you uninstall the option using the key code, click the **Key Code** option, and then set up the key code. When you uninstall the option using the key file, click the **Key File** option, and then set up the path for the key file name. Use Browse to set the path to a key file correctly. Click **OK**.
5. A message appears, click **Close**.
The Licenses list appears.
6. Confirm the **TC-DISTRIBUTED** character strings not on the **Licenses** list.

Un-installation of TCMD is now complete.

Enabling or disabling TCMD

TCMD can be set to "enable" or "disable" when it is installed. You can disable or re-enable it.

Prerequisites

- All of TCE or TrueCopy pairs must be deleted and the status of the volumes must be Simplex.
- All the remote path settings must be deleted.
- All the remote port CHAP secret settings must be deleted.

To enable or disable TCMD

1. In the Navigator 2 GUI, click the check box for the disk array, then click the **Show & Configure array** button.
2. In the tree view, click **Settings**, then click **Licenses**.
3. Select **TC-DISTRIBUTED** in the **Licenses** list.
4. Click **Change Status**. The Change License screen appears.



5. To *disable*, clear the **Enable: Yes** check box.
To *enable*, check the **Enable: Yes** check box.
 6. Click **OK**.
 7. A message appears, confirming that the feature is set. Click **Close**.
 8. The **Licenses** list screen appears. Confirm the **Status** of the **TC-DISTRIBUTED** is changed.
- Enabling or disabling of TCMD is now complete.

TrueCopy Modular Distributed setup

This chapter provides required information to set up your system for TrueCopy Modular Distributed. It includes:

Planning and design

- ☐ [Cautions and restrictions](#)
- ☐ [Recommendations](#)
- ☐ [Configuration guidelines](#)
- ☐ [Environmental conditions](#)

Setup procedures

- ☐ [Setting the remote path](#)
- ☐ [Deleting the remote path](#)
- ☐ [Setting the remote port CHAP secret](#)

Cautions and restrictions

Before using TCMD, confirm the cautions and restrictions described in the TCE [Operating system recommendations and restrictions on page 19-38](#) and the TrueCopy [Special problems and recommendations on page 15-25](#).

Precautions when writing from the host to the Hub array or Edge array

Be careful of the following points in the status where the TrueCopy pair is created for the Edge array by the Hub array.

- When performing the pair operation using Navigator 2:
 - When the TrueCopy pair status is Paired or Synchronizing, do not map the P-VOL of the Hub array to the host group. Write for the P-VOL of the TrueCopy pair in the Hub array causes an error.
 - -When the TrueCopy pair status is Split, the S-VOL of the Edge array can be mapped to the host group. However, when swapping from the S-VOL of the Edge array, do not map the S-VOL of the Edge array to the host group. If swapping is performed while mapping to the host group, the pair status may be Failure
- When performing the pair operation using CCI:
 - -Regardless of the TrueCopy pair status, map the S-VOL of the Edge array to the host group other than that the host belongs. If mapping to the same host group, the pair status may be PSUE.

Be careful of the following points in the status where the TrueCopy pair is created for the Hub array by the Edge array.

- When performing the pair operation using Navigator 2:
 - -When the TrueCopy pair status is Paired or Synchronizing, do not map the P-VOL of the Edge array to the host group. If you perform Write to the P-VOL of a TrueCopy pair in the Edge array, the pair status may become Failure.
 - -When the TrueCopy pair status is Split, the P-VOL of the Edge array can be mapped to the host group. However, when swapping from the S-VOL of the Hub array, do not map the S-VOL of the Edge array to the host group. If swapping is performed while mapping to the host group, the pair status may be Failure.
- When performing the pair operation using CCI:
 - -Regardless of the TrueCopy pair status, map the P-VOL of the Edge array to the host group other than that the host belongs. If mapping to the same host group, the pair status may be PSUE.

Setting the remote paths for each HUS in which TCMD is installed

When TCMD is installed, you will be able to set the Distributed mode to Hub or Edge. However, some combinations cannot set the remote path depending on the setting of the Distributed mode. [Table 24-1](#) shows the availability of connecting the remote path.

Table 24-1: Availability of connecting the remote path

Local Array	Remote Array		
	Hub	Edge	Normal (N/A)
Hub	Not available	Available	Not available
Edge	Available	Available	Available
Normal (N/A)	Not available	Available	Available

Setting the remote path: HUS 100 (TCMD install) and AMS2000/500/1000

Although Hitachi AMS500/1000 does not support TCMD, if the HUS100 series in which TCMD is installed is in the Edge mode, the remote path can be set. The AMS500/1000 with TCE cannot connect to the HUS with TCE and TCMD in HUB mode.

Although the Hitachi AMS2000 does not support the combination of TCE and TCMD, if the HUS100 in which TCE and TCMD are installed are in the Edge mode, the AMS2000 with TCE can connect to the HUS100. The AMS2000 with TCE cannot connect to the HUS with TCE and TCMD in HUB mode.

For the Hitachi AMS2000 on which TrueCopy and TCMD are installed, if connecting with the HUS100 series in which TrueCopy and TCMD are installed, the remote path can be set in the combination shown on [Table 24-1 on page 24-34](#) (it is the same whether the Hitachi AMS2000 is the local array or remote array). At this time, check that the firmware version of the AMS2000 to be connected is the version 08C0/A or more. If the firmware version is less than 08C0/A, the remote path cannot be set (You can set a remote path with HUS100 being a local array, but it is blocked after that.)

Important: When connecting the Hitachi AMS2000 set in the Hub mode and the HUS100 set in the Edge mode, only Fibre Channel can be used. When connecting the Hitachi AMS2000 set in the Edge mode and the HUS100 set in the Hub mode, Fibre Channel and iSCSI can be used.

Adding the Edge array in the configuration of the set TCMD

These precautions are related to CTG and cycle time, when using TCE. When adding the Edge array in the configuration of the set TCMD, you must assign new CTGs for pair creation in the Edge array and the Hub array. If the CTGs are used up to the maximum number, the Edge array cannot be added. Review the pair configuration and reduce the CTGs to be used.

The minimum value of the cycle time also increases depending on the increase of the number of CTGs. In accordance with the new configuration, review the cycle time of each array.

Check the number of CTGs in the Hub array in the configuration after the addition, and confirm that the cycle time is set to "the number of CTGs of the Hub array × 30 seconds" or more in all the Hub array and Edge arrays.

When the array whose cycle time is smaller than the minimum value exists, cycle time-out tends to occur due to the load. Furthermore, in the array whose cycle time is smaller than the minimum value, new pair creation and recreation, re-synchronizing, and swapping of the existing pairs cannot be performed.

Configuring TCMD adding an array to configuration (TCMD not used)

When you build a TCMD configuration based on the existing TrueCopy or TCE configuration, follow the steps below to add an array.

1. Split all the pairs of TrueCopy or TCE in the existing configuration.
2. Delete all the remote paths in the existing configuration.
3. Install the TCMD license in the arrays in the existing configuration and the array to be added.
4. Change the Distributed mode to Hub in the array to be Hub.
5. Create a remote path between the Hub and Edge arrays.

Changing an Array (part of a TCMD configuration with a different array)

When you change an array which is part of a TCMD configuration with a different array, ensure that you have deleted all the TrueCopy/TCE pairs and the remote path for the array before changing the array. If you try a pair operation or remote path setting when the pairs and the remote path remain, those operations may fail.

Precautions when setting the remote port CHAP secret

For the remote paths in the same array, the remote path whose CHAP secret is set to automatic input and the remote path whose CHAP secret is set to manual input cannot be mixed together. If setting the remote port CHAP in the configuration where the remote path whose CHAP secret is set to automatic input is already used, the remote path whose CHAP secret is set to automatic input cannot be connected. When setting the remote port CHAP, recreate the existing remote path by inputting the CHAP secret manually by using this procedure:

- Split all the pairs which use the target remote path.
- Delete the target remote path.
- Recreate the target remote path by inputting the CHAP secret manually.
- Resynchronize the split pair.

Recommendations

We recommend the copy pace be medium when creating and re-synchronizing pairs using TCMD.

If you create and resynchronize pairs for two or more Edge arrays from the Hub array at the same time, copy performance deteriorates and it takes more time to complete the copy. When creating and re-synchronizing pairs from the Hub array to two or more the Edge arrays, shift the time and execute it one by one.

Configuration guidelines

A system using TCMD is composed of various components, such as a Hub array, Edge array, P-VOL, S-VOL, and communication line. If there is one bottleneck on the performance among these components, the entire system performance is affected. Especially, many Edge arrays and a Hub array which performs copy processing by itself tend to become bottlenecks. To configure the system using TCMD, reducing the load to the Hub array becomes a key to maintain the performance balance of the entire system. [Figure 24-1](#) shows the example of the configuration of the system using TCMD.




Figure 24-1: TCMD configuration example

To reduce the load to the Hub array, it is necessary to consider where the bottleneck is in the entire system using TCMD. [Table 24-2 on page 24-8](#) shows the bottleneck points and effect on performance.

Table 24-2: System performance bottleneck points

Parameter	Contents	Bottleneck effect
Bandwidth	Line bandwidth connecting Hub array and Edge array	When the line bandwidth connecting the Hub array and the Edge array is a low-speed line, the line bandwidth on the Hub array side becomes a bottleneck, and the copy performance of the entire systems may deteriorate. In the low-speed line width environment, it is necessary to adjust the line bandwidth to avoid a remote path bottleneck on the Hub array side.
RAID group configuration	RAID group configuration of Hub array	When the line bandwidth connecting the Hub array and the Edge array is a high-speed line, the drive becomes a bottleneck depending on the RAID group configuration on the Hub array side, and the copy performance of the entire system may deteriorate. It is necessary to review the RAID group configuration to avoid a drive bottleneck on the Hub array side.
Drive performance	Drive type of Hub array	When the line bandwidth connecting the Hub array and the Edge array is a high-speed line, the drive becomes a bottleneck depending on the drive performance on the Hub array side, and the copy performance of the entire system may deteriorate. It is necessary to adopt a high-performance (SAS or SSD/FMD) drive to avoid a drive bottleneck on the Hub side.
Back-end performance	Back-end performance of Hub array	When the line bandwidth connecting the Hub array and the Edge array is a high-speed line, the back-end becomes a bottleneck depending on the back-end performance on the Hub array side, and the copy performance of the entire system may deteriorate. It is necessary to make the array on the Hub array side a high-performance model (HUS150) to avoid a backend bottleneck.
Copy performance	Hub array and Edge array	When no bottleneck is in the entire system environment, if there is a problem on the copy performance between the Hub array and the Edge array, check the copy environment for each array, referring to Planning volumes on page 19-36 and Pair assignment on page 15-2 .
Cycle time	Hub array and Edge array	When the cycle time is short in the Hub array or the Edge array, when using TCE, the copy transfer amount increases and a performance bottleneck may occur in the Hub array side. It is necessary to adjust the cycle time to avoid a performance bottleneck.

Environmental conditions

Acquire the environmental conditions information for what the system using TCMD needs in advance. The necessary information is:

- Line bandwidth value used for remote path
- Information of RAID group configuration used in the system
- Types of drives created by the above-mentioned RAID group
- Connection configuration of Hub array and Edge array

Based on the provided information, check if the environment of the system using TCMD is in the recommended environment for TCMD in [Figure 24-2](#). When it satisfies the recommended environment, two or more copies between the Hub array and the Edge array can be executed at the same time. When it does not satisfy the recommended environment, bottlenecks may occur in the Hub array. Reduce the load to the Hub array side by shifting the copy time, increasing the cycle time, or performing other actions suggested in [Figure 24-2](#).




Figure 24-2: TCMD environment

Create the RAID group configuration on the Hub array side by dividing it for each Edge array as shown in [Figure 24-3 on page 24-11](#)




Figure 24-3: RAID group configuration on the Hub array

5. Enter the bandwidth number into the **Bandwidth** field. Select **Over 1000,0Mbps** in the **Bandwidth** for over 1000,0 Bbps network bandwidth. When connecting the array directly to the other array, set the bandwidth according to the transfer rate.

Specify the value of the network bandwidth so that each remote path can use the bandwidth. When remote path 0 and remote path 1 use the same network, set the value of the half of the bandwidth that the remote path can use.

6. Select the local port number from the **Remote Path 0** and **Remote Path 1** drop-down list.

Local Port: Select the port number (0A and 1A) that connected to the remote path.

7. Click **OK**.
8. A message appears. Click **Close**.

Setting of the remote path is now complete.

To set up the remote path for the iSCSI array

1. Connect the array in which you want to set to the Hub array, and select the **Remote Path** icon in the Setup tree view of the **Replication** tree view. The **Remote Path** list appears

2. Click **Create Path**. The **Create Remote Path** screen appears.

Create Remote Path Help

Remote Path Properties

Define the settings for a remote path to the remote array. If CHAP authentication for iSCSI is enabled, enter the CHAP secret in the appropriate fields below. A secret can consist of 12 to 32 characters (alphanumeric characters or special symbols: !, ", #, \$, %, &, ' , (,), *, +, -, ., /, :, ;, <, =, >, ? , [, \, ^, _ , { , | , } , ~ or space).

* Interface Type: ☐ Fibre ☒ iSCSI

* Remote Array ID:

Remote Path Name : ☒ Use default value for Remote Path Name ☐ Enter Remote Path Name Manually

From 0 to 32 (alphanumeric characters, "%", "*", "+", "-", ".", "/", "w", "@", "_", ":", "[", or "]")
Disabled word "N/A" or "----"

* Bandwidth: ☒ Over 1000.0Mbps ☐ Manually Mbps
Form 0.2 to 999.9

CHAP secret: ☒ Automatic ☐ Manual

Remote Path 0:

* Local Port:

* Remote Port IP Address:

* TCP Port No. of Remote Port:

* CHAP Secret:

* CHAP Secret (Retype):

Remote Path 1:

* Local Port:

* Remote Port IP Address:

* TCP Port No. of Remote Port:

* CHAP Secret:

* CHAP Secret (Retype):

* Required field

OK Cancel Apply

3. Select **iSCSI** in the **Interface Type**.
4. Enter the remote array ID number in the **Remote Array ID**.
5. Specify the naming to the remote path name.
Use default value for Remote Path Name: The Remote Path Name named to Array_Remote Array ID.
Enter Remote Path Name Manually: Enter the characters strings the displaying characters.
6. Enter the bandwidth number to the **Bandwidth**. Select **Over 1000,0Mbps** in the Bandwidth for over 1000,0 Mbps network bandwidth. When connecting the array directly to the other array, set 1000 to the **Bandwidth**.

**NOTES:**

1. Specify the value of the network bandwidth so that each remote path can use. When remote path 0 and remote path 1 use the same network, set the value of the half of the bandwidth that the remote path can use.
 2. The bandwidth to input into the text box affects the setting of the time-out time. It does not limit the bandwidth which the remote pass uses.
-
7. When the **CHAP secret** specifies to the remote port, select manually.
 8. Specify the following items for the **Remote Path 0** and **Remote Path 1**:
 - **Local Port:** Select the port number that connected to the remote path. The IPv4 or IPv6 format can be used to specify the IP address.
 - **Remote Port IP Address:** Specify the remote port IP address that connected to the remote path.
 9. When the **CHAP secret** specifies to the remote port, enter the specified characters to the **CHAP Secret**.
 10. Click **OK**.
 11. A message appears. Click **Close**.
- Setting of the remote path is now complete.
- Repeat step 2 to 9 to set the remote path for the number of Edge arrays.

Deleting the remote path

When the remote path becomes unnecessary, delete the remote path.

Prerequisites

- To delete the remote patch, change all the TrueCopy pairs or all the TCE pairs in the array to the Simplex or Split status
- Do not do a pair operation for a TrueCopy/TCE pair when the remote path for the pair is not set up, because then the pair operation may not complete correctly.



NOTE: When performing the planned shutdown of the remote array, the remote path should not necessarily be deleted. Change all the TrueCopy pairs or all the TCE pairs in the array to the Split status, and then perform the planned shutdown of the remote array. After restarting the array, perform the pair resynchronization. However, when the Warning notice to the Failure Monitoring Department at the time of the remote path blockade or the notice by the SNMP Agent Support Function or the E-mail Alert Function is not desired, delete the remote path, and then turn off the power of the remote array.

To delete the remote path

1. Connect the Hub array, and select the **Remote Path** icon in the **Setup** tree view in the **Replication** tree. The **Remote Path** list appears.

2. Select the remote path you want to delete in the **Remote Path** list and click **Delete Path**.
3. A message appears. Click Close.

Deletion of the remote path is now complete.

Setting the remote port CHAP secret


For iSCSI array, the remote path can use the CHAP secret. Set the CHAP secret mode to the remote array to be the connection destination of the remote path. If you set the CHAP secret to the remote array, you can prevent creation of the remote path from the array that the same character string is not set as the CHAP secret.



NOTE: If setting the remote port CHAP in the array, the remote path whose CHAP secret is set to automatic input cannot be connected for the array. When setting the remote port CHAP secret while using the remote path whose CHAP secret is set to automatic input, see [Adding the Edge array in the configuration of the set TCMD](#) and recreate the remote path.

To set the remote port CHAP secret:

1. Connect to the remote array and click the **Remote Path** icon in the **Setup** tree in the **Replication** tree.




The **Remote Path** list appears.

2. Click the **Remote Port CHAP** tab and click **Add Remote Port CHAP**



The **Add Remote Port CHAP** screen appears.



3. Enter the array ID of the local array in **Local Array ID**.
4. Enter the CHAP secret to be set to each remote path in **Remote Path 0** and **Remote Path 1**. Enter it twice for confirmation.
5. Click **OK**.
6. The confirmation message appears. Click **Close**.

The setting of the remote port CHAP secret is completed.

Using TrueCopy Modular Distributed

This chapter provides procedures for performing basic TCMD operations using the Navigator 2 GUI. For CLI instructions, see the Appendix.

- ❑ [Configuration example: centralized backup using TCE](#)
- ❑ [Perform the aggregation backup](#)
- ❑ [Data delivery using TrueCopy Remote Replication](#)
- ❑ [Create a pair in data delivery configuration](#)
- ❑ [Executing the data delivery](#)
- ❑ [Setting the distributed mode](#)

Configuration example: centralized backup using TCE

You can perform an aggregation backup of the master data dispersed to multiple sites in one site. [Figure 25-1](#) shows the configuration example.




Figure 25-1: TCM configuration

Perform the aggregation backup

1. Copy the master data to the Hub array using TCE.
2. Wait for the pair status becomes **Paired**.
3. Perform the aggregation backup for all master data in the local site sequentially.
4. Create the backup by using SnapShot for the master data in the local site or the backup data in the backup site as needed.

The updated master data is copied to the backup site asynchronously as long as the pair status is **Paired**.

Refer to [Using TrueCopy Extended on page 20-1](#) for the TCE operations and refer to [Using Snapshot on page 10-1](#) for the SnapShot operations.

Data delivery using TrueCopy Remote Replication

You can perform distributing master data on the local site to multiple remote sites. In data delivery, TrueCopy, ShadowImage, SnapShot, and TCMD are used together.

Refer to the Hitachi Unified Storage TrueCopy Remote Replication User's Guide for the TrueCopy operations, the Hitachi Unified Storage ShadowImage in-system replication User's Guide for the ShadowImage operations, and the Hitachi Unified Storage Copy-on-write SnapShot User's Guide for the SnapShot operations.




Figure 25-2: Configuration of data delivery

Creating data delivery configuration

In a data delivery configuration, you need to build a cascade configuration to disperse data is needed to transfer master data to each delivery target array.

The delivery source array to deliver master data to multiple arrays needs pairs configured as follows (the detailed procedures are described later).




Figure 25-3: Data delivery configuration

Configure the delivery source array as follows.

Table 25-1: Data delivery configuration specifications

Parameter	Data delivery configuration specification (delivery source array)
User interface	- Navigator 2 : Version 23.50 or later: Used for creating volumes; setting remote path, command devices, and DMLU; and handling pairs; - CCI: used for the pair operations
Array type	HUS110/130/150 with firmware version 0935/A or later. HUS150 is highly recommended.
License	Licenses of TrueCopy, ShadowImage, SnapShot, and TCMD need to be installed.
Distributed mode	Set to Hub mode.
Remote Path	FC or iSCSI is available. Create bidirectional remote paths from the delivery source array to each delivery target array. It is required that 1.5 Mbps or more (100 Mbps or more is recommended) be guaranteed for each remote path. Two remote paths being set, the bandwidth must be 3.0 Mbps or more between the arrays.

Parameter	Data delivery configuration specification (delivery source array)
Volume	- Master data needs to be stored in a delivery source array. - If you want to deliver existing data in a delivery target array as master data, copy the data to a delivery source array in advance using TrueCopy. - In addition to a volume used for storing master data, a mirror volume is needed for temporarily storing data to be delivered - For one volume for master data, create one mirror volume the same size as that for master data. - A mirror volume can be a normal volume or a DP volume, but we recommend you to create one with the same volume type as that for master data. - We recommend you to create a mirror volume in a RAID group or DP pool other than that where a volume for master data belongs. - We recommend you create a mirror volume using 4D or more SSD/FMDs or SAS drives.
Command device	- This must be set when performing pair operations of CCI - Set them for both of the local and remote arrays.
DMLU	- This needs to be set to use a pair of TrueCopy. - Be sure to set this for both of the local and remote arrays
Pair structure	In a data delivery configuration, the following pairs are needed per volume for master data. - A ShadowImage pair where P-VOL is a volume for master data and S-VOL is a mirror volume. - SnapShot pairs where P-VOL is a mirror volume (the same number of pairs as that of delivery target arrays) - A TrueCopy pair where P-VOL is SnapShot V-VOL of the above and S-VOL is a volume in a delivery target array. In normal operation, pairs that are used for data delivery are Split. Data delivery is performed by pair resync.
Copy pace	The copy paces from a P-VOL to an S-VOL and vice versa can be adjusted in three stages.

The delivery target array needs a pair configured as follows (the detailed procedures are described later).




Figure 25-4: Delivery target pair configuration


Configure the delivery target array as follows.

Table 25-2: Data delivery configuration specification (delivery target array)


Parameter	Data delivery configuration specification (delivery target array)
User interface	- Navigator 2 : Version 23.50 or later: Used fused for the pair creation, for the setting of the remote paths, command devices, or DMLU and for the pair operations. - CCI: used for the pair operations
Array type	HUS110/130/150 with firmware version 0935/A or later. AMS2000 series with firmware version 08C0/A or later.
License	Licenses of TrueCopy, ShadowImage, SnapShot, and TCMD need to be installed.
Distributed mode	Set to Edge mode.
Remote Path	FC or iSCSI is available. Create bidirectional remote paths from the delivery source array to each delivery target array. It is required that 1.5 Mbps or more (100 Mbps or more is recommended) be guaranteed for each remote path. Two remote paths being set, the bandwidth must be 3.0 Mbps or more between the arrays.
Volume	- A delivery target volume is needed to receive delivered data. - For each set of master data, create a volume the same size as that for master data. - A delivery target volume can be a normal volume or a DP volume, but we recommend you to create one with the same volume type as that for master data. - A delivery target volume needs to be unmounted before data delivery because an access from a host causes an error. - A delivery target volume can be used in a cascade configuration of ShadowImage or SnapShot in a delivery target array.
Command device	- This must be set when performing pair operations of CCI - Set them for both of the local and remote arrays.
DMLU	- This needs to be set to use pairs of ShadowImage and TrueCopy. - Set the capacity of a volume based on a capacity to be used.
Copy pace	The copy paces from a P-VOL to an S-VOL and vice versa can be adjusted in three stages.

Create a pair in data delivery configuration


1. Create a mirror of the master data of the Hub array on the local site using ShadowImage.




2. After creating the mirror, split the ShadowImage pair on the local site.




3. Create a V-VOL for delivery using SnapShot by making the mirror as a P-VOL.




4. Execute Step 3 for the number of arrays on the remote site.




5. Split the SnapShot pair of the V-VOL for delivery.




6. Create a TrueCopy pair for the V-VOL for delivery and the volume on the remote site.



7. Execute Step 6 for the number of arrays on the remote site.




8. When copying is completed, split the TrueCopy pair.




Perform the above-mentioned operation for all master data in the local site sequentially.

Executing the data delivery


1. Resynchronize the ShadowImage pair of the master data and the mirror on the local site.




2. When the resynchronization is completed, split the ShadowImage pair of the master data and the mirror.




3. Resynchronize the SnapShot pair of the mirror and the V-VOL for delivery and split.




4. Execute Step 3 for the number of arrays on the remote site.




5. Resynchronize the V-VOL for delivery and the TrueCopy pair of the volume on the remote site.



6. Execute Step 5 for the number of arrays on the remote site.



7. When copying is completed, split the TrueCopy pair.



Perform the above operations to all the master data to be delivered.

Multiple sets of master data can be delivered simultaneously, but this places more workload. You should limit the number of configurations to be delivered simultaneously to two (two cascaded configurations). Each mirror volume used for simultaneous data delivery should belong to different RAID groups.

Master data is available for host access even during data delivery. In this case, the data at the time of ShadowImage pair split (when the above step 3 is completed) is delivered.

Setting the distributed mode

To set the remote paths between one array and two or more arrays using TCMD, set the Distributed mode to Hub for one array.

Before setting the Distributed mode, perform the following:

- Decide the configuration of the array that uses TCMD in advance, and check the array in which the Distributed mode is set to Hub and the array in which the Distributed mode remains as Edge.
- If you install the arrays in TCMD, the Distributed mode is set to Edge in the initial status.


Changing the Distributed mode to Hub from Edge

Prerequisites

- All the remote paths settings must be deleted.
- All the remote port CHAP secret settings must be deleted.
- All the remote pairs settings must be deleted (the status of all the volumes is Simplex).

To change the distributed mode to Hub from Edge

1. Connect the array want to set the Hub array, and select the **Remote Path** icon in the **Setup** tree view of the **Replication** tree view. The **Remote Path** screen appears.
2. Click **Change Distributed Mode**. The **Change Distributed Mode** dialog appears.



3. Select the **Hub** option and click **OK**.
4. A message appears, confirmation that the mode is changed. Click **Close**.

5. The Remote Path screen appears.

Remote Path

HUS110_91200026 > Replication > Setup > Remote Path

Change Distributed Mode Refresh Information Help

Summary

Local Array ID	91200026
Distributed Mode	Hub

Remote Path

Remote Array ID	Remote Path Name	Type	Path0Status	Path1Status	Bandwidth
No Object					

Create Path Recover Path 0 Recover Path 1 Edit Path Delete Path

6. Confirm that the Distributed Mode is Hub.

Changing the Distributed mode to Hub from Edge is now complete.

Changing the Distributed Mode to Edge from Hub

Prerequisites

- All the remote paths settings must be deleted.
- All the remote port CHAP secret settings must be deleted.

To change the distributed mode to Edge from Hub

1. Connect the array set to the Hub array, and select the **Remote Path** icon in the **Setup** tree view of the **Replication** tree view. The **Remote Path** screen appears
2. Click **Change Distributed Mode**. The **Change Distributed Mode** dialog appears.
3. Select the **Edge** option and click **OK**.
4. A message appears, confirmation that mode is changed. Click **Close**. The **Remote Path** screen appears.
5. Confirm the **Distributed Mode** is **Edge**

Changing the Distributed mode to Edge from Hub is now complete.

Troubleshooting TrueCopy Modular Distributed

This chapter provides information and instructions for troubleshooting and monitoring the TDE system.

□ [Troubleshooting](#)

Troubleshooting

For troubleshooting TCMD, use the same procedures as when troubleshooting TCE or TrueCopy. See [Monitoring and troubleshooting TrueCopy Extended on page 21-1](#), or [Monitoring and maintenance on page 16-2](#).

Cascading replication products

Cascading is connecting different types of replication program pairs, like ShadowImage with Snapshot, or ShadowImage with TrueCopy. It is possible to connect a local replication program pair with a local replication program pair and a local replication program pair with a remote replication program pair. Cascading different types of replication program pairs allows you to utilize the characteristics of both replication programs at the same time.

- ❑ [Cascading ShadowImage](#)
- ❑ [Cascading Snapshot](#)
- ❑ [Cascading TrueCopy Remote](#)
- ❑ [Cascading TCE](#)

Cascading ShadowImage

Cascading ShadowImage with Snapshot

Cascading a volume of Snapshot with a P-VOL of ShadowImage is supported only when the P-VOL of ShadowImage and a P-VOL of Snapshot are the same volume. Also, operations of the ShadowImage and Snapshot pairs are restricted depending on statuses of the pairs. See [Figure 27-1](#).




Figure 27-1: Cascading with a ShadowImage P-VOL

Restriction when performing restoration

When performing restoration, the pair status of the pairs must be different than those which make the restoration **Split**. While the ShadowImage pair is executing restoration, the V-VOLs of the cascaded Snapshot cannot be Read/Write. When the restoration is completed, Read/Write from/to all the V-VOLs will be possible again.




Figure 27-2: While restoring ShadowImage, the Snapshot V-VOL cannot be Read/Write

I/O switching function

The I/O switching function operates even in the configuration in which ShadowImage and Snapshot are cascaded. However, when cascading TrueCopy, if the I/O switching target pair and TrueCopy are cascaded, the I/O switching function does not operate.

Performance when cascading P-VOL of ShadowImage with Snapshot

When the P-VOL of ShadowImage and the P-VOL of Snapshot are cascaded, and when the ShadowImage pair status is any of **Paired**, **Paired Internally Synchronizing**, **Synchronizing** and **Split Pending**, and the Snapshot pair status is **Split**, the host I/O performance for the P-VOL deteriorates. Use ShadowImage in the **Split** status and, if needed, resynchronize the ShadowImage pair and acquire the backup.

Table 27-1 shows whether a read/write from/to a P-VOL of ShadowImage is possible or not in the case where a P-VOL of Snapshot and a Snapshot P-VOL are the same volume.

Failure in this table excludes a condition in which volume access is not possible (for example, volume blockage).

When one P-VOL configures a pair with one or more S-VOLs, decide which item is applied as the pair status of the P-VOL of ShadowImage with the following procedure:

- If all the pairs that the P-VOL configures are in **Split**, the item of **Split** is applied.
- If all the pairs that the P-VOL configures are **Split** or **Failure** status, **Split** is applied. However, when including the pair that became **Failure** during restore, the items of **Failure** (Restore) are applied.
- If a pair in the **Paired** status, the **Synchronizing** status, or the **Reverse Synchronizing** status is included in the pair that the P-VOL configures, **Paired**, **Synchronizing**, and **Reverse Synchronizing** is applied, respectively.
- When multiple Paired statuses and Synchronizing status exist in the pairs that the relevant P-VOL configures, if the respective statuses are all Readable, they are Readable. Moreover, if the respective statuses are all Writable, they are Writable.

Table 27-1: A Read / Write instruction to a ShadowImage P-VOL

Snapshot P_VOL	ShadowImage P_VOL							
	Paired (including paired internally synchronizing)	Syn-chroniz-ing	Reverse syn-chroniz-ing	Split	Split pend-ing	Failure	Failure (restore)	Failure (S_VOL switch)
Paired	YES	YES	YES	YES	YES	YES	NO	NO
Reverse Synchronizing	YES	YES	NO	YES	YES	YES	NO	NO
Split	YES	YES	YES	YES	YES	YES	YES	NO
Failure	YES	YES	YES	YES	YES	YES	YES	YES
Failure (Restore)	NO	NO	NO	YES	NO	YES	NO	NO

YES indicates a possible case
 NO indicates an unsupported case

Table 27-2 and Table 27-3 on page 27-6 shows pair status and operation when cascading Snapshot with ShadowImage. The shaded areas in the tables indicate unworkable combinations.

Table 27-2: ShadowImage pair operation when volume shared with P-VOL on ShadowImage and Snapshot ○

ShadowImage operation	Snapshot pair status				
	Paired	Reverse synchronizing	Split	Failure	Failure (restore)
Creating pairs	YES	NO	YES	YES	NO
Splitting pairs	YES		YES	YES	
Re-synchronizing pairs	YES	NO	YES	YES	NO
Restoring pairs	NO	NO	YES	YES	NO
Deleting pairs	YES	YES	YES	YES	YES

YES indicates a possible case
 NO indicates an unsupported case

Table 27-3: Snapshot pair operation when volume shared with P-VOL on ShadowImage and Snapshot)

Snapshot operation	ShadowImage pair status							
	Paired (including paired internally synchronizing)	Synchro nizing	Reverse synchro nizing	Split	Split pending	Failure	Failure (restore)	Failure (S_VOL switch)
Creating pairs	YES	YES	NO	YES	YES	YES	NO	NO
Splitting pairs	YES	YES		YES	YES	YES		NO
Re-synchroni zing pairs	YES	YES	NO	YES	YES	YES	NO	NO
Restoring pairs	NO	NO	NO	YES	NO	YES	NO	NO
Deleting pairs	YES	YES	YES	YES	YES	YES	YES	YES

YES indicates a possible case
 NO indicates an unsupported case

Cascading a ShadowImage S-VOL with Snapshot

Figure 27-3 shows the cascade of a ShadowImage S-VOL.




Figure 27-3: Cascading a ShadowImage S-VOL

Restrictions

- **Restriction of pair creation order.** When cascading a P-VOL of Snapshot with an S-VOL of ShadowImage, create a ShadowImage pair first. When a Snapshot pair is created earlier, delete the Snapshot pair once and create a pair using ShadowImage.
- **Restriction of Split Pending.** When the ShadowImage pair status is **Split Pending**, the Snapshot pair cannot be changed to the **Split** status. Execute it again after changing the ShadowImage pair status to other than **Split Pending**.
- **Changing the Snapshot pair to Split while copying ShadowImage.** When the Snapshot pair is changed to the **Split** status while the ShadowImage pair status is **Synchronizing** or **Paired Internally Synchronizing**, the V-VOL data of Snapshot cannot be guaranteed. This is because the status where the background copy of ShadowImage is operating is the V-VOL data of Snapshot.
- **Performing pair re-synchronization when the ShadowImage pair status is Failure.** If a pair is re-synchronized when the ShadowImage pair status is **Failure**, all data is copied from the P-VOL.

to the S-VOL of ShadowImage. When the Snapshot pair status is **Split**, all data of the P-VOL of Snapshot is saved to the V-VOL. Be careful of the free capacity of the data pool used by the V-VOL.

- **Performance at the time of cascading the S-VOL of Snapshot and ShadowImage.** When the S-VOL of ShadowImage and the P-VOL of Snapshot are cascaded, and when ShadowImage pair status is any of **Paired**, **Paired Internally Synchronizing**, **Synchronizing**, or **Split Pending**, and Snapshot pair status is **Split**, the host I/O performance for the P-VOL of ShadowImage deteriorates. Use ShadowImage in the **Split** status and, if needed, resynchronize the ShadowImage pair and acquire the backup.

Table 27-4 shows whether a read or write from or to an S-VOL of ShadowImage is possible when a P-VOL of Snapshot and an S-VOL of ShadowImage are the same volume.

Table 27-4: Read / Write instructions to an S-VOL of ShadowImage

Snapshot P_VOL status	ShadowImage S_VOL status							
	Paired (including Paired Internally Synchronizing)	Syn-chroniz-ing	Reverse Synchro-nizing	Split	Split Pend-ing	Failure	Failure (Restore)	Failure (S_VOL Switch)
Paired	R	R	R	R/W	R/W	R	R/W	R/W
Reverse Syn-chronizing	NO	NO	NO	R/W	NO	NO	NO	NO
Split	R	R	R	R/W	R/W	R	R/W	R/W
Failure	R	R	R	R/W	R/W	R	R/W	R/W
Failure (Restore)	NO	NO	NO	R/W	NO	R/W	NO	NO

R/W: Read/Write by a host is possible.

R: Read by a host is possible but write is unsupported .

NO indicates an unsupported case

R/W: Read/Write by a host is unsupported .

**NOTE: When using Snapshot with ShadowImage**

- **Failure** in this table excludes a condition in which volume access is not possible (for example, volume blockage).
- When one P-VOL configures a pair with one or more S-VOLs, decide which item is applied as the pair status of the P-VOL of the above-mentioned ShadowImage with the following procedure:
 - If all the pairs that the P-VOL configures are in the **Split** status, the item of **Split** is applied.
 - If all the pairs that the P-VOL configures are in the **Split** status or the **Failure** status, the item of **Split** is applied. However, when including the pair that became **Failure** during restore, the items of **Failure** (Restore) are applied.
 - If a pair in the **Paired** status, the **Synchronizing** status, or the **Reverse Synchronizing** status is included in the pair that the P-VOL concerned configures, the item of **Paired**, **Synchronizing**, and **Reverse Synchronizing** is applied, respectively.
 - When multiple Paired statuses and Synchronizing status exist in the pairs that the relevant P-VOL configures, if the respective statuses are all Readable, they are Readable. Moreover, if the respective statuses are all Writable, they are Writable.

Table 27-5: ShadowImage pair operation when volume shared with S-VOL on ShadowImage and P-VOL on Snapshot

ShadowImage operation	Snapshot pairs status				
	Paired	Reverse Synchronizing	Split	Failure	Failure (Restore)
Creating pairs	NO	NO	NO	NO	NO
Splitting pairs	YES		YES	YES	
Re-synchronizing pairs	YES	NO	YES	YES	NO
Restoring pairs	YES	NO	YES	YES	NO
Deleting pairs	YES	YES	YES	YES	YES

YES indicates a possible case

NO indicates an unsupported case

Table 27-6: Snapshot Pair Operation when volume shared with S-VOL on ShadowImage and P-VOL on Snapshot

Snapshot operation	ShadowImage pair status							
	Paired (including Paired Internally Synchronizing)	Synchronizing	Reverse Synchronizing	Split	Split Pending	Failure	Failure (Restore)	Failure (S_VOL Switch)
Creating pairs	YES	YES	YES	YES	NO	YES	YES	YES
Splitting pairs	YES	YES	YES	YES	NO	YES	YES	YES
Re-synchronizing pairs	YES	YES	YES	YES	NO	YES	YES	YES
Restoring pairs	NO	NO	NO	YES	NO	NO	NO	NO
Deleting pairs	YES	YES	YES	YES	YES	YES	YES	YES

YES indicates a possible case
NO indicates an unsupported case

Cascading restrictions with ShadowImage P-VOL and S-VOL

The P-VOL and the S-VOL of ShadowImage can cascade Snapshot at the same time as shown in [Figure 27-4](#). However, when operating ShadowImage in the status of **Paired**, **Paired Internally Synchronizing**, **Synchronizing**, or **Split Pending** and operating Snapshot in the **Split** status, the performance deteriorates significantly. Start the operation after advance verification




Figure 27-4: Simultaneous cascading restrictions with ShadowImage P-VOL and S-VOL

Cascading ShadowImage with TrueCopy

ShadowImage volumes can be cascaded with those of TrueCopy, as shown in the [Figure 27-5 on page 27-12](#) through [Figure 27-10 on page 27-17](#). ShadowImage P-VOL volumes can be cascaded with those of TrueCopy in the Split Pending or Paired Internally Synchronizing status. ShadowImage S-VOL volumes cannot be cascaded with those of TrueCopy in the Split Pending or Paired Internally Synchronizing status.

For details on concurrent use of TrueCopy, see [Cascading TCE on page 27-78](#).




Figure 27-5: Cascading a ShadowImage P-VOL with TrueCopy (P-VOL: S-VOL=1: 1)




Figure 27-6: Cascading a ShadowImage P-VOL with TrueCopy (P-VOL: S-VOL=1:3)

Cascading a ShadowImage S-VOL

A cascade of a ShadowImage S-VOL is used when making a backup on the remote side asynchronously. Because a backing up is made from a ShadowImage S-VOL to the remote side in this configuration, lowering of performance on the local side (a ShadowImage P-VOL) during the backing up can be minimized. Note that a TrueCopy pair must be placed in the Split status when re-synchronizing a ShadowImage volume on the local side.




Figure 27-7: Cascading a ShadowImage S-VOL with TrueCopy (P-VOL: S-VOL=1: 1)




Figure 27-8: Cascading a ShadowImage S-VOL with TrueCopy (P-VOL: S-VOL=1:3)

Cascading a ShadowImage P-VOL and S-VOL

Volumes of ShadowImage P-VOL can be cascaded with those of TrueCopy in the **Split Pending** or **Paired Internally Synchronizing** status. Volumes of ShadowImage S-VOL cannot be cascaded with those of TrueCopy in the **Split Pending** or **Paired Internally Synchronizing** status. See [Figure 27-9](#) and [Figure 27-10](#) on page 27-17.




Figure 27-9: Cascading a ShadowImage P-VOL and S-VOL




Figure 27-10: Cascading a ShadowImage P-VOL and S-VOL

Cascading restrictions on TrueCopy with ShadowImage and Snapshot

Snapshot can cascade ShadowImage and TrueCopy at the same time. However, since the performance may deteriorate, start the operation after advance verification. See [Chapter 12, TrueCopy Remote Replication theory of operation](#) for details on the concurrent use of TrueCopy.

Cascade restrictions of TrueCopy S-VOL with Snapshot V-VOL. In the configuration in which the P-VOL of ShadowImage and the P-VOL of Snapshot are cascaded as shown in [Figure 27-11 on page 27-18](#), and at the same time, in the configuration in which the V-VOL of Snapshot and the S-VOL of TrueCopy are cascaded, when the TrueCopy pair status is **Paired** or **Synchronizing**, ShadowImage cannot be restored. Change the TrueCopy pair status to **Split**, and then execute it again.




Figure 27-11: Cascade restrictions of TrueCopy S-VOL with Snapshot V-VOL

Cascading restrictions on TCE with ShadowImage

ShadowImage may be in concurrent use with TCE, however ShadowImage cannot be cascaded with TCE.

Cascading Snapshot

Cascading Snapshot with ShadowImage

Volumes of Snapshot can be cascaded with those of ShadowImage as shown in [Figure 27-12](#). For details, see [Cascading ShadowImage with Snapshot on page A-39](#).




Figure 27-12: Cascading Snapshot with ShadowImage

Cascading restrictions with ShadowImage P-VOL and S-VOL

When the firmware version of the disk array is 0920/B or more, the P-VOL and the S-VOL of ShadowImage can cascade Snapshot at the same time as shown in [Figure 27-4](#). However, when operating ShadowImage in the status of **Paired**, **Paired Internally Synchronizing**, **Synchronizing**, or **Split Pending** and operating Snapshot in the **Split** status as is, the performance deteriorates significantly. Start the operation after advance verification.




Figure 27-13: Simultaneous cascading restrictions with ShadowImage P-VOL and S-VOL

Cascading Snapshot with TrueCopy Remote

Volumes of Snapshot can be cascaded with those of TrueCopy as shown in [Figure 27-14](#). Because the cascade of Snapshot with TrueCopy lowers the performance, only use it when necessary.




Figure 27-14: Cascading a Snapshot P-VOL with TrueCopy

Cascading a Snapshot P-VOL

Snapshot volumes can be cascaded with those of TrueCopy, as shown in [Figure 27-15](#). For details on cascade use of TrueCopy, refer to [Cascading TrueCopy Remote](#) on page 27-29.




Figure 27-15: Cascade with a Snapshot P-VOL

Cascading a Snapshot V-VOL

A cascade of a Snapshot V-VOL is used when making a backup on the remote side asynchronously. Though the cascade of Snapshot can decrease an S-VOL (V-VOL) capacity differently from a cascade of ShadowImage, the performance on the local side (a P-VOL of Snapshot) is affected by the backup. Furthermore, SnapShot can create multiple V-VOLs for a P-VOL, but up to eight V-VOLs can be cascaded to TrueCopy. Note that a TrueCopy pair must be placed in the Split status when resync (giving the Snapshot instruction) a volume of Snapshot on the local side. See [Figure 27-16](#) and [Figure 27-17](#) on page 27-24.




Figure 27-16: Cascade with a Snapshot V-VOL




Figure 27-17: Cascade with a Snapshot V-VOL

Configuration restrictions on the Cascade of TrueCopy with Snapshot

Figure 27-18 shows an example of a configuration in which restrictions are placed on the cascade of TrueCopy with Snapshot. SnapShot can create multiple V-VOLs for a P-VOL, but up to eight V-VOLs can be cascaded to TrueCopy. Furthermore, when resynchronizing the local side SnapShot (giving the SnapShot instruction), you need to change the TrueCopy pairs to Split.




Figure 27-18: Configuration restrictions on the cascade of TrueCopy with Snapshot

Cascade restrictions on TrueCopy with ShadowImage and Snapshot

Snapshot can cascade ShadowImage and TrueCopy at the same time. However, since the performance may deteriorate, start the operation after advance verification. See [Cascading TrueCopy Remote on page 27-29](#) for details on the concurrent use of TrueCopy.

Cascade restrictions of TrueCopy S-VOL with Snapshot V-VOL. In the configuration in which the P-VOL of ShadowImage and the P-VOL of Snapshot are cascaded as shown in [Figure 27-19](#), and at the same time, in the configuration in which the V-VOL of Snapshot and the S-VOL of TrueCopy are cascaded, when the TrueCopy pair status is **Paired** or **Synchronizing**, ShadowImage cannot be restored. Change the TrueCopy pair status to **Split**, and then execute it again.




Figure 27-19: Cascade restrictions on TrueCopy S-VOL with Snapshot V-VOL

Cascading Snapshot with True Copy Extended

Volumes of TCE can be cascaded with those of Snapshot P-VOL as shown in [Figure 27-20](#) and [Figure 27-21](#). For details on concurrent use of TCE, refer to [Cascading TCE on page 27-78](#).




Figure 27-20: Cascading a Snapshot P-VOL with a TCE P-VOL




Figure 27-21: Cascading a Snapshot P-VOL with a TCE S-VOL

Restrictions on cascading TCE with Snapshot

Figure 27-22 shows an example of a configuration in which restrictions are placed on the cascade of TCE with Snapshot.




Figure 27-22: Restrictions on the cascade of TCE with Snapshot

Cascading TrueCopy Remote

TrueCopy can be cascaded with ShadowImage or Snapshot. It cannot be used with TrueCopy Extended Distance.

A cascaded system is a configuration in which a TrueCopy P-VOL or S-VOL is shared with a ShadowImage or Snapshot P-VOL or S-VOL. Cascading with ShadowImage reduces performance impact to the TrueCopy host and provides data for use with secondary applications.

Many but not all configurations, operations, and statuses between TrueCopy and ShadowImage or Snapshot are supported. See [Cascading ShadowImage on page 27-2](#), and [Cascading with Snapshot on page 27-78](#), for detailed information.

Cascading with ShadowImage

In a cascaded system the TrueCopy P-VOL or S-VOL is shared with a ShadowImage P-VOL or S-VOL. Cascading is usually done to provide a backup that can be used if the volumes are damaged or have inconsistent data. This section discusses the configurations, operations, and statuses that are allowed in the cascaded systems.

Cascade overview

TrueCopy's main function is to maintain a copy of the production volume in order to fully restore the P-VOL in the event of a disaster. A ShadowImage backup is another copy—of either the local production volume or the remote S-VOL. A backup ensures that the TrueCopy system:

- Has access to reliable data that can be used to stabilize inconsistencies between the P-VOL and S-VOL, which can result when a sudden outage occurs.
- Can complete the subsequent recovery of the production storage system.

Cascading TrueCopy with ShadowImage ensures the following:

- When ShadowImage is cascaded on the local side, TrueCopy operations can be conducted from the local ShadowImage S-VOL. In this case, the latency associated with the TrueCopy backup is lowered, improving host I/O performance.
- When ShadowImage is cascaded on the remote side, data in the ShadowImage S-VOL can be used as a backup for the TrueCopy S-VOL, which may be required in the event of failure during a TrueCopy resynchronization. The backup data is used to restore the TrueCopy S-VOL, if necessary, from which the local P-VOL can be restored.
- A full-volume copy can be used development, reporting, and so on.
- When both of TrueCopy and ShadowImage pairs are placed in the Paired status, performance of a host on the local side is lowered. It is recommended to make the status of a TrueCopy and ShadowImage pairs Split when host I/Os are done frequently.

Cascade configurations

Cascade configurations can consist of P-VOLs and S-VOLs, in both TrueCopy and ShadowImage.

The following sections show supported configurations.



NOTE:

1. Cascading TrueCopy with another TrueCopy system or TrueCopy Extended Distance is not supported.
 2. When a restore is done on ShadowImage, the TrueCopy pair must be split.
-

Configurations with ShadowImage P-VOLs

The ShadowImage P-VOL can be shared with the TrueCopy P-VOL or S-VOL.

Figure 27-23 and Figure 27-24 on page 27-32 shows cascade configurations where the ShadowImage P-VOL to S-VOL ratio is 1-to-1.

When a restore is executed on ShadowImage, the TrueCopy pair must be split.




Figure 27-23: Cascade using ShadowImage P-VOL (P-VOL:S-VOL = 1:1)




Figure 27-24: Cascade using ShadowImage P-VOL (P-VOL:S-VOL = 1:1)

Figure 27-25 and Figure 27-26 on page 27-34 shows cascade configurations where the ShadowImage P-VOL to S-VOL ratio is 1-to-3. I




Figure 27-25: Cascade using ShadowImage P-VOL (P-VOL:S-VOL = 1:3)




Figure 27-26: Cascade Using ShadowImage P-VOL (P-VOL:S-VOL = 1:3)

Configurations with ShadowImage S-VOLs

The ShadowImage S-VOL can be shared with the TrueCopy P-VOL only. The TrueCopy remote copy is created from the local side ShadowImage S-VOL. This results in an asynchronous TrueCopy system.



NOTE: a TrueCopy pair must be placed in the Split status when re-synchronizing a volume of ShadowImage on the local side.

Figure 27-27 shows cascade configurations where the ShadowImage P-VOL to S-VOL ratio is 1-to-1.




Figure 27-27: Cascade with ShadowImage S-VOL (P-VOL : S-VOL = 1 : 1)

Figure 27-28 shows cascade configurations where the ShadowImage P-VOL to S-VOL ratio is 1-to-3.




Figure 27-28: Cascade with ShadowImage S-VOL (P-VOL: S-VOL = 1:3)

Configurations with ShadowImage P-VOLs and S-VOLs

This sections shows configurations in which TrueCopy is cascaded with ShadowImage P-VOLs and S-VOLs.

The lower pair in [Figure 27-29](#) shows both ShadowImage P-VOL and S-VOL cascaded with the TrueCopy P-VOL; on the right-side one of the pairs is reversed due to a pair swap.




Figure 27-29: Swapped pair configuration

Figure 27-30 shows multiple cascade volumes. The right side configuration shows pairs that have been swapped.




Figure 27-30: Multiple swapped pair configuration

Cascading a TrueCopy P-VOL with a ShadowImage P-VOL

If by using TrueCopy, the P-VOL on ShadowImage is cascaded, when restores using ShadowImage is executed, TrueCopy must be in the Split status. If restore using ShadowImage is executed in the Synchronizing status or Paired status of TrueCopy, the data in the volumes for P-VOL that are cascaded using TrueCopy on the local side and the remote side cannot be assured of equality. See [Figure 27-31](#).




Figure 27-31: Cascading a TrueCopy P-VOL with a ShadowImage P-VOL (P-VOL: S-VOL=1: 1)



NOTE: When both of TrueCopy and ShadowImage pairs are placed in the Paired status, performance of a host on the local side is lowered. It is recommended to make the status of a TrueCopy and ShadowImage pairs **Split** when host I/Os are instructed frequently.

Volume shared with P-VOL on ShadowImage and P-VOL on TrueCopy

Table 27-7 shows whether a read/write from/to a P-VOL of ShadowImage on the local side is possible when a P-VOL of ShadowImage and a P-VOL of TrueCopy are the same volume.

Table 27-7: Read / Write instructions to a ShadowImage P-VOL on the local side

TrueCopy P-VOL		ShadowImage P-VOL						
		Paired (including Paired Internally Synchronizing)	Synchroni zing	Synchronizing (Restore)		Split (including Split Pending)	Failure	Failure (Restore)
				R / W	W			
Paired		R/W	R/W	NO	NO	R/W	R/W	NO
Synchronizing		R/W	R/W	NO	NO	R/W	R/W	NO
Split	R / W	R/W	R/W	R/W	W	R/W	R/W	Δ ~~R/W~~
Failure	R / W	R/W	R/W	Δ R/W	Δ W	R/W	Δ R/W	Δ ~~R/W~~
	R	R	R	NO	NO	R	Δ R	NO
	~~R/W~~	~~R/W~~	~~R/W~~	NO	NO	~~R/W~~	Δ ~~R/W~~	NO

R/W: Read/Write by a host is possible.

R: Read by a host is possible but write is not supported.

NO means an unsupported case

Δ indicates a case where a pair operation causes an error (a case that can occur as a result of a change of the pair status to Failure)

~~R/W~~: Read/Write by a host is not supported

W: Write by a host is possible but read is not supported



NOTE: *Failure* in this table excludes any condition where volume access is not possible (for example, volume blockage).

When one P-VOL configures a pair with one or more S-VOLs, decide which item is applied as the pair status of the P-VOL of the above-mentioned ShadowImage with the following procedure.

1. If all the pairs that the P-VOL concerned configures are in the *Split* status, the item of *Split* is applied.
2. If all the pairs that the P-VOL concerned configures are in the *Split* status or the *Failure* status, the item of *Split* is applied. However, when including the pair that became *Failure* during restore, the items of *Failure* (Restore) are applied.

3. If a pair in the *Paired* status, the *Synchronizing* status, or the *Reverse Synchronizing* status is included in the pair that the P-VOL concerned configures, the item of *Paired*, *Synchronizing*, and *Reverse Synchronizing* is applied, respectively. (Two or more pairs in the *Paired* status, the *Synchronizing* status, and the *Reverse Synchronizing* status are never included in the pair that the P-VOL concerned configures.)
4. When multiple *Paired* statuses and *Synchronizing* status exist in the pairs that the relevant P-VOL configures, if the respective statuses are all *Readable*, they are *Readable*. Moreover, if the respective statuses are all *Writable*, they are *Writable*.

Pair Operation restrictions for cascading TrueCopy/ShadowImage

Table [Table 27-8](#) to [Table 27-9](#) shows pair status and operation when cascading TrueCopy with ShadowImage. The shaded areas in the tables indicate unworkable combinations.

Table 27-8: TrueCopy Pair Operation when volume shared with P-VOL on TrueCopy and ShadowImage

TrueCopy operation	ShadowImage pair status					
	Paired including Paired internally Synchronizing	Synchronizing	Synchronizing (Restore)	Split including Split Pending	Failure	Failure Restore
Creating pairs	YES	YES	NO	YES	YES	NO
Splitting pairs	YES	YES		YES	YES	
Re-synchronizing pairs	YES	YES	NO	YES	YES	NO
Swapping pairs	YES	YES	NO	YES	YES	NO
Deleting pairs	YES	YES	YES	YES	YES	YES

YES indicates a possible case NO indicates an unsupported case

Table 27-9: ShadowImage pair operation when volume shared with P-VOL on TrueCopy and ShadowImage

ShadowImage Operation	True Copy Pair Status			
	Paired	Synchronizing	Split	Failure
Creating pairs	YES	YES	YES	YES
Splitting pairs	YES	YES	YES	YES
Re-synchronizing pairs	YES	YES	YES	YES
Restoring pairs	NO	NO	YES	NO
Deleting pairs	YES	YES	YES	YES

Cascading a TrueCopy S-VOL with a ShadowImage P-VOL

If by using TrueCopy, the P-VOL on ShadowImage is cascaded, when restores using ShadowImage is executed, TrueCopy must be in the Split status. If restore using ShadowImage is executed in the Synchronizing status or Paired status of TrueCopy, the data in the volumes for P-VOL that are cascaded using TrueCopy on the local side and the remote side cannot be assured of equality. See [Figure 27-32](#).




Figure 27-32: Cascading a TrueCopy S-VOL with a ShadowImage P-VOL



NOTE: When both of TrueCopy and ShadowImage pairs are placed in the Paired status, performance of a host on the local side is lowered. It is recommended to make the status of a TrueCopy and ShadowImage pairs **Split** when host I/Os are instructed frequently.

Volume shared with P-VOL on ShadowImage and S-VOL on TrueCopy

Table 27-10 shows whether or not a read/write from/to a P-VOL of ShadowImage on the remote side is possible in the case where a P-VOL of ShadowImage and a S-VOL of TrueCopy are the same volume.

Table 27-10: Read / Write instructions to a ShadowImage P-VOL on the remote side

TrueCopy S-VOL		ShadowImage P-VOL						
		Paired (including Paired Internally Synchronizing)	Synchroni zing	Synchronizing (Restore)		Split (including Split Pending)	Failure	Failure (Restore)
				R / W	W			
Paired		R	R	NO	NO	R	R	NO
Synchronizing		R	R	NO	NO	R	R	NO
Split	R / W	R/W	R/W	R/W	W	R/W	R/W	Δ R/W
	R	R	R	NO	NO	R	R	NO
Failure		R	R	NO	NO	R	Δ R	NO

R/W: Read/Write by a host is possible.

R: Read by a host is possible but write is not supported.

W: Write by a host is possible but read is not supported

NO means an unsupported case

Δ indicates a case where a pair operation causes an error (a case that can occur as a result of a change of the pair status to Failure)

~~R/W~~: Read/Write by a host is not supported



NOTE: *Failure* in this table excludes any condition where volume access is not possible (for example, volume blockage).

When one P-VOL configures a pair with one or more S-VOLs, decide which item is applied as the pair status of the P-VOL of the above-mentioned ShadowImage with the following procedure.

1. If all the pairs that the P-VOL concerned configures are in the *Split* status, the item of *Split* is applied.
2. If all the pairs that the P-VOL concerned configures are in the *Split* status or the *Failure* status, the item of *Split* is applied. However, when including the pair that became *Failure* during restore, the items of *Failure* (Restore) are applied.

3. If a pair in the *Paired* status, the *Synchronizing* status, or the *Reverse Synchronizing* status is included in the pair that the P-VOL concerned configures, the item of *Paired*, *Synchronizing*, and *Reverse Synchronizing* is applied, respectively. (Two or more pairs in the *Paired* status, the *Synchronizing* status, and the *Reverse Synchronizing* status are never included in the pair that the P-VOL concerned configures.)
4. When multiple *Paired* statuses and *Synchronizing* status exist in the pairs that the relevant P-VOL configures, if the respective statuses are all Readable, they are Readable. Moreover, if the respective statuses are all Writable, they are Writable.

Volume shared with TrueCopy S-VOL and ShadowImage P-VOL

Table [Table 27-11](#) to [Table 27-12](#) shows pair status and operation when cascading TrueCopy with ShadowImage. The shaded areas in the tables indicate unworkable combinations.

Table 27-11: TrueCopy Pair Operation when volume shared with S-VOL on TrueCopy and P-VOL on ShadowImage

TrueCopy Operation	ShadowImage pair status					
	Paired (including Paired internally Synchronizing)	Synchronizing	Synchronizing (Restore)	Split (including Split Pending)	Failure	Failure (Restore)
Creating pairs	YES	YES	NO	YES	YES	NO
Splitting pairs	YES	YES		YES	YES	
Re-synchronizing pairs	YES	YES	NO	YES	YES	NO
Swapping pairs	YES	YES	NO	YES	YES	NO
Deleting pairs	YES	YES	YES	YES	YES	YES

YES indicates a possible case
NO indicates an unsupported case

Table 27-12: ShadowImage pair operation when volume shared with S-VOL on TrueCopy and P-VOL on ShadowImage

ShadowImage operation	True Copy pair status			
	Paired	Synchronizing	Split	Failure
Creating pairs	YES	YES	YES	YES
Splitting pairs	YES	YES	YES	YES
Re-synchronizing pairs	YES	YES	YES	YES
Restoring pairs	NO	NO	YES*	NO
Deleting pairs	YES	YES	YES	YES
*When the S-VOL attribute is Read Only by pair splitting, cannot be restored.				

Cascading a TrueCopy P-VOL with a ShadowImage S-VOL

Cascade of an volume of TrueCopy with a ShadowImage S-VOL is supported only when the ShadowImage S-VOL and a TrueCopy P-VOL are the same volume. Also, operations of the ShadowImage and TrueCopy pairs are restricted depending on statuses of the pairs.

When cascading volumes of TrueCopy with a ShadowImage S-VOL, create a ShadowImage pair first. When a TrueCopy pair is created earlier, split the TrueCopy pair once and create a pair using ShadowImage.

When changing the status of a ShadowImage pair, the status of a TrueCopy pair must be Split or Failure. When changing the status of a TrueCopy pair, the status of a ShadowImage pair must be Split.

See [Figure 27-33 on page 27-47](#).




Figure 27-33: Cascading a TrueCopy P-VOL with a ShadowImage S-VOL



NOTE: A cascade of a ShadowImage S-VOL is used when making a backup on the remote side asynchronously. Because a backup is made from an S-VOL of ShadowImage to the remote side in this configuration, lowering of performance on the local side (a P-VOL of ShadowImage) during the backing up can be minimized. Note that a TrueCopy pair must be placed in the Split status when re-synchronizing a volume of ShadowImage on the local side.

Volume shared with S-VOL on ShadowImage and P-VOL on TrueCopy

Table 27-13 shows whether or not a read/write from/to a P-VOL of ShadowImage on the remote side is possible in the case where a S-VOL of ShadowImage and a P-VOL of TrueCopy are the same volume.

Table 27-13: Read / Write instructions to an S-VOL of ShadowImage on the local side

TrueCopy P-VOL		Shadow Image S-VOL							
		Paired (including Paired Internally Synchronizing)	Synchronizing	Synchronizing (Restore)		Split	Split Pending	Failure	Failure (Restore)
Paired		NO	NO	NO		R/W	R/W	NO	NO
Synchronizing		NO	NO	NO		R/W	R/W	NO	NO
Split	R/W	R	R	R		R/W	R/W	Δ R	Δ ~~R/W~~
Failure	R/W	R	R	R		R/W	R/W	Δ R	Δ ~~R/W~~
	R	R	R	R		R	R	Δ R	Δ ~~R/W~~
	~~R/W~~	~~R/W~~	~~R/W~~	~~R/W~~	~~R/W~~	~~R/W~~	~~R/W~~	Δ ~~R/W~~	Δ ~~R/W~~

R/W: Read/Write by a host is possible.

R: Read by a host is possible but write is not supported.

NO means an unsupported case

Δ indicates a case where a pair operation causes an error (a case that can occur as a result of a change of the pair status to Failure)

~~R/W~~: Read/Write by a host is not supported



NOTE: *Failure* in this table excludes any condition where volume access is not possible (for example, volume blockage).

Volume Shared with TrueCopy P-VOL and ShadowImage S-VOL

Table [Table 27-14](#) to [Table 27-15](#) shows pair status and operation when cascading TrueCopy with ShadowImage. The shaded areas in the tables indicate unworkable combinations.

Table 27-14: TrueCopy pair operation when volume shared with P-VOL on TrueCopy and S-VOL on ShadowImage

TrueCopy operation	ShadowImage pair status						
	Paired (including Paired internally Synchronizing)	Synchronizing	Synchronizing (Restore)	Split	Split Pending	Failure	Failure (Restore)
Creating pairs	NO	NO	NO	YES	YES	NO	NO
Splitting pairs				YES	YES		
Re-synchronizing pairs	NO	NO	NO	YES	YES	NO	NO
Restoring pairs	NO	NO	NO	YES	NO	NO	NO
Deleting pairs	YES	YES	YES	YES	YES	YES	YES

YES indicates a possible case
NO indicates an unsupported case

Table 27-15: ShadowImage pair operation when volume shared with P-VOL on TrueCopy and S-VOL on ShadowImage

ShadowImage operation	True Copy pair status			
	Paired	Synchronizing	Split	Failure
Creating pairs	NO	NO	NO	NO
Splitting pairs			YES	YES
Re-synchronizing pairs	NO	NO	YES	YES
Restoring pairs	NO	NO	YES	YES
Deleting pairs	YES	YES	YES	YES

Volume Shared with S-VOL on TrueCopy and ShadowImage

Table [Table 27-16](#) to [Table 27-17](#) shows pair status and operation when cascading TrueCopy with ShadowImage.

Table 27-16: TrueCopy pair operation when volume shared with S-VOL on TrueCopy and S-VOL on ShadowImage

TrueCopy operation	ShadowImage pair status						
	Paired including Paired internally Synchronizing	Synchronizing	Synchronizing (Restore)	Split	Split Pending	Failure	Failure (Restore)
Creating pairs	NO	NO	NO	NO	NO	NO	NO
Splitting pairs				YES		YES	
Re-synchronizing pairs				YES		NO	
Swapping pairs				YES		NO	
Deleting pairs				YES		YES	

Table 27-17: ShadowImage pair operation when volume shared with S-VOL on TrueCopy and S-VOL on ShadowImage

ShadowImage operation	True Copy pair status			
	Paired	Synchronizing	Split	Failure
Creating pairs	NO	NO	NO	NO
Splitting pairs			YES	
Re-synchronizing pairs	NO	NO	NO	NO
Restoring pairs	NO	NO	NO	NO
Deleting pairs	YES	YES	YES	YES

YES indicates a possible case
NO indicates an unsupported case

Cascading TrueCopy with ShadowImage P-VOL and S-VOL 1:1

Cascade with a ShadowImage P-VOL (P-VOL: S-VOL=1:1)




Figure 27-34: Cascade with a ShadowImage P-VOL

Cascade with a ShadowImage S-VOL (P-VOL: S-VOL=1:1)




Figure 27-35: Cascade with a ShadowImage S-VOL

Simultaneous cascading of TrueCopy with ShadowImage

Simultaneous cascade with a P-VOL and an S-VOL of ShadowImage (P-VOL: S-VOL=1:1)




Figure 27-36: Cascade with a P-VOL and a S-VOL of ShadowImage

Cascading TrueCopy with ShadowImage P-VOL and S-VOL 1:3




Figure 27-37: Cascade with a ShadowImage P-VOL

Cascade with a ShadowImage S-VOL (P-VOL: S-VOL=1:3)




Figure 27-38: Cascade with a ShadowImage S-VOL

Simultaneous cascading of TrueCopy with ShadowImage

Cascading with a ShadowImage P-VOL and S-VOL (P-VOL: S-VOL=1:3)




Figure 27-39: Cascading with a P-VOL and a S-VOL of ShadowImage

Swapping when cascading TrueCopy and ShadowImage Pairs

In a cascade configuration of TrueCopy and ShadowImage pairs, P-VOL data can be restored through the use of backup data on the remote side (a ShadowImage S-VOL on the remote side). To restore data using backup data on the remote side, it is required to perform swapping.

The following task description explains the procedure for the swapping, using an example of a cascade configuration of TrueCopy and ShadowImage pairs.

Suppose that the system is usually operated in the cascade configuration of TrueCopy and ShadowImage pairs (Figure 27-40). At a certain time, entire data on the local side became invalid because a failure occurred in the local array (Figure 27-41). In this situation, a swapping must be performed in order to restore data using the backup data on the remote side (remote array).




Figure 27-40: Cascade configuration of TrueCopy and ShadowImage




Figure 27-41: Cascade configuration of TrueCopy and ShadowImage

To perform a swap:

1. Perform restoration from an S-VOL of ShadowImage on the remote side to a P-VOL.
2. Split the ShadowImage pair on the remote side after completing the restoration.
3. Split the ShadowImage pair on the local side.
4. Perform a swapping for a TrueCopy pair that straddles the local and remote array.
5. Split the TrueCopy pair after completing the swapping.
6. Perform a swapping again for a TrueCopy pair that straddles the local and remote array.
7. Split the TrueCopy pair after completing the swapping.
8. Perform restoration of a ShadowImage pair on the local side. Here, the host I/O can be resumed.
9. Return to the usual operation after completing the restoration.

Creating a backup with ShadowImage**To back up TrueCopy volume using ShadowImage**

1. Split the TrueCopy pair. At this time, remaining data buffers on the host are flushed to the P-VOL, then copied to the S-VOL. This ensures a complete and consistent copy.
2. Create the ShadowImage pair.
3. Split the ShadowImage pair.
4. Re-synchronize the TrueCopy pair.

Figure 27-42 shows the backup process.




Figure 27-42: Backup operations

For disaster recovery using a ShadowImage backup, see [Resynchronizing the pair on page 15-20](#).

Cascading with Snapshot

In a cascaded system, the TrueCopy P-VOL or S-VOL is shared with a Snapshot P-VOL or S-VOL. Cascading is usually done to provide a backup that can be used if the volumes are damaged or have inconsistent data. This appendix discusses the configurations, operations, and statuses that are allowed in the cascaded system.

Cascade overview

Snapshot is cascaded with TrueCopy to:

- Make a backup of the TrueCopy S-VOL on the remote side
- Pair the Snapshot V-VOL on the local side with the TrueCopy S-VOL. This results in an asynchronous TrueCopy pair.
- Provide any other traditional use for Snapshot

While a Snapshot V-VOL is smaller than a ShadowImage S-VOL would be, performance when cascading with Snapshot is lower than it would be by cascading with ShadowImage.

This section provides the following:

- Supported cascade configurations
- TrueCopy and Snapshot operations allowed
- The combined TrueCopy and Snapshot statuses allowed
- The combined statuses that allow read/write
- Best Practices

Cascade configurations

Cascade configurations can consist of P-VOLs and S-VOLs, in both TrueCopy and Snapshot.

The following sections show supported configurations.



NOTE:

1. Cascading TrueCopy with another TrueCopy system or TrueCopy Extended Distance is not supported.
2. When a restore is done on ShadowImage, the TrueCopy pair must be split.

In Configurations 2 and 4, the Snapshot cascade backs up data on the remote side and manages generation on the remote side.

Configurations with Snapshot P-VOLs

The Snapshot P-VOL can be cascaded with the TrueCopy P-VOL or S-VOL. [Figure 27-43](#) and [Figure 27-44 on page 27-61](#) shows supported configurations.




Figure 27-43: Cascade with Snapshot P-VOL




Figure 27-44: Cascade with Snapshot P-VOL

Cascading with a Snapshot V-VOL

The Snapshot V-VOL can be cascaded with the TrueCopy P-VOL only. Only one V-VOL in a Snapshot pair may be cascaded.

Figure 27-45 shows supported configurations.




Figure 27-45: Examples of cascade with Snapshot V-VOL

Cascading a TrueCopy P-VOL with a Snapshot P-VOL

When restore using Snapshot is executed, TrueCopy must be in the Split status. If restore using Snapshot is executed in the Synchronizing status or Paired status of TrueCopy, the data in the volumes for P-VOL that are cascaded using TrueCopy on the local side and the remote side cannot be assured of equality.




Figure 27-46: Cascade connection of TrueCopy P-VOL with Snapshot P-VOL

Volume shared with P-VOL on Snapshot and P-VOL on TrueCopy

Table 27-18 shows whether a read/write from/to a Snapshot P-VOL on the local side is possible or not in the case where a Snapshot P-VOL and a TrueCopy P-VOL are the same volume.

Table 27-18: A Read / Write instruction to a Snapshot P-VOL on the local side

TrueCopy P-VOL		Snapshot P-VOL				
		Paired	Synchronizing (Restore)	Split	Failure	Failure (Restore)
Paired		R/W	NO	R/W	R/W	NO
Synchronizing		R/W	NO	R/W	R/W	NO
Split	R / W	R/W	R/W	R/W	R/W	Δ ~~R/W~~
Failure	R / W	R/W	Δ R/W	R/W	Δ R/W	Δ ~~R/W~~
	R	R	NO	R	Δ R	NO
	~~R/W~~	~~R/W~~	NO	~~R/W~~	Δ ~~R/W~~	NO

R/W: Read/Write by a host is possible.

R: Read by a host is possible but write is not supported.

NO means an unsupported case

Δ indicates a case where a pair operation causes an error (a case that can occur as a result of a change of the pair status to Failure)

~~R/W~~: Read/Write by a host is not supported



NOTE: *Failure* in this table excludes any condition where volume access is not possible (for example, volume blockage).

V-VOLs number of Snapshot

V-VOLs of up to 1,024 generations can be made even when the Snapshot P-VOL is cascaded with the TrueCopy P-VOL in the same way as when no cascade connection is made.

Table 27-19: TrueCopy pair operation when volume shared with P-VOL on TrueCopy and P-VOL on Snapshot

TrueCopy operation	Snapshot pair status				
	Paired	Reverse Synchronizing	Split	Failure	Failure (Restore)
Creating pairs	YES	NO	YES	YES	NO
Splitting pairs	YES	NO	YES	YES	NO
Re-synchronizing pairs	YES	NO	YES	YES	NO
Restoring pairs	YES	NO	YES	YES	NO
Deleting pairs	YES	YES	YES	YES	YES

YES indicates a possible case

NO indicates an unsupported case

Table 27-20: Snapshot pair operation when volume shared with P-VOL on TrueCopy and P-VOL on Snapshot

Snapshot operation	True Copy pair status			
	Paired	Synchronizing	Split	Failure
Creating pairs	YES	YES	YES	YES
Splitting pairs	YES	YES	YES	YES
Re-synchronizing pairs	YES	YES	YES	YES
Restoring pairs	NO	NO	YES	NO
Deleting pairs	YES	YES	YES	YES

Cascading a TrueCopy S-VOL with a Snapshot P-VOL

If a TrueCopy pair is placed in the Paired status when a Snapshot pair is in the Split status, the performance of a host on the local side is lowered more. It is recommended to make the status of a TrueCopy pair Split in a period of time when host I/Os are instructed frequently.




Figure 27-47: Cascading a TrueCopy S-VOL with a Snapshot P-VOL

Volume Shared with Snapshot P-VOL and TrueCopy S-VOL

Table 27-21 shows whether a read/write from/to a Snapshot P-VOL on the remote side is possible or not in the case where a Snapshot P-VOL and a TrueCopy S-VOL are the same volume.

Table 27-21: A Read / Write instruction to a Snapshot P-VOL on the local side

TrueCopy P-VOL		Snapshot P-VOL				
		Paired	Synchronizing (Restore)	Split	Failure	Failure (Restore)
Paired		R	NO	R	R	NO
Synchronizing		R	NO	R	R	NO
Split	R/W	R/W	R/W	R/W	R/W	Δ R/W
	R	R	NO	R	R	NO
Failure		R	NO	R	R	NO

R/W: Read/Write by a host is possible.

R: Read by a host is possible but write is not supported.

NO means an unsupported case

Δ indicates a case where a pair operation causes an error (a case that can occur as a result of a change of the pair status to Failure)

~~R/W~~: Read/Write by a host is not supported



NOTE: *Failure* in this table excludes any condition where volume access is not possible (for example, volume blockage).

V-VOLs number of Snapshot

V-VOLs of up to 1,024 generations can be made even in the case where the Snapshot P-VOL is cascaded with the TrueCopy S-VOL in the same way as in the case where no cascade connection is made.

Table 27-22: TrueCopy Pair Operation when Volume Shared with TrueCopy S-VOL and Snapshot P-VOL

TrueCopy operation	Snapshot pair status				
	Paired	Reverse Synchronizing	Split	Failure	Failure (Restore)
Creating pairs	YES	NO	YES	YES	NO
Splitting pairs	YES	NO	YES	YES	NO
Re-synchronizing pairs	YES	NO	YES	YES	NO
Swapping pairs	YES	NO	YES	YES	NO
Deleting pairs	YES	YES	YES	YES	YES

YES indicates a possible case

NO indicates an unsupported case

Table 27-23: Snapshot pair operation when volume shared with TrueCopy S-VOL and Snapshot P-VOL

Snapshot operation	True Copy pair status				
	Paired	Synchronizing	Split	Failure	Take-over
Creating pairs	YES	YES	NO	YES	YES
Splitting pairs	YES	YES	YES	YES	YES
Re-synchronizing pairs	YES	YES	YES	YES	YES
Restoring pairs	NO	NO	YES*	YES	NO
Deleting pairs	YES	YES	YES	YES	YES
* When the S-VOL attribute is Read Only by pair splitting, cannot be restored.					

YES indicates a possible case

NO indicates an unsupported case

Cascading a TrueCopy P-VOL with a Snapshot V-VOL

A cascade of a Snapshot V-VOL is used when making a backup on the remote side asynchronously. Though the cascade of Snapshot can decrease an S-VOL (V-VOL) capacity differently from a cascade of ShadowImage, the performance on the local side (a P-VOL of Snapshot) is affected by the backup. Snapshot can make two or more V-VOLs for a P-VOL, however, only a single V-VOL can be cascaded with TrueCopy. Note that a TrueCopy pair must be placed in the Split status when re syncing (giving the Snapshot instruction) a volume of Snapshot on the local side.




Figure 27-48: Cascading a TrueCopy S-VOL with a Snapshot P-VOL)

Table 27-24: TrueCopy pair operation when volume shared with TrueCopy P-VOL and Snapshot V-VOL

TrueCopy operation	Snapshot pair status				
	Paired	Reverse Synchronizing	Split	Failure	Failure (Restore)
Creating pairs	NO	NO	YES	NO	NO
Splitting pairs			YES		
Re-synchronizing pairs	NO	NO	YES	NO	NO
Swapping pairs	NO	NO	YES	NO	NO
Deleting pairs	YES	YES	YES	YES	YES

YES indicates a possible case NO indicates an unsupported case

Table 27-25: Snapshot pair operation when volume shared with TrueCopy P-VOL and Snapshot V-VOL

Snapshot operation	Truecopy pair status			
	Paired	Reverse Synchronizing	Split	Failure
Creating pairs	NO	NO	NO	NO
Splitting pairs			YES	YES
Re-synchronizing pairs	NO	NO	YES	YES
Swapping pairs	NO	NO	YES	YES
Deleting pairs	YES	YES	YES	YES

YES indicates a possible case
NO indicates an unsupported case

Transition of statuses of TrueCopy and Snapshot pairs

A cascade of an volume of TrueCopy with a Snapshot V-VOL is supported only when the Snapshot V-VOL and a TrueCopy P-VOL are the same volume. Also, operations of the Snapshot and TrueCopy pairs are restricted depending on statuses of the pairs.

When cascading volumes of TrueCopy with a Snapshot V-VOL, create a Snapshot pair first. When a TrueCopy pair is created earlier, split the TrueCopy pair once and create a pair using Snapshot

When changing a status of a Snapshot pair, a status of a TrueCopy pair must be Split or Failure. When changing a status of a TrueCopy pair, a status of a Snapshot pair must be Split.

Table 27-26 shows whether a read/write from/to a Snapshot V-VOL on the local side is possible or not in the case where a Snapshot V-VOL and a TrueCopy P-VOL are the same volume.

Table 27-26: A Read / Write instruction to a Snapshot V-VOL on the Local Side

TrueCopy P-VOL		Snapshot S-VOL				
		Paired	Synchronizing (Restore)	Split	Failure	Failure (Restore)
Paired		NO	NO	R/W	NO	NO
Synchronizing		NO	NO	R/W	NO	NO
Split	R/W	~~R/W~~	~~R/W~~	R/W	Δ , ~~R/W~~	Δ , ~~R/W~~
	R	~~R/W~~	~~R/W~~	R	Δ , ~~R/W~~	Δ , ~~R/W~~
Failure	R/W	~~R/W~~	~~R/W~~	R/W	Δ , ~~R/W~~	Δ , ~~R/W~~
	R	~~R/W~~	~~R/W~~	R	Δ , ~~R/W~~	Δ , ~~R/W~~
	~~R/W~~	~~R/W~~	~~R/W~~	~~R/W~~	Δ , ~~R/W~~	Δ , ~~R/W~~

R/W: Read/Write by a host is possible.

R: Read by a host is possible but write is unsupported .

NO indicates an unsupported case

Δ indicates a case where a pair operation causes an error (a case that can occur as a result of a change of the pair status to Failure)

~~R/W~~: Read/Write by a host is unsupported .



NOTE: *Failure* in this table excludes any condition where volume access is not possible (for example, volume blockage).

Swapping when cascading a TrueCopy pair and a Snapshot pair

In a cascade configuration of TrueCopy and Snapshot pairs, P-VOL data can be restored through the use of backup data on the remote side (a Snapshot S-VOL on the remote side). To restore data using backup data on the remote side, it is required to perform swapping.

The following task description explains the procedure for the swapping, using an example of a cascade configuration of TrueCopy and Snapshot pairs.

Suppose that the system is usually operated in the cascade configuration of TrueCopy and Snapshot pairs (Figure 27-49). At a certain time, entire data on the local side became invalid because a failure occurred in the local array (Figure 27-50). In this situation, a swapping must be performed in order to restore data using the backup data on the remote side (remote array).




Figure 27-49: Cascade configuration of TrueCopy and Snapshot




Figure 27-50: Cascade configuration of TrueCopy and Snapshot

To perform a swap:

1. Perform restoration from an S-VOL of Snapshot on the remote side to a P-VOL.
2. Split the Snapshot pair on the remote side after completing the restoration.
3. Split the Snapshot pair on the local side.
4. Perform a swapping for a TrueCopy pair that straddles the local and remote array.
5. Split the TrueCopy pair after completing the swapping.
6. Perform a swapping again for a TrueCopy pair that straddles the local and remote array.
7. Split the TrueCopy pair after completing the swapping.
8. Perform restoration of a Snapshot pair on the local side. Here, the host I/O can be resumed.
9. Return to the usual operation after completing the restoration.

Table 27-27: TrueCopy pair operation when volume shared with TrueCopy S-VOL and Snapshot V-VOL

TrueCopy operation	Snapshot pair status				
	Paired	Reverse Syn-chronizing	Split	Failure	Failure (Restore)
Creating pairs	NO	NO	NO	NO	NO
Splitting pairs			YES		
Re-synchronizing pairs	NO	NO	YES	NO	NO
Swapping pairs	NO	NO	YES	NO	NO
Deleting pairs	NO	NO	YES	YES	NO

YES indicates a possible case
NO indicates an unsupported case

Table 27-28: Snapshot pair operation when volume shared with TrueCopy S-VOL and Snapshot V-VOL

Snapshot operation	True Copy pair status				
	Paired	Synchronizing	Split	Failure	Take-over
Creating pairs	NO	NO	NO	NO	NO
Splitting pairs			YES	YES	NO
Re-synchronizing pairs	NO	NO	NO	NO	NO
Restoring pairs	NO	NO	NO	NO	NO
Deleting pairs	YES	YES	YES	YES	YES

YES indicates a possible case
NO indicates an unsupported case

Creating a backup with Snapshot

To back up TrueCopy volume using Snapshot

1. Split the TrueCopy pair. At this time, remaining data buffers on the host are flushed to the P-VOL, then copied to the S-VOL. This ensures a complete and consistent copy.
2. Create the Snapshot pair.
3. Split the Snapshot pair.
4. Re-synchronize the TrueCopy pair.

Figure 27-51 shows the backup process.




Figure 27-51: Backup operations

For disaster recovery using a Snapshot backup, see [Resynchronizing the pair on page 15-20](#).

When to create a backup

When a pair is synchronizing, each write to the P-VOL is also sent to the remote S-VOL. The host does not send a write until confirmation is received for the previous write is copied to the S-VOL. This latency impacts host I/O. The best time to synchronize a pair is during off-peak hours. [Figure 27-52](#) illustrates host I/O when a pair is split and when a pair is synchronizing.




Figure 27-52: I/O performance impact during resynchronization

Cascading with ShadowImage and Snapshot

TrueCopy can cascade ShadowImage and Snapshot. When ShadowImage pair status is any of **Paired**, **Paired Internally Synchronizing**, **Synchronizing**, and **Split Pending**, and at the same time, Snapshot pair status is **Split**, and at the same time, TrueCopy pair status is **Paired**, the host I/O performance for the P-VOL deteriorates. Use ShadowImage or TrueCopy in the **Split** status and, if needed, resynchronize the pair and acquire the backup.

Cascade restrictions of TrueCopy with Snapshot and ShadowImage

TrueCopy can cascade Snapshot and ShadowImage at the same time. However, since the performance may deteriorate, start the operation after advance verification.

Cascade restrictions of TrueCopy S-VOL with Snapshot V-VOL

In the configuration in which the P-VOL of ShadowImage and the P-VOL of Snapshot are cascaded as shown below, and also, in the configuration in which the V-VOL of Snapshot and the S-VOL of TrueCopy are cascaded, when the TrueCopy pair status is **Paired** or **Synchronizing**, ShadowImage cannot be restored. Change the TrueCopy pair status to **Split**, and then execute it again.




Figure 27-53: Cascade restrictions of TrueCopy S-VOL with Snapshot V-VOL

Cascading restrictions

Figure 27-54 shows TrueCopy cascade connections that are not available.




Figure 27-54: Cascade connection restrictions

Concurrent use of TrueCopy and ShadowImage or Snapshot

By using TrueCopy and ShadowImage concurrently, the volume in the array that the TrueCopy function used is duplicated. Even when the TrueCopy function is in progress, the host I/O operation to the volume is guaranteed. In addition, it is possible to replace ShadowImage with Snapshot. When ShadowImage is replaced with Snapshot, the S-VOL capacity can be decreased but the performance is lowered. Therefore, make a selection accordingly.

Cascading TCE

TCE P-VOLs and S-VOLs can be cascaded with Snapshot P-VOLs. This section discusses the supported configurations, operations, and statuses.

Cascading with Snapshot

V-VOLs number of Snapshot

V-VOLs of up to 1,024 generations can be made even in the case where the Snapshot

P-VOL is cascaded with the TCE P-VOL and TCE S-VOL in the same way as in the case where no cascade connection is made.

DP pool

In case of cascading Snapshot P-VOL with TCE P-VOL or TCE S-VOL, the DP pool that the Snapshot pair uses and the one that the TCE pair uses must be the same. That is, if a Snapshot pair is cascaded to a TCE P-VOL, the DP pool number specified when creating the Snapshot pair must be the same as the one specified for the local array when creating the TCE pair. If a Snapshot pair is cascaded to a TCE S-VOL, the DP pool number specified when creating the Snapshot pair must be the same as the one specified for the remote array when creating the TCE pair.

Cascading a TCE P-VOL with a Snapshot P-VOL

When combining the backup operation by Snapshot in the local array and the backup to the remote array by TCE, cascade the P-VOLs of TCE and Snapshot. as shown in [Figure 27-55](#).




Figure 27-55: Cascading a TCE P-VOL with a Snapshot P-VOL

When cascading the P-VOLs of TCE and Snapshot, the pair operation has the following restrictions according to each pair status of TCE and Snapshot. [Table 27-30 on page 27-80](#) shows execution conditions of the TCE pair operation and [Table 27-31 on page 27-81](#) shows those of the Snapshot pair operation.

As for the volume (P-VOL of the local side Snapshot) shared between TCE and Snapshot, the availability of Read/Write is decided by combination of the statuses of the TCE pair and Snapshot pair. [Table 27-29 on page 27-80](#) shows the availability of Read/Write for the P-VOL of the local side Snapshot.

The restoration of the Snapshot pair cascaded with the TCE P-VOL can be done only when the status of the TCE pair is Simplex, Split, or Pool Full.



NOTE: When the target volume of the TCE pair is the Snapshot P-VOL and the Snapshot pair status becomes Reverse Synchronizing or the Snapshot pair status becomes Failure during restore, you cannot execute pair creation or pair resynchronization of TCE. Therefore, it is required to recover the Snapshot pair.

Table 27-29: Read / Write instructions to a Snapshot P-VOL on the local side

TCE P-VOL	Snapshot P-VOL Status					
	Paired	Synchronizing (Restore)	Split	Threshold over	Failure	Failure (Restore)
Paired	R/W	NO	R/W	R/W	R/W	NO
Synchronizing	R/W	NO	R/W	R/W	R/W	NO
Split	R/W	R/W	R/W	R/W	R/W	Δ , ~~R/W~~
Pool Full	R/W	R/W	R/W	R/W	R/W	Δ , ~~R/W~~
Failure	R/W	Δ R/W	R/W	R/W	Δ R/W	Δ , ~~R/W~~

R/W: Read/Write by a host is possible.

R: Read by a host is possible but write is not supported.

NO indicates an unsupported case

Δ indicates a case where a pair operation causes an error (a case that can occur as a result of a change of the pair status to Failure)

~~R/W~~: Read/Write by a host is not supported



NOTE: Failure in this table excludes a condition in which access of a volume is not possible (for example, volume blockage).

Table 27-30: TCE Pair Operation when Volume Shared with P-VOL on TCE and P-VOL on Snapshot

TCE operation	Snapshot-VOL status				
	Paired	Synchronizing (Restore)	Split	Failure	Failure (Restore)
Creating pairs	YES	NO	YES	YES	NO
Splitting pairs	YES	NO	YES	YES	NO
Re-synchronizing pairs	YES	NO	YES	YES	NO
Restoring pairs	YES	NO	YES	YES	NO
Deleting pairs	YES	YES	YES	YES	YES

YES indicates a possible case

NO indicates an unsupported case

Table 27-31: Snapshot pair operation when volume shared with P-VOL on TCE and P-VOL on Snapshot

Snapshot operation	TCE P-VOL status			
	Paired	Synchronizing	Split	Failure
Creating pairs	YES	YES	YES	YES
Splitting pairs	YES	YES	YES	YES
Re-synchronizing pairs	YES	YES	YES	YES
Restoring pairs	NO	NO	YES	NO
Deleting pairs	YES	YES	YES	YES

YES indicates a possible case
NO indicates an unsupported case

Cascading a TCE S-VOL with a Snapshot P-VOL

In the remote array, cascade the S-VOL of TCE and the P-VOL of Snapshot to remain the backup of the S-VOL of TCE. as shown in [Figure 27-56](#).




Figure 27-56: Cascading a TCE S-VOL with a Snapshot P-VOL

When cascading the S-VOL of TCE and the P-VOL of Snapshot, the pair operation has the following restrictions according to each pair status of TCE and Snapshot. [Table 27-33 on page 27-84](#) shows execution conditions of the TCE pair operation and [Table 27-34 on page 27-84](#) shows those of the Snapshot pair operation.

As for the volume (P-VOL of remote side Snapshot) shared between TCE and Snapshot, the availability of Read/Write is decided by combination of the status of the TCE pair and Snapshot pair. [Table 27-32 on page 27-83](#) shows the availability of Read/Write for the P-VOL of the remote side Snapshot.

When restoring the S-VOL of TCE and the cascaded Snapshot, it is required to change the TCE status to Simplex or Split. If the status is Takeover, you can execute the restore. However, the restore is not possible in the Busy status which is in the middle of the restoration processing of the S-VOL from the DP pool.

In the Busy status where the S-VOL of TCE is in the middle of the restoration processing from the DP pool, Read/Write from/to the S-VOL of TCE and the V-VOL of the cascaded Snapshot is not possible.



NOTE:

1. When the target volume of the TCE pair is the Snapshot P-VOL and the Snapshot pair status becomes Reverse Synchronizing or the Snapshot pair status becomes Failure during restore, you cannot execute pair creation or pair resynchronization of TCE. Therefore, it is required to recover the Snapshot pair
 2. If the restoration of the data from the DP pool fails due to failure occurrence while the TCE pair status is Busy, the Snapshot pair status becomes Failure. It does not recover unless you delete the TCE pair and Snapshot pair and create the pairs again.
 3. Failure in this table excludes a condition in which access of a volume is not possible (for example, volume blockage).
-

Table 27-32: A Read / Write instruction to a Snapshot P-VOL on the remote side

TCE S-VOL		Snapshot P-VOL					
		Paired	Synchronizing (Restore)	Split	Threshold over	Failure	Failure (Restore)
Paired		R	NO	R	R	R	NO
Synchronizing		R	NO	R	R	R	NO
Split	R / W	R/W	R/W	R/W	R/W	R/W	Δ , ~~R/W~~
	R	R	NO	R	R	R	NO
Inconsistent		Δ , ~~R/W~~	NO	Δ , ~~R/W~~	Δ , ~~R/W~~	Δ , ~~R/W~~	NO
Take Over		R/W	R/W	R/W	R/W	R/W	Δ , ~~R/W~~
Paired internally busy		R/W	NO	R/W	R/W	R/W	NO
Busy		R/W	NO	R/W	R/W	R/W	NO
Pool Full		R	NO	R	R	Δ , R	NO

R/W: Read/Write by a host is possible.

R: Read by a host is possible but write is not supported.

NO indicates an unsupported case

Δ indicates a case where a pair operation causes an error (a case that can occur as a result of a change of the pair status to Failure)

~~R/W~~: Read/Write by a host is not supported

Table 27-33: TCE pair operation when volume shared with S-VOL on TCE and P-VOL on Snapshot

TCE operation	Snapshot P-VOL status				
	Paired	Synchronizing (Restore)	Split	Failure	Failure (Restore)
Creating pairs	YES	NO	YES	YES	NO
Splitting pairs	YES	NO	YES	YES	NO
Re-synchronizing pairs	YES	NO	YES	YES	NO
Restoring pairs	YES	NO	YES	YES	NO
Deleting pairs	YES	YES	YES	YES	YES

YES indicates a possible case
NO indicates an unsupported case

Table 27-34: Snapshot pair operation when volume shared with S-VOL on TCE and P-VOL on Snapshot

Snapshot operation	TCE S-VOL Status							
	Paired	Synchronizing	Split RW mode	Split R mode	Inconsistent	Take over	Busy	Paired internal ly busy
Creating pairs	YES	YES	YES	YES	YES	YES	YES	YES
Splitting pairs	NO ¹	NO	YES	YES	NO	YES	YES	YES
Re-synchronizing pairs	YES	YES	YES	YES	NO	YES	YES	YES
Restoring pairs	NO	NO	YES ²	NO	NO	YES	NO	NO
Deleting pairs	YES	YES	YES	YES	YES	YES	YES	YES

Notes:

1. Split pair is available only when the conditions for execution described in "1. Issue a pair split to Snapshot pairs on the remote array using HSNM2 or CCI" in [Snapshot cascade configuration local and remote backup operations on page 27-85](#). are met.
2. When the S-VOL attribute is Read Only by pair splitting, cannot be restore

YES indicates a possible case
NO indicates an unsupported case

Snapshot cascade configuration local and remote backup operations

TCE provides feature functions for supporting backup operations. These feature functions are explained below.

Local snapshot creation function: Both asynchronous remote copy operations by TCE and snapshot operation by Snapshot can be performed together. Up to 1,024 snapshots of the P-VOL of a TCE pair can be created in the local array. When the host issues a command for creation of a snapshot of the P-VOL (1), the local array retains the P-VOL data at the time as a snapshot (2). See [Figure 27-57 on page 27-85](#).




Figure 27-57: Local Snapshot creation function

Remote snapshot creation command function: In a remote snapshot, the snapshot of application data in the remote array is obtained. Up to 1,024 snapshots per TCE pair can be acquired. This function enables remote backup using the asynchronous remote copy function that hitherto was an off-site backup operation in which a tape medium was physically transported. More reliable backup operation is achieved by automating the work.

There are two ways to create a remote snapshot:

1. Issue a pair split to Snapshot pairs on the remote array using HSNM2 or CCI

When the S-VOLs of TCE are cascaded with the P-VOLs of Snapshot and the TCE pairs and Snapshot pairs are all in Paired state, you can directly issue a pair split to Snapshot pairs on the remote array. That pair split will not be executed immediately but will be reserved for the Snapshot group (the Snapshot pair remains in Paired state) when TCE is transferring its replication data (differential data) to the remote array. And once TCE has transferred all replication data in the TCE group, the pair split that has been reserved for the Snapshot group will be actually executed, which changes the Snapshot pair to Split state. See [Figure 27-58 on page 27-86](#).




Figure 27-58: Remote snapshot creation 1

Normally, when the S-VOLs of TCE are cascaded with the P-VOLs of Snapshot, pair split to the Snapshot pairs is rejected while the TCE pair state is being Paired. However, when all of the following conditions are met, a pair split to the Snapshot pairs will be accepted even while TCE pair state is Paired, which means you can obtain remote snapshots. Please note that you have to issue a pair split to Snapshot pairs by group. If you issue a pair split by pair, the pair split will be rejected even when the conditions are met.

See [Figure 27-59 on page 27-87](#).

The conditions for execution are:

- The S-VOLs of TCE are cascaded with P-VOLs of a Snapshot group that will receive pair split.
- In the above cascade configuration, all Snapshot pairs in the Snapshot group that will receive pair split are cascaded with TCE pairs.
- In the above cascade configuration, the number of Snapshot pairs in the Snapshot group is the same as the one of TCE pairs in the TCE pair group.
- All Snapshot pairs in the Snapshot group that will receive pair split are in Paired state.
- All TCE pairs in the TCE group that are cascaded with Snapshot are in Paired state. Paired:split and Paired:delete do not meet the condition. Also, the condition is not met when a "pairsplit -mscas" command is being executed from CCI.




Figure 27-59: Pair split available or not available

After a pair split has been reserved for a Snapshot group, if one of the following occurs as a result of executing a pair operation to Snapshot or TCE or the pair state being changed, all pairs in the Snapshot group for which the pair split has been reserved will change to Failure state.

- The conditions for executing mentioned above are not met.
- TCE cycle copy fails or stops
- Temporal consistency in the Snapshot group is not ensured.

Specific examples of pair operations and pair changes are listed below that can change all pairs in the Snapshot group to Failure state.

- Some pairs in the Snapshot group change to Failure state.
- Pairs in the TCE group change to Pool Full state or inconsistency state.

- TCE pair creation is executed and a new pair is created for the TCE group.
- Pair resync is executed to TCE pairs by pair or by group after a problem makes the TCE pairs change to Failure state.
- On the local array, pair split or pair deletion is executed to the TCE pairs by pair (when pair split or pair deletion is executed by group on the local array, the Snapshot pairs will change to Split state once the TCE pairs have been split or deleted).
- On the remote array, pair deletion is executed to the TCE pairs by pair or by group.
- On the remote array, forced takeover is executed to the TCE pairs by pair or by group.
- Planned shutdown is performed on the remote array or the remote array is down due to a problem (because the cycle copy stops, all pairs in the Snapshot group change to Failure state when the remote array is recovered).

Other cautions on this feature.

- A reserved pair split to Snapshot can time out in case that cycle time takes too long to complete or cycle time does not complete due to a problem.
- This feature can be executed when both local array and remote array are HUS 100 series.
- After a pair split has been reserved for Snapshot, if online firmware replacement is performed on the remote array, the reserved pair split can time out. Do not perform online firmware replacement on the remote array after a pair split has been reserved.
- When you issue a pair split to a Snapshot group using CCI, you need to set a value (on the second time scale) of about two times the cycle time to -t option. Here is an example for a cycle time of 3600 seconds:
pairsplit -g ss -t 7200

2. Issue "pairsplit -mscas" command of CCI to TCE pairs from the local array.

The remote Snapshot creation command function makes a local host on which applications are running issue a command to split a snapshot cascading from S-VOL of the remote array. The data determined for the remote snapshot is the P-VOL data at the time that the local array receives the split request. See [Figure 27-60 on page 27-89](#).

When the host issues a remote snapshot creation command to the P-VOL (1), the local array performs in-band communication by using the remote line, and requests the creation of a remote snapshot (2). The remote array creates a snapshot of the S-VOL according to the command (3). This communication (2) is executed after the P-VOL data is determined at the time when the split command was issued to the P-VOL and the determined P-VOL data is reflected onto the S-VOL.

By commanding creation of a remote snapshot from the local host, the timing of the I/O stop of the application and the snap shot creation is synchronized, and the consistent backup data can be performed.

Even while remote snapshot processing is in progress, a TCE pair status keeps being Paired and the S-VOL continues to be updated. When a combination of TrueCopy and ShadowImage is used for remote backup, several commands need to be used and the pair status cannot keep being Paired. Many procedures, such as suspending and re-synchronizing the ShadowImage pair and the TrueCopy pair, are therefore required. This limits creating backup data once every several hours. TCE simplifies backup operation because only one command is required. In addition, the backup frequency can be several seconds to several minutes.




Figure 27-60: Remote snapshot creation 2

Naming function: This function adds a human-readable character string of ASCII 31 characters to a remote snapshot. Because a snapshot can be identified by a character string rather than a volume number, a snapshot which includes files to restore can be easily found from several generations and it reduces the risk of operator error.

Time management function: Array manages at which time a remote snapshot holds P-VOL data. This function simplifies snapshot aging, meaning finding and deleting old snapshots. The managed time is the time on the local array. In an array, two controllers have independent clock, respectively. Therefore, when the time management of the remote snapshot is used, it must be set as that there is no big difference in the time of both controllers. It is recommended to use NTP (Network Time Protocol) to adjust clocks among the controllers.

TCE with Snapshot cascade restrictions

Cascade connection with TCEs is not available.




Figure 27-61: Cascade connection restrictions

ShadowImage In-system Replication reference information

This appendix includes:

- ❑ [ShadowImage general specifications](#)
- ❑ [Operations using CLI](#)
- ❑ [Operations using CCI](#)
- ❑ [I/O switching mode feature](#)

ShadowImage general specifications

Table A-1 lists and describes the external specifications for ShadowImage.

Table A-1: External specifications for ShadowImage

Item	Specification
Applicable model	For dual configuration only.
Host interface	Fibre Channel or iSCSI.
Number of pairs	HUS 130/HUS 150: 2,047 (maximum) HUS 110: 1,023 (maximum) Note: When a P-VOL is paired with eight S-VOLs, the number of pairs is eight.
Command devices	Required for CCI. - Maximum: 128 per disk array - Volume size: 33MB or greater
Unit of pair management	Volumes are the target of ShadowImage pairs, and are managed per volume.
Pair structure (number of S-VOLs per P-VOL)	1 P-VOL:8 S-VOLs
Differential Management LU (DMLU)	- The DMLU size must be greater than or equal to 10 GB. The recommended size is 64 GB. The minimum DMLU size is 10 GB, the maximum size is 128GB. - The stripe size is 64KB minimum, 256KB maximum. - If you are using a merged volume for DMLU, it is necessary that each sub volume capacity is more than 1GB in average. - There is only one DMLU. Redundancy is necessary because a secondary DMLU is not available - A SAS drive and RAID 1+0 is recommended for performance.
RAID level	- P-VOL: RAID 0 (2D to 16D), RAID 1+0 (2D+2D to 8D+8D), RAID 5 (2D+1P to 15D+1P), RAID 6 (2D+2P to 28D+2P), RAID 1 (1D+1D) ("with redundancy" recommended) - S-VOL: RAID 0 (2D to 16D), RAID 1+0 (2D+2D to 8D+8D), RAID 5 (2D+1P to 15D+1P), RAID 6 (2D+2P to 28D+2P), RAID 1 (1D+1D) ("with redundancy" recommended)
Combination of RAID groups	P-VOL and S-VOL should be paired on different RAID groups. The number of data disks does not have to be the same.
Size of P-VOL and S-VOL	P-VOL = S-VOL. The max volume size is 128 TB.
Types of drive for the P-VOL and S-VOL	If the drive types are supported by the disk array, they can be set for the P-VOL and S-VOL. Assign a volume consisting of SAS or SSD/FMD drives to a P-VOL.

Table A-1: External specifications for ShadowImage

Item	Specification
Consistency Group (CTG) number	CTG per disk array: 1,024/array (maximum) HUS 130/HUS 150: 2,047 pairs/CTG (maximum) HUS 110: 1,023 pairs/CTG (maximum)
MU number	This is used for specifying a pair in CCI. For ShadowImage pairs, the value from 0 to 39 can be specified.
Mixing ShadowImage and non-ShadowImage	Mixing volumes (P-VOL and S-VOL) of ShadowImage and volumes of non-ShadowImage are available within the disk array. However, note that there may be some effects to the performance. The performance decreases when re-synchronizing pairs operation is in priority during resynchronization (even if the volumes are non-ShadowImage).
Concurrent use with TrueCopy	Yes. When the firmware of the disk array is less than 0920/B, the cascade connection cannot be executed with the ShadowImage pair including the DP-VOL created by Dynamic Provisioning. See Cascading ShadowImage with TrueCopy on page 27-11 for more information.
Concurrent use of TCE	ShadowImage and TCE can be used together at the same time, but a cascade between ShadowImage and TCE is not supported.
Concurrent use of Dynamic Tiering	The DP-VOL created by Dynamic Provisioning can be used as a ShadowImage P-VOL or S-VOL. For more details, see Concurrent use of Dynamic Provisioning on page 4-12 .
Concurrent use of Dynamic Provisioning	The DP volume of the DP pool whose tier mode is enabled in Dynamic Tiering can be used as a P-VOL and an S-VOL of ShadowImage. For more details, see Concurrent use of Dynamic Tiering on page 4-16 .
Concurrent use of Snapshot	Snapshot and ShadowImage can be used together at the same time. The number of CTG at the time of using Snapshot and ShadowImage together is limited to the maximum of 1,024 combining that of Snapshot and ShadowImage.
Formatting, growing/shrinking, deleting during Coupling (RAID group, P-VOL, S-VOL).	Not available. However, when pair status is failure (S-VOL Switch), a P-VOL can be formatted. When pair status is Simplex, you can grow or shrink volumes.
Concurrent use of Volume Migration	Yes, however a P-VOL, an S-VOL, and a reserved volume of Volume Migration cannot be specified as a ShadowImage P-VOL. The maximum number of the pairs and the number of pairs whose data can be copied in the background is limited when ShadowImage is used together with Volume Migration.

Table A-1: External specifications for ShadowImage

Item	Specification
Concurrent use of Cache Residency Manager	Yes, however the volume specified for Cache Residency (volume cache residence) cannot be used as a P-VOL, S-VOL.
Concurrent use of Cache Partition Manager	Yes
Concurrent use of SNMP Agent	Yes. Traps are sent following when occurs. Pair status changes to Failure.
Concurrent use of Data Retention Utility	Yes. However, when S-VOL Disable is set for n volume, the volume cannot be used in a ShadowImage pair. When S-VOL Disable is set for a volume that is already a S-VOL, no suppression of the pair takes place, unless the pair status is split.
Concurrent use of Power Saving/Power Saving Plus	Yes. However, when a P-VOL or S-VOL is included in a RAID group in which Power Saving/Power Saving Plus is enabled, the only ShadowImage pair operation that can be performed are pair split and the pair release.
Concurrent use of unified volume	Yes.
Concurrent use of LUN Manager	Yes.
Concurrent use of Password Protection	Yes.
ShadowImage I/O switching function	Yes. DP-VOLs can be used for a P-VOL or an S-VOL of ShadowImage. For details, see I/O switching mode feature on page A-34 .
Load balancing function	The load balancing function applies to a ShadowImage pair. When the load balancing function is activated for a ShadowImage pair, the ownership of the P-VOL and S-VOL changes to the same controller. When the pair state is Synchronizing or Reverse Synchronizing, the ownership of the pair will change across the cores but not across the controllers.
Maximum supported capacity value of S-VOL (TB)	See Calculating maximum capacity on page 4-19 for details.
License	ShadowImage must be installed using the key code.
Management of volumes while using ShadowImage	Formatting and deleting volumes are not available. When formatting and deleting volumes, split ShadowImage pair(s) using the pairsplit command.
Restriction for formatting the volumes	Do not execute ShadowImage operations while formatting the volume. Formatting takes priority and the ShadowImage operations will be suspended.

Table A-1: External specifications for ShadowImage

Item	Specification
Restriction during RAID group expansion	A RAID group with a ShadowImage P-VOL or S-VOL can be expanded only when the pair status is Simplex or Split.
DMLU	The DMLU is an exclusive volume for storing the differential data at the time when the volume is copied.
Failures	When a failure of the copy operation from P-VOL to S-VOL occurs, ShadowImage will suspend the pair and the status changes to failure. If a volume failure occurs, ShadowImage suspends the pair. If a drive failure occurs, the ShadowImage pair status is not affected because of the RAID architecture.
Reduction of memory	The memory cannot be reduced when ShadowImage, Snapshot, or TrueCopy are enabled. Reduce memory after disabling the functions.

Operations using CLI

This topic describes basic Navigator 2 CLI procedures for performing ShadowImage operations.

- ❑ [Installing and uninstalling ShadowImage](#)
- ❑ [ShadowImage operations](#)
- ❑ [Creating ShadowImage pairs that belong to a group](#)
- ❑ [Splitting ShadowImage pairs that belong to a group](#)
- ❑ [Sample back up script for Windows](#)



NOTE: For additional information on the commands and options used in this appendix, see the *Hitachi Unified Storage Command Line Interface Reference Guide*.

Installing and uninstalling ShadowImage

If ShadowImage was purchased when the order for Hitachi Unified Storage was placed, then ShadowImage came bundled with the system and no installation is necessary. Proceed to [Enabling or disabling ShadowImage on page A-8](#).

If you purchased ShadowImage on an order separate from your Hitachi Unified Storage, it must be installed before enabling. A key code or key file is required



NOTE: Before installing/uninstalling ShadowImage, verify that the array is operating in a normal state. If a failure such as a controller blockade has occurred, installation/un-installation cannot be performed.

Installing ShadowImage

To install ShadowImage, the key code or key file provided with the optional feature is required. You can obtain it from the download page on the HDS Support Portal, <https://portal.hds.com>

To install ShadowImage

1. From the command prompt, register the array in which the ShadowImage is to be installed, then connect to the array.
2. Execute the auopt command to install ShadowImage. For example:

```
% auopt -unit subsystem-name -lock off -licensefile license-
file-path\license-file-name
No.  Option Name
   1  ShadowImage In-system Replication
Please specify the number of the option to unlock.
When you unlock two or more options, partition the numbers
given in the list with space(s). When you unlock all options,
input 'all'. Input 'q', then break.
The number of the option to unlock. (number/all/q [all]): 1
Are you sure you want to unlock the option?
(y/n [n]): y

Option Name                                     Result
ShadowImage In-system Replication              Unlock

The process was completed.
%
```

3. Execute the auopt command to confirm whether ShadowImage has been installed.

```
% auopt -unit array-name -refer
Option Name          Type      Term      Status
  Reconfigure Memory Status
SHADOWIMAGE          Permanent ---      Enable
      N/A
%
```

ShadowImage is installed and the status is "Enable". Installation of ShadowImage is now complete.

Uninstalling ShadowImage

To uninstall ShadowImage, the key code provided with the optional feature is required. Once uninstalled, ShadowImage cannot be used again until it is installed using the key code or key file.

To uninstall ShadowImage:

1. All ShadowImage pairs must be released (the status of all volumes are Simplex) before uninstalling ShadowImage.
2. From the command prompt, register the array in which the ShadowImage is to be uninstalled, then connect to the array.
3. Execute the auopt command to uninstall ShadowImage. For example:

```
% auopt -unit subsystem-name -lock on -keycode downloaded-48
characters-key code
Are you sure you want to lock the option?
(y/n [n]): y
The option is locked.
%
```

4. Execute the auopt command to confirm whether ShadowImage has been uninstalled. For example:

```
% auopt -unit subsystem-name -refer
DMEC002015: No information displayed.
%
```

Uninstalling ShadowImage is now complete.

Enabling or disabling ShadowImage

Once ShadowImage is installed, it can be enabled or disabled.

The following describes the enabling/disabling procedure.

1. If you are disabling ShadowImage, all pairs must be released (the status of all volumes are Simplex).
2. From the command prompt, register the array in which the status of the feature is to be changed, then connect to the array.
3. Execute the auopt command to change the status (enable or disable). The following is an example of changing the status from enable to disable. If you want to change the status from disable to enable, enter enable after the -st option, as shown in the example below.

```
% auopt -unit subsystem-name -option SHADOWIMAGE -st disable
Are you sure you want to disable the option? (y/n [n]): y
The option has been set successfully.
%
```

4. Execute the auopt command to confirm whether the status has been changed. For example:

```
% auopt -unit array-name -refer
Option Name          Type          Term          Status
      Reconfigure Memory Status
SHADOWIMAGE          Permanent ---          Disable
      N/A
%
```

Enabling or disabling ShadowImage is now complete.

Setting the DMLU

The DMLU is an exclusive logical unit for storing the differential data while the volume is being copied. The DMLU in the disk array is treated in the same way as the other logical units. However, a logical unit that is set as the DMLU is not recognized by a host (it is hidden).

When the DMLU is not set, it must be created. Set a logical unit with a size of 10 GB minimum as the DMLU.

Prerequisites

- LUs for the DMLUs must be set up and formatted.
- DMLU size must be at least 10 GB.
- One DMLU is needed but two are recommended, with the second used as a backup.
- For RAID considerations, see the bullet on DMLU in [Cascading ShadowImage with TrueCopy on page 27-11](#).



NOTE: When either pair of ShadowImage, TrueCopy, or Volume Migration exist and when only one DMLU is set, the DMLU cannot be removed.

To set up DMLU

1. From the command prompt, register the array on which you want to create the DMLU and connect to that array.
2. Execute the `audmlu` command to create a DMLU. This command first displays volumes that can be assigned as DMLUs and later creates a DMLU. For example:

```
% audmlu -unit array-name -availablelist
Available Logical Units
LUN Capacity  RAID Group DP Pool  RAID Level  Type  Status
      0      10.0 GB      0      N/A    5( 4D+1P)  SAS   Normal
%
% audmlu -unit array-name -set -lu 0
Are you sure you want to set the DM-LU? (y/n [n]): y
The DM-LU has been set successfully.
%
```

3. To release an already set DMLU, specify the `-rm` in the `audmlu` command. For example:

```
% audmlu -unit array-name -rm 0
Are you sure you want to release the DM-LU? (y/n [n]): y
The DM-LU has been released successfully.
%
```

To add DMLU

1. To add an already set DMLU capacity, specify the **-chgsz** and **-size** options in the audmlu command.

```
% audmlu -unit array-name -chgsz -size capacity after
adding -rg RAID group number
Are you sure you want to add the capacity of the DM-LU?
(y/n [n]): y
The capacity of DM-LU has been added successfully.
%
```

The -rg option can be specified only when the DMLU is a normal volume. Select a RAID group which meets the following conditions:

- The drive type and the combination are the same as the DMLU
- A new volume can be created
- A sequential free area for the capacity to be expanded exists

Setting the ShadowImage I/O switching mode

The following procedure explains how to set the ShadowImage I/O switching mode to ON. For more information, see [I/O switching mode feature on page A-34](#).

To set the ShadowImage I/O Switching Mode

1. From the command prompt, register the array on which you want to set the ShadowImage I/O Switching Mode. Connect to the array.
2. Execute the `ausystemparam` command.

When you want to reset the ShadowImage I/O Switching Mode, enter `disable` following the `-set -ShadowImageIOSwitch` option. For example:

```
% ausystemparam -unit array-name -set -ShadowImageIOSwitch
enable
Are you sure you want to set the system parameter? (y/n [n]): y
The system parameter has been set successfully.
%
```

3. Execute the `ausystemparam` command to verify that the ShadowImage I/O Switching Mode has been set. For example:

```
% ausystemparam -unit array-name -refer
Options
Turbo LU Warning = OFF
:
:
ShadowImage I/O Switch Mode = ON
:
:
Operation if the Processor failures Occurs = Reset a Fault
:
:
%
```



NOTE: When turning off the I/O Switching Mode, pair status must be other than Failure (S-VOL Switch) and Synchronizing (S-VOL Switch).

Setting the system tuning parameter

This setting limits the number of times processing is executed for flushing the dirty data in the cache to the drive at the same time.

To set the Dirty Data Flush Number Limit of a system tuning parameters:

1. From the command prompt, register the array on which you want to set a system tuning parameters and connect to that array.
2. Execute the `ausystuning` command to set the system tuning parameters.

Example:

```
% ausystuning -unit array-name -set dirtylimit enable
Are you sure you want to set the system tuning parameter? (y/n [n]): y
Changing Dirty Data Flush Number Limit may have performance impact when local re
plication is enabled and time out may occur if I/O load is heavy. Please change
the setting when host I/O load is light.
Do you want to continue processing? (y/n [n]): y
The system tuning parameter has been set successfully.
%
```

ShadowImage operations

The `aureplicationlocal` command operates ShadowImage pair. To refer the `aureplicationlocal` command and its options, type in `aureplicationlocal -help` at the command prompt.

Confirming pairs status

To confirm the ShadowImage pairs, use the `aureplicationlocal -refer` command.

1. From the command prompt, register the array on which you want to confirm the ShadowImage pairs. Connect to the array.
2. Execute the `aureplicationlocal` command to confirm the ShadowImage pairs, as shown in the example below.

```
% aureplicationlocal -unit subsystem-name -refer -si
Pair Name          LUN Pair LUN Status
    Copy Type      Group
SI_LU1020_LU1021    1020    1021 Paired(
    0%)              S
ShadowImage  ---:Ungrouped
%
```

Creating ShadowImage pairs

The following procedure explains how to create on pair. To create pairs in a group, refer to [Creating ShadowImage pairs that belong to a group on page A-17](#).

To create the ShadowImage pairs, use the `aureplicationlocal -create` command.

1. From the command prompt, register the array on which you want to create the ShadowImage pairs. Connect to the array.
2. Execute the `aureplicationlocal` command to create the ShadowImage pairs.

When you want to automatically split the pair immediately after creation is completed, create the pair specifying the `-compsplit` option. In this case, the pair status immediately after pair creation becomes Split Pending.

In the following example, the P-VOL LUN is 1020 and the S-VOL LUN is 1021.

```
% aureplicationlocal -unit subsystem-name -create -si -pvol
1020 -svol 1021
Are you sure you want to create pairs "SI_LU1020_LU1021"?
(y/n [n]): y
The pair has been created successfully.
%
```

3. Verify the pair status, as shown in the example below.

```
% aureplicationlocal -unit subsystem-name -refer -si
Pair Name                LUN Pair LUN Status
Copy Type    Group
SI_LU1020_LU1021        1020    1021 Synchronizing(
40%)                ShadowImage ---:Ungrouped
%
```

The ShadowImage pair is created.

Splitting ShadowImage pairs

To split the ShadowImage pairs, use the `aureplicationlocal -split` command.

1. From the command prompt, register the array on which you want to split the ShadowImage pairs. Connect to the array.
2. Execute the `aureplicationlocal` command to split the ShadowImage pairs.

When you want to split Quick Mode, split the pair specifying the **-quick** option. In this case, the pair status immediately after pair splitting becomes Split Pending.

In the following example, the P-VOL LUN is 1020 and the S-VOL LUN is 1021.

```
% aureplicationlocal -unit subsystem-name -split -si -pvol
    1020 -svol 1021
Are you sure you want to split pairs?
(y/n [n]): y
The pair has been split successfully.
%
```

3. Verify the pair status as shown in the example below.

```
% aureplicationlocal -unit subsystem-name -refer -si
Pair Name          LUN Pair LUN Status
    Copy Type    Group
SI_LU1020_LU1021    1020    1021 Split(100%)
    ShadowImage  ---:Ungrouped
%
```

The ShadowImage pair is split.

Re-synchronizing ShadowImage pairs

To re-synchronize the ShadowImage pairs, use the `aureplicationlocal -resync` command.

1. From the command prompt, register the array on which you want to re-synchronize the ShadowImage pairs. Connect to the array.
2. Execute the `aureplicationlocal` command to re-synchronize the ShadowImage pairs.

When you want to re-synchronize Quick Mode, re-synchronize the pair specifying the **-quick** option. In this case, the pair status immediately after pair re-synchronizing becomes Paired Internally Synchronizing.

In the following example, the P-VOL LUN is 1020 and the S-VOL LUN is 1021.

```
% aureplicationlocal -unit subsystem-name -resync -si -pvol
1020 -svol 1021
Are you sure you want to re-synchronize pairs?
(y/n [n]): y
The pair has been re-synchronized successfully.
%
```

3. Verify the pair status as shown in the example below.

```
% aureplicationlocal -unit subsystem-name -refer -si
Pair Name          LUN Pair LUN Status
Copy Type      Group
SI_LU1020_LU1021      1020    1021 Synchronizing(
40%)                ShadowImage ---:Ungrouped
%
```

The ShadowImage pair is resynchronized.

Restoring the P-VOL

To restore the ShadowImage pairs, use the `aureplicationlocal -restore` command.

1. From the command prompt, register the array on which you want to restore the ShadowImage pairs. Connect to the array.
2. Execute the `aureplicationlocal` command to restore the ShadowImage pairs.

In the following example, the P-VOL LUN is 1020 and the S-VOL LUN is 1021.

```
% aureplicationlocal -unit subsystem-name -restore -si -pvol
1020 -svol 1021
Are you sure you want to restore pairs?
(y/n [n]): y
The pair has been restored successfully.
%
```

3. Verify the pair status as shown in the example below.

```
% aureplicationlocal -unit subsystem-name -refer -si
Pair Name          LUN Pair LUN Status
  Copy Type    Group
SI_LU1020_LU1021      1020      1021  Reverse
  synchronizing( 40%) ShadowImage ---:Ungrouped
%
```

The ShadowImage pair is restored.

Deleting ShadowImage pairs

To delete the ShadowImage pairs, use the `aureplicationlocal -simplex` command.

1. From the command prompt, register the array on which you want to release the ShadowImage pairs. Connect to the array.
2. Execute the `aureplicationlocal` command to release the ShadowImage pairs.

In the following example, the P-VOL LUN is 1020 and the S-VOL LUN is 1021.

```
% aureplicationlocal -unit subsystem-name -simplex -si -pvol
    1020 -svol 1021
Are you sure you want to release pairs?
(y/n [n]): y
The pair has been released successfully.
%
```

3. Verify the pair status as shown in the example below.

```
% aureplicationlocal -unit subsystem-name -refer -si
DMEC002015: No information displayed.
%
```

The ShadowImage pair is deleted.

Editing pair information

You can change the pair name, group name, and/or copy pace.

To change the pair information:

1. From the command prompt, register the array on which you want to change the ShadowImage pair information. Connect to the array.
2. Execute the `aureplicationlocal` command to change the ShadowImage pair information.

In the following example, the P-VOL LUN is 1020 and the S-VOL LUN is 1021.

```
% aureplicationlocal -unit subsystem-name -chg -si -pace slow
-pvol 1020 -svol 1021
Are you sure you want to change pair information?
(y/n [n]): y
The pair information has been changed successfully.
%
```

The ShadowImage pair information is changed.

Creating ShadowImage pairs that belong to a group

To create multiple ShadowImage pairs that belong to a group:

1. Create the first pair that belongs to a group specifying an unused group number for the new group with the `-gno` option. The new group has been created and in this group, the new pair has been created too.

```
% aureplicationlocal -unit array-name -create -si -pvol 1020
-svol 1021 -gno 20
Are you sure you want to create pairs "SI_LU1020_LU1021"?
(y/n [n]): y
The pair has been created successfully.
%
```

2. Add the name to the group if necessary using command to change the pair information.

```
% aureplicationlocal -unit array-name -chg -si -gno 20 -
newgname group-name
Are you sure you want to change pair information?
(y/n [n]): y
The pair information has been changed successfully.
%
```

3. Create the next pair that belongs to the created group specifying the number of the created group with `-gno` option.
4. By repeating the step 3, the multiple pairs that belong to the same group can be created.



NOTE: You cannot use the options of the group number specification and automatic split after pair creation at the same time. To create two or more pairs that utilize the group by using Quick Mode, create all pairs belonging to the group, specify the quick option, and execute the split by group unit.

Splitting ShadowImage pairs that belong to a group

To split two or more ShadowImage pairs that belong to a group:

1. Execute the **aureplicationlocal** command to split the ShadowImage pairs. Display the status of the pairs belonging to the group to be the target, and split the pairs after checking that all pairs are in the split-able status.

```
% aureplicationlocal -unit array-name -refer -si
Pair Name          LUN Pair LUN Status
Copy Type          Group
SI_LU1000_LU1003   1000      1003 Paired(100%)
ShadowImage 0:
SI_LU1001_LU1004   1001      1004 Paired(100%)
ShadowImage 0:
SI_LU1002_LU1005   1002      1005 Paired(100%)
ShadowImage 0:
%
% aureplicationlocal -unit array-name -si -split -gno 0
Are you sure you want to split pair?
(y/n [n]): y
The pair has been split successfully.
%
```



NOTE: If a pair that cannot be split is mixed in the specified group, the pair split in the group unit does not operate. when this occurs, an error as a response to the pair split operation may or may not be displayed. Also, the pair split-able status differs depending on whether Quick Mode is used or not. Therefore, check that all the pairs belonging to the group to be the pair split target are in the following statuses according to each case.

When using Quick Mode:

- Paired
- Paired Internally Synchronizing
- Synchronizing

When not using the Quick Mode:

- Paired
- Paired Internally Synchronizing

2. Verify the pair status executing the **aureplicationlocal** command.

```
% aureplicationlocal -unit array-name -refer -si
Pair Name          LUN Pair LUN Status
Copy Type          Group
SI_LU1000_LU1003   1000      1003 Split(100%)
ShadowImage 0:
SI_LU1001_LU1004   1001      1004 Split(100%)
ShadowImage 0:
SI_LU1002_LU1005   1002      1005 Split(100%)
ShadowImage 0:
%
```

The ShadowImage pair is split.

Sample back up script for Windows

This section provides sample script for backing a volume on Windows Server.

```
echo off
REM Specify the registered name of the arrays
set UNITNAME=Array1
REM Specify the group name (Specify "Ungroup" if the pair
    doesn't belong to any group)
set G_NAME=Ungrouped
REM Specify the pair name
set P_NAME=SI_LU0001_LU0002
REM Specify the directory path that is mount point of P-VOL
    and S-VOL
set MAINDIR=C:\main
set BACKUPDIR=C:\backup
REM Specify GUID of P-VOL and S-VOL
PVOL_GUID=xxxxxxxx-xxxx-xxxx-xxxx-xxxxxxxxxxxxxx
SVOL_GUID=yyyyyyyy-yyy-yyy-yyy-yyy-yyyyyyyyyyyy

REM Unmounting the S-VOL
pairdisplay -x umount %BACKUPDIR%
REM Re-synchronizing pair (Updating the backup data)
aureplicationlocal -unit %UNITNAME% -si -resync -pairname
    %P_NAME% -gname %G_NAME%
aureplicationmon -unit %UNITNAME% -evwait -si -pairname
    %P_NAME% -gname %G_NAME% -st paired -pvol

REM Unmounting the P-VOL
pairdisplay -x umount %MAINDIR%
REM Splitting pair (Determine the backup data)
aureplicationlocal -unit %UNITNAME% -si -split -pairname
    %P_NAME% -gname %G_NAME%
aureplicationmon -unit %UNITNAME% -evwait -si -pairname
    %P_NAME% -gname %G_NAME% -st split -pvol
REM Mounting the P-VOL
pairdisplay -x mount %MAINDIR% Volume{%PVOL_GUID%}

REM Mounting the S-VOL
pairdisplay -x mount %BACKUPDIR% Volume{%SVOL_GUID%}
< The procedure of data copy from C:\backup to backup
    appliance>
```



NOTE: For Windows Server environments, The CCI mount/unmount commands must be used when mounting/un-mounting a volume.

Operations using CCI

This topic describes basic CCI procedures for setting up and performing ShadowImage operations.

- ❑ [Setting up CCI](#)
- ❑ [ShadowImage operations using CCI](#)
- ❑ [Pair, group name differences in CCI and Navigator 2](#)

Setting up CCI

CCI is used to display ShadowImage volume information, create and manage ShadowImage pairs, and issue commands for replication operations. CCI resides on the UNIX/Windows management host and interfaces with the arrays through dedicated volumes. CCI commands can be issued from the UNIX/Windows command line or using a script file.

The following sub-topics describe necessary set up procedures for CCI for ShadowImage.

Setting the command device

The command devices and LU mapping setting is used Navigator 2.

To designate command devices

1. From the command prompt, register the array to which you want to set the command device. Connect to the array.
2. Execute the aucmddev command to set a command device. First, display volumes to be assignable command device, and later set a command device. When you want to use the protection function of CCI, enter enable following the -dev option.

The following example specifies LUN 2 for command device 1.

```
% aucmddev -unit disk array-name -availablelist
Available Logical Units
  LUN  Capacity  RAID Group  DP Pool  RAID Level  Type  Status
    2    35.0 MB      0         N/A      6( 9D+2P)  SAS  Normal
    3    35.0 MB      0         N/A      6( 9D+2P)  SAS  Normal
%
% aucmddev -unit disk array-name -set -dev 1 2
Are you sure you want to set the command devices?
(y/n [n]): y
The command devices have been set successfully.
%
```

3. Execute the aucmddev command to verify that the command device has been set. For example:

```
% aucmddev -unit disk array-name -refer
Command Device LUN  RAID Manager Protect
1              2   Disable
%
```



NOTE: To set the alternate command device function or to avoid data loss and disk array downtime, designate two or more command devices. For details on alternate Command Device function, refer to the *Hitachi Unified Storage Command Control Interface Installation and Configuration Guide*

4. The following example releases a command device:

```
% aucmddev -unit disk array-name -rm -dev 1
Are you sure you want to release the command devices?
(y/n [n]): y
This operation may cause CCI, which is accessing this command
device, to freeze.
Stop the CCI, which is accessing this command device, before
performing this operation.
Are you sure you want to release the command devices? (y/n [n]): y
The specified command device will be released.
Are you sure you want to execute? (y/n [n]): y
The command devices have been released successfully.
%
```

5. To change an already set command device, release the command device, then change the volume number. The following example specifies LUN 3 for command device 1.

```
% aucmddev -unit disk array-name -set -dev 1 3
Are you sure you want to set the command devices?
(y/n [n]): y
The command devices have been set successfully.
%
```

Setting LU mapping

If using iSCSI, use the autargetmap command instead of the auhgmap command used with fibre channel.

To set up LU mapping

1. From the command prompt, register the disk array to which you want to set the LU Mapping, then connect to the disk array.
2. Execute the auhgmap command to set the LU Mapping. The following is an example of setting LUN 0 in the disk array to be recognized as 6 by the host. The port is connected via target group 0 of port 0A on controller 0.

```
% auhgmap -unit disk array-name -add 0 A 0 6 0
Are you sure you want to add the mapping information?
(y/n [n]): y
The mapping information has been set successfully.
%
```

3. Execute the auhgmap command to verify that the LU Mapping is set. For example:

```
% auhgmap -unit disk array-name -refer
Mapping mode = ON
Port  Group  H-LUN    LUN
  0A      0      6      0
%
```

Defining the configuration definition file

The configuration definition file describes the system configuration. It is required to make CCI operational. The configuration definition file is a text file created and/or edited using any standard text editor. It can be defined from the PC where CCI software is installed.

A sample configuration definition file, HORCM_CONF, is included with the CCI software. It should be used as the basis for creating your configuration definition file(s). The system administrator should copy the sample file, set the necessary parameters in the copied file, and place the copied file in the proper directory. For details on configuration definition file, refer to the *Hitachi Unified Storage Command Control Interface Installation and Configuration Guide*.

The configuration definition file can be automatically created using the mkconf command tool. For details on the mkconf command, refer to the *Hitachi Unified Storage Command Control Interface Installation and Configuration Guide*. However, the parameters, such as poll(10ms) must be set manually (see step 4 below).

To define the configuration definition file

The following is an example that manually defines the configuration definition file. The system is configured with two instances within the same Windows host.

1. On the host where CCI is installed, verify that CCI is not running. If CCI is running, shut it down using the horcmshutdown command.
2. In the command prompt, make two copies of the sample file (horcm.conf). For example:

```
c:\HORCM\etc> copy \HORCM\etc\horcm.conf \WINDOWS\horcm0.conf  
c:\HORCM\etc> copy \HORCM\etc\horcm.conf \WINDOWS\horcm1.conf
```

3. Open horcm0.conf using the text editor.
4. In the **HORCM_MON** section, set the necessary parameters.



NOTE: A value more than or equal to 6000 must be set for poll(10ms). Specifying the value incorrectly may cause resource contention in the internal process, resulting the process temporarily suspending and pausing the internal processing of the disk array.

5. In the **HORCM_CMD** section, specify the physical drive (command device) on the disk array. [Figure A-1](#) and [Figure A-2](#) show examples of the horcm0.conf file in which the ShadowImage P-VOL-to-S-VOL ratio is 1:1 and 1:3, respectively.

```

horcm0.conf - Notepad
File Edit Search Help

HORCM_MON
#ip_address      service      poll(10ms)      tineout(10ms)
XXXXXXXXX        5000          12000           3000

HORCM_CMD
#dev_name        dev_name        dev_name
\\.\PHYSICALDRIVE1

HORCM_DEV
#dev_group      dev_name      port#      TargetID      LU#      MU#
VG01            oradb1        CL1-A      1             1       0

HORCM_INST
#dev_group      ip_address      service
VG01            XXXXXXXXX      5001

```

Figure A-1: Horcm0.conf example 1 (P-VOL: S-VOL=1: 1)

```

horcm0.conf - Notepad
File Edit Search Help

HORCM_MON
#ip_address      service      poll(10ms)      tineout(10ms)
XXXXXXXXX        5000          12000           3000

HORCM_CMD
#dev_name        dev_name        dev_name
\\.\PHYSICALDRIVE1

HORCM_DEV
#dev_group      dev_name      port#      TargetID      LU#      MU#
VG01            oradb1        CL1-A      1             1       0
VG02            oradb2        CL1-A      1             1       1
VG03            oradb3        CL1-A      1             1       2

HORCM_INST
#dev_group      ip_address      service
VG01            XXXXXXXXX      5001

```

Figure A-2: Horcm0.conf example 2 (P-VOL: S-VOL=1: 3)

```

horcm0.conf - Notepad
File Edit Format View Help

HORCM_MON
#ip_address      service      poll(10ms)      tineout(10ms)
XXXXXXXXX        5000          12000           3000

HORCM_CMD
#dev_name        dev_name        dev_name
\\.\PHYSICALDRIVE1

HORCM_DEV
#dev_group      dev_name      port#      TargetID      LU#      MU#
VG01            oradb1_SI     CL1-A      1             1       0
VG02            oradb2_SS     CL1-A      1             2       1

HORCM_INST
#dev_group      ip_address      service
VG01            XXXXXXXXX      5001
VG02            XXXXXXXXX      5001

```

Figure A-3: Horcm0.conf example (cascading ShadowImage S-VOL with Snapshot P-VOL)

6. Save the configuration definition file and use the horcmstart command to start CCI.

7. Execute the raidscan command and write down the target ID displayed in the execution result.
8. Shut down CCI and then open the configuration definition file again.
9. In the **HORCM_DEV** section, set the necessary parameters. For the target ID, set the ID of the raidscan result you wrote down. Also, the item MU# must be added after the LU#.
10. In the **HORCM_INST** section, set the necessary parameters, and then save (overwrite) the file.
11. Repeat Steps 3 to 10, using [Figure A-4](#) to [Figure A-6](#) on page A-26 for examples.

```

horcm1.conf - Notepad
File Edit Search Help

HORCM_MON
#ip_address      service      poll(10ms)      timeout(10ms)
XXXXXXXXX        5001          12000           3000

HORCM_CMD
#dev_name        dev_name        dev_name
\\.\PHYSICALDRIVE1

HORCM_DEV
#dev_group      dev_name      port#      TargetID      LU#      MU#
UG01            oradb1        CL1-A      1              2        0

HORCM_INST
#dev_group      ip_address      service
UG01            XXXXXXXXX      5000
  
```

Figure A-4: Horcm1.conf example 3 (P-VOL: S-VOL=1: 1)

```

horcm1.conf - Notepad
File Edit Search Help

HORCM_MON
#ip_address      service      poll(10ms)      timeout(10ms)
XXXXXXXXX        5001          12000           3000

HORCM_CMD
#dev_name        dev_name        dev_name
\\.\PHYSICALDRIVE1

HORCM_DEV
#dev_group      dev_name      port#      TargetID      LU#      MU#
UG01            oradb1        CL1-A      1              2        0
UG02            oradb2        CL1-A      1              3        0
UG03            oradb3        CL1-A      1              4        0

HORCM_INST
#dev_group      ip_address      service
UG01            XXXXXXXXX      5000
  
```

Figure A-5: Horcm1.conf example 4 (P-VOL: S-VOL=1: 3)

```

horcm1.conf - Notepad
File Edit Format View Help
HORCM_MON
#ip_address      service      poll(10ms)  timeout(10ms)
xxxxxxxxx        5001          12000       3000

HORCM_CMD
#dev_name        dev_name        dev_name
\\.\PHYSICALDRIVE1

HORCM_DEV
#dev_group       dev_name       port#       TargetID     LU#     MU#
VG01             oradb1_SI      CL1-A       1            2       0
VG02             oradb2_SS      CL1-A       1            3       0

HORCM_INST
#dev_group       ip_address      service
VG01             xxxxxxxx        5000
VG02             xxxxxxxx        5000

```

Figure A-6: Horcm1.conf example (cascading ShadowImage S-VOL with Snapshot P-VOL)

12. Enter the following example lines in the command prompt to verify the connection between CCI and the disk array.



NOTE: Volumes of ShadowImage can be cascaded with those of Snapshot. There is no distinction between ShadowImage pairs and Snapshot pairs on the configuration definition file of CCI. Therefore, the configuration definition file when cascading the P-VOL of ShadowImage and the P-VOL of Snapshot can be defined the same as the one shown in [Figure A-2 on page A-24](#) and [Figure A-5 on page A-25](#). Moreover, the configuration definition file when cascading the S-VOL of ShadowImage and the P-VOL of Snapshot can be defined the same as the one shown in [Figure A-3 on page A-24](#) and [Figure A-6 on page A-26](#)

For details on the configuration definition file, refer to *Hitachi Unified Storage Command Control Interface Installation and Configuration Guide*.

```

C:\>cd horcm\etc

C:\HORCM\etc>echo hd1-3 | .\inraid
Harddisk 1 -> [ST] CL1-A Ser =91100174 LDEV = 0 [HITACHI ] [DF600F-CM ]
Harddisk 2 -> [ST] CL1-A Ser =91100174 LDEV = 1 [HITACHI ] [DF600F ]
                HORC = SMPL HOMRCF[MU#0 = SMPL MU#1 = NONE MU#2 = NONE]
                RAID6[Group 1-0] SSID = 0x0000
Harddisk 3 -> [ST] CL1-A Ser =91100174 LDEV = 2 [HITACHI ] [DF600F ]
                HORC = SMPL HOMRCF[MU#0 = SMPL MU#1 = NONE MU#2 = NONE]
                RAID6[Group 2-0] SSID = 0x0000

C:\HORCM\etc>

```

Setting the environment variable

To perform ShadowImage operations, you must set the environment variable for the execution environment. The following describes an example in which two instances are configured within the same Windows host.

1. Set the environment variable for each instance. Enter the following from the command prompt:

```
C:\HORCM\etc>set HORCMINST=0
```

2. To enable ShadowImage, the environment variable must be set as follows:


```
C:\HORCM\etc>set HORCC_MRCF=1
```

3. Execute the horcmstart script, and then execute the pairedisplay command to verify the configuration, as shown in the following example:

```
C:\HORCM\etc>horcmstart 0 1
starting HORCM inst 0
HORCM inst 0 starts successfully.
starting HORCM inst 1
HORCM inst 1 starts successfully.

C:\HORCM\etc>pairedisplay -g VG01
group  PairVOL(L/R) (Port#,TID, LU-M) ,Seq#,LDEV#.P/S,Status, Seq#,P-LDEV# M
VG01   oradb1(L)    (CL1-A , 1,  1-0 )91100174    1.SMPL -----,----- ---- -
VG01   oradb1(R)    (CL1-A , 1,  2-0 )91100174    2.SMPL -----,----- ---- -
```

CCI setup for performing ShadowImage operations is now complete.

ShadowImage operations using CCI

Pair operation using CCI are shown in Figure A-7.




Figure A-7: ShadowImage pair status transitions

Confirming pair status

Table A-2 shows the related CCI and Navigator 2 GUI pair status.

Table A-2: CCI / Navigator 2 GUI pair status

Description	CCI	Navigator 2
Status where a pair is not created.	SMPL	Simplex
Initial copy or resynchronization copy is in execution.	COPY	Synchronizing
Status that the copying is completed and the contents written to the P-VOL is reflected in the S-VOL.	PAIR	Paired
Status that copy is not completed but the pair split in the quick mode is accepted.	PAIR(IS)	Paired Internally Synchronizing
Status that the written contents are managed as differential data by split.	PSUS/ SSUS	Split
Status that the written contents by the quick split is managed as the differential data.	PSUS(SP)/ COPY	Split Pending
Status that the differential data is copied from the S-VOL to the P-VOL for restoration.	RCPY	Reverse Synchronizing
Status that suspends copying forcibly when a failure occurs.	PSUE	Failure or Failure(R)

To confirm ShadowImage pairs



NOTE: The following are assigned automatically to a pair name:
SS_LUXXXX_LUYYYY_ZZZZZ

XXXX: LU number of P-VOL (four digits with 0)

YYYY: LU number of V-VOL (four digits with 0)

ZZZZZ: 5 digits of optional number

For the example below, the group name in the configuration definition file is VG01.

1. Execute the `pairdisplay` command to verify the pair status and the configuration. For example:

```
C:\HORCM\etc>pairdisplay -g VG01
Group PairVol(L/R) (Port#,TID, LU-M) ,Seq#,LDEV#.P/S,Status, Seq#,P-LDEV# M
VG01  oradb1(L)   (CL1-A , 1, 1-0 )91100174 1.P-VOL PAIR,91100174 2 -
VG01  oradb1(R)   (CL1-A , 1, 2-0 )91100174 2.S-VOL PAIR,----- 1 -
```

The pair status is displayed. For details on the `pairdisplay` command and its options, refer to the *Hitachi Unified Storage Command Control Interface Installation and Configuration Guide*.



NOTE: CCI displays PSUE for both Failure and Failure(R).

Creating pairs (paircreate)

To create ShadowImage pairs

1. Execute the pairdisplay command to verify that the status of the ShadowImage volumes is SMPL. The following example specifies the group name in the configuration definition file as VG01.

```
C:\HORCM\etc>pairdisplay -g VG01
Group  PairVol(L/R) (Port#,TID, LU-M) ,Seq#,LDEV#.P/S,Status, Seq#,P-LDEV# M
VG01   oradb1(L)   (CL1-A , 1,  1-0 )91100174    1.SMPL -----,----- ---- -
VG01   oradb1(R)   (CL1-A , 1,  2-0 )91100174    2.SMPL -----,----- ---- -
```

2. Execute the paircreate command, then execute the pairevtwait command to verify that the status of each volume is PAIR. When using the paircreate command, the -c option is copying pace, which can vary between 1-15. 6-10 (medium) is recommended. 1-5 is a slow pace, which is used when I/O performance must be prioritized. 11-15 is a fast pace, which is used when copying is prioritized. The following example shows the paircreate and pairevtwait commands.

```
C:\HORCM\etc>paircreate -g VG01 -vl -c 15
C:\HORCM\etc>pairevtwait -g VG01 -s pair -t 300 10
pairevtwait : Wait status done.
```

3. Execute the pairdisplay command to verify the pair status and the configuration. For example:

```
C:\HORCM\etc>pairdisplay -g VG01
Group  PairVol(L/R) (Port#,TID, LU-M) ,Seq#,LDEV#.P/S,Status, Seq#,P-LDEV# M
VG01   oradb1(L)   (CL1-A , 1,  1-0 )91100174    1.P-VOL PAIR,91100174    2  -
VG01   oradb1(R)   (CL1-A , 1,  2-0 )91100174    2.S-VOL PAIR,-----    1  -
```

Pair creation using a consistency group

A consistency group insures that the data in two or more S-VOLs included in a group are of the same time. For more information, see [Consistency group \(CTG\) on page 2-18](#).

To create a pair using a consistency group

1. Execute the `pairdisplay` command to verify that the status of the ShadowImage volumes is SMPL. In the following example, the group name in the configuration definition file is VG01.

```
C:\HORCM\etc>pairdisplay -g VG01
Group  PairVol(L/R) (Port#,TID, LU-M) ,Seq#,LDEV#.P/S,Status, Seq#,P-LDEV# M
VG01   oradb1(L)   (CL1-A , 1, 1-0 )91100174    1.SMPL -----,----- ---- -
VG01   oradb1(R)   (CL1-A , 1, 2-0 )91100174    2.SMPL -----,----- ---- -
VG01   oradb2(L)   (CL1-A , 1, 3-0 )91100174    3.SMPL -----,----- ---- -
VG01   oradb2(R)   (CL1-A , 1, 4-0 )91100174    4.SMPL -----,----- ---- -
VG01   oradb3(L)   (CL1-A , 1, 5-0 )91100174    5.SMPL -----,----- ---- -
VG01   oradb3(R)   (CL1-A , 1, 6-0 )91100174    6.SMPL -----,----- ---- -
```

2. Execute the `paircreate -m grp` command, then execute the `pairevtwait` command to verify that the status of each volume is PAIR.

```
C:\HORCM\etc>paircreate -g VG01 -vl -m grp
C:\HORCM\etc>pairevtwait -g VG01 -s pair -t 300 10
pairevtwait : Wait status done.
```

3. Execute the `pairdisplay` command to verify the pair status and the configuration. For example:

```
C:\HORCM\etc>pairdisplay -g VG01
Group  PairVol(L/R) (Port#,TID, LU-M) ,Seq#,LDEV#.P/S,Status, Seq#,P-LDEV# M
VG01   oradb1(L)   (CL1-A , 1, 1-0 )91100174    1.P-VOL PAIR,91100174    2 -
VG01   oradb1(R)   (CL1-A , 1, 2-0 )91100174    2.S-VOL PAIR,-----    1 -
VG01   oradb2(L)   (CL1-A , 1, 3-0 )91100174    3.P-VOL PAIR,91100174    4 -
VG01   oradb2(R)   (CL1-A , 1, 4-0 )91100174    4.S-VOL PAIR,-----    3 -
VG01   oradb3(L)   (CL1-A , 1, 5-0 )91100174    5.P-VOL PAIR,91100174    6 -
VG01   oradb3(R)   (CL1-A , 1, 6-0 )91100174    6.S-VOL PAIR,-----    5 -
```

Splitting pairs (pairsplit)

To split ShadowImage pairs

1. Execute the pairsplit command to split the ShadowImage pair in the PAIR status. In the following example, the group name in the configuration definition file is VG01

```
C:\HORCM\etc>pairsplit -g VG01
```

2. Execute the pairdisplay command to verify the pair status and the configuration.

```
C:\HORCM\etc>pairdisplay -g VG01
Group   PairVol(L/R) (Port#,TID,LU-M) ,Seq#,LDEV#.P/S,Status, Seq#,P-LDEV# M
VG01    oradb1(L)   (CL1-A , 1, 1-0 )91100174    1.P-VOL PSUS,91100174    2  -
VG01    oradb1(R)   (CL1-A , 1, 2-0 )91100174    2.S-VOL SSUS,-----    1  -
```

When it is required to split two or more S-VOLs included in a group at the same time and to assure that data of the same time are stored in the S-VOLs, use the CTG. In order to use the CTG, create a pair adding the -m grp option with the paircreate command.

Resynchronizing pairs (pairresync)

To resynchronize ShadowImage pairs

1. Execute the pairresync command to resynchronize the ShadowImage pair, then execute the pairevtwait command to verify that the status of each volume is PAIR. When using the -c option (copy pace), see the explanation in [Creating pairs \(paircreate\) on page A-30](#). For the following example, the group name in the configuration definition file is VG01.

```
C:\HORCM\etc>pairresync -g VG01 -c 15
C:\HORCM\etc>pairevtwait -g VG01 -s pair -t 300 10
pairevtwait : Wait status done.
```

2. Execute the pairdisplay command to verify the pair status and the configuration. For example:

```
C:\HORCM\etc>pairdisplay -g VG01
Group   PairVol(L/R) (Port#,TID, LU-M) ,Seq#,LDEV#.P/S,Status, Seq#,P-LDEV# M
VG01    oradb1(L)   (CL1-A , 1, 1-0 )91100174    1.P-VOL PAIR,91100174    2  -
VG01    oradb1(R)   (CL1-A , 1, 2-0 )91100174    2.S-VOL PAIR,-----    1  -
```

Releasing pairs (pairsplit -S)

To release the ShadowImage pair and change the status to SMPL

1. Execute the pairdisplay command to verify that the status of the ShadowImage pair is PAIR. In the following example, the group name in the configuration definition file is VG01.

```
C:\HORCM\etc>pairdisplay -g VG01
Group  PairVol(L/R) (Port#,TID, LU-M) ,Seq#,LDEV#.P/S,Status, Seq#,P-LDEV# M
VG01   oradb1(L)   (CL1-A , 1, 1-0 )91100174    1.P-VOL PAIR,91100174    2  -
VG01   oradb1(R)   (CL1-A , 1, 2-0 )91100174    2.S-VOL PAIR,-----    1  -
```

2. Execute the pairsplit (pairsplit -S) command to release the ShadowImage pair.

```
C:\HORCM\etc>pairsplit -g VG01 -S
```

3. Execute pairdisplay command to verify that the pair status changed to SMPL. For example:

```
C:\HORCM\etc>pairdisplay -g VG01
Group  PairVol(L/R) (Port#,TID, LU-M) ,Seq#,LDEV#.P/S,Status, Seq#,P-LDEV# M
VG01   oradb1(L)   (CL1-A , 1, 1-0 )91100174    1.SMPL -----,-----    -  -
VG01   oradb1(R)   (CL1-A , 1, 2-0 )91100174    2.SMPL -----,-----    -  -
```

Pair, group name differences in CCI and Navigator 2

Pairs and groups that were created using CCI will be displayed differently when status is confirmed in Navigator 2.

- Pairs created with CCI and defined in the configuration definition file display unnamed in Navigator 2.
- Groups defined in the configuration definition file are also different in Navigator 2.
- Pairs defined in a group on the configuration definition file using CCI are displayed in Navigator 2 as ungrouped.

For information about how to manage a group defined on the configuration definition file as a CTG, see the *Hitachi Unified Storage Command Control Interface Installation and Configuration Guide*.

I/O switching mode feature

This topic provides a description, specifications, and setup instructions for the I/O Switching Mode feature.

- ❑ [I/O Switching Mode feature operating conditions](#)
- ❑ [Specifications](#)
- ❑ [Recommendations](#)
- ❑ [Enabling I/O switching mode](#)
- ❑ [Recovery from a drive failure](#)

I/O Switching Mode feature operating conditions

When ShadowImage I/O Switching function operates during a drive double failure (triple failures for RAID 6), pair status is changed to "Failure (S-VOL Switch)" and the object of a host I/O instruction is switched from a P-VOL to an S-VOL. Since a report is made as if sent from a P-VOL, a job can continue without interruption. In a configuration where one P-VOL is paired with more than one S-VOLs, I/O is handed to a S-VOL only if the S-VOL which has the smallest logical number is in the Paired status.

This feature only operates under the following conditions:

- Turn on ShadowImage I/O Switching mode through Navigator 2. For details, see [Create a pair on page 5-6](#).
- It operates only when the pair is in the Paired status.
- When one P-VOL configures a pair with one or more S-VOLs, it operates. In the case, the S-VOL which has the smallest logical number becomes the target of I/O changing. If the status of the pair is not Paired, I/O changing is not performed even if the statuses of the other pairs is Paired.

Also, an S-VOL cannot be created newly from the P-VOL whose I/O is switched to the S-VOL.

- In the ShadowImage I/O Switching function, DP-VOLs created by



NOTE: If a P-VOL or S-VOL of ShadowImage exists in a RAID group in which a double (triple failures for RAID 6) drive failure occurred when the ShadowImage I/O Switching mode is turned on, all the volumes in the RAID group become unformatted irrespective of whether they are ShadowImage pairs or not.

Dynamic Provisioning can be used for a P-VOL or an S-VOL of ShadowImage.

Figure A-8 illustrates I/O Switching Mode.




Figure A-8: I/O Switching Mode function

The I/O Switching feature activates when a drive double failure (triple failures for RAID 6) occurs. At that time, the pair's status is changed to "Failure (S-VOL Switch)", and host read/write access is automatically transferred from a P-VOL to an S-VOL. When one P-VOL configures a pair with one or more S-VOLs, switch to the S-VOL that has the smallest volume number.



NOTE: When I/O Switching is activated, all LUs in the associated RAID group become unformatted, whether or not they are in a ShadowImage pair.

Specifications

Table A-3 shows specifications for the ShadowImage I/O Switching Mode function.

Table A-3: I/O Switching Mode specifications

Item	Specification
Preconditions for operation	- The ShadowImage I/O Switching mode must be turned on. - Pair status must be PAIR. - The ShadowImage I/O switching target pair and TrueCopy must not cascade.
Scope of application	All ShadowImage pairs that satisfy the preconditions for operation. In ShadowImage I/O Switching function, DP-VOLs can be used for a P-VOL or an S-VOL of ShadowImage.
Access to a P-VOL	Execution of host I/O continues after a drive failure because a report is sent to the host from the S-VOL as if from the P-VOL.
Access to an S-VOL	An I/O instruction issued to an S-VOL results in an error.
Display of the status	In Navigator 2: - When an object of a host I/O is switched to an S-VOL, the pairstatus is displayed as Failure (S-VOL Switch). - When executing the re-synchronizing instruction by Failure (S-VOL Switch), restoration operates and Reverse Synchronizing (S-VOL Switch) is displayed. In CCI (pairedisplay command): - Even when host I/O is switched to an S-VOL, the pair status is displayed as PSUE. However, when the <code>pairmon - allsnd -nowait</code> command is issued, the code (internal code of the pair status) is displayed as 0x08. - After an object of a host I/O is switched to an S-VOL, the <code>pairresync</code> command is executed, the pair status is displayed as RCPY.
Formatting	Quick formatting can only be performed when the pair status is PSUE (S-VOL Switch).
Notes	The <code>pairsplit</code>, <code>pairresync -restore</code> or <code>pairsplit - S</code> commands cannot be performed when the status is Failure (S-VOL Switch) or Reverse Synchronizing (S-VOL Switch).

Recommendations

- Locate P-VOLs and S-VOLs in respective RAID groups. When both are located in the same RAID group, they can become unformatted in the event of a drive failure.
- When a pair is in the Paired status, as it must be for the I/O Switching Mode, performance is lower than when the pair is Split. Hitachi recommends assigning a volume that uses SAS or SSD/FMD drives to an S-VOL to assure best performance results.

Enabling I/O switching mode

To use CLI, see [Setting the ShadowImage I/O switching mode on page A-11](#).

Use the following procedure to enable I/O switching mode which displays a Navigator 2 applet screen. If your system does not display the screen (it takes a minute or two to appear), see the GUI online Help for Advanced Settings.

To enable I/O Switching Mode

1. In Navigator 2, select the subsystem in which ShadowImage is to be operated, then click **Show & Configure disk array**.
2. In the tree view, select the **Advanced Settings** icon.
3. Click **Open Advanced Settings**. The Array Unit screen displays, as shown in [Figure A-9](#). This may take a few minutes.

If you have problems with screen, see the following section.




Figure A-9: Array Unit screen

4. Select **Configuration Settings**, then click **Set**. The Configuration Settings screen displays.
5. Click the **System Parameter** tab.
6. Click the **ShadowImage I/O Switch Mode** check box, then click **Apply**.
7. On the confirmation message, click **OK**.
8. Click **Close** on the **Configuration Settings** page.
9. Click **Close** on the subsequent message screen.



NOTE: When disabling I/O Switching Mode, pair statuses must be other than Failure (S-VOL Switch) and Synchronizing (S-VOL Switch).

About the array unit screen

This screen is an applet connected to the SNM2 Server. After 20 minutes elapses while displaying this applet screen, automatic logoff occurs. Therefore, when your operation is completed, close the screen.

If the applet screen does not display, the login to the SNM2 Server may have failed. In this case, the applet screen cannot be displayed again. The code: 0x000000000000b045 or "DMEG800003: The error occurred in connecting RMI server." is displayed on the applet screen. Take the following actions.

- Close the Web browser, stop the SNM2 Server, restart it, then navigate to the disk array.
- Close the Web browser; confirm the SNM2 Server is started. If it has stopped, start it and display the screen of the disk array that you want to operate.
- Return to the disk array screen after 20 minutes elapsed and display the screen of the disk array you want to operate.

Recovery from a drive failure

When the I/O Switching Mode feature is used, recovery from a drive failure on which the P-VOL is located can be undertaken during the newly-established host read/write operations to the S-VOL. This section provides the basic procedure for recovery.

To recover from a drive failure after I/O Switching Mode is activated

1. After a drive double failure occurs (triple failure for RAID 6) in a P-VOL and I/O from the host is transferred via the I/O Switching Mode feature to the S-VOL, have the P-VOL drives replaced.
2. When one P-VOL configures a pair with one or more S-VOLs, delete the pairs other than the Paired in which the I/O switching to the S-VOL was operated.
3. When a drive double failure occurs (triple failures for RAID 6) in a P-VOL that is DP-VOL, re initialize the DP pool.
4. When the drives have been replaced, perform quick formatting of the P-VOL.
5. Perform a reverse-resync of the pair, copying S-VOL data to the P-VOL. Performance is lowered during the reverse-resync, but host I/O can be continued. When resynchronization is completed, the pair status becomes Paired.
6. When one P-VOL configures a pair with one or more S-VOLs, create a pair in the original pair configuration.

Copy-on-Write Snapshot reference information

This appendix includes:

- ❑ [Snapshot specifications](#)
- ❑ [Operations using CLI](#)
- ❑ [Operations using CCI](#)
- ❑ [Setting the command device for raidcom command](#)
- ❑ [Using Snapshot with Cache Partition Manager](#)

Snapshot specifications

Table B-1 lists external specifications for Snapshot.

Table B-1: General Specifications

Item	Specification
Host interface	Fibre Channel or iSCSI.
Number of pairs	HUS 150/HUS 130/HUS 110: 100,000 (maximum) Note: When one P-VOL pairs with 1,024 V-VOLs, the number of pairs is 1,024.
Cache Memory	HUS 150: 8, 16 GB per controller HUS 130: 8 GB per controller HUS 110: 4 GB per controller
Command devices	Required for CCI. - Maximum: 128 per disk array - Volume size: 33MB or greater
Unit of pair management	Volumes are the target of Snapshot pairs, and are managed per volume.
Number of V-VOLs per P-VOL	1: 1,024
RAID level	RAID 1+0 (2D+2D to 8D+8D) RAID 5 (2D+1P to 15D+1P) RAID 6 (2D+2P to 28D+2P) RAID 1 (1D+1D)
Combination of RAID levels	All combinations are supported. The number of data disks may be different.
Volume size	Volumes for the V-VOL must be equal in size to the P-VOL. The max volume size is 128 TB.
Drive types for the P-VOL and data pool	If the drive types are supported by the disk array, they can be set for the P-VOL and data pool. SAS drives, SAS7.2K drives, or SSD/FMD drives are recommended. A DP-VOL cannot be a P-VOL.
Consistency Group (CTG) number	Max 1,024/array HUS 150/HUS 130: 2,046 pairs/CTG (maximum) HUS 110: 1,022 pairs/CTG (maximum)
MU Number	This is used for specifying a pair in CCI. For SnapShot pairs, the value from 0 to 1032 can be specified.
Consumed capacity of DP pool	Snapshot stores their replication data and management information into a DP pool. For details , see DP pool consumption on page 9-7 .
Differential management	When the status of P-VOL and V-VOL is Split, write operations received individually will be managed as the differential data of the P-VOL and the V-VOL. When one P-VOL configures a pair with more than one V-VOL, the difference is managed for each pair.
Data pool	HUS 150/HUS 130: Max 64/array (DP pool number is 0 to 63) HUS 110: Max 50/array (DP pool number is 0 to 49)

Table B-1: General Specifications (Continued)

Item	Specification
Access to the DP pool from a host	The DP pool is not recognizable from the host.
Expansion of DP pool capacity	Expansion of the DP is possible. The capacity is expanded through an addition of RAID groups to the DP pool. The extension of a DP pool can be made while a pair that uses the DP pool is created. However, the RAID group with different drive types cannot be mixed.
Max supported capacity of P-VOL and data pool	The supported capacity of Snapshot is limited based on P-VOL and data pool size. For details, see Requirements and recommendations for Snapshot Volumes on page 9-14 .
Reduction of data pool capacity	Possible only when all the pairs that use the data pool have been deleted.
Unifying, growing and shrinking of a volume assigned to a data pool	No.
Formatting, deleting, growing, shrinking of a volume in a pair: Deleting RAID group in a pair:	No.
Pairing with an expanded volume	Only P-VOL can be expanded
Formatting or expanding V-VOL	No.
Pairing with unified volume	When the disk array firmware version is less than 0920/B the capacity of each volume before the unification must be 1 GB or larger.
Deletion of the V-VOL	Only possible when P-VOL and V-VOL are in simplex status and not paired.
Swap V-VOL for P-VOL	No.
Load balancing	The load balancing works for the P-VOL, but it does not work for the V-VOL.
Restriction during RAID group expansion	A RAID group with a Snapshot P-VOL or V-VOL can be expanded only when the pair status is Simplex or Paired.
Initial copy when creating a pair	Not necessary.
Re-synchronizing	Not necessary.
Restoration (re-synchronizing V-VOL to P-VOL)	Possible.
Pair deleting	Possible (When pair deleted, V-VOL data is annulled).
Pair splitting	Always splitting.
Concurrent use with ShadowImage	Snapshot and ShadowImage can be used at the same time on the same disk array. If Snapshot is used concurrently with ShadowImage, CTGs limited to 1,024.

Table B-1: General Specifications (Continued)

Item	Specification
Snapshot use with expanded volumes	Yes.
Concurrent use with TrueCopy and TCE	TrueCopy can be cascaded with Snapshot. TCE can be cascaded with a Snapshot P-VOL. See Cascading Snapshot with TrueCopy Remote on page 27-21 for more information.
Concurrent use of Dynamic Provisioning	Because Snapshot uses the DP pools to work, Dynamic Provisioning is necessary. A DP-VOL can be used for the P-VOL of Snapshot. See Concurrent use of Dynamic Provisioning on page 9-26 .
Concurrent use of Dynamic Tiering	The DP volume of the DP pool whose tier mode is enabled in Dynamic Tiering can be used as a P-VOL of SnapShot. Furthermore, the DP pool whose tier mode is enabled in Dynamic Tiering can be specified as the replication data DP pool and the management area DP pool. For details, see Concurrent use of Dynamic Tiering on page 9-28 .
Concurrent use with LUN Manager	Yes.
Concurrent use with Password Protection	Yes.
Concurrent use of Volume Migration	Yes, however a Volume Migration P-VOL, S-VOL, and Reserved volume cannot be specified as a Snapshot P-VOL.
Concurrent use of SNMP Agent Support Function	Available. The SNMP Agent Support Function notifies users of a event happening when the pair status changes to Threshold Over as the usage rate of the DP pool exceeds the Replication Depletion Alert threshold value as well as when the pair status changes to Failure as the usage rate exceeds the Replication Data Released threshold or some failures occur on Snapshot
Concurrent use of Cache Residency Manager	Yes, however the volume specified for Cache Residency (volume cache residence) cannot be used as a P-VOL, V-VOL, or data pool.
Concurrent use of Cache Partition Manager	Yes. Cache partition information is initialized when Snapshot is installed. Data pool volume segment size must be the default size (16 kB) or less. See Setting the command device for raidcom command on page B-36 .
Concurrent use of SNMP Agent	Yes. Traps are sent when a failure occurs. Pair status changes to Threshold over or Failure.
Concurrent use of Data Retention Utility	Yes, but note following: • When S-VOL Disable is set for an volume, the volume cannot be used in a Snapshot pair. • When S-VOL Disable is set for a volume that is already a V-VOL, no suppression of the pair takes place, unless the pair status is split. • When S-VOL Disable is set for a P-VOL, restoration of the P-VOL is suppressed.

Table B-1: General Specifications (Continued)

Item	Specification
Concurrent use of Power Saving/Power Saving Plus	Yes. However, when a P-VOL is included in a RAID group in which Power Saving/Power Saving Plus is enabled, the only Snapshot pair operation that can be performed are the pair split and the pair release.
Potential effect caused by a P-VOL failure	The V-VOL relies on P-VOL data, therefore a P-VOL failure results in a V-VOL failure also.
Potential effect caused by installation of the Snapshot function	When the firmware version of the disk array is less than 0920/B, reboot is required to acquire data pool resources.
Requirement for Snapshot installation	Reboot is required to acquire pool resources.
Potential effect at the time of one controller blockade	One controller blockade does not affect the V-VOL data.
Treatment when exceeding the replication threshold value of DP pool usage rate.	Pair status to be changed. Returns warning to CCI. Also, the E-mail Alert Function and SNMP Agent Support Function will work to notify you of the event happening. When the usage rate of the DP pool exceeds the Replication Data Released threshold, the pair status changes to Failure. (The threshold value can be set per user)
Action to be taken when the limit of usable POOL capacity is exceeded	When data pool usage is 100%, statuses of all the V-VOLs using the POOL become failure.
Reduction of memory	Memory cannot be reduced when Snapshot, ShadowImage, TrueCopy, or TCE are enabled. Reduce memory after disabling the functions.

Operations using CLI

This section describes Storage Navigator 2 Command Line Interface (CLI) procedures for Snapshot enabling, configuration and copy operations.

- ❑ [Installing and uninstalling Snapshot](#)
- ❑ [Enabling or disabling Snapshot](#)
- ❑ [Operations for Snapshot configuration](#)
- ❑ [Setting the system tuning parameter \(optional\)](#)
- ❑ [Performing Snapshot operations](#)
- ❑ [Sample back up script for Windows](#)



NOTE: For additional information on the commands and their options used in this appendix, see the *Hitachi Unified Storage Command Line Interface Reference Guide*

Installing and uninstalling Snapshot

Installation instructions are provided for Navigator 2 GUI.

Important prerequisite information

- A key code or key file is required to install or uninstall. If you do not have the key file or code, you can obtain it from the download page on the HDS Support Portal, <https://portal.hds.com>
- Before installing or uninstalling Snapshot, verify that the Storage system is operating in a normal state. Installation/un-installation cannot be performed if a failure has occurred.
- Hitachi recommends changing TrueCopy or TCE pair status to Split before installing Snapshot on the remote array.
- When Snapshot is used together with TCE, restarting the array by the function that is installed after the function that was installed first is not required. The restart was done by the function that was installed first in order to ensure the resource for the data pool in the cache memory.



NOTE: If you install or uninstall Snapshot because of a spin-down instruction when Power Saving, the spin down may fail if the instruction is received immediately after the array restarts. If the spin-down fails, perform the spin-down again. Check that the spin-down instruction has not been issued or has been completed (no RAID group in the Power Saving Status of **Normal(Command Monitoring)** exists) before installing or uninstalling Snapshot.

Installing Snapshot

Snapshot cannot usually be selected (locked) when first using the array. To make Snapshot available, you must install the Snapshot and make its function selectable (unlocked).

To install Snapshot

1. From the command prompt, register the array in which Snapshot is to be installed, then connect to the array.
2. Execute the auopt command to install Snapshot. For example:

```
% auopt -unit array-name -lock off -keycode manual-attached-keycode
Are you sure you want to unlock the option? (y/n [n]): y
The option is unlocked.
A DP pool is required to use the installed function.
Create a DP pool before you use the function.
```

3. Execute the auopt command to confirm whether Snapshot has been installed. For example:

```
% auopt -unit array-name -refer
```

Option	Name	Type	Term	Status
	Reconfigure Memory Status			
SNAPSHOT		Permanent	---	Enable
	N/A			

```
%
```

Snapshot is installed and Status is **Enable**.

Snapshot installation is complete.

Snapshot requires the DP pool of Hitachi Dynamic Provisioning (HDP). If HDP is not installed, install HDP.

Uninstalling Snapshot

Once uninstalled, Snapshot cannot be used again until it is installed using the key code or key file. Once uninstalled, Snapshot cannot be used (locked) until it is again unlocked using the key code or key file.

Prerequisites

- The key code or key file provided with the optional feature is required to uninstall Snapshot.
- Snapshot pairs must be released and their status returned to Simplex.
- The replication data is deleted after the pair deletion is completed. The replication data deletion may be operated in the background at the time of the pair deletion. Check that the DP pool capacity is recovered after the pair deletion. If it is recovered, the replication data has been deleted.
- All Snapshot volumes (V-VOL) must be deleted.
- For additional prerequisites, see [Important prerequisite information on page B-7](#).



NOTE: If you uninstall Snapshot because of a spin-down instruction when Power Saving, the spin down may fail if the instruction is received immediately after the array restarts. If the spin-down fails, perform the spin-down again. Check that the spin-down instruction has not been issued or has been completed (no RAID group in the Power Saving Status of **Normal(Command Monitoring)** exists) before uninstalling Snapshot.

To uninstall Snapshot

1. From the command prompt, register the array in which the Snapshot is to be uninstalled, then connect to the array.
2. Execute the auopt command to uninstall Snapshot. For example:

```
% auopt -unit array-name -lock on -keycode manual-attached-keycode
Are you sure you want to lock the option? (y/n [n]): y
The option is locked.
```

3. Execute the auopt command to confirm whether Snapshot has been uninstalled. For example:

```
% auopt -unit array-name -refer
DMEC002015: No information displayed.
%
```

Snapshot uninstall is complete.

Enabling or disabling Snapshot

Snapshot is bundled with the Adaptable array. You must enable it before using.

Prerequisites

The following conditions must be satisfied in order to disable Snapshot:

- All Snapshot pairs must be released (that is, the status of all volumes are Simplex).
- The replication data is deleted after the pair deletion is completed. The replication data deletion may be operated in the background at the time of the pair deletion. Check that the DP pool capacity is recovered after the pair deletion. If it is recovered, the replication data has been deleted.
- All Snapshot volumes (V-VOL) must be deleted.



NOTE: If you enable/disable Snapshot because of a spin-down instruction when Power Saving, the spin down may fail if the instruction is received immediately after the array restarts. If the spin-down fails, perform the spin-down again. Check that the spin-down instruction has not been issued or has been completed (no RAID group in the Power Saving Status of **Normal(Command Monitoring)** exists) before disabling or enabling Snapshot.

To enable or disable Snapshot Data

1. To enable Snapshot using CLI, from the command prompt, register the array in which the status of the feature is to be changed, then connect to the array.
2. Execute the auopt to change the status (enable or disable).

Following is an example of changing the status from enable to disable. If you want to change the status from disable to enable, enter enable after the -st option.

```
% auopt -unit array-name -option SNAPSHOT -st disable
Are you sure you want to disable the option? (y/n [n]): y
The option has been set successfully.
%
```

3. Execute auopt to confirm whether the status has been changed. For example:

```
% auopt -unit array-name -refer
Option Name          Type      Term      Status
Reconfigure Memory Status
SNAPSHOT              Permanent ---      Disable
N/A
%
```

Snapshot Enable/Disable is complete.

Operations for Snapshot configuration

Setting the DP pool

For instructions to set a DP pool, refer to the *Hitachi Unified Storage Dynamic Provisioning Configuration Guide*.

Setting the replication threshold (optional)

To set the Depletion Alert and/or the Replication Data Released of replication threshold:

1. From the command prompt, execute the `audpool` command, change the Depletion Alert and/or the Replication Data Released of replication threshold.

Example: This example shows changing the Depletion Alert threshold.

```
% audpool -unit array-name -chg -dppoolno 0 -repdepletion_alert 50
Are you sure you want to change the DP pool attribute? (y/n [n]): y
DP pool attribute changed successfully.
%
%
```

2. Execute the `audpool` command to confirm the DP pool attribute.

```

% audppool -unit array-name -refer -detail -dppoolno 0 -t
DP Pool : 0
RAID Level          : 6(6D+2P)
Page Size           : 32MB
Stripe Size         : 256KB
Type                : SAS
Status              : Normal
Reconstruction Progress : N/A
Capacity
  Total Capacity    : 8.9 TB
  Consumed Capacity
    Total           : 2.2 TB
    User Data       : 0.7 TB
    Replication Data : 0.4 TB
    Management Area : 0.5 TB
  Needing Preparation Capacity : 0.0 TB
DP Pool Consumed Capacity Alert
  Early Alert       : 40%
  Depletion Alert   : 50%
  Notifications Active : Enable
Over Provisioning Threshold
  Warning           : 100%
  Limit             : 130%
  Notifications Active : Enable
Replication Threshold
  Replication Depletion Alert : 50%
  Replication Data Released   : 95%
Defined LU Count          : 0
DP RAID Group
  DP RAID Group  RAID Level  Capacity  Capacity  Percent
                49    6(6D+2P)    8.9 TB    2.2 TB    24%
Drive Configuration
  DP RAID Group  RAID Level Unit  HDU  Type  Capacity  Status
                49    6(6D+2P)   0    0  SAS    300GB  Standby
                49    6(6D+2P)   0    1  SAS    300GB  Standby
                :
                :
Logical Unit
  Consumed          Stripe Cache      Pair Cache
Number
  LU Capacity      Capacity Consumed % Size  Partition Partition
Status of Paths
%

```

Setting the V-VOL (optional)

Since the Snapshot volume (V-VOL) is automatically created at the time of the pair creation, it is not always necessary to set V-VOL. However, you may create the V-VOL before the pair creation and perform the pair creation with the created V-VOL.

Prerequisites

- When deleting the V-VOL, the pair state must be Simplex.

To set the V-VOL:

1. From the command prompt, register the array to which you want to set the V-VOL, then connect to the array.
2. Execute the `aureplicationvvol` command create a V-VOL. For example:

```
% aureplicationvvol -unit array-name -add -lu 1000 -size 1
Are you sure you want to create the Snapshot logical unit 1000?
(y/n[n]): y
The Snapshot logical unit has been successfully created.
%
```

3. To delete an existing Snapshot logical unit, refer to the following example of deleting Snapshot logical unit 1000. When deleting the V-VOL, the pair state must be Simplex.

```
% aureplicationvvol -unit array-name -rm -lu 1000
Are you sure you want to delete the Snapshot logical unit 1000?
(y/n[n]): y
The Snapshot logical unit has been successfully deleted.
%
```

Setting the system tuning parameter (optional)

This setting limits the number of times processing is executed for flushing the dirty data in the cache to the drive at the same time.

To set the Dirty Data Flush Number Limit of a system tuning parameters:

1. From the command prompt, register the array on which you want to set a system tuning parameters and connect to that array.
2. Execute the `ausystuning` command to set the system tuning parameters.

Example:

```
% ausystuning -unit array-name -set -dtynumlimit enable
Are you sure you want to set the system tuning parameter? (y/n [n]): y
Changing Dirty Data Flush Number Limit may have performance impact when local re
placation is enabled and time out may occur if I/O load is heavy. Please change
the setting when host I/O load is light.
Do you want to continue processing? (y/n [n]): y
The system tuning parameter has been set successfully.
%
```

Performing Snapshot operations

The `aureplicationlocal` command operates Snapshot pair. To refer the `aureplicationlocal` command and its options, type in `aureplicationlocal -help` at the command prompt.

Creating Snapshot pairs using CLI



NOTE: When not specifying the pair name (`-pairname pair_name`), the following are assigned automatically:
When specifying the logical unit number the V-VOL: `SS_LUXXXX_LUYYYY`
When not specifying the logical unit number the V-VOL:
`SS_LUXXXX_LUNONE_nnnnnnnnnnnnnnn`
XXXX: LU number of P-VOL (four digits with 0)
YYYY: LU number of V-VOL (four digits with 0)
nnnnnnnnnnnnnnnn: 14-digit optional numbers
("Year/Month/Day/hour/minute/second/millisecond" at the time of the pair creation)

To create Snapshot pairs using CLI:

1. From the command prompt, register the array to which you want to create the Snapshot pair, then connect to the array.
2. Execute the `aureplicationlocal` command to create a pair.
First, display the volumes to be assigned to a P-VOL, and then create a pair. Refer to the following example:

```
% aureplicationlocal -unit array-name -ss -availablelist
-pvol
Available Logical Units
  LUN  Capacity RAID Group DP Pool RAID Level  Type Status
  100   30.0 GB           0   N/A  6( 9D+2P) SAS  Normal
  200   35.0 GB           0   N/A  6( 9D+2P) SAS  Normal
%
% aureplicationlocal -unit array-name -ss -create -pvol 200
-svol 1001 -compsplit
Are you sure you want to create pair "SS_LU0200_LU1001"?
(y/n[n]): y
The pair has been created successfully.
%
```

3. Execute the `aureplicationlocal` command to verify that the pair has been created. Refer to the following example.

```
% aureplicationlocal -unit array-name -ss -refer
Pair name          LUN Pair LUN Status
Copy Type      Group
SS_LU0200_LU1001      200 1001   Split(100%)
Snapshot        ---:Ungrouped
%
```

The Snapshot pair is created.

Splitting Snapshot Pairs

To split the Snapshot pairs:

1. From the command prompt, register the array to which you want to split the Snapshot pair, then connect to the array.
2. Execute the `aureplicationlocal` command to split the pair.

In the following example, the P-VOL LUN is 200 and the S-VOL LUN is 1001.

```
% aureplicationlocal -unit array-name -ss -resync -pvol 200
-svol 1001
Are you sure you want to split pair?
(y/n[n]): y
The split of pair has been required.
%
```

3. Execute `aureplicationlocal` to verify the pair.

```
% aureplicationlocal -unit array-name -ss -refer
Pair name          LUN Pair LUN Status
Copy Type      Group
SS_LU0200_LU1001      200 1001   Split(100%)
Snapshot        ---:Ungrouped
%
```

The Snapshot pair is split.

Re-synchronizing Snapshot Pairs

To re-synchronize Snapshot pairs:

1. From the command prompt, register the array to which you want to re-synchronize the Snapshot pair, then connect to the array.
2. Execute the `aureplicationlocal` command to re-synchronize the pair.

In the following example, the P-VOL LUN is 200 and the S-VOL LUN is 1001.

```
% aureplicationlocal -unit array-name -ss -resync -pvol 200 -
svol 1001
Are you sure you want to re-synchronize pair? (y/n [n]): y
The re-synchronizing of pair has been required.
%
```

3. Execute `aureplicationlocal` to verify the pair.

```
% aureplicationlocal -unit array-name -ss -refer
Pair name                LUN  Pair LUN  Status
Copy Type              Group
SS_LU0200_LU1001        200    1001  Synchronizing(
40%)
Snapshot                ---:Ungrouped
%
```

The Snapshot pair is resynchronized.

Restoring V-VOL to P-VOL using CLI

To restore the V-VOL to the P-VOL using CLI:

1. From the command prompt, register the array to which you want to restore the Snapshot pair, then connect to the array.
2. Execute the `aureplicationlocal` command restore the pair.

First, display the pair status, and then restore the pair. Refer to the following example.

```
% aureplicationlocal -unit array-name -ss -refer
Pair name          LUN Pair LUN Status
Copy Type      Group
SS_LU0200_LU1001      200 1001    Split(100%)
Snapshot        ---:Ungrouped
%
% aureplicationlocal -unit array-name -ss -restore -pvol 200
-svol 1001
Are you sure you want to restore pair?
(y/n[n]): y
The pair has been restored successfully.
%
```

3. Execute `aureplicationlocal` to restore the pair. Refer to the following example.

```
% aureplicationlocal -unit array-name -ss -refer
Pair name          LUN Pair LUN Status
Copy Type      Group
SS_LU0200_LU1001      200 1001    Paired( 40%)
Snapshot        ---:Ungrouped
%
```

V-VOL to P-VOL is restored.

Deleting Snapshot pairs

To delete the Snapshot pair and change the status to Simplex using CLI:

1. From the command prompt, register the array to which you want to delete the Snapshot pair, then connect to the array.
2. Execute the `aureplicationlocal` command to delete the pair. Refer to the following example.

```
% aureplicationlocal -unit array-name -ss -simplex -pvol 200
-svol 1001
Are you sure you want to release pair?
(y/n[n]): y
The pair has been released successfully.
%
```

3. Execute `aureplicationlocal` to confirm the deleted pair. Refer to the following example.

```
% aureplicationlocal -unit array-name -ss -refer
DMEC002015: No information is displayed.
%
```

The Snapshot pair is deleted.

Changing pair information

You can change the pair name, assignment/deprivation of volume number to secondary volume, group name, and/or copy pace.

1. From the command prompt, register the array to which you want to change the Snapshot pair information, then connect to the array.
2. Execute the `aureplicationlocal` command change the pair information. This is an example of changing a copy pace.

```
% aureplicationlocal -unit array-name -ss -chg -pace slow -pvol 200
-svol 1001
Are you sure you want to change pair information?
(y/n[n]): y
The pair information has been changed successfully.
%
```

3. Execute the `aureplicationlocal` command to assign the volume number to a secondary volume.

```
% aureplicationlocal -unit array-name -ss -chg
-pairname SS_LU2000_LUNNONE_20110320180000 -gno
0 -svol 2002
Are you sure you want to change pair information? (y/n [n]): y
The pair information has been changed successfully.
%
```

4. Execute the `aureplicationlocal` command deprive the volume number of the secondary volume.

```
% aureplicationlocal -unit array-name -ss -chg
-pairname SS_LU2000_LU_2002 -gno 0 -svol notallocate
Are you sure you want to change pair information? (y/n [n]): y
The pair information has been changed successfully.
%
```

The Snapshot pair information is changed.

Creating multiple Snapshot pairs that belong to a group using CLI

To create multiple Snapshot pairs that belong to a group using CLI:

1. Create the first pair specifying an unused group number for the new group with the `-gno` option. Refer to the following example.

```
% aureplicationlocal -unit array-name -ss -create -pvol 200
-svol 1001 -gno 20
Are you sure you want to create pair "SS_LU0200_LU1001"?
(y/n[n]): y
The pair has been created successfully.
%
```

The new group has been created, and in this group, the new pair has been created too.

2. Add the name of the group by specifying the group name with the `-newgname` option to change the pair information. Refer to the following example.


```
% aureplicationlocal -unit array-name -ss -chg -gno 20 -
newgname group-name
Are you sure you want to change pair information?
(y/n[n]): y
The pair information has been changed successfully.
%
```

3. Create the next pair that belongs to the created group specifying the number of the created group with `-gno` option.
Snapshot pairs that share the same P-VOL must use same Data Pool.
4. By repeating the step 3, the multiple pairs that belong to the same group can be created.

Sample back up script for Windows

This section provides sample script for backing a volume on Windows Server.

```
echo off
REM Specify the registered name of the arrays
set UNITNAME=Array1
REM Specify the group name (Specify "Ungroup" if the pair
doesn't belong to any group)
set G_NAME=Ungrouped
REM Specify the pair name
set P_NAME=SS_LU0001_LU0002
REM Specify the directory path that is mount point of P-VOL
and V-VOL
set MAINDIR=C:\main
set BACKUPDIR=C:\backup
REM Specify GUID of P-VOL and V-VOL
PVOL_GUID=xxxxxxxx-xxxx-xxxx-xxxx-xxxxxxxxxxxxxx
SVOL_GUID=yyyyyyyyy-yyyy-yyyy-yyyy-yyyyyyyyyyyyyy

REM Unmounting the V-VOL
pairdisplay -x umount %BACKUPDIR%
REM Re-synchronizing pair (Updating the backup data)
aureplicationlocal -unit %UNITNAME% -ss -resync -pairname
%P_NAME% -gname %G_NAME%
aureplicationmon -unit %UNITNAME% -evwait -ss -pairname
%P_NAME% -gname %G_NAME% -st paired -pvol

REM Unmounting the P-VOL
pairdisplay -x umount %MAINDIR%
REM Splitting pair (Determine the backup data)
aureplicationlocal -unit %UNITNAME% -ss -split -pairname
%P_NAME% -gname %G_NAME%
aureplicationmon -unit %UNITNAME% -evwait -ss -pairname
%P_NAME% -gname %G_NAME% -st split -pvol
REM Mounting the P-VOL
pairdisplay -x mount %MAINDIR% Volume{%PVOL_GUID%}

REM Mounting the V-VOL
pairdisplay -x mount %BACKUPDIR% Volume{%SVOL_GUID%}
< The procedure of data copy from C:\backup to backup
appliance>
```



NOTE: In case Windows Server is used, the CCI mount command must be used when mounting/un-mounting a volume. Also, the GUID, which is displayed by the mountvol command, is needed as an argument to use mount command of CCI.

Operations using CCI

This topic describes basic CCI procedures for setting up and performing Snapshot operations.

- ❑ [Setting up CCI](#)
- ❑ [Performing Snapshot operations](#)
- ❑ [Pair and group name differences in CCI and Navigator 2](#)

Setting up CCI

The following sub-sections describe necessary set up procedures for CCI for Snapshot.

Setting the command device

The Command Device is a dedicated logical volume on the disk array that functions as the interface to CCI software.

Prerequisite

The Command Device must be defined in the HORCM_CMD section of the configuration definition file for the CCI instance on the attached host. Up to 128 Command Devices can be designated for the disk array.

The command devices and LU mapping setting use Navigator 2.



- When pair operation using CCI, the P-VOL and V-VOL, whose mapping information is not set for the port that has been set in the configuration definition file, cannot be paired. When you do not want them recognized by a host, map to a port that is not connected to the host or to a host group in which no host has been registered using LUN Manager.
- Volumes set for Command Devices must be recognized by the host. The Command Device volume size must be greater than or equal to 33 MB.

To set up a command devices

1. From the command prompt, register the disk array to which you want to set the command device. Connect to the disk array.
2. Execute the aucmddev command to set a command device. First, display the volumes to be assignable command device, and later set a command device. When you want to use the protection function of CCI, enter enable following the -dev option.

The following example specifies LU 200 for command device 1.

```
% aucmddev -unit disk array-name -availablelist
Available Logical Units
  LUN  Capacity RAID Group DP Pool RAID Level  Type Status
    2   35.0 MB      0    N/A   6( 9D+2P) SAS  Normal
    3   35.0 MB      0    N/A   6( 9D+2P) SAS  Normal
%
% aucmddev -unit disk array-name -set -dev 1 200
Are you sure you want to set the command devices?
(y/n [n]): y
The command devices have been set successfully.
%
```

3. Execute the aucmddev command to verify that the command device has been set. For example:

```
% aucmddev -unit disk array-name -refer
Command Device LUN RAID Manager Protect
1                200 Disable
%
```



NOTE: To set the alternate command device function or to avoid data loss and disk array downtime, designate two or more command devices. For details on alternate Command Device function, refer to the *Hitachi Unified Storage Command Control Interface Installation and Configuration Guide*.

4. The following example releases a command device:

```
% aucmddev -unit disk array-name -rm -dev 1
Are you sure you want to release the command devices?
(y/n [n]): y
This operation may cause the CCI, which is accessing to this
command device, to freeze.
Please make sure to stop the CCI, which is accessing to this
command device, before performing this operation.
Are you sure you want to release the command devices? (y/n [n]): y
The specified command device will be released.
Are you sure you want to execute? (y/n [n]): y
The command devices have been released successfully.
%
```

5. To change an already set command device, release the command device, then change the volume number. The following example specifies LU 201 for command device 1.

```
% aucmddev -unit disk array-name -set -dev 1 201
Are you sure you want to set the command devices?
(y/n [n]): y
The command devices have been set successfully.
%
```

Setting LU Mapping information

The Mapping Information is specified using Navigator 2. If no mapping is done on the P-VOLs and V-VOLs specified in the CCI configuration files when the Mapping mode is enabled, which means the hosts cannot recognize the P-VOLs and V-VOLs, no pair operation can be done on the P-VOLs and V-VOLs. Use LUN Manager if you want to hide the volumes from the hosts.

Prerequisite

For iSCSI, use the autargetmap command instead of the auhgmmap command.

To set up LU Mapping

1. From the command prompt, register the disk array to which you want to set the LU Mapping, then connect to the disk array.
2. Execute the auhgmmap command to set the LU Mapping. The following is an example of setting LU 0 in the disk array to be recognized as 6 by the host. The port is connected via target group 0 of port 0A on controller 0.

```
% auhgmap -unit disk array-name -add 0 A 0 6 0
Are you sure you want to add the mapping information?
(y/n [n]): y
The mapping information has been set successfully.
%
```

3. Execute the auhgmap command to verify that the LU Mapping is set. For example:

```
% auhgmap -unit disk array-name -refer
Mapping mode = ON
Port  Group  H-LUN
   0A      0      6      0
%
```

Defining the configuration definition file

The configuration definition file is a text file created and/or edited using any standard text editor. It can be defined from the PC where CCI software is installed. It is required to make CCI operational.

A sample configuration definition file, HORCM_CONF, is included with the CCI software. It should be used as the basis for creating your configuration definition file(s). The system administrator should copy the sample file, set the necessary parameters in the copied file, and place the copied file in the proper directory. For details on configuration definition file, refer to the *Hitachi Unified Storage Command Control Interface Installation and Configuration Guide*.

The configuration definition file can be automatically created using the mkconf command tool. For details on the mkconf command, refer to the *Hitachi Unified Storage Command Control Interface Installation and Configuration Guide*. However, the parameters, such as poll(10ms) must be set manually (see step 4 below).

Example for manually defining the configuration definition file

The following describes an example for manually defining the configuration definition file when the system is configured with two instances within the same Windows host.

The P-VOL and V-VOLs are conceptually diagrammed in [Figure B-1](#).




Figure B-1: P-VOL and V-VOLs

1. On the host where CCI is installed, verify that CCI is not running. If CCI is running, shut it down using the horcmshutdown command.

2. In the command prompt, make two copies of the sample file (horcm.conf). For example:

```
c:\HORCM\etc> copy \HORCM\etc\horcm.conf \WINDOWS\horcm0.conf
c:\HORCM\etc> copy \HORCM\etc\horcm.conf \WINDOWS\horcm1.conf
```

3. Open horcm0.conf using the text editor.
4. In the **HORCM_MON** section, set the necessary parameters.

Important: A value more than or equal to 6000 must be set for poll(10ms). Specifying the value incorrectly may cause resource contention in the internal process, resulting in the process temporarily suspending and pausing the internal processing of the disk array. See the *Hitachi Unified Storage Command Control Interface Installation and Configuration Guide* for more information.

5. In the **HORCM_CMD** section, specify the physical drive (command device) on the disk array. For example:




Figure B-2: Horcm0.conf example (P-VOL; S-VOL=1: 3)




Figure B-3: Horcm0.conf example (cascading ShadowImage S-VOL with Snapshot P-VOL)

6. Set the necessary parameters in the **HORCM_LDEV** section, then in the **HORCM_INST** section.
7. Save the configuration definition file.
8. Repeat Steps 3 to 7 for the horcm1.conf file. Example:




Figure B-4: Horcm1.conf example (P-VOL: S-VOL=1: 3)




Figure B-5: Horcm1.conf example (cascading ShadowImage S-VOL with Snapshot P-VOL)

9. Enter the following example lines in the command prompt to verify the connection between CCI and the disk array:

```
C:\>cd HORCM\etc

C:\HORCM\etc>echo hd1-7 | .\in RAID
Harddisk 1 -> [ST] CL1-A Ser =91100123 LDEV = 200 [HITACHI ] [DF600F-CM      ]
Harddisk 2 -> [ST] CL1-A Ser =91100123 LDEV =   2 [HITACHI ] [DF600F ]
                HORC = SMPL  HOMRCF[MU#0 = SMPL MU#1 = NONE MU#2 = NONE]
                RAID6RAID5[Group  2-  0] SSID = 0x0000
Harddisk 3 -> [ST] CL1-A Ser =91100123 LDEV =   3 [HITACHI ] [DF600F ]
                HORC = SMPL  HOMRCF[MU#0 = SMPL MU#1 = NONE MU#2 = NONE]
                RAID6RAID5[Group  3-  0] SSID = 0x0000
Harddisk 4 -> [ST] CL1-A Ser =91100123 LDEV =   2 [HITACHI ] [DF600F ]
                HORC = SMPL  HOMRCF[MU#0 = NONE MU#1 = SMPL MU#2 = NONE]
                RAID6RAID5[Group  2-  1] SSID = 0x0000
Harddisk 5 -> [ST] CL1-A Ser =91100123 LDEV =   4 [HITACHI ] [DF600F ]
                HORC = SMPL  HOMRCF[MU#0 = NONE MU#1 = SMPL MU#2 = NONE]
                RAID6RAID5[Group  4-  0] SSID = 0x0000
Harddisk 6 -> [ST] CL1-A Ser =91100123 LDEV =   2 [HITACHI ] [DF600F ]
                HORC = SMPL  HOMRCF[MU#0 = NONE MU#1 = NONE MU#2 = SMPL]
                RAID6RAID5[Group  2-  2] SSID = 0x0000
Harddisk 7 -> [ST] CL1-A Ser =91100123 LDEV =   5 [HITACHI ] [DF600F ]
                HORC = SMPL  HOMRCF[MU#0 = NONE MU#1 = NONE MU#2 = SMPL]
                RAID6RAID5[Group  5-  0] SSID = 0x0000

C:\HORCM\etc>
```

For more information on the configuration definition file, refer to the *Hitachi Unified Storage Command Control Interface Installation and Configuration Guide*.

Setting the environment variable

To perform ShadowImage operations, you must set the environment variable. The following describes an example in which two instances are configured within the same host.

1. Set the environment variable for each instance. Enter the following from the command prompt:

```
C:\HORCM\etc>set HORCMINST=0
```

2. To enable Snapshot, the environment variable must be set as follows:

```
C:\HORCM\etc>set HORCC_MRCF=1
```

- Execute the horcmstart script, and then execute the pairedisplay command to verify the configuration, as shown in the following example:

```
:\HORCM\etc>horcmstart 0 1
starting HORCM inst 0
HORCM inst 0 starts successfully.
starting HORCM inst 1
HORCM inst 1 starts successfully.

C:\HORCM\etc>pairedisplay -g VG01
group   PairVOL(L/R) (Port#,TID,LU-M) ,Seq#,LDEV#.P/S,Status, Seq#,P-LDEV# M
VG01    oradb1(L)   (CL1-A , 1, 2-0 )91100123 2.SMPL ----,----- ---- -
VG01    oradb1(R)   (CL1-A , 1, 3-0 )91100123 3.SMPL ----,----- ---- -
VG01    oradb2(L)   (CL1-A , 1, 2-1 )91100123 2.SMPL ----,----- ---- -
VG01    oradb2(R)   (CL1-A , 1, 4-0 )91100123 4.SMPL ----,----- ---- -
VG01    oradb3(L)   (CL1-A , 1, 2-2 )91100123 2.SMPL ----,----- ---- -
VG01    oradb3(R)   (CL1-A , 1, 5-0 )91100123 5.SMPL ----,----- ---- -
```

Performing Snapshot operations

Pair operation using CCI are shown in [Figure B-6](#).




Figure B-6: Snapshot pair status for CCI

Confirming pair status

Table B-2 shows the related CCI and Navigator 2 GUI pair status.

Table B-2: CCI / Navigator 2 GUI pair status

CCI	Navigator 2	Description
SMPL	Simplex	Status where a pair is not created.
PAIR	Paired	Status that exists in order to give interchangeability with ShadowImage.
RCPY	Reverse Synchronizing	Status in which the backup data retained in the V-VOL is being restored to the P-VOL.
PSUS/ SSUS	Split	Status in which the P-VOL data at the time of the pair splitting is retained in the V-VOL.
PFUS	Threshold Over	Status in which the used rate of DP pool reaches the threshold of Replication Depletion Alert.
PSUE	Failure	Status that suspends copying forcibly when a failure occurs.

To confirm Snapshot pairs

For the example below, the group name in the configuration definition file is VG01.

1. Execute the `pairstatus` command to verify the pair status and the configuration. For example:

```
C:\HORCM\etc>pairstatus -g VG01
Group   PairVOL(L/R) (Port#,TID,LU-M) ,Seq#,LDEV#.P/S,Status, Seq#,P-LDEV# M
VG01    oradb1(L)   (CL1-A , 1, 2-0 )91100123      2.P-VOL PSUS,-----  ---- -
VG01    oradb1(R)   (CL1-A , 1, 3-0 )91100123      3.S-VOL SSUS,-----  ---- -
```

The pair status is displayed. For details on the `pairstatus` command and its options, refer to *Hitachi Unified Storage Command Control Interface Installation and Configuration Guide*.

Pair create operation



NOTE: The following are assigned to pair names automatically:

SS_LUXXXX_LUYYYY_ZZZZZ

XXXX: LU number of P-VOL (four digits with 0)

YYYY: LU number of V-VOL (four digits with 0)

ZZZZ: 5 digits of optional number

To create Snapshot pairs

In the examples below, the group name in the configuration definition file is VG01.

1. Execute `pairstatus` to verify that the status of the Snapshot volumes is SMPL.
2. Execute `paircreate`; then execute `pairevtwait` to verify that the status of each volume is PSUS. For example:

```
C:\HORCM\etc>paircreate -split -g VG01 -d oradb1 -vl
C:\HORCM\etc>pairevtwait -g VG01 -s psus -t 300 10
pairevtwait : Wait status done.
```

3. Execute `pairdisplay` to verify the pair status and the configuration. For example:

```
C:\HORCM\etc>pairdisplay -g VG01
group  PairVOL(L/R) (Port#,TID,LU-M) ,Seq#,LDEV#.P/S,Status, Seq#,P-LDEV# M
VG01   oradb1(L)   (CL1-A , 1, 2-0 )91100123 2.P-VOL PSUS,----- ---- -
VG01   oradb1(R)   (CL1-A , 1, 3-0 )91100123 3.S-VOL SSUS,----- ---- -
VG01   oradb2(L)   (CL1-A , 1, 2-1 )91100123 2.SMPL -----,----- ---- -
VG01   oradb2(R)   (CL1-A , 1, 4-0 )91100123 4.SMPL -----,----- ---- -
VG01   oradb3(L)   (CL1-A , 1, 2-2 )91100123 2.SMPL -----,----- ---- -
VG01   oradb3(R)   (CL1-A , 1, 5-0 )91100123 5.SMPL -----,----- ---- -
```



NOTE: Using CCI, the same DP pool is used for the replication data DP pool and the management area DP pool. The different DP pools cannot be specified respectively. You can use HSNM2 for pair creation to specify separate DP pools for the replication data DP pool and the management area DP pool.

Pair creation using a consistency group

A consistency group insures that the data in two or more S-VOLs included in a group are of the same time. For more information, see [Consistency Groups \(CTG\) on page 7-11](#).

To create a pair using a consistency group

1. Execute the `pairdisplay` command to verify that the status of the Snapshot volumes is SMPL. In the following example, the group name in the configuration definition file is VG01.
2. Execute `paircreate -m grp`; then, execute `pairevtwait` to verify that the status of each volume is PAIR. For example:

```
C:\HORCM\etc>paircreate -g VG01 -vl -m grp
C:\HORCM\etc>pairevtwait -g VG01 -s pair -t 300 10
pairevtwait : Wait status done.
```

3. Execute `pairsplit`; then, execute `pairevtwait` to verify that the status of each volume is PSUS. For example:

```
C:\HORCM\etc>pairsplit -g VG01
C:\HORCM\etc>pairevtwait -g VG01 -s psus -t 300 10
pairevtwait : Wait status done.
```

4. Execute `pairdisplay` to verify the pair status and the configuration. For example:

```
C:\HORCM\etc>pairdisplay -g VG01
group  PairVOL(L/R) (Port#,TID,LU-M) ,Seq#,LDEV#.P/S,Status, Seq#,P-LDEV# M
VG01   oradb1(L)   (CL1-A , 1, 2-0 )91100123 2.P-VOL PSUS,----- ---- -
VG01   oradb1(R)   (CL1-A , 1, 3-0 )91100123 3.S-VOL SSUS,----- ---- -
VG01   oradb2(L)   (CL1-A , 1, 2-1 )91100123 2.P-VOL PSUS,----- ---- -
VG01   oradb2(R)   (CL1-A , 1, 4-0 )91100123 4.S-VOL SSUS,----- ---- -
VG01   oradb3(L)   (CL1-A , 1, 2-2 )91100123 2.P-VOL PSUS,----- ---- -
VG01   oradb3(R)   (CL1-A , 1, 5-0 )91100123 5.S-VOL SSUS,----- ---- -
```



NOTE: When using the consistency group, the -m grp option is required. However, the -split option and the -m grp option cannot be used at the same time.

Pair Splitting

To split the Snapshot pairs:

1. For example, if the group name in the configuration definition file is VG01: Change the status to PSUS using pairsplit.

```
C:\HORCM\etc>pairevtwait -g VG01 -s pair -t 300 10
pairevtwait : Wait status done.
:\HORCM\etc>pairsplit -g VG01 -d oradb1
```

2. .Execute pairdisplay to update the pair status and the configuration.
For example:

```
C:\HORCM\etc>pairdisplay -g VG01
group  PairVOL(L/R) (Port#,TID,LU-M) ,Seq#,LDEV#.P/S,Status, Seq#,P-LDEV# M
VG01   oradb1(L)   (CL1-A , 1, 2-0 )91100123 2.P-VOL PSUS,----- ---- -
VG01   oradb1(R)   (CL1-A , 1, 3-0 )91100123 3.S-VOL SSUS,----- ---- -
VG01   oradb2(L)   (CL1-A , 1, 2-1 )91100123 2.SMPL ----,----- ---- -
VG01   oradb2(R)   (CL1-A , 1, 4-0 )91100123 4.SMPL ----,----- ---- -
VG01   oradb3(L)   (CL1-A , 1, 2-2 )91100123 2.SMPL ----,----- ---- -
VG01   oradb3(R)   (CL1-A , 1, 5-0 )91100123 5.SMPL ----,----- ---- -
```

The Snapshot pair is split.

Re-synchronizing Snapshot pairs

To re-synchronize the Snapshot pairs:

1. For example, if the group name in the configuration definition file is VG01:

Change the PSUS status of the Snapshot pair to PAIR status using pairresync.

```
C:\HORCM\etc>pairresync -g VG01 -d oradb1
```

2. Execute pairdisplay to update the pair status and the configuration.
For example:

```
C:\HORCM\etc>pairdisplay -g VG01
group  PairVOL(L/R) (Port#,TID,LU-M) ,Seq#,LDEV#.P/S,Status, Seq#,P-LDEV# M
VG01   oradb1(L)   (CL1-A , 1, 2-0 )91100123 2.P-VOL PSUS,----- ---- -
VG01   oradb1(R)   (CL1-A , 1, 3-0 )91100123 3.S-VOL SSUS,----- ---- -
VG01   oradb2(L)   (CL1-A , 1, 2-1 )91100123 2.SMPL ----,----- ---- -
VG01   oradb2(R)   (CL1-A , 1, 4-0 )91100123 4.SMPL ----,----- ---- -
VG01   oradb3(L)   (CL1-A , 1, 2-2 )91100123 2.SMPL ----,----- ---- -
VG01   oradb3(R)   (CL1-A , 1, 5-0 )91100123 5.SMPL ----,----- ---- -
```

The Snapshot pair is re-synchronized.

Restoring a V-VOL to the P-VOL

To restore the V-VOL to the P-VOL

In the examples below, the group name in the configuration definition file is VG01.

1. Execute `pairresync -restore` to restore the V-VOL to the P-VOL. For example:

```
C:\HORCM\etc>pairresync -restore -g VG01 -d oradb1 -c 15
```

2. Execute `pairstatus` to display pair status and the configuration. For example:

```
C:\HORCM\etc>pairstatus -g VG01
group  PairVOL(L/R) (Port#,TID,LU-M) ,Seq#,LDEV#.P/S,Status, Seq#,P-LDEV# M
VG01   oradb1(L)   (CL1-A , 1,  2-0 )91100123 2.P-VOL RCCOPY,-----  ----  -
VG01   oradb1(R)   (CL1-A , 1,  3-0 )91100123      3.S-VOL RCCOPY,-----  ----  -
VG01   oradb2(L)   (CL1-A , 1,  2-1 )91100123      2.SMPL  ----,-----  ----  -
VG01   oradb2(R)   (CL1-A , 1,  4-0 )91100123      4.SMPL  ----,-----  ----  -
VG01   oradb3(L)   (CL1-A , 1,  2-2 )91100123      2.SMPL  ----,-----  ----  -
VG01   oradb3(R)   (CL1-A , 1,  5-0 )91100123      5.SMPL  ----,-----  ----  -
```

3. Execute the `pairsplit` command. Pair status becomes PSUS. For example:

```
C:\HORCM\etc>pairsplit -g VG01 -d oradb1
```

The V-VOL is restored to the P-VOL.

Deleting Snapshot pairs

To delete the Snapshot pair and change status to SMPL

In the examples below, the group name in the configuration definition file is VG01.

1. Execute the `pairstatus` command to verify that the status of the Snapshot pair is PSUS or PSUE. For example:

```
C:\HORCM\etc>pairstatus -g VG01
group  PairVOL(L/R) (Port#,TID,LU-M) ,Seq#,LDEV#.P/S,Status, Seq#,P-LDEV# M
VG01   oradb1(L)   (CL1-A , 1, 2-0 )91100123 2.P-VOL PSUS,----- ---- -
VG01   oradb1(R)   (CL1-A , 1, 3-0 )91100123 3.S-VOL SSUS,----- ---- -
VG01   oradb2(L)   (CL1-A , 1, 2-1 )91100123 2.SMPL ----,----- ---- -
VG01   oradb2(R)   (CL1-A , 1, 4-0 )91100123 4.SMPL ----,----- ---- -
VG01   oradb3(L)   (CL1-A , 1, 2-2 )91100123 2.SMPL ----,----- ---- -
VG01   oradb3(R)   (CL1-A , 1, 5-0 )91100123 5.SMPL ----,----- ---- -
```

2. Execute the `pairsplit -S` command to delete the Snapshot pair. For example:

```
C:\HORCM\etc>pairsplit -S -g VG01 -d oradb1
```

3. Execute the `pairstatus` command to verify that the pair status changed to SMPL. For example:

```
C:\HORCM\etc>pairstatus -g VG01
group  PairVOL(L/R) (Port#,TID,LU-M) ,Seq#,LDEV#.P/S,Status, Seq#,P-LDEV# M
VG01   oradb1(L)   (CL1-A , 1, 2-0 )9110012391100123 2.SMPL ----,----- ---- -
VG01   oradb1(R)   (CL1-A , 1, 3-0 )91100123 3.SMPL ----,----- ---- -
VG01   oradb2(L)   (CL1-A , 1, 2-1 )91100123 2.SMPL ----,----- ---- -
VG01   oradb2(R)   (CL1-A , 1, 4-0 )91100123 4.SMPL ----,----- ---- -
VG01   oradb3(L)   (CL1-A , 1, 2-2 )91100123 2.SMPL ----,----- ---- -
VG01   oradb3(R)   (CL1-A , 1, 5-0 )91100123 5.SMPL ----,----- ---- -
```

Pair and group name differences in CCI and Navigator 2

Pairs and groups that were created using CCI will be displayed differently when status is confirmed in Navigator 2.

- Pairs created with CCI and defined in the configuration definition file display unnamed in Navigator 2.
- Groups defined in the configuration definition file are also different in Navigator 2.
- Pairs defined in a group on the configuration definition file using CCI are displayed in Navigator 2 as ungrouped.

For information about how to manage a group defined on the configuration definition file as a CTG, see the *Hitachi Unified Storage Command Control Interface Installation and Configuration Guide*.

Performing Snapshot operations using raidcom

You can use the `raidcom` command other than the `paircreate/pairsplit/pairresync` command for the pair operation using CCI. For more details on the `raidcom` command, refer to the *Hitachi Unified Storage Command Control Interface Installation and Configuration Guide*.

Figure B-7 shows a pair operation using the `raidcom` command.




Figure B-7: Snapshot pair status transitions using raidcom command

Setting the command device for raidcom command

See [Setting the command device on page B-22](#).

Creating the configuration definition file for raidcom command

Refer to [Defining the configuration definition file on page B-24](#), step 5, in the **HORCM_CMD** section, to specify the array's physical drive (command device). Only one Configuration Definition File is required (only the **HORCM_CMD** is necessary).

Setting the environment variable for raidcom command

Refer to [Setting the environment variable on page B-27](#) and perform step 1 and step 3 to execute the horcmstart script.

Creating a snapshotset and registering a P-VOL

You must register the P-VOL and the DP pool to be used in the snapshotset before creating the snapshot data. If the specified snapshotset does not exist, a snapshotset is created.



NOTE: The following are assigned automatically to pair names:
When not specifying the logical unit number the V-VOL:

SS_LUXXXX_LUNONE_ZZZZZ

XXXX: LU number of P-VOL (four digits with 0)

ZZZZZ: 5 digits of optional number

1. The snapshotset name of the snapshot set at the registration destination is "snap1", the number of the DP pool to be used is "50" and the volume number of the P-VOL to be registered to "snap1" is "10". Execute the raidcom add snapshotset command and register the P-VOL and the DP pool in the snapshotset. To enable the CTG mode of the snapshotset, add -snap_mode CTG.

```
C:\HORCM\etc>raidcom and snapshotset -ldev_id 10 -pool 50 -snapshot_name
snap1
                        -snap_mode CTG
```



NOTE: Using CCI, the same DP pool is used for the replication data DP pool and the management area DP pool. The different DP pools cannot be specified respectively. You can use HSNM2 for pair creation to specify separate DP pools for the replication data DP pool and the management area DP pool.

2. Execute the raidcom get snapshotset command and check that the creation of the snapshotset and the registration of the P-VOL and the DP pool are executed. (Check that STAT is changed to PAIR).

```
C:\HORCM\etc>raidcom get snapshotset -snapshot_name snap1
Snapshot_name  P/S  STAT  Serial#  LDEV#  MU#  P-LDEV#  PID    %  MODE  SPLT-TIME
snap1          P-VOL PAIR  93000007  10   1010   -      50     50  G---  -
```

Creating a Snapshot data

The process of creating the snapshot data (which is the duplication of the P-VOL) is :

1. The snapshot data of the P-VOL whose volume number registered in the snapshotset of the snapshot set name "snap1" is 10 will be created. Execute the `raidcom modify snapshotset -snapshot_data create` command and create the snapshot data.

```
C:\HORCM\etc>raidcom modify snapshotset -ldev_id 10 -snapshot_name snap1
                        -create snapshot
```

2. Execute the `raidcom get snapshotset` command and check that the snapshot data is created. (Check that STAT is changed to PSUS).

```
C:\HORCM\etc>raidcom get snapshotset -snapshot_name snap1
```

Snapshot_name	P/S	STAT	Serial#	LDEV#	MU#	P-LDEV#	PID	%	MODE	SPLT-TIME
snap1	P-VOL	PSUS	93000007	10	1010	-	50	50	G---	4F677A10

3. When multiple P-VOLs are registered in the same snapshotset, you can create the snapshot data of the multiple P-VOLs at once by setting the operation target of the `raidcom modify snapshotset -snapshot_data create` command in the snapshotset. The point that you can operate the multiple snapshot data in the same snapshotset at once is similar for the `raidcom modify snapshotset -snapshot_data resync` or the `raidcom modify snapshotset -snapshot_data restore` command.

Example of creating the Snapshot data of the multiple P-VOLs

1. The snapshot data of two P-VOLs whose volume numbers registered in the snapshotset of the snapshotset name "snap1" are 10 and 20 are created first. Register two P-VOLs in the snapshotset.

```
C:\HORCM\etc>raidcom and snapshotset -ldev_id 10 -pool 50 -snapshot_name snap1
                        -snap_mode CTG
C:\HORCM\etc>raidcom and snapshotset -ldev_id 50 -pool 50 -snapshot_name snap1
                        -snap_mode CTG
```

2. Execute the `raidcom modify snapshotset -snapshot_data create` command and create the snapshot data of two P-VOLs at once.

```
C:\HORCM\etc>raidcom modify snapshotset -snapshot_name snap1 -
snapshot_data create
```

3. Execute the `raidcom get snapshotset` command and check that two snapshot data are created. (Check that STAT is changed to PSUS).

```
C:\HORCM\etc>raidcom get snapshotset -snapshot_name snap1
```

Snapshot_name	P/S	STAT	Serial#	LDEV#	MU#	P-LDEV#	PID	%	MODE	SPLT-TIME
snap1	P-VOL	PSUS	93000007	10	1010	-	50	50	G---	4F677A10
snap1	P-VOL	PSUS	93000007	20	1011	-	50	50	G---	4F677A10

Discarding Snapshot data

1. The snapshot data of the P-VOL whose volume number registered in the snapshotset of the snapshotset name "snap1" is 10 will be discarded. Execute the `raidcom modify snapshotset -snapshot_data resync` command and discard the snapshot data.

```
C:\HORCM\etc>raidcom modify snapshotset -ldev_id 10 -snapshot_name snap1
                        -snapshot_data resync
```

2. Execute the `raidcom get snapshotset` command and check that the snap data is discarded. (Check that STAT is changed to PAIR.)

```
C:\HORCM\etc>raidcom get snapshotset -snapshot_name snap1
Snapshot_name  P/S  STAT  Serial#  LDEV#  MU#  P-LDEV#  PID    %  MODE  SPLT-TIME
snap1          P-VOL PAIR  93000007  10  1010  -        50    50  G---  -
```

Restoring Snapshot data

1. The snapshot data of the P-VOL whose volume number registered in the snapshotset of the snapshotset name "snap1" is 10 will be restored. Execute the `raidcom modify snapshotset -snapshot_data restore` command and restore the snapshot data.

```
C:\HORCM\etc>raidcom modify snapshotset -ldev_id 10 -snapshot_name snap1
-snapshot_data restore
```

2. Execute the `raidcom get snapshotset` command and check that the snap data is restored. (Check that STAT is changed to RCPY. When the restoration is completed, STAT is changed to PAIR.)

```
C:\HORCM\etc>raidcom get snapshotset -snapshot_name snap1
Snapshot_name  P/S  STAT  Serial#  LDEV#  MU#  P-LDEV#  PID    %  MODE  SPLT-TIME
snap1          P-VOL RCPY  93000007  10  1010  -        50    50  G---  -
C:\HORCM\etc>raidcom get snapshotset -snapshot_name snap1
Snapshot_name  P/S  STAT  Serial#  LDEV#  MU#  P-LDEV#  PID    %  MODE  SPLT-TIME
snap1          P-VOL PAIR  93000007  10  1010  -        50    50  G---  -
```

Changing the Snapshotset name

1. The snapshotset name of the snapshotset to which the snapshot data whose volume number of the P-VOL is 10 and MU# is 1010 belongs is changed to "snap2". Execute the `raidcom modify snapshotset -snapshot_data rename` command and change the snapshotset name.

```
C:\HORCM\etc>raidcom modify snapshotset -ldev_id 10 -mirror_id 1010 -
snapshot_name snap2
-snapshot_data rename
```

2. Execute the `raidcom get snapshotset` command and check that the snapshot set name is changed.

```
C:\HORCM\etc>raidcom get snapshotset -ldev_id 10
Snapshot_name  P/S  STAT  Serial#  LDEV#  MU#  P-LDEV#  PID    %  MODE  SPLT-TIME
snap2          P-VOL PAIR  93000007  10  1010  -        50    50  G---  -
```

Volume number mapping to the Snapshot data

It is necessary to map the volume number to the snapshot data to make the host recognize the snapshot data.

1. The volume number 30 is mapped to the snapshot data of the P-VOL whose volume number registered in the snapshotset of the snapshot set name "snap1" is 10. Execute the `raidcom map snapshotset` command and map the volume number to the snapshot data.

```
C:\HORCM\etc>raidcom map snapshotset -ldev_id 10 30 -snapshot_name snap1
```

2. Execute the `raidcom get snapshotset` command and check that the volume number is mapped to the snapshot data. (Check that the volume number mapped to P-LDEV# is displayed.).

```
C:\HORCM\etc>raidcom get snapshotset -snapshot_name snap1
Snapshot_name  P/S  STAT  Serial#  LDEV#  MU#  P-LDEV#  PID    %  MODE  SPLT-TIME
snap1         P-VOL PSUS  93000007  10   1010   30      50    50  G---  4F677A10
```

Volume number un-mapping of the Snapshot data

1. The volume number 30 mapped to the snapshot data will be unmapped. Execute the `raidcom unmap snapshotset` command and un-map the volume number mapped to the snapshot data.

```
C:\HORCM\etc>raidcom unmap snapshotset -ldev_id 30
```

2. Execute the `raidcom get snapshotset` command and check that the volume number mapped to the snapshot data is unmapped. (Check that P-LDEV# becomes -).

```
C:\HORCM\etc>raidcom get snapshotset -snapshot_name snap1
Snapshot_name  P/S  STAT  Serial#  LDEV#  MU#  P-LDEV#  PID    %  MODE  SPLT-TIME
snap1         P-VOL PSUS  93000007  10   1010   -       50    50  G---  4F677A10
```

Changing the volume assignment number of the Snapshot data

1. The volume number mapped to the snapshot data is 30 and the snapshot set to which the assignment destination snapshot data of the volume number belongs is "snap2". Execute the `raidcom replace snapshotset` command and change the assignment of the volume number of the snapshot data.



NOTE: The assignment of the volume number can only be changed between the snapshot data of the same P-VOL

```
C:\HORCM\etc>raidcom replace snapshotset -ldev_id 30 -snapshot_name snap2
```

2. Execute the `raidcom get snapshotset` command and check that the assignment of the volume number of the snapshot data is changed. (Check Snapshot_name and P-LDEV# and check that the volume number is assigned to the target snapshot data.)

```
C:\HORCM\etc>raidcom get snapshotset -snapshot_name snap1
Snapshot_name  P/S  STAT  Serial#  LDEV#  MU#  P-LDEV#  PID    %  MODE  SPLT-TIME
snap1         P-VOL PSUS  93000007  10   1010   -       50    50  G---  4F677A10
snap2         P-VOL PSUS  93000007  20   1010   30      50    50  G---  4F677A10
```

Deleting the snapshotset

Once the snapshotset is deleted, all the snapshot data belonging to the relevant snapshotset is also deleted.

1. The snapshotset of the snapshotset name "snap1" will be deleted.
Execute the raidcom delete snapshotset command and delete the snapshotset...

```
C:\H0RCM\etc>raidcom delete snapshotset -snapshot_name snap1
```

2. Execute the raidcom get snapshotset command and check that the snapshot data is deleted.

```
C:\H0RCM\etc>raidcom get snapshotset
Snapshot_name  P/S  STAT  Serial#  LDEV#  MU#  P-LDEV#  PID    %  MODE  SPLT-TIME
-             -    -     -        -      -    -        -    -   -    -
```

Using Snapshot with Cache Partition Manager

This topic provides special instructions for using Snapshot with Cache Partition Manager. Snapshot uses a part of the cache area to manage internal resources. Because of this, cache capacity for Cache Partition Manager decreases.

In this case, make sure that cache partition information is initialized, which results in the following efficiencies:

- Logical units are moved to the master partitions on the side of the default owner controller.
- Sub-partitions are deleted and the size of the each master partition is reduced to half of the user data area after the installing Snapshot.

Figure B-8 shows partitions before Snapshot is installed; Figure B-9 shows them with Snapshot.




Figure B-8: Cache partitions with Cache Partition Manager




Figure B-9: Cache partitions when Snapshot installed with Cache Partition Manager



TrueCopy Remote Replication reference information

This appendix includes:

- ☐ [TrueCopy specifications](#)
- ☐ [Operations using CLI](#)
- ☐ [Operations using CCI](#)

TrueCopy specifications

Table C-1 lists external specifications for TrueCopy.

Table C-1: External specifications for TrueCopy

Parameter	TrueCopy requirement
License	Key code required for installation. TrueCopy Remote and TrueCopy Extended cannot coexist, and have different licenses.
User Interface	- Navigator 2 GUI and/or CLI: Setup and pair operations. - CCI: Pair operations. Certain related operations only available using CCI.
Command device	- Required for CCI, one per disk array. - Up to 128 allowed per disk array. - Must be 65,538 blocks or more (1 block = 512 bytes) (33 M bytes or more).
Host Interface	Fibre Channel or iSCSI
Controller configuration	Configuration of dual controller is required.
Remote path	- The interface type of two remote paths between disk arrays must be the same, fibre channel or iSCSI. - One remote path per controller is required, with a total of two between the disk arrays in a the dual controller configuration.
Port modes	Initiator and target intermix mode. One port may be used for host I/O and TrueCopy at the same time.
License	A key code is required for installation. TrueCopy Remote and TrueCopy Extended cannot coexist, and have different licenses.
Bandwidth supported	1.5 M bps or more (100 M bps or more is recommended.) Low transfer rate results in greater time for TrueCopy operations and reduced host I/O performance.
DMLU	- Minimum size: 10 GB - Maximum size: <=128 GB - One DMLU is required - Be sure to set the DMLU for both the local and remote arrays.
Unit of pair management	Volumes are the target of TrueCopy pairs, and are managed per volume.
Maximum number of volumes in which pair can be created	- HUS 110: 2,046 - HUS 130/HUS 150: 4,094 - The maximum number of volumes when different types of arrays are combined is that of the array whose maximum number of volumes is smaller
Pair structure	One copy (S-VOL) per P-VOL.

Table C-1: External specifications for TrueCopy (Continued)

Parameter	TrueCopy requirement
Supported RAID level	- RAID 1 (1D+1D), RAID 5 (2D+1P to 15D+1P) - RAID 1+0 (2D+2D to 8D+8D) - RAID 6 (2D+2P to 28D+2P)
Combination of RAID levels	All combinations supported. The number of data disks does not have to be the same.
Size of volume	1:1, P-VOL = S-VOL. The max volume size is 128 TB.
Drive Types for P-VOL/S-VOL	If the drive types are supported by the disk array, they can be set for a P-VOL and an S-VOL. However, SAS drives or SSD/FMD drives are recommended, especially for the P-VOL. When a pair is created using two volumes configured by the SAS7.2K drives, requirements for using the SAS7.2K drives may differ.
Supported capacity value of P-VOL and S-VOL	The capacity for TrueCopy is limited. See Calculating supported capacity on page 14-19 .
Consistency Groups (CTG) supported	- Up to 256 CTGs per disk array. - Maximum number of pairs one CTG can manage is: 2,046 for HUS 110, 4,094 for HUS 130/HUS 150.
Management of volumes while using TrueCopy	A TrueCopy pair must be deleted before the following operations: - Deletion of the pair's RAID group - Deletion of a volume - Deletion of the DMLU - Formatting, growing, or shrinking a volume
Restrictions during volume formatting	A TrueCopy pair cannot be created by specifying a formatting volume.
Restrictions during RAID group expansion	A RAID group with a TrueCopy P-VOL or S-VOL can be expanded only when the pair status is Simplex or Split.
Pair creation of a unified volume	A TrueCopy pair can be created by specifying a unified volume. However, unification of volumes or release of the unified volume cannot be done for the paired volumes.
Failures	When a failure of the copy operation from P-VOL to S-VOL occurs, TrueCopy will suspend the pair (Failure). If a volume failure occurs, TrueCopy suspend the pair. If a drive failure occurs, the TrueCopy pair status is not affected because of the RAID architecture.
Concurrent use of Data Retention Utility	Yes, but note following: - When S-VOL Disable is set for an volume, the volume cannot be used in a pair. - When S-VOL Disable is set for a volume that is already an S-VOL, no suppression of the pair takes place, unless the pair status is split.
Concurrent use of SNMP Agent	Yes. A trap is transmitted when: - A failure occurs in the remote path - Pair status changes to Failure.
Concurrent use of Volume Migration	Yes, but, a Volume Migration P-VOL, S-VOL, and Reserved volume of cannot be specified as a TrueCopy P-VOL or an S-VOL
Concurrent use of TCE	No.

Table C-1: External specifications for TrueCopy (Continued)

Parameter	TrueCopy requirement
Concurrent use of ShadowImage	Yes. TrueCopy can be used together with ShadowImage and cascaded with ShadowImage.
Concurrent use of Snapshot	Yes. TrueCopy can be used together with Snapshot and cascaded with Snapshot.
Concurrent use of Power Saving /Power Saving Plus	Yes. However, when a P-VOL or an S-VOL is included in a RAID group for which the Concurrent use of Power Saving /Power Saving Plus is specified, only a TrueCopy pair split and the pair delete can be performed.
Concurrent use of Dynamic Provisioning	Available. For more details, see Concurrent use of Dynamic Provisioning on page 14-14 .
Concurrent use of Dynamic Tiering	Available. Formore details, see Concurrent use of Dynamic Tiering on page 14-17 .
Load balancing function	The Load balancing function applies to a TrueCopy pair.
Reduction of memory	Memory cannot be reduced when the ShadowImage, Snapshot, TrueCopy, or Volume Migration function are enabled. Reduce memory after disabling the functions.
Remote Copy over iSCSI in the WAN environment	We recommend using TrueCopy in the WAN environment of MTU1500 or more. However, if TCE needs to be implemented in the WAN environment of less than MTU1500, change the maximum segment size (MSS) of the WAN router to a value less than 1500, and then create a remote path. The data length transmitted from TCE of Hitachi Unified Storage changes to the specified value less than 1500. When creating a remote path without changing the MSS value or not creating the remote path after changing the MSS value again, a data transfer error occurs because TCE of HUS100 transmits the MTU1500 data. To change the MSS value, request the customer or the WAN router provider.

Operations using CLI

This section describes CLI procedures for setting up and performing TrueCopy operations.

- ❑ [Installation and setup](#)
- ❑ [Pair operations](#)
- ❑ [Sample scripts](#)



NOTE: For additional information on the commands and options used in this appendix, see the *Hitachi Unified Storage Command Line Interface Reference Guide*.

Installation and setup

This section provides installation/uninstalling, enabling/disabling, and setup procedures.

TrueCopy is an extra-cost option and must be installed using a key code or file. Obtain it from the download page on the HDS Support Portal, <http://support.hds.com>. For prerequisites, see the GUI-based instructions in [Installation procedures on page 13-3](#).

Installing

TrueCopy cannot be installed if more than 239 hosts are connected to a port on the array.

To install TrueCopy

1. From the command prompt, register the array in which the TrueCopy is to be installed, and then connect to the array.
2. Execute the `auopt` command to install TrueCopy. For example:

```
% auopt -unit array-name -lock off -keycode manual-attached-keycode
Are you sure you want to unlock the option? (y/n [n]): y
When Cache Partition Manager is enabled, if the option using data
pool will be enabled the default cache partition information will be
restored.
Do you want to continue processing? (y/n [n]): y
The option is unlocked.
%
```

3. Execute the `auopt` command to confirm whether TrueCopy has been installed. For example:

```
% auopt -unit array-name -refer
Option Name          Type      Term      Status
    Refigure Memory Status
TRUECOPY              Permanent ---      Enable
    N/A
%
```

TrueCopy is installed and enabled. Installation is complete.

Enabling or disabling

TrueCopy can be disabled or enabled. When TrueCopy is first installed it is automatically enabled.

Prerequisites for disabling

- TrueCopy pairs must be released (the status of all volumes must be Simplex).
- The remote path must be released.
- TrueCopy cannot be enabled if more than 239 hosts are connected to a port on the array.

To enable /disable TrueCopy

1. From the command prompt, register the array in which the status of the feature is to be changed, and then connect to the array.
2. Execute the auopt command to change TrueCopy status (enable or disable).

The following is an example of changing the status from enable to disable. If you want to change the status from disable to enable, enter enable after the -st option.

```
% auopt -unit array-name -option TRUCCOPY -st disable
Are you sure you want to disable the option? (y/n [n]): y
The option has been set successfully.
%
```

3. Execute the auopt command to confirm that the status has been changed. For example:

```
% auopt -unit array-name -refer
Option Name          Type      Term      Status
      Reconfigure Memory Status
TRUCCOPY              Permanent ---      Disable
N/A
%
```

Uninstalling

To uninstall TrueCopy, the key code provided for optional features is required.

Prerequisites for uninstalling

- All TrueCopy pairs must be released (the status of all volumes must be Simplex).
- The remote path must be released.

To uninstall TrueCopy

1. From the command prompt, register the array in which the TrueCopy is to be uninstalled, and then connect to the array.
2. Execute the auopt command to uninstall TrueCopy. For example:

```
% auopt -unit array-name -lock on -keycode manual-attached-keycode
Are you sure you want to lock the option? (y/n [n]): y
The option is locked.
%
```

3. Execute the auopt command to confirm that TrueCopy is uninstalled. For example:

```
% auopt -unit array-name -refer
DMEC002015: No information displayed.
%
```

Setting the Differential Management Logical Unit

The DMLU must be set up before TrueCopy copies can be made. Please see the prerequisites under [Differential Management LU \(DMLU\) on page 12-4](#) before proceeding.

To set up the DMLU

1. From the command prompt, register the array to which you want to set the DMLU. Connect to the array.
2. Execute the `audmlu` command. This command first displays volumes that can be assigned as DMLUs and then creates a DMLU. For example:

```
% audmlu -unit array-name -availablelist
Available Logical Units
  LUN Capacity    RAID Group  DP Pool RAID Level  Type Status
    0   10.0 GB         0      (N/A   5( 4D+1P)  SAS  Normal
%
% audmlu -unit array-name -set -lu 0
Are you sure you want to set the DM-LU? (y/n [n]): y
The DM-LU has been set successfully.
%
```

Release a DMLU

There is this restriction when either pair of ShadowImage, Volume Migration, or TrueCopy exists.

- The DMLU cannot be released.

To release a TrueCopy DMLU

Use the following example:

```
% audmlu -unit array-name -rm -lu 0
Are you sure you want to release the DM-LU? (y/n [n]): y
The DM-LU has been released successfully.
%
```

Adding a DMLU capacity

To add a DMLU capacity that has been previously set, specify the `-chgsz` and `-size` options in the `audmlu` command.

The `-rg` option can be specified only when the DMLU is a normal volume.

Select a RAID group which meets the following conditions:

- The drive type and the combination are the same as the DMLU
- A new volume can be created
- A sequential free area for the capacity to be expanded exists

Example:

Use the following example:

```
% audmlu -unit array-chgsize -size capacity after adding -rg RAID
group number
Are you sure you want to add the capacity of DM-LU?
(y/n [n]): y
The capacity of DM-LU has been added successfully.
%
```

Setting the remote port CHAP secret

For iSCSI, the remote path can employ a CHAP secret. Set the CHAP secret mode on the remote array. For more information on the CHAP secret, see [Adding or changing the remote port CHAP secret \(iSCSI only\) on page 14-23](#).

The setting procedure of the remote port CHAP secret is shown below.

1. From the command prompt, register the array in which you want to set the remote path, and then connect to the array.
2. Execute the aurmtpath command with the -set option and perform the CHAP secret of the remote port. The input example and the result are shown below. For example:

```
% aurmtpath -unit array-name -set -target -local 91200027
-secret
Are you sure you want to set the remote path information?
(y/n[n]): y
Please input Path 0 Secret.
Path 0 Secret:
Re-enter Path 0 Secret:
Please input Path 1 Secret.
Path 1 Secret:
Re-enter Path 1 Secret:
The remote path information has been set successfully.
%
```

The setting of the remote port CHAP secret is completed.

Setting the remote path

Data is transferred from the local to the remote array over the remote path.

Please review Prerequisites in GUI instructions in [Remote path requirements on page 14-33](#) before proceeding.

To set up the remote path

1. From the command prompt, register the array in which you want to set the remote path, and then connect to the array.
2. The following shows an example of referencing the remote path status where remote path information is not yet specified.

Fibre Channel example:

```
% aurmtpath -unit array-name -refer
Initiator Information
Local Information
  Array ID      : 91200026
  Distributed Mode : N/A

Path Information
  Interface Type      : ---
  Remote Array ID     : ---
  Remote Path Name    : ---
  Bandwidth [0.1 Mbps] : ---
  iSCSI CHAP Secret   : ---

Path      Status      Local      Remote      Remote Port      TCP Port No. of
0      Undefined      ---      ---      IP Address      Remote Port
1      Undefined      ---      ---      ---      ---

%
```

iSCSI example:

```
% aurmtpath -unit array-name -refer
Initiator Information
Local Information
  Array ID      : 91200027
  Distributed Mode : N/A

Path Information
  Interface Type      : ---
  Remote Array ID     : ---
  Remote Path Name    : ---
  Bandwidth [0.1 Mbps] : ---
  iSCSI CHAP Secret   : ---

Path      Status      Local      Remote      Remote Port      TCP Port No. of
0      Undefined      ---      ---      IP Address      Remote Port
1      Undefined      ---      ---      ---      ---

Target Information
Local Array ID      :

%
```

3. Execute the aurmtpath command to set the remote path.

Fibre Channel example:

```
% aurmtpath -unit array-name -set -remote 91200027 -band 15
               -path0 0A 0A -path1 1A 1B
Are you sure you want to set the remote path information?
(y/n[n]): y
The remote path information has been set successfully.
%
```

iSCSI example:

```
% aurmtpath -unit array-name -set -initiator -remote 91200027 -secret disable
               -path0 0B -path0_addr 192.168.1.201 -band 100
               -path1 1B -path1_addr 192.168.1.209
Are you sure you want to set the remote path information?
(y/n[n]): y
The remote path information has been set successfully.
%
```

4. Execute the aurmtpath command to confirm whether the remote path has been set. For example:

Fibre Channel example:

```
% aurmtpath -unit array-name -refer
Initiator Information
Local Information
  Array ID      : 91200026
  Distributed Mode : N/A

Path Information
  Interface Type      : FC
  Remote Array ID     : 91200027
  Remote Path Name    : Array 91200027
  Bandwidth [0.1 Mbps] : 15
  iSCSI CHAP Secret   : N/A

Path Status Local Remote Port TCP Port No. of
0 Normal 0A 0A N/A Remote Port
1 Normal 1A 1B N/A Remote Port
%
```

iSCSI example:

```
% aurmtpath -unit array-name -refer
```

Initiator Information

Local Information

Array ID : 91200026

Distributed Mode : N/A

Path Information

Interface Type : iSCSI

Remote Array ID : 91200027

Remote Path Name : Array 91200027

Bandwidth [0.1 Mbps] : 100

iSCSI CHAP Secret : Disable

Path	Status	Local	Remote	Port IP Address	TCP Port No. of Remote Port
0	Normal	0B	N/A	192.168.0.201	3260
1	Normal	1B	N/A	192.168.0.209	3260

Target Information

Local Array ID : 91200026

```
%
```

Deleting the remote path

When the remote path becomes unnecessary, delete the remote path.

Prerequisites

- The pair status of the volumes using the remote path to be deleted must be Simplex or Split
- Do not do a pair operation for a TrueCopy pair when the remote path for the pair is not set up, because the pair operation may not complete correctly.



NOTE: When performing the planned shutdown of the remote array, the remote path should not necessarily be deleted. Change all the TrueCopy pairs in the array to the Split status, and then perform the planned shutdown of the remote array. After restarting the array, perform the pair resynchronization. However, when the Warning notice to the Failure Monitoring Department at the time of the remote path blockade or the notice by the SNMP Agent Support Function or the E-mail Alert Function is not desired, delete the remote path, and then turn off the power of the remote array.

To delete the remote path

1. From the command prompt, register the array in which you want to delete the remote path, and then connect to the array.
2. Execute the aurmtpath command to delete the remote path. For example:

```
% aurmtpath -unit array-name -rm -remote 91200027
Are you sure you want to delete the remote path information?
(y/n[n]): y
The remote path information has been deleted successfully.
%
```

3. Execute the aurmtpath command to confirm that the path is deleted. For example:

```
% aurmtpath -unit array-name -refer
Initiator Information
Local Information
  Array ID      : 91200026
  Distributed Mode : N/A

  Path Information
    Interface Type      : ---
    Remote Array ID     : ---
    Remote Path Name    : ---
    Bandwidth [0.1 Mbps] : ---
    iSCSI CHAP Secret   : ---

  Path      Status      Local      Remote Port      TCP Port No. of
  Port      Port
    0      Undefined    ---      ---      ---      ---
    1      Undefined    ---      ---      ---      ---
%
```


Pair operations

The following sections describe the CLI procedures and commands for performing TrueCopy operations.

The `aureplicationremote` command operates TrueCopy pair. To refer to the `aureplicationremote` command and its options, type in `aureplicationremote -help`.

Confirm that the state of the remote path is **Normal** before doing a pair operation. If you do a pair operation when a remote path is not set up or its state is **Diagnosing** or **Blocked**, the pair operation may not be completed correctly.

Displaying status for all pairs

To display all pair status

1. From the command prompt, register the array to which you want to display the status of paired logical volumes. Connect to the array.
2. Execute the `aureplicationremote -refer` command. For example:

```
% aureplicationremote -unit local array-name -refer
Pair name          Local LUN  Attribute  Remote LUN  Status
Copy Type          Group Name
TC_LU0000_LU0000   0  P-VOL    0  Paired(100%)
TrueCopy           0:
TC_LU0001_LU0001   1  P-VOL    1  Paired(100%)
TrueCopy           0:
%
```

Displaying detail for a specific pair

To display pair details

1. From the command prompt, register the array to which you want to display the status and other details for a pair. Connect to the array.
2. Execute the `aureplicationremote -refer -detail` command to display the detailed pair status. For example:

```
% aureplicationremote -unit local array-name -refer -detail -pvol 0 -svol 0
                        -locallun pvol -remote 91200027
Pair Name              : TC_LU0000_LU0000
Local Information
  LUN                  : 0
  Attribute             : P-VOL
  DP Pool
    Replication Data   : N/A
    Management Area    : N/A
Remote Information
  Array ID             : 91200027
  Path Name            : Array_91200027
  LUN                  : 0
  Capacity              : 50.0 GB
  Status                : Paired(100%)
  Copy Type            : TrueCopy
  Group Name           : ---:Ungrouped
  Consistency Time     : N/A
  Difference Size      : N/A
  Copy Pace            : Prior
  Fence Level          : Never
  Previous Cycle Time  : N/A
%
```

Creating a pair

See prerequisite information under [Creating pairs on page 15-3](#) before continuing.

To create a pair

1. From the command prompt, register the local array in which you want to create pairs, and then connect to the array.
2. Execute the `aureplicationremote -refer -availablelist` command to display volumes available for copy as the P-VOL. For example:

```
% aureplicationremote -unit local array-name -refer -availablelist -tc -pvol
Available Logical Units
LUN Capacity  RAID Group  DP Pool  RAID Level  Type Status
   0   10.0 GB         0     N/A       6(9D+2P)  SAS  Normal
%
```

3. Execute the `aureplicationremote -refer -availablelist` command to display volumes on the remote array that are available as the S-VOL. For example:

```
% aureplicationremote -unit remote array-name -refer -availablelist -tc -svol
Available Logical Units
LUN Capacity  RAID Group  DP Pool  RAID Level  Type Status
   0   10.0 GB         0     N/A       6(9D+2P)  SAS  Normal
%
```

4. Specify the volumes to be paired and create a pair using the `aureplicationremote -create` command. For example:

```
% aureplicationremote -unit local array-name -create -tc -pvol 2 -svol 2
                        -remote xxxxxxxx
Are you sure you want to create pair "TC_LU0002_LU0002"?
(y/n [n]): y
The pair has been created successfully.
%
```

Creating pairs belonging to a group

To create multiple pairs that belongs to a group

1. Create a pair that belongs to a group by specifying a group number for the -gno command option. The new group and new pair are created at the same time. For example:

```
% aureplicationremote -unit local array-name -create -tc -pvol
2000 -svol 2002
                        -gno 20 -remote xxxxxxxx
Are you sure you want to create pair "TC_LU2000_LU2002"?
(y/n [n]): y
The pair has been created successfully.
%
```

2. Create new pairs as needed and assign them to the group.

Splitting a pair

To split a pair

1. From the command prompt, register the local array in which you want to split pairs, and then connect to the array.
2. Execute the aureplicationremote -split command to split the specified pair. For example:

```
% aureplicationremote -unit local array-name -split -tc -pvol 2000 -svol 2002
                        -remote xxxxxxxx
Are you sure you want to split the pair?
(y/n [n]): y
The pair has been split successfully.
%
```

Resynchronizing a pair

To resynchronize a pair

1. From the command prompt, register the local array in which you want to re-synchronize pairs, and then connect to the array.
2. Execute the aureplicationremote -resync command to re-synchronize the specified pair. For example:

```
% aureplicationremote -unit local array-name -resync -tc -pvol
2000 -svol 2002
                        -remote xxxxxxxx
Are you sure you want to re-synchronize pair?
(y/n [n]): y
The pair has been re-synchronized successfully.
%
```

Swapping a pair

Please review the Prerequisites in [Swapping pairs on page 15-11](#).

To swap the pairs, the remote path must be set to the local array from the remote array.

To swap a pair

1. From the command prompt, register the remote array in which you want to swap pairs, and then connect to the array.
2. Execute the `aureplicationremote -swaps` command to swap the specified pair. For example:

```
% aureplicationremote -unit remote array-name -swaps -tc -svol
2002
Are you sure you want to swap pair?
(y/n [n]): y
The pair has been swapped successfully.
%
```

Deleting a pair

To delete a pair

1. From the command prompt, register the local array in which you want to delete pairs, and then connect to the array.
2. Execute the `aureplicationremote -simplex` command to delete the specified pair. For example:

```
% aureplicationremote -unit local array-name -simplex -tc
-local lun pvol
                        -pvol 2000 -svol 2002 -remote xxxxxxxx
Are you sure you want to release pair?
(y/n [n]): y
The pair has been released successfully.
%
```

3. Verify the pair status. For example:

```
% aureplicationremote -unit local array-name -refer
DMEC002015: No information displayed.
%
```

4. When executing the pair deletion in the batch file or the script, insert a five-second wait before executing the next processing step. An example batch file with a five-second wait is:

```
ping 127.0.0.1 -n 5 > nul
```

- Pair creation of TrueCopy which specified the volume specified as the S-VOL of the deleted pair
- Pair creation of Volume Migration which specified the volume specified as the S-VOL of the deleted pair
- Deletion of the volume specified as the S-VOL of the deleted pair
- Shrinking of the volume specified as the S-VOL of the deleted pair
- Removing of the DMLU
- Expanding capacity of the DMLU

Changing pair information

You can change the pair name, group name, and/or copy pace.

To change pair information

1. From the command prompt, register the local array on which you want to change the TrueCopy pair information, and then connect to the array.
2. Execute the `aureplicationlocal -chg` command to change the TrueCopy pair information. In the following example, the copy pace is changed from normal to slow.

```
% aureplicationremote -unit local array-name -tc -chg -pace slow
                        -locallun pvol -pvol 2000 -svol 2002 -remote
xxxxxxx
Are you sure you want to change pair information?
(y/n [n]): y
%
```

Sample scripts

This section provides sample CLI scripts for executing a backup and for monitoring pair status.

Backup script

The following example provides sample script commands for backing up a volume on a Windows Server.

```
echo off
REM Specify the registered name of the arrays
set UNITNAME=Array1
REM Specify the group name (Specify "Ungroup" if the pair doesn't
belong to any group)
set G_NAME=Ungrouped
REM Specify the pair name
set P_NAME=TC_LU0001_LU0002
REM Specify the directory path that is mount point of P-VOL and
S-VOL
set MAINDIR=C:\main
set BACKUPDIR=C:\backup
REM Specify GUID of P-VOL and S-VOL
PVOL_GUID=xxxxxxxx-xxxx-xxxx-xxxx-xxxxxxxxxxxxxx
SVOL_GUID=yyyyyyyy-yyy-yyy-yyy-yyy-yyyyyyyyyyyy

REM Unmounting the S-VOL
pairstatus -x umount %BACKUPDIR%
REM Re-synchronizing pair (Updating the backup data)
aureplicationremote -unit %UNITNAME% -tc -resync -pairname
%P_NAME% -gname %G_NAME%
aureplicationmon -unit %UNITNAME% -evwait -tc -pairname %P_NAME%
-gname %G_NAME% -st paired -pvol

REM Unmounting the P-VOL
pairstatus -x umount %MAINDIR%
REM Splitting pair (Determine the backup data)
aureplicationremote -unit %UNITNAME% -tc -split -pairname
%P_NAME% -gname %G_NAME%
aureplicationmon -unit %UNITNAME% -evwait -tc -pairname %P_NAME%
-gname %G_NAME% -st split -pvol
REM Mounting the P-VOL
pairstatus -x mount %MAINDIR% Volume{%PVOL_GUID%}

REM Mounting the S-VOL
pairstatus -x mount %BACKUPDIR% Volume{%SVOL_GUID%}
< The procedure of data copy from C:\backup to backup appliance>
```



NOTE: When Windows Server is used, the CCI mount command is required when mounting or un-mounting a volume. The GUID, which is displayed by mountvol command, is needed as an argument when using the mount command. For more information, see the *Hitachi Unified Storage Command Line Interface Reference Guide*.

Pair-monitoring script

The following is a sample script for monitoring two TrueCopy pairs (TC_LU0001_LU0002 and TC_LU0003_LU0004). The script includes commands for informing the user when pair failure occurs. The script is re-activated after several minutes. The array must be registered.

```
echo OFF
REM Specify the registered name of the arrays
set UNITNAME=Array1
REM Specify the name of target group (Specify "Ungroup" if the
pair doesn't belong to any group)
set G_NAME=Ungrouped
REM Specify the name of target pair
set P1_NAME=TC_LU0001_LU0002
set P2_NAME=TC_LU0003_LU0004
REM Specify the value to inform "Failure"
set FAILURE=14

REM Checking the first pair
:pair1
aureplicationmon -unit %UNITNAME% -evwait -tc -pairname %P1_NAME%
-gname %G_NAME% -nowait
if errorlevel %FAILURE% goto pair1_failure
goto pair2
:pair1_failure
<The procedure for informing a user>*

REM Checking the second pair
aureplicationmon -unit %UNITNAME% -evwait -tc -pairname %P2_NAME%
-gname %G_NAME% -nowait
if errorlevel %FAILURE% goto pair2_failure
goto end
:pair2_failure
<The procedure for informing a user>*

:end
```

Operations using CCI

This topic describes CCI procedures for setting up and performing TrueCopy operations, and provides examples of TrueCopy commands using the Windows Server.

- ❑ [Setting up CCI](#)
- ❑ [Preparing for CCI operations](#)
- ❑ [Pair operations](#)

Setting up CCI

CCI is used to display TrueCopy volume information, create and manage TrueCopy pairs, and issue commands for replication operations. CCI resides on the UNIX/Windows management host and interfaces with the disk arrays through dedicated logical volumes. CCI commands can be issued from the UNIX/Windows command line or using a script file.

When the operating system of the host is Windows Server, CCI is required to mount or un-mount the volume.

Preparing for CCI operations

The following must be set up for TrueCopy CCI operations

- Command device
- LU mapping
- Environment variable
- Configuration definition file

Enter CCI commands from the command prompt on the host where CCI is installed.

Setting the command device

CCI interfaces with Hitachi Unified Storage through the *command device*, a dedicated logical volume located on the local and remote arrays. CCI software on the host issues read/write commands to the command device. The command device is accessed as a raw device (no file system, no mount operation). It is dedicated to CCI communications and cannot be used by any other applications.

- The command device is defined in the HORCM_CMD section of the configuration definition file for the CCI instance on the attached host.
- 128 command devices can be designated for the array.
- Logical units used as command devices must be recognized by the host.
- The command device must be 33 MB or greater.



If a command device fails, all commands are terminated. CCI supports an alternate command device function, in which two command devices are specified within the same array, to provide a backup. For details on the alternate command device function, refer to the *Hitachi Unified Storage Command Control Interface Installation and Configuration Guide*.

To designate a command device

1. From the command prompt, register the array to which you want to set the command device, and then connect to the array.
2. Execute the `aucmddev` command to set a command device. When this command is run, logical units that can be assigned as a command device display, then the command device is set. To use the CCI protection function, enter `enable` following the `-dev` option. The following is an example of specifying LUN 2 for command device 1. First, display volumes to be assignable command device, and later set a command device.

```
% aucmddev -unit array-name -availablelist
Available Logical Units
  LUN  Capacity RAID Group DP Pool RAID Level  Type Status
    2   35.0 MB         0  NA      6(9D+2P)  SAS  Normal
    3   35.0 MB         0  NA      6(9D+2P)  SAS  Normal
%
% aucmddev -unit array-name -set -dev 1 2
Are you sure you want to set the command devices?
(y/n [n]): y
The command devices have been set successfully.
%
```

3. Execute the `aucmddev` command to verify that the command device is set. For example:

```
% aucmddev -unit array-name -refer
Command Device LUN  RAID Manager Protect
1              2   Disable
%
```

4. To release a command device, follow the example below, in which command device 1 is released.

```
% aucmddev -unit array-name -rm -dev 1
Are you sure you want to release the command devices?
(y/n [n]): y
This operation may cause the CCI, which is accessing to this
command device, to freeze.
Please make sure to stop the CCI, which is accessing to this
command device, before performing this operation.
Are you sure you want to release the command devices? (y/n [n]): y
The specified command device will be released.
Are you sure you want to execute? (y/n [n]): y
The command devices have been released successfully.
%
```

5. To change a command device, first release it, then change the volume number. The following example of specifies LU 3 for command device 1.

```
% aucmddev -unit array-name -set -dev 1 3
Are you sure you want to set the command devices?
(y/n [n]): y
The command devices have been set successfully.
%
```

Setting LU mapping

For iSCSI, use the `autargetmap` command instead of the `auhgmap` command.

To set up LU Mapping

1. From the command prompt, register the array to which you want to set the LU Mapping, then connect to the array.
2. Execute the `auhgmap` command to set the LU Mapping. The following is an example of setting LUN 0 in the array to be recognized as 6 by the host. The port is connected via target group 0 of port 0A on controller 0.

```
% auhgmap -unit array-name -add 0 A 0 6 0
Are you sure you want to add the mapping information?
(y/n [n]): y
The mapping information has been set successfully.
%
```

3. Execute the `auhgmap` command to verify that the LU Mapping is set. For example:

```
% auhgmap -unit array-name -refer
Mapping mode = ON
Port   Group   H-LUN   LUN
  0A   000:000     6     0
%
```

Defining the configuration definition file

The Configuration Definition file describes the system configuration for CCI. It must be set by the user. The configuration definition file is a text file created and/or edited using any standard text editor. It can be defined from the PC where CCI software is installed.

A sample configuration definition file, `HORCM_CONF`, is included with the CCI software. It should be used as the basis for creating your configuration definition file(s). The system administrator should copy the sample file, set the necessary parameters in the copied file, and place the copied file in the proper directory. For more information on configuration definition file, refer to the *Hitachi Unified Storage Command Control Interface Installation and Configuration Guide*.

The configuration definition file can be automatically created using the `mkconf` command tool. For more information on the `mkconf` command, refer to the *Hitachi Unified Storage Command Control Interface Installation and Configuration Guide*. However, the parameters, such as `poll(10ms)` must be set manually (see Step 4 below).

To define the configuration definition file

The following example defines the configuration definition file with two instances on the same Windows host.

1. On the host where CCI is installed, verify that CCI is not running. If CCI is running, shut it down using the `horcmshutdown` command.
2. From the command prompt, make two copies of the sample file (`horcm.conf`). For example:

```
c:\HORCM\etc> copy \HORCM\etc\horcm.conf \WINDOWS\horcm0.conf
c:\HORCM\etc> copy \HORCM\etc\horcm.conf \WINDOWS\horcm1.conf
```

3. Open **horcm0.conf** using the text editor.
4. In the `HORCM_MON` section, set the necessary parameters.

Important: A value more than or equal to 6000 must be set for `poll(10ms)`. Specifying the value incorrectly may cause resource contention in the internal process, resulting the process temporarily suspending and pausing the internal processing of the array.

5. In the HORCM_CMD section, specify the physical drive (command device) on the array. [Figure C-1](#) shows an example of the horcm0.conf file.

```

horcm0.conf - Notepad
File Edit Search Help

HORCM_MON
#ip_address      service      poll(10ms)      timeout(10ms)
XXXXXXXXX        5000            12000           3000

HORCM_CMD
#dev_name        dev_name        dev_name
\\.\PHYSICALDRIVE1

HORCM_DEV
#dev_group       dev_name       port#           TargetID        LU#           MU#
UG01            oradb1         CL1-A           1               1

HORCM_INST
#dev_group       ip_address     service
UG01            XXXXXXXX      5001

```

Figure C-1: Horcm0.conf example

6. Save the configuration definition file and use the horcmstart command to start CCI.
7. Execute the raidscan command; in the result, note the target ID.
8. Shut down CCI and open the configuration definition file again.
9. In the HORCM_DEV section, set the necessary parameters. For the target ID, enter the ID of the raidscan result. For MU#, do not set a parameter.
10. In the HORCM_INST section, set the necessary parameters, and then save (overwrite) the file.
11. Repeat steps 3 to 10 for the **horcm1.conf** file. [Figure C-2](#) shows an example of the horcm1.conf file.




Figure C-2: Horcm1.conf example

12. Enter the following in the command prompt to verify the connection between CCI and the array:

```

C:\>cd horcm\etc

C:\horcm\etc>echo hd1-3 | .\inqraid
Harddisk 1 -> [ST] CL1-A Ser =9000174 LDEV = 0 [HITACHI ] [DF600F-CM ]
Harddisk 2 -> [ST] CL1-A Ser =9000174 LDEV = 1 [HITACHI ] [DF600F ]
                HORC = SMPL HOMRCF[MU#0 = NONE MU#1 = NONE MU#2 = NONE]
                RAID5[Group 1-0] SSID = 0x0000
Harddisk 3 -> [ST] CL1-A Ser =85000175 LDEV = 2 [HITACHI ] [DF600F ]
                HORC = SMPL HOMRCF[MU#0 = NONE MU#1 = NONE MU#2 = NONE]
                RAID5[Group 2-0] SSID = 0x0000

C:\horcm\etc>
  
```

Setting the environment variable

The environment variable must be set up for the execution environment. The following describes an example in which two instances (0 and 1) are configured on the same Windows Server.

1. Set the environment variable for each instance. Enter the following from the command prompt:

```
C:\HORCM\etc>set HORCMINST=0
```

2. Execute the horcmstart script, and then execute the pairedisplay command to verify the configuration. For example:

```
C:\HORCM\etc>horcmstart 0 1
starting HORCM inst 0
HORCM inst 0 starts successfully.
starting HORCM inst 1
HORCM inst 1 starts successfully.

C:\HORCM\etc>pairedisplay -g VG01
group  PairVOL(L/R) (Port#,TID, LU) ,Seq#,LDEV#.P/S,Status,Fence, Seq#,P-LDEV# M
VG01   oradb1(L)    (CL1-A , 1,  1 )91200174    1.SMPL ---- -,----- -
VG01   oradb1(R)    (CL1-A , 1,  2 )91200175    2.SMPL ---- -,----- -
```

Pair operations

This section provides information and instructions for performing TrueCopy operations.

Multiple CCI requests and order of execution

Commands issued from CCI are generally executed in the order they are received in the volumes. However, when more than five commands are issued per controller, the execution order changes.




Figure C-3: Copy order

For example, VOL 0, 1, 2, 3, 5, 6, 7, 8 begin multi-transmission and start the copy operation almost at the same time. VOL#4 starts copying when one of the volumes (0 to 3) completes the operation. VOL#9 starts copying when one of the volumes (5 to 8) completes the operation.

Operations and pair status

Each TrueCopy operation requires a specific pair status. [Figure C-4](#) shows the relationship between status and operations.

Confirm that the state of the remote path is **Normal** before doing a pair operation. If you do a pair operation when a remote path is not set up or its state is **Diagnosing** or **Blocked**, the pair operation may not be completed correctly.




Figure C-4: TrueCopy pair status transitions

Confirming pair status

Table C-2 shows the related CCI and Navigator 2 GUI pair status.

Table C-2: CCI / Navigator 2 GUI pair status

CCI	Navigator 2	Description
SMPL	Simplex	Status where a pair is not created.
COPY	Synchronizing	Initial copy or resynchronization copy is in execution.
PAIR	Paired	Status that the copying is completed and the contents written to the P-VOL is reflected in the S-VOL.
PSUS/SSUS	Split	Status that the written contents are managed as differential data by split.
SSWS	Takeover	Takeover
PSUE	Failure	Status that suspends copying forcibly when a failure occurs.

To confirm TrueCopy pairs

For the example below, the group name in the configuration definition file is VG01.

1. Execute the `pairdisplay` command to verify the pair status and the configuration. For example:

```
c:\HORCM\etc>pairdisplay -g VG01
Group  PairVol(L/R) (Port#,TID, LU) ,Seq#,LDEV#.P/S,Status,Fence, Seq#,P-LDEV# M
VG01   oradb1(L)  (CL1-A , 1, 1 )91200174    1.P-VOL COPY Never ,91200175    2  -
VG01   oradb1(R)  (CL1-A , 1, 2 )91200175    2.S-VOL COPY Never ,-----    1  -
```

The pair status is displayed. For details on the `pairdisplay` command and its options, refer to *Hitachi Unified Storage Command Control Interface Installation and Configuration Guide*.

Creating pairs (paircreate)

In the examples, VG01 is the group name in the configuration definition file.

To create TrueCopy pairs



NOTE: A pair created using CCI and defined in the configuration definition file appear unnamed in the Navigator 2 GUI.

1. Execute the `pairdisplay` command to verify that the status of the possible volumes to be copied is SMPL. For example:

```
C:\HORCM\etc>pairdisplay -g VG01
Group  PairVol(L/R) (Port#,TID, LU) ,Seq#,LDEV#.P/S,Status,Fence, Seq#,P-LDEV# M
VG01   oradb1(L)  (CL1-A , 1, 1 )91200174    1.SMPL -----,-----    -
VG01   oradb1(R)  (CL1-A , 1, 2 )91200175    2.SMPL -----,-----    -
```

2. Execute the `paircreate` command. The `-c` option (medium) is recommended when specifying copying pace. See [Copy pace on page 15-4](#) for more information.
3. Execute the `pairevtwait` command to verify that the status of each volume is PAIR. The following example shows the `paircreate` and `pairevtwait` commands.

```
C:\HORCM\etc>paircreate -g VG01 -f never -vl -c 10
C:\HORCM\etc>pairevtwait -g VG01 -s pair -t 300 10
pairevtwait : Wait status done.
```

4. Execute the `pairstatus` command to verify pair status and the configuration. For example:

```
C:\HORCM\etc>pairstatus -g VG01
Group PairVol(L/R) (Port#,TID, LU) ,Seq#,LDEV#.P/S,Status,Fence, Seq#,P-LDEV# M
VG01 oradb1(L) (CL1-A , 1, 1 )91200174 1 .P-VOL COPY Never ,91200175 2 -
VG01 oradb1(R) (CL1-A , 1, 2 )91200175 2.S-VOL COPY Never ,----- 1 -
```

Creating pairs that belong to a group

In the examples, VG01 is the group name in the configuration definition file.

The examples below show how to create a pair using the consistency group (CTG).



NOTE: Consistency groups created using CCI and defined in the configuration definition file are not seen in the Navigator 2 GUI. Also, pairs assigned to groups using CCI appear ungrouped in the Navigator 2 GUI.

1. Execute the `pairstatus` command to verify that volume status is SMPL. For example:

```
C:\HORCM\etc>pairstatus -g VG01
Group PairVol(L/R) (Port#,TID, LU) ,Seq#,LDEV#.P/S,Status,Fence, Seq#,P-LDEV# M
VG01 oradb1(L) (CL1-A , 1, 1 )91200174 1.SMPL -----,----- ---- -
VG01 oradb1(R) (CL1-A , 1, 2 )91200175 2.SMPL -----,----- ---- -
VG01 oradb2(L) (CL1-A , 1, 3 )91200174 3.SMPL -----,----- ---- -
VG01 oradb2(R) (CL1-A , 1, 4 )91200175 4.SMPL -----,----- ---- -
VG01 oradb3(L) (CL1-A , 1, 5 )91200174 5.SMPL -----,----- ---- -
VG01 oradb3(R) (CL1-A , 1, 6 )91200175 6.SMPL -----,----- ---- -
```

2. Execute the `paircreate -fg` command, then execute the `pairevtwait` command to verify that the status of each volume is PAIR. For example:

```
C:\HORCM\etc>paircreate -g VG01 -f never -vl -m fg -c 10
C:\HORCM\etc>pairevtwait -g VG01 -s pair -t 300 10
pairevtwait : Wait status done.
```

3. Execute the `pairstatus` command to verify the pair status and the configuration. For example:

```
C:\HORCM\etc>pairstat -g VG01
Group PairVol(L/R) (Port#,TID, LU) ,Seq#,LDEV#.P/S,Status,Fence, Seq#,P-LDEV# M
VG01 oradb1(L) (CL1-A , 1, 1 )91200174 1.P-VOL COPY Never ,91200175 2 -
VG01 oradb1(R) (CL1-A , 1, 2 )91200175 2.S-VOL COPY Never ,----- 1 -
VG01 oradb2(L) (CL1-A , 1, 3 )91200174 3.P-VOL COPY Never ,91200175 4 -
VG01 oradb2(R) (CL1-A , 1, 4 )91200175 4.S-VOL COPY Never ,----- 3 -
VG01 oradb3(L) (CL1-A , 1, 5 )91200174 5.P-VOL COPY Never ,91200175 6 -
VG01 oradb3(R) (CL1-A , 1, 6 )91200175 6.S-VOL COPY Never ,----- 5 -
```

Splitting pairs (pairsplit)

In the examples, VG01 is the group name in the configuration definition file.

Two or more pairs can be split at the same time if they are in the same consistency group.

To split pairs

1. Execute the pairsplit command to split the TrueCopy pair in the PAIR status. For example:

```
C:\HORCM\etc>pairsplit -g VG01
```

2. Execute the pairstat command to verify the pair status and the configuration. For example:

```
c:\horcm\etc>pairstat -g VG01
Group PairVol(L/R) (Port#,TID, LU) ,Seq#,LDEV#.P/S,Status,Fence, Setup-LDEV# M
VG01 oradb1(L) (CL1-A , 1, 1 )91200174 1.P-VOL PSUS Never ,91200175 2 -
VG01 oradb1(R) (CL1-A , 1, 2 )91200175 2.S-VOL SSUS Never ,----- 1 -
```

Resynchronizing pairs (pairresync)

In the examples, VG01 is the group name in the configuration definition file.

To resynchronize pairs

1. Execute the pairresync command. Enter between 1 to 15 for copy pace, 1 being slowest (and therefore best I/O performance), and 15 being fastest (and therefore lowest I/O performance). A medium value is recommended.
2. Execute the pairevtwait command to verify that the status of each volume is PAIR. The following example shows the pairresync and the pairevtwait commands.

```
C:\HORCM\etc>pairresync -g VG01 -c 10
C:\HORCM\etc>pairevtwait -g VG01 -s pair -t 300 10
pairevtwait : Wait status done.
```

3. Execute the pairstat command to verify the pair status and the configuration. For example:

```
c:\horcm\etc>pairstat -g VG01
Group PairVol(L/R) (Port#,TID, LU) ,Seq#,LDEV#.P/S,Status,Fence, Seq#,P-LDEV# M
VG01 oradb1(L) (CL1-A , 1, 1 )91200174 1.P-VOL PAIR NEVER ,91200175 2 -
VG01 oradb1(R) (CL1-A , 1, 2 )91200175 2.S-VOL PAIR NEVER ,----- 1 -
```

Suspending pairs (pairsplit -R)

In the examples, VG01 is the group name in the configuration definition file.

To suspend pairs

1. Execute the pairdisplay command to verify that the pair to be suspended is in PAIR status. For example:

```
c:\horcm\etc>pairdisplay -g VG01
Group  PairVol(L/R) (Port#,TID, LU) ,Seq#,LDEV#.P/S,Status,Fence, Seq#,P-LDEV# M
VG01   oradb1(L)   (CL1-A , 1,  1 )91200174    1.P-VOL PAIR NEVER ,91200175    2  -
VG01   oradb1(R)   (CL1-A , 1,  2 )91200175    2.S-VOL PAIR NEVER ,-----    1  -
```

2. Execute the pairsplit -R command to split the pair. For example:

```
C:\HORCM\etc>pairsplit -g VG01 -R
```

3. Execute the pairdisplay command to verify that the P-VOL pair status changed to PSUE. For example:

```
c:\horcm\etc>pairdisplay -g VG01
Group  PairVol(L/R) (Port#,TID, LU) ,Seq#,LDEV#.P/S,Status,Fence, Seq#,P-LDEV# M
VG01   oradb1(L)   (CL1-A , 1,  1 )91200174    1.P-VOL PSUE NEVER ,91200175    2  -
VG01   oradb1(R)   (CL1-A , 1,  2 )91200175    2.SMPL ----- ,-----    -
```

Releasing pairs (pairsplit -S)

In the examples, VG01 is the group name in the configuration definition file.

To release pairs and change status to SMPL

1. Execute the pairsplit -S command to release the pair. For example:

```
C:\HORCM\etc>pairsplit -g VG01 -S
```

2. Execute the pairdisplay command to verify that the pair status changed to SMPL. For example:

```
c:\horcm\etc>pairdisplay -g VG01
Group  PairVol(L/R) (Port#,TID, LU) ,Seq#,LDEV#.P/S,Status,Fence, Seq#,P-LDEV# M
VG01   oradb1(L)   (CL1-A , 1,  1 )91200174    1.SMPL ----- ,-----    -
VG01   oradb1(R)   (CL1-A , 1,  2 )91200175    2.SMPL ----- ,-----    -
```

Mounting and unmounting a volume

When Windows 2000 Server and Windows Server 2008 are used, the CCI mount command is required to mount or un-mount a volume. The GUID, which is displayed by mountvol command, is needed as an argument when using the mount command. For more information, see the *Hitachi Unified Storage Command Control Interface Installation and Configuration Guide*.



TrueCopy Extended Distance reference information

This appendix contains:

- ❑ [TCE system specifications](#)
- ❑ [Operations using CLI](#)
- ❑ [Operations using CCI](#)
- ❑ [Initializing Cache Partition when TCE and Snapshot are installed](#)
- ❑ [Wavelength Division Multiplexing \(WDM\) and dark fibre](#)

TCE system specifications

Table D-1 describes the TCE specifications.

Table D-1: TCE specifications

Parameter	TCE Specification
User interface	• Navigator 2 GUI: used for the setting of DP pool, remote paths, or command devices, and for the pair operations. • Navigator 2 CLI • CCI: used for the pair operations
Controller configuration	Configuration of dual controller is required.
Cache memory	• HUS 110: 4 GB/controller • HUS 130: 8 GB/controller • HUS 150: 8, 16 GB/controller
Host interface	Fibre Channel or iSCSI (cannot mix)
Remote path	• One remote path per controller is required—totaling two for a pair. • The interface type of multiple remote paths between local and remote arrays must be the same.
Number of hosts when remote path is iSCSI	Maximum number of connectable hosts per port: 239.
DP pool	• HUS 150/HUS 130, up to 64 DP pools can be specified for one array. • The DP pools must set in each local array and remote array.
Port modes	Initiator and target intermix mode. One port may be used for host I/O and TCE at the same time.
Bandwidth	• Minimum: 1.5 M. • Recommended: 100M or more. • When low bandwidth is used: - The time limit for execution of CCI commands and host I/O must be extended. - Response time for CCI commands may take several seconds.
License	Entry of the key code enables TCE to be used. TrueCopy and TCE cannot coexist and the licenses to use them are different from each other.
Command device (CCI only)	• Required for CCI. • Minimum size: (33 MB; 65,538 blocks (1 block = 512 bytes)) • Must be set up on local and remote arrays. • Maximum # allowed per array: 128
Unit of pair management	• Volumes are the target of TCE pairs, and are managed per volume

Table D-1: TCE specifications (Continued)

Parameter	TCE Specification
Maximum # of volumes that can be used for TCE pairs	- HUS 110: 2,046 volume - HUS 150/HUS 130: 4,094 volume The maximum number of volumes when different types of arrays are combined is that of the array whose maximum number of volumes is smaller
Pair structure	One S-VOL per P-VOL.
Supported RAID level	- RAID 1 (1D+1D), RAID 5 (2D+1P to 15D+1P) - RAID 1+0 (2D+2D to 8D+8D) - RAID 6 (2D+2P to 28D+2P)
Combination of RAID levels	Local RAID level can be different than remote level. The number of data disks does not have to be the same.
Size of volumes	The volume size must always be P-VOL = S-VOL. The max volume size is 128 TB..
Types of drive for P-VOL, S-VOL, and DP pool	If the drive types are supported by the array, they can be set for a P-VOL, an S-VOL, and DP pools. SAS, SAS7.2K, or SSD/FMD drives are recommended. Set all configured volumes using the same drive type.
Supported capacity value of P-VOL and S-VOL	Capacity is limited.
Copy pace	User adjustable rate that data is copied to remote array. See the copy pace step on page 20-5 for more information.
Consistency Group (CTG)	- Maximum allowed: 64 - Maximum number of pairs allowed per consistency group: - HUS 110: 2,046 - HUS 150/HUS 130: 4,094
Management of volumes while using TCE	For all the P-VOLs and S-VOLs creating pairs, deletion of a RAID group, deletion of a volume, deletion of DP pool, volume formatting, and growing or shrinking of a volume cannot be done. When performing any of these operations, perform it after deleting the TCE pairs
RAID group deleting, volume deleting, and formatting for a paired P-VOL or S-VOL	RAID group deleting, volume deleting, DP pool deleting, and formatting for a paired P-VOL or S-VOL cannot be performed.
Pair creation using unified volumes	- A TCE pair can be created using a unified volume. - When array firmware is less than 0920/B, the size of each volume making up the unified volume must be 1 GB or larger. - When the array firmware is 0920/B or later, there are no restrictions on the volumes making up the unified volume. - Volumes that are already in a P-VOL or S-VOL cannot be unified. - Unified Volumes that are in a P-VOL or S-VOL cannot be released.

Table D-1: TCE specifications (Continued)

Parameter	TCE Specification
Restriction during RAID group expansion	A RAID group in which a TCE P-VOL or DP pool exists can be expanded only when pair status is Simplex or Split. If the TCE DP pool is shared with Snapshot, the Snapshot pairs must be in Simplex or Paired status.
Pair creation of a unified volume	A TCE pair can be created by specifying a unified volume. However, unification of volumes or release of the unified volumes cannot be done for the paired volumes
Unified volume for DP pool	Not allowed.
Differential data	When pair status is Split, data sent to the P-VOL and S-VOL are managed as differential data.
Host access to a DP pool	A DP pool volume is hidden from a host.
Expansion of DP pool capacity	A DP pool capacity can be expanded by adding a RAID group to the DP pool. The capacity can be expanded even when the DP pool is being used by pairs. However, the RAID group with different drive types cannot be mixed
Reduction of DP pool capacity	Yes. The pairs associated with a DP pool must be deleted before the DP pool can be reduced. The capacity can be reduced by adding the necessary capacity or RAID groups again after deleting all RAID groups set to the DP pool once.
Failures	- When the copy operation from P-VOL to S-VOL fails, TCE suspends the pair (Failure). Because TCE copies data to the remote S-VOL regularly, data is restored to the S-VOL from the update immediately before the occurrence of the failure. - A drive failure does not affect TCE pair status because of the RAID architecture.
Depletion of a DP pool	When the usage rate of the DP pool in local array reaches the Replication Data Released threshold or the usage rate of the DP pool in remote array reaches the Replication Data Released threshold, the status of any pair becomes "Pool Full" and the P-VOL data becomes unable to update the S-VOL data
Cycle time	If necessary, the cycle time, which updates an S-VOL using the differential data when the pair status is "Paired," can be changed. The default cycle time is 300 seconds; and the cycle can be specified by the second up to 3,600 seconds. The shortest limit value that can be set is calculated as number of CTGs of local array or remote array × 30 seconds.
Assuring the order in which the data transferred to the S-VOL is written	The differential data is transferred from the P-VOL to the S-VOL in a cycle specified by a user. Therefore, order of the transferred data, which is to be reflected on the S-VOL, in each cycle, is assured

Table D-1: TCE specifications (Continued)

Parameter	TCE Specification
Array restart at TCE installation	The array is restarted after installation to set the DP pool, unless it is also used by Snapshot. Then there is no restart.
TCE use with TrueCopy	Not allowed.
TCE use with Snapshot	Snapshot can be cascaded with TCE or used separately. Only a Snapshot P-VOL can be cascaded with TCE.
TCE use with ShadowImage	Although TCE can be used at the same time as a ShadowImage system, it cannot be cascaded with ShadowImage.
TCE use with LUN Expansion	When firmware version is less than 0920/B, it is not allowed to create a TCE pair using unified volumes, which unify the volume with 1 GB or less capacity.
TCE use with Data Retention Utility	Allowed. - When S-VOL Disable is set for an volume, a pair cannot be created using the volume as the S-VOL. - S-VOL Disable can be set for an volume that is currently an S-VOL, if pair status is Split.
TCE use with Cache Residency Manager	Available. However, a volume specified by Cache Residency Manager cannot be specified as a P-VOL or an S-VOL.
TCE use with Cache Partition Manager	- TCE can be used together with Cache Partition Manager. - Make the segment size of volumes to be used as a TCE DP pool no larger than the default, (16 KB). See Initializing Cache Partition when TCE and Snapshot are installed on page D-43 for details on initialization.
TCE use with SNMP Agent	Allowed. A trap is transmitted for the following: - Remote path failure. - Threshold value of the DP pool is exceeded. - Actual cycle time exceeds the default or user-specified value. - Pair status changes to: - Pool Full - Failure. - Inconsistent because the DP pool is full or because of a failure.
TCE use with Volume Migration	Allowed. However, a Volume Migration P-VOL, S-VOL, or Reserved volume cannot be used as a TCE P-VOL or S-VOL.
Concurrent use of TrueCopy	Not available.
Concurrent use of TCMD	By using TCMD together with TCE, you can set the remote paths among nine arrays and can create TCE pairs. For more detail, see TrueCopy Modular Distributed overview on page 22-2 .
Concurrent use of ShadowImage	Though TCE can be used together with ShadowImage, it cannot be cascaded with ShadowImage.

Table D-1: TCE specifications (Continued)

Parameter	TCE Specification
Concurrent use of Snapshot	TCE can be used together with Snapshot and cascaded with only a Snapshot P-VOL. Also, the number of volumes that can be paired is limited to the maximum number or less depending on the number of Snapshot P-VOL's
Concurrent use of Power Saving/ Power Saving Plus	Available. However, when the P-VOL or the S-VOL is included in a RAID group for which the Power Saving/Power Saving Plus has been specified, no pair operation can be performed except the pair split and the pair delete.
Concurrent use of Dynamic Provisioning	Available. For details, see Concurrent use of Dynamic Provisioning on page 19-45 .
Concurrent use of Dynamic Tiering	Available. For details, see Concurrent use of Dynamic Tiering on page 19-49 .
Reduction of memory	Reduce memory only after disabling TCE.
Load balancing function	The Load balancing function applies to a TCE pair. When the pair state is Paired and cycle copy is being performed, the load balancing function does not work.
Reduction of memory	The memory cannot be reduced when the ShadowImage, SnapShot, TCE, or Volume Migration function is validated. Make the reduction after invalidating the function.
Remote Copy over iSCSI in the WAN environment	We recommend using TrueCopy in the WAN environment of MTU1500 or more. However, if TCE needs to be implemented in the WAN environment of less than MTU1500, change the maximum segment size (MSS) of the WAN router to a value less than 1500, and then create a remote path. The data length transmitted from TCE of Hitachi Unified Storage changes to the specified value less than 1500. When creating a remote path without changing the MSS value or not creating the remote path after changing the MSS value again, a data transfer error occurs because TCE of HUS100 transmits the MTU1500 data. To change the MSS value, request the customer or the WAN router provider.

Operations using CLI

This section describes CLI procedures for setting up and performing TCE operations.

- ❑ [Installation and setup](#)
- ❑ [Pair operations](#)
- ❑ [Procedures for failure recovery](#)
- ❑ [Sample script](#)



NOTE: For additional information on the commands and options in this appendix, see the *Hitachi Unified Storage Command Line Interface Reference Guide*.

Installation and setup

The following sections provide installation/uninstalling, enabling/disabling, and setup procedures using CLI.

TCE is an extra-cost option and must be installed using a key code or file. Obtain it from the download page on the HDS Support Portal, <https://portal.hds.com>. See [Installation procedures on page 18-3](#) for prerequisites information.

Before installing or uninstalling TCE, verify the following:

- The array must be operating in a normal state. Installation and un-installation cannot be performed if a failure has occurred.
- Make sure that a spin-down operation is not in progress when installing or uninstalling TCE.
- The array may require a restart at the end of the installation procedure. If Snapshot is already enabled, no restart is necessary. If restart is required, it can be done when prompted, or at a later time.
- TCE cannot be installed if more than 239 hosts are connected to a port on the array.

Installing

To install TCE

1. From the command prompt, register the array in which the TCE is to be installed, and then connect to the array.
2. Execute the `auopt` command to install TCE. For example:

```
% auopt -unit array-name -lock off -keycode manual-attached-keycode
Are you sure you want to unlock the option? (y/n [n]): y
The option is unlocked.
A DP pool is required to use the installed function.
Create a DP pool before you use the function.

%
%
```

3. Execute the `auopt` command to confirm whether TCE has been installed. For example:

```
% auopt -unit array-name -refer
Option Name          Type      Term      Status
      Reconfigure Memory Status
TC-EXTENDED          Permanent ---      Enable
      N/A

%
```

TCE is installed and **Status** is **Enabled**. Installation of TCE is now complete.



NOTE: TCE needs the DP pool of Dynamic Provisioning. If Dynamic Provisioning is not installed, install Dynamic Provisioning.

Enabling and disabling

TCE can be disabled or enabled. When TCE is first installed it is automatically enabled.

Prerequisites for disabling

- TCE pairs must be released (the status of all volumes must be Simplex).
- The remote path must be deleted, unless TrueCopy continues to be used.
- TCE cannot be enabled if more than 239 hosts are connected to a port on the array.

To enable / disable TCE

1. From the command prompt, register the array in which the status of the feature is to be changed, and then connect to the array.
2. Execute the auopt command to change TCE status (enable or disable).

The following is an example of changing the status from enable to disable. If you want to change the status from disable to enable, enter enable after the -st option.

```
% auopt -unit array-name -option TC-EXTENDED -st disable
Are you sure you want to disable the option? (y/n [n]): y
The option has been set successfully.

%
```

3. Execute the auopt command to confirm that the status has been changed. For example:

```
% auopt -unit array-name -refer
Option Name          Type      Term      Status
      Reconfigure Memory Status
TC-EXTENDED          Permanent ---      Disable
      Reconfiguring(10%)

%
```

Enabling or disabling TCE is now complete.

Un-installing TCE

To uninstall TCE, the key code or key file provided with the optional feature is required. Once uninstalled, TCE cannot be used (locked) until it is again installed using the key code or key file.

Prerequisites for uninstalling

- TCE pairs must be released (the status of all volumes must be Simplex).
- The remote path must be released, unless TrueCopy continues to be used.

To uninstall TCE

1. From the command prompt, register the array in which the TCE is to be uninstalled, and then connect to the array.
2. Execute the `auopt` command to uninstall TCE. For example:

```
% auopt -unit array-name -lock on -keycode manual-attached-keycode
Are you sure you want to lock the option? (y/n [n]): y
The option is locked.
.%
```

3. Execute the `auopt` command to confirm that TCE is uninstalled. For example:

```
% auopt -unit array-name -refer
DMEC002015: No information displayed.
%
```

Uninstalling TCE is now complete

Setting the DP pool

For instructions to set a DP pool, refer to the *Hitachi Unified Storage Dynamic Provisioning Configuration Guide*.

Setting the replication threshold

To set the Depletion Alert and/or the Replication Data Released of replication threshold:

1. From the command prompt, execute the `audpool` command, change the Depletion Alert and/or the Replication Data Released of replication threshold.

Example: This example shows changing the Depletion Alert threshold.

```
% audpool -unit array-name -chg -dppoolno 0 -repdepletion_alert 50
Are you sure you want to change the DP pool attribute? (y/n [n]): y
DP pool attribute changed successfully.
%
%
```

2. Execute the `audpool` command to confirm the DP pool attribute.


```

% audppool -unit array-name -refer -detail -dppoolno 0 -t
DP Pool : 0
RAID Level           : 6(6D+2P)
Page Size            : 32MB
Stripe Size          : 256KB
Type                 : SAS
Status               : Normal
Reconstruction Progress : N/A
Capacity
  Total Capacity      : 8.9 TB
  Consumed Capacity
    Total             : 2.2 TB
    User Data         : 0.7 TB
    Replication Data   : 0.4 TB
    Management Area    : 0.5 TB
    Needing Preparation Capacity : 0.0 TB
DP Pool Consumed Capacity Alert
  Early Alert         : 40%
  Depletion Alert     : 50%
  Notifications Active : Enable
Over Provisioning Threshold
  Warning             : 100%
  Limit               : 130%
  Notifications Active : Enable
Replication Threshold
  Replication Depletion Alert : 50%
  Replication Data Released   : 95%
Defined LU Count       : 0
DP RAID Group
  DP RAID Group    RAID Level    Capacity    Capacity    Percent
                49      6(6D+2P)      8.9 TB      2.2 TB      24%
Drive Configuration
  DP RAID Group    RAID Level Unit   HDU   Type   Capacity   Status
                49      6(6D+2P)    0     0   SAS    300GB   Standby
                49      6(6D+2P)    0     1   SAS    300GB   Standby
                :
                :
Logical Unit
  Consumed          Stripe Cache    Pair Cache
Number
  LU pacity    Capacity Consumed % Size    Partition Partition
Status of Paths
%

```

Setting the cycle time

Cycle time is the time between updates to the remote copy when the pair is in Paired status. The default is 300 seconds. You can set cycle time between 30 to 3600 seconds. Set the cycle time for each array. The shortest value that can be set is calculated as number of CTGs of local array or remote array × 30 seconds.

Copying may take a longer than the cycle time, depending on the amount of the differential data or low bandwidth.

To set the cycle time

1. From the command prompt, register the array to which you want to set the cycle time, and then connect to the array.
2. Execute the `autruecopyopt` command to confirm the existing cycle time. For example:

```
% autruecopyopt -unit array-name -refer
Cycle Time[sec.]      : 300
Cycle OVER report     : Disable
%
```

3. Execute the `autruecopyopt` command to set the cycle time. The cycle time is 300 seconds by default and can be specified within a range from 30 to 3600 seconds. For example:

```
% autruecopyopt -unit array-name -set -cycletime 300
Are you sure you want to set the TrueCopy options? (y/n [n]): y
The TrueCopy options have been set successfully.
%
```

Setting mapping information

The following is the procedure for mapping information. *For iSCSI*, use the `autargetmap` command in place of `auhgmap`.

1. From the command prompt, register the array to which you want to set the mapping information, and then connect to the array.
2. Execute the `auhgmap` command to set the mapping information. The following example defines LU 0 in the array to be recognized as 6 by the host. The port is connected via host group 0 of port 0A on controller 0.

```
% auhgmap -unit array-name -add 0 A 0 6 0
Are you sure you want to add the mapping information? (y/n [n]): y
The mapping information has been set successfully.
%
```

3. Execute the `auhgmap` command to verify that the mapping information has been set. For example:

```
% auhgmap -unit array-name -refer
Mapping mode = ON
Port  Group  H-LUN    LUN
  0A  000:000    6      0
%
```

Setting the remote port CHAP secret

iSCSI systems only. The remote path can employ a CHAP secret. Set the CHAP secret mode on the remote array. For more information on the CHAP secret, see [Setting the cycle time on page 19-56](#).

The setting procedure of the remote port CHAP secret is shown below.

1. From the command prompt, register the array in which you want to set the remote port CHAP secret, and then connect to the array.
2. Execute the `aurmtpath` command with the `-set` option and perform the CHAP secret of the remote port. The input example and the result are shown below. For example:

```
% aurmtpath -unit array-name -set -target -local 9120027
-secret
Are you sure you want to set the remote path information?
(y/n[n]): y
Please input Path 0 Secret.
Path 0 Secret:
Re-enter Path 0 Secret:
Please input Path 1 Secret.
Path 1 Secret:
Re-enter Path 1 Secret:
The remote path information has been set successfully.
%
```

The setting of the remote port CHAP secret is completed.

Setting the remote path

Data is transferred from the local to the remote array over the remote path.

Please review Prerequisites in [Setting the remote path on page 19-58](#) before proceeding.

To set up the remote path

1. From the command prompt, register the array in which you want to set the remote path, and then connect to the array.
2. The following shows an example of referencing the remote path status where remote path information is not yet specified.

Fibre Channel example:

```
% aurmtpath -unit array-name -refer
Initiator Information
Local Information
  Array ID      : 91200026
  Distributed Mode : N/A

Path Information
  Interface Type      :
  Remote Array ID     :
  Remote Path Name    :

  Bandwidth [0.1 M] :
  iSCSI CHAP Secret  :

  Path  Status      Local      Remote Port      TCP Port No. of
      0  Undefined   ---      ---      IP Address      Remote Port
      1  Undefined   ---      ---      ---              ---
%
```

iSCSI example:

```
% aurmtpath -unit array-name -refer
Initiator Information
Local Information
  Array ID      : 91200026
  Distributed Mode : N/A

Path Information
  Interface Type      : FC
  Remote Array ID     : 91200027
  Remote Path Name    : N/A
  Bandwidth [0.1 M] : 15
  iSCSI CHAP Secret  : N/A

  Path  Status      Local      Remote Port      TCP Port No. of
      0  Undefined   ---      ---      IP Address      Remote Port
      1  Undefined   ---      ---      ---              ---

Target Information
  Local Array ID      :
%
```

3. Execute the aurmtpath command to set the remote path.

Fibre Channel example:

```
% aurmtpath -unit array-name -set -remote 91200027 -band 15
               -path0 0A 0A -path1 1A 1B
Are you sure you want to set the remote path information?
(y/n[n]): y
The remote path information has been set successfully.
%
```

iSCSI example:

```
% aurmtpath -unit array-name -set -initiator -remote 91200027
-secret disable
               -path0 0B -path0_addr 192.168.1.201 -band 100
               -path1 1B -path1_addr 192.168.1.209
Are you sure you want to set the remote path information?
(y/n[n]): y
The remote path information has been set successfully.
%
```

4. Execute the aurmtpath command to confirm whether the remote path has been set. For example:

Fibre Channel example:

```
% aurmtpath -unit array-name -refer
Initiator Information
Local Information
  Array ID      : 91200026
  Distributed Mode : N/A

Path Information
Interface Type      : FC
Remote Array ID     : 91200027
Remote Path Name    : N/A
Bandwidth [0.1 M]  : 15
iSCSI CHAP Secret   : N/A

      Path   Status      Local      Remote Port      TCP Port No. of
      Port                                Remote IP Address  Remote
      0   Normal        0A        0A   N/A              N/A
      1   Normal        1A        1B   N/A              N/A
%
```

iSCSI example:

```
% aurmtpath -unit array-name -refer
```

Initiator Information

Local Information

```
Array ID      : 91200026  
Distributed Mode : N/A
```

Path Information

```
Interface Type      : iSCSI  
Remote Array ID     : 91200027  
Remote Path Name    : N/A  
Bandwidth [0.1 M]  : 100  
iSCSI CHAP Secret   : Disable
```

Path	Status	Local	Remote	Remote Port IP Address	TCP Port No. of Remote Port
0	Normal	0B	N/A	192.168.0.201	3260
1	Normal	1B	N/A	192.168.0.209	3260

Target Information

```
Local Array ID      : 91200026
```

```
%
```

Deleting the remote path

When the remote path becomes unnecessary, delete the remote path.

Prerequisites

- To delete the remote path, change all the TCE pairs in the array to **Simplex** or **Split** status.
- Do not do a pair operation for a TCE pair when the remote path for the pair is not set up, because then the pair operation may not complete correctly.



NOTE: When performing the planned shutdown of the remote array, the remote path should not necessarily be deleted. Change all the TCE pairs in the array to the Split status, and then perform the planned shutdown of the remote array. After restarting the array, perform the pair resynchronization. However, when the Warning notice to the Failure Monitoring Department at the time of the remote path blockade or the notice by the SNMP Agent Support Function or the E-mail Alert Function is not desired, delete the remote path, and then turn off the power of the remote array.

To delete the remote path

1. From the command prompt, register the array in which you want to delete the remote path, and then connect to the array.
2. Execute the aurmtpath command to delete the remote path. For example:

```
% aurmtpath -unit array-name -rm -remote 91200027
Are you sure you want to delete the remote path information?
(y/n[n]): y
The remote path information has been deleted successfully.
%
```

3. Execute the aurmtpath command to confirm that the path is deleted. For example:

```
% aurmtpath -unit array-name -refer
Initiator Information
Local Information
  Array ID      : 91200026
  Distributed Mode : N/A

Path Information
  Interface Type      :
  Remote Array ID     :
  Remote Path Name    :
  Bandwidth [0.1 M]  :
  iSCSI CHAP Secret  :

Path  Status      Local      Remote      Remote Port      TCP Port No. of
      0  Undefined  ---      ---      IP Address      Remote Port
      1  Undefined  ---      ---      ---              ---
%
```

Pair operations

The following sections describe the CLI procedures and commands for performing TCE operations.

The `aureplicationremote` command operates TCE pair. To refer to the `aureplicationremote` command and its options, type in `aureplicationremote -help`.

Confirm that the state of the remote path is **Normal** before doing a pair operation. If you do a pair operation when a remote path is not set up or its state is **Diagnosing** or **Blocked**, the pair operation may not be completed correctly.

Displaying status for all pairs

To display all pair status

1. From the command prompt, register the array to which you want to display the status of paired logical volumes. Connect to the array.
1. Execute the `aureplicationremote -refer` command. For example:

% aureplicationremote -unit local array-name -refer					
Pair name		Local LUN	Attribute	Remote LUN	Status
	Copy Type		Group Name		
TCE_LU0000_LU0000		0	P-VOL	0	Paired(100
%)	TrueCopy Extended Distance	0:			
TCE_LU0001_LU0001		1	P-VOL	1	Paired(100
%)	TrueCopy Extended Distance	0:			
%					

Displaying detail for a specific pair

To display pair details

1. From the command prompt, register the array to which you want to display the status and other details for a pair. Connect to the array.

2. Execute the `aureplicationremote -refer -detail` command to display the detailed pair status. For example:

```
% aureplicationremote -unit local array-name -refer -detail -pvol 0 -svol 0
                        -locallun pvol -remote 91200027
Pair Name               : TCE_LU0000_LU0000
Local Information
  LUN                   : 0
  Attribute              : P-VOL
  DP Pool
    Replication Data    : 0
    Management Area     : 0
Remote Information
  Array ID              : 91200027
  Path Name             : N/A
  LUN                   : 0
  Capacity              : 50.0 GB
  Status                : Paired(100%)
  Copy Type             : TrueCopy Extended Distance
  Group Name            : 0:
  Consistency Time      : 2011/07/29 11:09:34
  Difference Size        : 2.0 MB
  Copy Pace             : ---
  Fence Level           : N/A
  Previous Cycle Time    : 504 sec.
%
```

Creating a pair

See prerequisite information under [Creating the initial copy on page 20-2](#) before proceeding.

To create a pair

1. From the command prompt, register the local array in which you want to create pairs, and then connect to the array.
2. Execute the `aureplicationremote -refer -availablelist` command to display volumes available for copy as the P-VOL. For example:

```
% aureplicationremote -unit local array-name -refer -availablelist -tce -
pvol
Available Logical Units
  LUN  Capacity RAID Group DP Pool RAID Level  Type Status
    2   50.0 GB           0   N/A    6( 9D+2P) SAS  Normal
%
```

3. Execute the `aureplicationremote -refer -availablelist` command to display volumes on the remote array that are available as the S-VOL. For example:

```
% aureplicationremote -unit remote array-name -refer -availablelist -tce -
pvol
Available Logical Units
  LUN  Capacity RAID Group DP Pool RAID Level  Type Status
    2   50.0 GB           0   N/A    6( 9D+2P) SAS  Normal
%
```

4. Specify the volumes to be paired and create a pair using the `aureplicationremote -create` command. For example:

```
% aureplicationremote -unit local array-name -create -tce -pvol 2 -svol 2
                        -remote xxxxxxxx -gno 0 -remotepoolno 0
Are you sure you want to create pair "TCE_LU0002_LU0002"?
(y/n [n]): y
The pair has been created successfully.
%
```

Splitting a pair

A pair split operation on a pair belonging to a group results in all pairs in the group being split.

To split a pair

1. From the command prompt, register the local array in which you want to split pairs, and then connect to the array.
2. Execute the `aureplicationremote -split` command to split the specified pair. For example:

```
% aureplicationremote -unit local array-name -split -tce -localvol 2 -
remotevol 2
                        -remote xxxxxxxx -locallun pvol
Are you sure you want to split pair? (y/n [n]): y
The split of pair has been required.
%
```

Resynchronizing a pair

To resynchronize a pair

1. From the command prompt, register the local array in which you want to re-synchronize pairs, and then connect to the array.
2. Execute the `aureplicationremote -resync` command to re-synchronize the specified pair. For example:

```
% aureplicationremote -unit local array-name -resync -tce -pvol 2 -svol 2
                        -remote xxxxxxxx
Are you sure you want to re-synchronize pair?
(y/n [n]): y
The pair has been re-synchronized successfully.
%
```

Swapping a pair

Please review the Prerequisites in [Swapping pairs on page 20-9](#).

To swap the pairs, the remote path must be set to the local array from the remote array.

To swap a pair

1. From the command prompt, register the remote array in which you want to swap pairs, and then connect to the array.

2. Execute the `aureplicationremote -swaps` command to swap the specified pair. For example:

```
% aureplicationremote -unit remote array-name -swaps -tce -gno 1
Are you sure you want to swap pair?
(y/n [n]): y
The pair has been swapped successfully.
%
```

Deleting a pair

To delete a pair

1. From the command prompt, register the local array in which you want to delete pairs, and then connect to the array.
2. Execute the `aureplicationremote -simplex` command to delete the specified pair. For example:

```
% aureplicationremote -unit local array-name -simplex -tce -locallun pvol
                        -pvol 2 -svol 2 -remote xxxxxxxx
Are you sure you want to release pair?
(y/n [n]): y
The pair has been released successfully.
%
```

Changing pair information

You can change the pair name, group name, and/or copy pace.

1. From the command prompt, register the local array on which you want to change the TCE pair information, and then connect to the array.
2. Execute the `aureplicationlocal -chg` command to change the TCE pair information. In the following example, change the copy pace from normal to slow.

```
% aureplicationremote -unit local array-name -tce -chg -pace slow
                        -locallun pvol -pvol 2000 -svol 2002 -remote xxxxxxxx
Are you sure you want to change pair information?
(y/n [n]): y
The pair information has been changed successfully.
%
```

Monitoring pair status

To monitor pair status

1. From the command prompt, register the local array on which you want to monitor pair status, and then connect to the array.
2. Execute the `aureplicationmon -evwait` command. For example:

```
% aureplicationmon -unit local array-name -evwait -tce -st simplex -gno 0
                  -waitmode backup
Simplex Status Monitoring...
Status has been changed to Simplex.
%
```

Confirming consistency group (CTG) status

You can display information about a consistency group using the `aureplicationremote` command. The information is displayed in a list.

To display consistency group status

1. From the command prompt, register the local array on which you want to view consistency group status, and then connect to the array.
2. Execute the `aureplicationremote -unit unit_name -refer - groupinfo` command. For example:

Descriptions of the consistency group information that displays is shown in [Table D-2](#).

Table D-2: CTG information

Displayed item	Contents
CTG No.	CTG number
Lapsed Time	The Lapsed Time after the current cycle is started is displayed (in hours, minutes, and seconds).
Remaining Difference Size	The size of the residual differential data to be transferred in the current cycle is displayed. The size of the differential data in the pair information shows a total size of the data that have not been transferred and thus remains in the local array, whereas the size of the remaining differential data does not include the size of the data to be transferred in the following cycle. Therefore, the size of the remaining differential data does not coincide with the total size of the differential data of the pairs included in the CTG.
Transfer Rate	The Transfer Rate of the current cycle is displayed (KB/s). During a period from the start of the cycle to the execution of the copy operation or a waiting period from completion of the copy operation to the start of the next cycle, "---" is displayed. While the Transfer Rate to be output is being calculated, "Calculating" is displayed.
Prediction Time of Transfer Completion	The predicted time when the data transfer is completed (Prediction Time of Transfer Completion) for each cycle of the CTG is displayed (in hours, minutes, and seconds). If the predicted time when the data transfer is completed (Prediction Time of Transfer Completion) cannot be calculated because it is maximized temporarily, "99:59:59" is displayed. During a waiting period from completion of the cyclic operation till the start of the next cycle, "Waiting" is displayed. While the predicted time to be output is being calculated, "Calculating" is displayed.

Procedures for failure recovery

Displaying the event log

When a failure occurs, you can learn useful information from the event log. The contents of the event log include the time when an error occurred, an error message, and an error detail code.

To display the event log

1. Register the array to confirm the event log on the command prompt.
2. Execute the `auinformsg` command and confirm the event log. For example:

```
% auinformsg -unit array-name
Controller 0/1 Common
12/18/2007 11:32:11 C0 IB1900 Remote copy failed(CTG-00)
12/18/2007 11:32:11 C0 IB1G00 Pair status changed by the
error(CTG-00)
:
12/18/2007 16:41:03 00 I10000 Subsystem is ready

Controller 0
12/17/2007 18:31:48 00 RBE301 Flash program update end
12/17/2007 18:31:08 00 RBE300 Flash program update start

Controller 1
12/17/2007 18:32:37 10 RBE301 Flash program update end
12/17/2007 18:31:49 10 RBE300 Flash program update start
%
```

The event log was displayed. When searching the specified messages or error detail codes, store the output result in the file and use the search function of the text editor as shown below.

```
% auinformsg -unit array-name>infomsg.txt
%
```

Reconstructing the remote path

To reconstruct the remote path

1. Register the array to reconstruct the remote path on the command prompt.
2. Execute the `aurmtpath` command with the `-reconst` option and enable the remote path status. For example:

```
% aurmtpath -unit array-name -reconst -remote 91200027 -path0
Are you sure you want to reconstruct the remote path?
(y/n [n]): y
The reconstruction of remote path has been required.
Please check "Status" as -refer option.
%
```

Sample script

The following example provides sample script commands for backing up a volume on a Windows Server.

```
echo off
REM Specify the registered name of the arrays
set UNITNAME=Array1
REM Specify the group name (Specify "Ungroup" if the pair doesn't belong to any
group)
set G_NAME=Ungrouped
REM Specify the pair name
set P_NAME=TCE_LU0001_LU0002
REM Specify the directory path that is mount point of P-VOL and S-VOL
set MAINDIR=C:\main
set BACKUPDIR=C:\backup
REM Specify GUID of P-VOL and S-VOL
PVOL_GUID=xxxxxxxx-xxxx-xxxx-xxxx-xxxxxxxxxxxxxx
SVOL_GUID=yyyyyyyy-yyy-yyy-yyy-yyy-yyyyyyyyyyyy

REM Unmounting the S-VOL
pairedisplay -x umount %BACKUPDIR%
REM Re-synchronize pair (Updating the backup data)
aureplicationremote -unit %UNITNAME% -tce -resync -pairname %P_NAME% -gno 0
aureplicationmon -unit %UNITNAME% -evwait -tce -pairname %P_NAME% -gno 0 -st
paired -pvol

REM Unmounting the P-VOL
pairedisplay -x umount %MAINDIR%
REM Splitting pair (Determine the backup data)
aureplicationremote -unit %UNITNAME% -tce -split -pairname %P_NAME% -gname
%G_NAME%
aureplicationmon -unit %UNITNAME% -evwait -tce -pairname %P_NAME% -gname
%G_NAME% -st split -pvol
REM Mounting the P-VOL
pairedisplay -x mount %MAINDIR% Volume{%PVOL_GUID%}

REM Mounting the S-VOL
pairedisplay -x mount %BACKUPDIR% Volume{%SVOL_GUID%}
< The procedure of data copy from C:\backup to backup appliance>
```

When Windows Server is used, the CCI mount command is required when mounting or un-mounting a volume. The GUID, which is displayed by the Windows mountvol command, is needed as an argument when using the mount command. For more information, refer to the *Hitachi Unified Storage Command Control Interface Installation and Configuration Guide*.

Operations using CCI

This section describes CCI procedures for setup and performing TCE operations, and provides examples of TCE commands using the Windows Server.

- ❑ [Setting the command device](#)
- ❑ [Setting mapping information](#)
- ❑ [Defining the configuration definition file](#)
- ❑ [Setting the environment variable](#)

Setup

The following sections provide procedures for setting up CCI for TCE.

Setting the command device

The command device is used by CCI to conduct operations on the array.

- Volumes used as command devices must be recognized by the host.
- The command device must be 33 MB or greater.
- Assign multiple command devices to different RAID groups to avoid disabled CCI functionality in the event of drive failure.



If a command device fails, all commands are terminated. CCI supports an alternate command device function, in which two command devices are specified within the same array, to provide a backup. For details on the alternate command device function, refer to the *Hitachi Unified Storage Command Control Interface Installation and Configuration Guide*.

To designate a command device

1. From the command prompt, register the array to which you want to set the command device, and then connect to the array.
2. Execute the aucmddev command to set a command device. When this command is run, LUNS that can be assigned as a command device display, then the command device is set. To use the CCI protection function, enter enable following the -dev option. The following is an example of specifying LUN 2 for command device 1.


```
% aucmddev -unit array-name -availablelist
Available Logical Units
  LUN  Capacity RAID Group DP Pool  RAID Level  Type Status
    2   35.0 MB          0   N/A    6( 9D+2P)  SAS  Normal
    3   35.0 MB          0   N/A    6( 9D+2P)  SAS  Normal
%
% aucmddev -unit array-name -set -dev 1 2
Are you sure you want to set the command devices?
(y/n [n]): y
The command devices have been set successfully.
%
```

3. Execute the aucmddev command to verify that the command device is set. For example:

```
% aucmddev -unit array-name -refer
Command Device LUN  RAID Manager Protect
1              2   Disable
%
```

4. To release a command device, follow the example below, in which command device 1 is released.

```
% aucmddev -unit array-name -rm -dev 1
Are you sure you want to release the command devices?
(y/n [n]): y
This operation may cause the CCI, which is accessing to this
command device, to freeze.
Please make sure to stop the CCI, which is accessing to this
command device, before performing this operation.
Are you sure you want to release the command devices? (y/n
[n]): y
The specified command device will be released.
Are you sure you want to execute? (y/n [n]): y
The command devices have been released successfully.
%
```

5. To change a command device, first release it, then change the volume number. The following example of specifies LUN 3 for command device 1.

```
% aucmddev -unit array-name -set -dev 1 3
Are you sure you want to set the command devices?
(y/n [n]): y
The command devices have been set successfully.
%
```

Setting mapping information

For iSCSI, use the autargetmap command instead of the auhgmap command.

To set up LU Mapping

1. From the command prompt, register the array to which you want to set the LU Mapping, then connect to the array.
2. Execute the auhgmap command to set the mapping information. The following is an example of setting LUN 0 in the array to be recognized as 6 by the host. The port is connected via target group 0 of port 0A on controller 0.

```
% auhgmap -unit array-name -add 0 A 0 6 0
Are you sure you want to add the mapping information?
(y/n [n]): y
The mapping information has been set successfully.
%
```

3. Execute the auhgmap command to verify that the LU Mapping is set. For example:

```
% auhgmap -unit array-name -refer
Mapping mode = ON
Port  Group                      H-LUN  LUN
  0A   000:G000                   6      0
%
```

Defining the configuration definition file

The configuration definition file describes system configuration. It is required to make CCI operational. The configuration definition file is a text file created and/or edited using any standard text editor. It can be defined from the PC where CCI software is installed.

A sample configuration definition file, HORCM_CONF, is included with the CCI software. It should be used as the basis for creating your configuration definition file(s). The system administrator should copy the sample file, set the necessary parameters in the copied file, and place the copied file in the proper directory. For more information on configuration definition file, refer to the *Hitachi Unified Storage Command Control Interface Installation and Configuration Guide*.

The configuration definition file can be automatically created using the mkconf command tool. For more information on the mkconf command, refer to the *Hitachi Unified Storage Command Control Interface Installation and Configuration Guide*. However, the parameters, such as poll(10ms) must be set manually (see Step 4 below).

To define the configuration definition file

The following example defines the configuration definition file with two instances on the same Windows host.

1. On the host where CCI is installed, verify that CCI is not running. If CCI is running, shut it down using the horcmshutdown command.
2. From the command prompt, make two copies of the sample file (horcm.conf). For example:

```
c:\HORCM\etc> copy \HORCM\etc\horcm.conf
\WINDOWS\horcm0.conf
c:\HORCM\etc> copy \HORCM\etc\horcm.conf
\WINDOWS\horcm1.conf
```

3. Open **horcm0.conf** using the text editor.
4. In the HORCM_MON section, set the necessary parameters.
Important: A value more than or equal to 6000 must be set for poll(10ms). Specifying the value incorrectly may cause resource contention in the internal process, resulting the process temporarily suspending and pausing the internal processing of the array.
5. In the HORCM_CMD section, specify the physical drive (command device) on the array. [Figure D-1](#) shows an example of the horcm0.conf file.




Figure D-1: Horcm0.conf example

6. Save the configuration definition file and use the horcmstart command to start CCI.
7. Execute the raidscan command; in the result, note the target ID.
8. Shut down CCI and open the configuration definition file again.
9. In the HORCM_DEV section, set the necessary parameters. For the target ID, enter the ID of the raidscan result. For MU#, do not set a parameter.
10. In the HORCM_INST section, set the necessary parameters, and then save (overwrite) the file.
11. Repeat Steps 3 to 10 for the **horcm1.conf** file. [Figure D-2](#) shows an example of the horcm1.conf file.

```

HORCM_MON
#ip_address      service      poll(10ms)      timeout(10ms)
XXXXXXXXXX       5001         12000          3000

HORCM_CMD
#dev_name        dev_name        dev_name
\\.\PHYSICALDRIVE1

HORCM_DEV
#dev_group       dev_name       port#          TargetID       LU#          MU#
UG01            oradb1         CL1-A          1              2

HORCM_INST
#dev_group       ip_address      service
UG01            XXXXXXXX       5000

```

Figure D-2: Horcm1.conf example

12. Enter the following in the command prompt to verify the connection between CCI and the array:

```

C:\>cd horcm\etc

C:\HORCM\etc>echo hd1-3 | .\inraid
Harddisk 1 -> [ST] CL1-A Ser =91200174 LDEV = 0 [HITACHI ] [DF600F-CM      ]
Harddisk 2 -> [ST] CL1-A Ser =91200174 LDEV = 1 [HITACHI ] [DF600F      ]
                HORC = SMPL  HOMRCF[MU#0 = SMPL  MU#1 = NONE  MU#2 = NONE]
                RAID5[Group 1-0] SSID = 0x0000
Harddisk 3 -> [ST] CL1-A Ser =91200175 LDEV = 2 [HITACHI ] [DF600F      ]
                HORC = SMPL  HOMRCF[MU#0 = NONE  MU#1 = NONE  MU#2 = NONE]
                RAID5[Group 2-0] SSID = 0x0000

C:\HORCM\etc>

```

Setting the environment variable

The environment variable must be set up for the execution environment. The following describes an example in which two instances (0 and 1) are configured on the same Windows Server.

1. Set the environment variable for each instance. Enter the following from the command prompt:

```
C:\HORCM\etc>set HORCMINST=0
```

2. Execute the horcmstart script, and then execute the pairedisplay command to verify the configuration. For example:

```

C:\HORCM\etc>horcmstart 0 1
starting HORCM inst 0
HORCM inst 0 starts successfully.
starting HORCM inst 1
HORCM inst 1 starts successfully.

```

```

C:\HORCM\etc>pairstat -g VG01

```

group	PairVOL(L/R)	(Port#,TID, LU)	Seq#	LDEV#	P/S	Status	Fence	Seq#	P-LDEV#	M
VG01	oradb1(L)	(CL1-A , 1, 1)	91200174	1	SMPL	----	-----	-----	----	-
VG01	oradb1(R)	(CL1-A , 1, 2)	91200175	2	SMPL	----	-----	-----	----	-

Pair operations

This section provides CCI procedures for performing TCE pairs operations. In the examples provided, the group name defined in the configuration definition file is VG01.



NOTE: A pair created using CCI and defined in the configuration definition file appear unnamed in the Navigator 2 GUI. Consistency groups created using CCI and defined in the configuration definition file are not seen in the Navigator 2 GUI. Also, pairs assigned to groups using CCI appear ungrouped in the Navigator 2 GUI.

Checking pair status

To check TCE pair status

1. Execute the `pairstatus` command to display the pair status and the configuration. For example:

```
C:\H0RCM\etc>pairstatus -g VG01
Group PairVol(L/R) (Port#,TID, LU) ,Seq#,LDEV#.P/S,Status,Fence,
Seq#,P-LDEV# M
vg01 oradb1(L) (CL1-A, 1, 1)91200174 1.P-VOL PAIR ASYNC
,91200175 2 -
vg01 oradb1(R) (CL1-B, 2, 2)91200175 2.S-VOL PAIR ASYNC , -
---- 1 -
```

The pair status is displayed. For details on the `pairstatus` command and its options, refer to the *Hitachi Unified Storage Command Control Interface Installation and Configuration Guide*.

CCI and Navigator 2 GUI pair statuses are described in [Table D-3](#).

Table D-3: Pair status descriptions

CCI	Navigator 2	Description
SMPL	Simplex	Status where a pair is not created.
COPY	Synchronizing	Initial copy or resynchronization copy is in execution.
PAIR	Paired	Status where copy is completed and update copy between pairs started.
PSUS/SSUS	Split	Update copy between pairs stopped by split.
PFUS	Pool Ful	Status that updating copy from the P-VOL to the S-VOL cannot continue due to too much use of the DP pool.
SSWS	Takeover	Takeover
SSUS	Inconsistent	Status that updating copy from the P-VOL to the S-VOL cannot continue due to the S-VOL failure.
PSUE	Failure	Update copy between pairs stopped by failure occurrence.

Creating a pair (paircreate)

To create a pair

1. Execute the pairdisplay command to verify that the status of the possible volumes to be copied is SMPL. The group name in the example is VG01.

```
C:\HORCM\etc>pairdisplay -g VG01
Group  PairVol(L/R) (Port#,TID, LU) ,Seq#,LDEV#.P/S,Status,Fence, Seq#,P-
LDEV# M
VG01   oradb1(L)   (CL1-A , 1,  1 )90000174      1.SMPL -----,---
--    ---- -
VG01   oradb1(R)   (CL1-A , 1,  2 )90000175      2.SMPL -----,---
--    ---- -
```

2. Execute the paircreate command. The -c option (medium) is recommended when specifying copying pace. See [Changing copy pace on page 21-16](#) for more information.
3. Execute the pairevtwait command to verify that the status of each volume is PAIR. The following example shows the paircreate and pairevtwait commands. For example:

```
C:\HORCM\etc>paircreate -g VG01 -f async -jp 0 -js 0 -vl -c
10
C:\HORCM\etc>pairevtwait -g VG01 -s pair -t 300 10
pairevtwait : Wait status done.
```

4. Execute the pairdisplay command to verify pair status and the configuration. For example:

```
c:\HORCM\etc>pairdisplay -g VG01
Group  PairVol(L/R) (Port#,TID, LU) ,Seq#,LDEV#.P/S,Status,Fence, Seq#,P-
LDEV# M
VG01   oradb1(L)   (CL1-A , 1,  1 )90000174      1.P-VOL PAIR Never ,90000175
2 -
VG01   oradb1(R)   (CL1-A , 1,  2 )90000175      2.S-VOL PAIR Never ,----
-      1 -
```



NOTE: When using CCI, the same DP pool is used for the replication data DP pool and the management area DP pool. The different DP pools cannot be specified respectively. You can use HSNM2 for pair creation to specify separate DP pools for the replication data DP pool and the management area DP pool.

Splitting a pair (pairsplit)

Two or more pairs can be split at the same time if they are in the same consistency group.

To split a pair

1. Execute the pairsplit command to split the TCE pair in the PAIR status. The group name in the example is VG01.

```
C:\HORCM\etc>pairsplit -g VG01
```

2. Execute the `pairstatus` command to verify the pair status and the configuration. For example:

```
c:\horcm\etc>pairstatus -g VG01
Group   PairVol(L/R) (Port#,TID, LU) ,Seq#,LDEV#.P/S,Status,Fence, Seq#,P-
LDEV# M
VG01    oradb1(L)  (CL1-A , 1, 1 )90000174    1.P-VOL PSUS ASYNC ,90000175
2 -
VG01    oradb1(R)  (CL1-A , 1, 2 )90000175    2.S-VOL SSUS ASYNC ,-----
1 -
```

Resynchronizing a pair (`pairresync`)

To resynchronize TCE pairs

1. Execute the `pairresync` command. Enter between 1 to 15 for copy pace, 1 being slowest (and therefore best I/O performance), and 15 being fastest (and therefore lowest I/O performance). A medium value is recommended.
2. Execute the `pairevtwait` command to verify that the status of each volume is PAIR. The following example shows the `pairresync` and the `pairevtwait` commands. The group name in the example is VG01.

```
C:\HORCM\etc>pairresync -g VG01 -c 15
C:\HORCM\etc>pairevtwait -g VG01 -s pair -t 300 10
pairevtwait : Wait status done.
```

3. Execute the `pairstatus` command to verify the pair status and the configuration. For example:

```
c:\horcm\etc>pairstatus -g VG01
Group   PairVol(L/R) (Port#,TID, LU) ,Seq#,LDEV#.P/S,Status,Fence, Seq#,P-LDEV# M
VG01    oradb1(L)  (CL1-A , 1, 1 )91200174    1.P-VOL PAIR ASYNC ,91200175    2 -
VG01    oradb1(R)  (CL1-A , 1, 2 )91200175    2.S-VOL PAIR ASYNC ,-----    1 -
```

Suspending pairs (`pairsplit -R`)

To suspend pairs

1. Execute the `pairstatus` command to verify that the pair to be suspended is in PAIR status. The group name in the example is VG01.

```
c:\horcm\etc>pairstatus -g VG01
Group   PairVol(L/R) (Port#,TID, LU) ,Seq#,LDEV#.P/S,Status,Fence, Seq#,P-LDEV# M
VG01    oradb1(L)  (CL1-A , 1, 1 )91200174    1.P-VOL PAIR ASYNC ,90000175    2 -
VG01    oradb1(R)  (CL1-A , 1, 2 )91200175    2.S-VOL PAIR ASYNC ,-----    1 -
```

2. Execute the `pairsplit -R` command to split the pair. For example:

```
C:\HORCM\etc>pairsplit -g VG01 -R
```

3. Execute the `pairstatus` command to verify that the pair status changed to SMPL. For example:


```
c:\horcm\etc>pairdisplay -g VG01
Group  PairVol(L/R) (Port#,TID, LU) ,Seq#,LDEV#.P/S,Status,Fence, Seq#,P-LDEV# M
VG01   oradb1(L)   (CL1-A , 1, 1 )91200174      1.P-VOL PSUE ASYNC ,91200175      2  -
VG01   oradb1(R)   (CL1-A , 1, 2 )91200175      2.S-VOL ----- ,----- -----
```

Releasing pairs (pairsplit -S)

To release pairs and change status to SMPL

1. Execute the pairsplit -S command to release the pair. The group name in the example is VG01.

```
C:\H0RCM\etc>pairsplit -g VG01 -S
```

2. Execute the pairdisplay command to verify that the pair status changed to SMPL. For example:

```
c:\horcm\etc>pairdisplay -g VG01
Group  PairVol(L/R) (Port#,TID, LU) ,Seq#,LDEV#.P/S,Status,Fence, Seq#,P-LDEV# M
VG01   oradb1(L)   (CL1-A , 1, 1 )91200174      1.SMPL ----- ,----- -----
VG01   oradb1(R)   (CL1-A , 1, 2 )91200175      2.SMPL ----- ,----- -----
```

Splitting TCE S-VOL/Snapshot V-VOL pair (pairsplit -mscas)

The pairsplit -mscas command splits a Snapshot pair that is cascaded with an S-VOL of a TCE pair. The data to be split is the P-VOL data of the TCE pair at the time when the pairsplit -mscas command is accepted.

CCI adds a human-readable character string of ASCII 31 characters to a remote snapshot. Because a snapshot can be identified by a character string rather than an volume number, it can be used for discrimination of the Snapshot volumes of many generations.

Requirements

- Cascade configuration of TCE and Snapshot pairs is required.
- This command is issued to TCE; however, the pair to be split is the Snapshot pair cascaded with the TCE S-VOL.
- This command can only be issued for the TCE consistency group (CTG). It cannot be issued directly to a pair.
- The TCE pair must be in PAIR status; the Snapshot pair must be in either PSUS or PAIR status.
- When both TCE and Snapshot pairs are in PAIR status, any pair split command directly to the Snapshot pair, other than the pairsplit command with the -mscas option, cannot be executed.

Restrictions

- The operation cannot be issued when the TCE S-VOL is in Synchronizing or Paired status from a remote host.
- When even a single pair that is under the end operation (delete? synchronizing?) exists, the command cannot be executed.
- When even a single pair that is under the splitting operation exists, the command cannot be executed.

- When the pairsplit -mscas command is being executed for even a single Snapshot pair that is cascaded with a pair in the specified CTG, the command cannot be executed. The pairsplit -mscas processing is continued unless it becomes Failure or Pool Full. The processing is started from the continuation at the time of the next start even if the main switch of the primary array is turned off during the processing.

Also, review the -mscas restrictions in [Miscellaneous troubleshooting on page 21-33](#). Also see [Figure D-3](#).




Figure D-3: Cascade configuration example

To split the TCE S-VOL / Snapshot V-VOL

In the example, the group name is ora. Group names of the cascaded Snapshot pairs are o0 and o1.

- Execute the pairsplit -mscas command to the TCE pair. The status must be PAIR. For example:

```
c:\horcm\etc>pairsplit -g ora -mscas Split-Marker 1
```

- Verify that the status of the TCE pair is still PAIR by executing the pairedisplay command. The group in the example is ora.

```
c:\horcm\etc>pairedisplay -g ora
Group  PairVol(L/R) (Port#,TID, LU) ,Seq#,LDEV#.P/S,Status,Fence, Seq#,P-LDEV# M
ora    oradb1(L)   (CL1-A , 1, 1 )91200174    1.PAIR -----,----- -
ora    oradb1(R)   (CL1-B , 1, 2 )91200175    2.PAIR -----,----- -
```

- Confirm that the Snapshot Pair is split using the indirect or direct methods.
 - For the indirect method, execute the pairsyncwait command to verify that the P-VOL data has been transferred to the S-VOL. For example:

```
c:\horcm\etc>pairsyncwait -g ora -t 10000
UnitID  CTGID   Q-Marker   Status   Q-Num
      0      3  00101231ef Done        2
```

The status may not display for one cycle after the command is issued.

Q-Marker counts up one by executing the `pairsplit -mscas` command. For the direct method, execute the `pairevtwait` command. For example:

```
c:\horcm\etc>pairevtwait -g o1 -s psus -t 300 10
pairevtwait : Wait status done.
```

Verify that the cascaded Snapshot pair is split by executing the `pairstatus -v smk` command. The group in the example below is o1.

```
c:\HORCM\etc>pairstatus -g o1 -v smk
Group   PairVol(L/R) Serial# LDEV# P/S Status UTC-TIME -----Split-
Maker-----
o1      URA_000(L)   91200175    2 P-VOL PSUS  -      -
o1      URA_000(R)   91200175    3 S-VOL SSUS  123456ef Split-Marker
(R)     90000175     3 S-VOL SSUS  123456ef Split-Marker
```

The TCE pair is released. For details on the `pairsplit` command, the `-mscas` option, and `pairsyncwait` command, refer to the *Hitachi Unified Storage Command Control Interface (CCI) Reference Guide*.

Confirming data transfer when status is PAIR

When the TCE pair is in the PAIR status, data is transferred in regular cycles to the S-VOL. However, the P-VOL data that was settled as S-VOL data must be checked, as well as when the S-VOL data was settled.

When you execute the `pairsyncwait` command, any succeeding commands must wait until the P-VOL data at the time of the cycle update is reflected in S-VOL data.

For more information, please refer to the *Hitachi Unified Storage Command Control Interface Installation and Configuration Guide*.

Pair creation/resynchronization for each CTG

In the pair creation/resynchronization performed with a specification of a certain group, renewal of the cycle is started from a pair for which the initial copy is completed first and the status of the pair above is changed to PAIR. A pair, initial copy for which is not completed before the first cycle renewal, renews the cycle taking the next occasion for the renewal. Therefore, the time when the status of the each pair is changed to PAIR may differ from the other one by the cycle time length.




Figure D-4: Pair creation/resynchronization for each CTG-1

In the pair creation/resynchronization newly performed for each CTG, the time for renewing the cycle is decided by the pair for which the initial copy is completed first. The pair, the initial copy for which is completed later than the first completion of the initial copy, employs the renewed cycle starting from the cycle after next at the earliest.

When pair creation or resynchronization is performed for a group, the new cycle time begins for any pair in the group that is in PAIR status. A pair whose initial copy is not complete is not updated in the current update cycle, but will update during the next cycle. Cycle time is determined according to the first pair to complete the initial copy.




Figure D-5: Pair creation/resynchronization for each CTG-2

When a pair is newly added to the CTG, the pair is synchronized with the existing cycle timing. In the example, the pair is synchronized with the existing cycle from the cycle 3 and its status is changed to PAIR from the cycle 4.

- When the paircreate or pairresync command is executed, the pair undergoes the differential copy in the COPY status, undergoes the cyclic copy once, and then placed in the PAIR status. When a new pair is added to a CTG, which is already placed in the PAIR status, by the paircreate or pairresync command, the copy operation halts until the time of the existing cyclic copy after the differential copy is completed. Further, it is not placed in the PAIR status until the first cyclic copy is completed after it begins to act in time to the cycle. Therefore, the pair synchronization rate displayed by Navigator 2 or CCI may be 100% or not changed when the pair status is COPY.
- When you want to confirm the time from the stop of the copy operation to the start of the cyclic copy, check the start of the next cycle by displaying the predicted time of completing the copy using Navigator 2. For the procedure for displaying the predicted time of completing the copy, refer to section 5.2.7.

Response time of pairsplit command

A response time of a pairsplit command depends on a pair status and an option. [Table D-4](#) summarizes a response time for each CCI command.

In a case of splitting and deleting a pair with PAIR status, a completion of a processing takes time depending on the amount of differential data at P-VOL.

In a case of creating a remote snapshot, CCI command returns immediately but a completion of creating a snapshot depending on the amount of differential data at P-VOL. In order to check the completion, see Split-Marker of a remote snapshot is updated or a creation time of a snapshot is updated.



NOTE: Only -g option is valid. The -d option is not accepted. If there are pairs which status is not PAIR, in a CTG, a command cannot be accepted. All S-VOLs with PAIR status need to have corresponding cascading V-VOLs and MU# of these Snapshot pairs must match the MU# specified in a pairsplit -mscas command option.

Table D-4: Response time of CCI commands

Command	Options	Status	Response	Next Status	Remarks
pairsplit	-S Delete pair	PAIR	Depend on differential data	SMPL	S-VOL data consistency guaranteed
		COPY	Immediate	SMPL	No S-VOL data consistency
		Others	Immediate	SMPL	No S-VOL data consistency
	-R Delete pair	PAIR	Immediate	SMPL (S-VOL only)	No S-VOL data consistency
		COPY	Immediate	SMPL (S-VOL only)	No S-VOL data consistency Can not be executed for SSWS(R) status
		Others	Immediate	SMPL (S-VOL only)	No S-VOL data consistency Can not be executed for SSWS(R) status
	-mascas Create remote snapshot (See note)	PAIR	Immediate	No change	A completion time depends on the amount of differential data. A completion can be check by Split-Marker and a creation time. Cycle updating process stops during creating a remote snapshot.
		Others	—	—	—
	Others Split pair	PAIR	Depend on differential data	PSUS	S-VOL data consistency guaranteed
		COPY	Immediate	PSUS	S-VOL data consistency guaranteed
		Others	Immediate	No change	S-VOL data consistency guaranteed

Table D-5: TCE pair statuses and relationship to takeover

Object Volume		CCI Commands		
Attribute	Status	Paircurchk Result	SVOL_Takeover	
			Data Consistency	Next Status
SMPL	-	To be confirmed	No	SMPL
P-VOL	-	To be confirmed	No	-
S-VOL	COPY	Inconsistent	No	COPY
	PAIR	To be analyzed	CTG	SSWS
	PSUS	Suspected	Pair	SSWS
	PSUS(N)	Suspected	No	PSUS(N)
	PFUS	Suspected	CTG	SSWS
	PSUE	Suspected	CTG	SSWS
	SSWS	Suspected	Pair	SSWS

Responses of paircurchk.

- To be confirmed: The object volume is not an S-VOL. Check is required.
- Inconsistent: There is no write order guarantee of an S-VOL because an initial copy or a resync copy is on going or because of S-VOL failures. So SVOL_Takeover cannot be executed.
- To be analyzed: Mirroring consistency cannot be determined just from a pair status of an S-VOL. However TCE does not support mirroring consistency, this result always shows that S-VOL has data consistency across a CTG not depending on a pair status of a P-VOL.
- Suspected: There is no mirroring consistency of an S-VOL. If a pair status is PSUE or PFUS, there is data consistency across a CTG. If a pair status is PSUS or SSWS, there is data consistency for each pair in a CTG. In a case of PSUS(N), there is no data consistency.
- Data consistency after SVOL_Takeover and its response
 - CTG: Data consistency across a CTG is guaranteed.
 - Pair: Data consistency of each pair is guaranteed.
 - **No**: No data consistency of each pair.
 - Good: Response of takeover is normal.
 - NG: Response of takeover is an error. If a pair status of an S-VOL is PSUS, the pair status is changed to SSWS even if the response is an error.

See the *Hitachi Unified Storage Command Control Interface Installation and Configuration Guide* for more details about horctakeover.

Pair, group name differences in CCI and Navigator 2

Pairs and groups that were created using CCI will be displayed differently when status is confirmed in Navigator 2.

- Pairs created with CCI and defined in the configuration definition file display unnamed in Navigator 2.
- Pairs defined in a group on the configuration definition file are displayed in Navigator 2 as ungrouped.

For information about how to manage a group defined on the configuration definition file as a CTG, see the *Hitachi Unified Storage Command Control Interface Installation and Configuration Guide*.

TCE and Snapshot differences

Table D-6 summaries differences between TCE and Snapshot

Table D-6: TCE, Snapshot behaviors

Condition	TCE	Snapshot
Replication threshold over	TCE does not refer to the Replication Depletion Alert threshold over value. When the usage rate of the DP pool exceeds the Replication Data Released threshold, the pair status becomes Pool Full.	When the usage rate of the DP pool exceeds the Replication Depletion Alert threshold, a Snapshot pair in Split status changes to in Threshold over status When the usage rate of the DP pool exceeds the Replication Data Released threshold, the pair status becomes Failure.
DP pool full at local	Pair status of P-VOL with Paired status changes to Pool Full. Pair status of P-VOL with Synchronizing status changes to Failure.	Pair status changes to Failure.
DP pool full at remote	Pair status of P-VOL changes to Failure. Pair status of S-VOL with Paired status changes to Pool Full. Pair status of S-VOL with Synchronizing status changes to Inconsistent.	Pair status changes to Failure.
Data consistency when DP pool full	S-VOL data stays consistent at consistency- group level.	V-VOL data is invalid.
How to recover from Failure	Resync the pair.	Delete then recreate the pair.

Table D-6: TCE, Snapshot behaviors

Condition	TCE	Snapshot
Failures	Failures at local: P-VOL changes to Failure. S-VOL does not change. Data consistency is ensured if pair status of S-VOL is Paired. Failures at remote: P-VOL changes to Failure. S-VOL changes to Inconsistent. No data consistency for S-VOL.	Pair status changes to Failure and V-VOL data is invalid.
Number of consistency groups supported	64 * CTG number for Snapshot and CTG number for TCE are independent. Snapshot supports 1,024 and TCE supports 64.	1,024


Initializing Cache Partition when TCE and Snapshot are installed

TCE and Snapshot use part of the cache to manage internal resources, causing a reduction in the cache capacity used by Cache Partition Manager.

Cache partition information should be initialized as follows, when TCE or Snapshot are installed after Cache Partition Manager is installed:

- All the volumes should be moved to the master partitions on the side of the default owner controller.
- All the sub-partitions must be deleted and the size of each master partition should be reduced to half of the user data area after installation of TCE or Snapshot.

Figure D-6 shows an example of Cache Partition Manager usage. Figure D-7 shows an example where TCE/Snapshot is installed when Cache Partition Manager already in use.

**Figure D-6: Cache Partition Manager usage**




Figure D-7: TCE or Snapshot installation with Cache Partition Manager

On the remote array, Synchronize Cache Execution mode should be turned off to avoid TCE remote path failure.

Wavelength Division Multiplexing (WDM) and dark fibre

This topic discusses WDM and dark fibre, which are used to extend Fibre Channel remote paths.

The integrity of a light wavelength remains intact when it is combined with other light wavelengths. Light wavelengths can be combined together in a transmission by multiplexing several optical signals on a dark fibre.

Wavelength Division Multiplexing uses this technology to increase the amount of data that can be transported across distances in a dark fibre extender.

- WDM signifies the multiplexing of several channels of the optical signal.
- Dense Wavelength Division Multiplexing (DWDM) signifies the multiplexing of several *dozen* channels of the optical signal.

Figure D-8 shows an illustration of WDM.




Figure D-8: Wavelength division multiplexing

WDM has the following characteristics:

- Response time is extended with WDM. This deterioration is made up by increasing the Fibre Channel BB-Credit (the number of buffer) without waiting for the response. This requires a switch.

If the array is connected directly to an WDM extender without a switch, BB-Credit is 4 or 8. If the array is connected with a switch (Brocade), BB-Credits are 16 and can hold up to 10 km on the standard scale. BB-Credits can be increased to a maximum of 60. By adding the Extended Fabrics option to a switch, BB-Credits can hold up to 100 km.

- For short distances (within several dozen kilometers), both signals of IN and OUT can be transmitted via one dark fiber.

- For long distances (more than several dozen kilometers), an optical amplifier is required to amplify the wavelength between two extenders to prevent attenuation through a fiber. Therefore, dark fibers are required to prepare for IN and OUT respectively. This is illustrated in [Figure D-9](#).




Figure D-9: Dark Fiber with WDM

- The WDM function can also be multiplexed in one dark fiber for G Ethernet.
- If switching is executed during a dark fibre failure, data transfer must be moved to another path, as shown in [Figure D-10](#).




Figure D-10: Dark Fiber failure

- It is recommended that a second line be set up for monitoring. This allows monitoring to continue if a failure occurs in the dark fiber.




Figure D-11: Line for monitoring

TrueCopy Modular Distributed reference information

This appendix contains:

- ❑ [TCMD system specifications](#)
- ❑ [Operations using CLI on page E-5](#)

TCMD system specifications

Table E-1 describes specifications of TCE expanded by TCMD.

Table E-1: Specification of TCE (Installed TCMD)

Parameter	TCMD specification
User interface	- Navigator 2 GUI and CLI: used for the setting of DP pool, remote paths, or command devices, and for the pair operations. - CCI: used for the pair operations
Controller configuration	Configuration of dual controller is required.
Host interface	Fibre Channel or iSCSI (cannot mix). For iSCSI environments, the HUS 100 firmware must be upgraded to V2.0B (0920/B, SNM2 Version 22.02) at a minimum.
Remote path	- Fibre Channel or iSCSI. - One remote path per controller is necessary and a total of two remote paths are necessary between the arrays because of the dual controller configuration. - You can set 16 (two for each array) remote paths to the maximum of eight Edge arrays in the Hub array. - The Fibre Channel remote path and the iSCSI remote path can coexist in one Hub array. However, the interface type of two remote paths between the arrays must be the same.
Port modes	Initiator and target intermix mode. One port may be used for host I/O and TCE at the same time.
Range of supported transfer rate	- It is required that 1.5 M bps or more (100 M bps or more is recommended) is guaranteed for each remote path - The remote path to be set two, the bandwidth must be 3.0 M bps or more between the arrays. - Further, when the range of the transfer rate is narrow, response to a CCI command may take several seconds
License	Entry of the key code enables TCE to be used. When using TCMD, it is further required to enter the key code of TCMD. TrueCopy and TCE cannot coexist and the licenses to use them are different from each other.
Command device (CCI only)	- This must be set when performing pair operations of CCI - Up to 128 command devices per array can be set - It is necessary to set 65,538 blocks or more (1 block = 512 bytes) (33 M bytes or more). - Set them for both of the local and remote arrays.
Unit of pair management	Volumes are the target of TCE pairs, and are managed per volume
Maximum # of volumes that can be used for pairs	- HUS 110: 2,046 volume - HUS 150/HUS 130: 4,094 volume The maximum number of volumes when different types of arrays are combined is that of the array whose maximum number of volumes is smaller - When using TCMD, the maximum number of volumes that can create pairs between two or more Edge arrays and Hub array is the maximum number of volumes of the types of the Hub array.
Pair structure	One S-VOL per P-VOL.

Table E-1: Specification of TCE (Installed TCMD) (Continued)

Parameter	TCMD specification
Supported RAID level	- RAID 1 (1D+1D), RAID 5 (2D+1P to 15D+1P) - RAID 1+0 (2D+2D to 8D+8D) - RAID 6 (2D+2P to 28D+2P)
Combination of RAID levels	There is no need to use RGs whose RAID levels and numbers of drives are same between the P-VOL and the S-VOL.
Size of pair volumes	Volume size of the P-VOL and S-VOL must be equal—identical block counts.
Types of drive for P-VOL and S-VOL	If the drive types are supported by the array, they can be set for a P-VOL and an S-VOL. It is recommended to set a volume configured by the SAS drives or the SSD/FMD to a P-VOL .
Copy pace	The copy paces from a P-VOL to an S-VOL and vice versa can be adjusted in three stages.
Consistency Group (CTG)	- Maximum allowed: 64 for any array models . - A pair with one local destination array can belong to one CTG.
Cycle time	- The cycle time to update the differential data from a P-VOL to an S-VOL when the pair status is Paired can be changed as needed - The default is 300 seconds and the maximum of 3,600 seconds can be set by the second. - The lowest value that can be set becomes number of CTGs residing on the Hub array × 30 seconds. - Set the cycle time more than or equal to that of the Hub array in the Edge array

Table E-2 describes specifications of TCE expanded by TCMD.

Table E-2: Specification of TrueCopy (Installed TCMD)

Parameter	TCMD specification
User interface	- Navigator 2 : used for the setting of DP pool, remote paths, or command devices, and for the pair operations. - CCI: used for the pair operations
Controller configuration	Configuration of dual controller is required.
Host interface	Fibre Channel or iSCSI (cannot mix).
Remote path	- Fibre Channel or iSCSI. - One remote path per controller is necessary and a total of two remote paths are necessary between the arrays because of the dual controller configuration. - You can set 16 (two for each array) remote paths to the maximum of eight Edge arrays in the Hub array. - The Fibre Channel remote path and the iSCSI remote path can coexist in one Hub array. However, the interface type of two remote paths between the arrays must be the same.
Port modes	One port is usable for the host I/O and the copy of TrueCopy at the same time.
Range of supported transfer rate	- It is required that 1.5 M bps or more (100 M bps or more is recommended) is guaranteed for each remote path - The remote path to be set two, the bandwidth must be 3.0 M bps or more between the arrays. - Further, when the range of the transfer rate is narrow, response to a CCI command may take several seconds
License	Entry of the key code enables TrueCopy to be used. When using TCMD, it is further required to enter the key code of TCMD. TrueCopy and TCE cannot coexist and the licenses to use them are different from each other. When using TrueCopy and TCMD together, the volume constructing the remote pair cannot be mounted directly to the host. For connecting the host, it is required to enter a key code of ShadowImage.
Command device (CCI only)	- This must be set when performing pair operations of CCI - Up to 128 command devices per array can be set. - It is necessary to set 65,538 blocks or more (1 block = 512 bytes) (33 M bytes or more). - Set them for both of the local and remote arrays.
DMLU	- This needs to be set to use a pair of TrueCopy. - Be sure to set this for both of the local and remote arrays
Unit of pair management	Volumes are the target of TrueCopy pairs, and are managed per volume

Table E-2: Specification of TrueCopy (Installed TCMD) (Continued)

Parameter	TCMD specification
Maximum # of volumes that can be used for pairs	- HUS 110: 2,046 volume - HUS 150/HUS 130: 4,094 volume The maximum number of volumes when different types of arrays are combined is that of the array whose maximum number of volumes is smaller - When using TCMD, the maximum number of volumes that can create pairs between two or more Edge arrays and Hub array is the maximum number of volumes of the types of the Hub array.
Pair structure	One S-VOL per P-VOL.
Supported RAID level	- RAID 1 (1D+1D), RAID 5 (2D+1P to 15D+1P) - RAID 1+0 (2D+2D to 8D+8D) - RAID 6 (2D+2P to 28D+2P)
Combination of RAID levels	All combinations supported. The number of data disks does not have to be the same.
Size of pair volumes	Volume size of the P-VOL and S-VOL must be equal—identical block counts.
Types of drive for P-VOL and S-VOL	If the drive types are supported by the array, they can be set for a P-VOL and an S-VOL. It is recommended to set a volume configured by the SAS drives or the SSD/FMD to a P-VOL .
Copy pace	The copy paces from a P-VOL to an S-VOL and vice versa can be adjusted in three stages.
Consistency Group (CTG)	- Maximum allowed: 256 for any array models . - A pair with one local destination array can belong to one CTG.

Operations using CLI

This section describes CLI procedures for setting up and performing TCMD operations.

- ☐ [Installation and uninstalling](#)
- ☐ [Enabling and disabling](#)
- ☐ [Setting the Distributed Mode](#)
- ☐ [Setting the remote port CHAP secret](#)
- ☐ [Setting the remote path](#)
- ☐ [Deleting the remote path](#)



NOTE: For additional information on the commands and options in this appendix, see the *Hitachi Unified Storage Command Line Interface Reference Guide*.

Installation and uninstalling

Since TCMD is an extra-cost option, TCMD cannot usually be selected (locked) when first using the array. To make TCMD available, you must install TCMD and make its function selectable (unlocked).

TCMD can be installed from Navigator 2. This section describes the installation/un-installation procedures performed by using Navigator 2 via the Command Line Interface (CLI).

Before installing or uninstalling TCMD, verify the following:

- The array must be operating in a normal state. Installation and un-installation cannot be performed if a failure has occurred. If a failure such as a controller blockade has occurred, installation/un-installation cannot be performed.
- To install TCMD, TCE must be installed and enabled.
- To install TCMD, the key code or key file provided with the optional feature is required

Installing TCMD

Since TCMD is an extra-cost option, TCMD cannot usually be selected (locked) when first using the array. To make TCMD available, you must install TCMD and make its function selectable (unlocked).

TCMD can be installed from Navigator 2. This section describes the installation/un-installation procedures performed by using Navigator 2 via the Command Line Interface (CLI).



NOTE: TCMD can be installed from Navigator 2. This section describes the installation/un-installation procedures performed by using Navigator 2 via the Command Line Interface (CLI).

NOTE: To install TCMD, TCE or TrueCopy must be installed and its status to be valid

To install TCMD

1. From the command prompt, register the array in which the TCMD is to be installed, and then connect to the array.
2. Execute the `auopt` command to install TCMD. For example:

```
% auopt -unit array-name -lock off -keycode manual-attached-keycode
Are you sure you want to unlock the option? (y/n [n]): y
The option is unlocked.
%
%
```

3. Execute the `auopt` command to confirm whether TCMD has been installed. For example:

% auopt -unit array-name -refer			
Option Name	Type	Term	Status
Reconfigure	Memory	Status	
TC-EXTENDED	Permanent	---	Enable
N/A			
TC-DISTRIBUTED	Permanent	---	Enable
N/A			
%			

TCMD is installed and **Status** is **Enabled**. Installation of TCMD is now complete.

Un-installingTCMD

To uninstall TCMD, the key code or key file provided with the optional feature is required. Once uninstalled, TCMD cannot be used (locked) until it is again installed using the key code or key file.

Prerequisites for uninstalling

- All of TCE or TrueCopy pairs pairs must be released (the status of all volumes must be Simplex).
- All the remote path settings must be deleted.
- All the remote port CHAP secret settings must be deleted.

To uninstall TCMD

1. From the command prompt, register the array in which the TCMD is to be uninstalled, and then connect to the array.
2. Execute the auopt command to uninstall TCMD. For example:

```
% auopt -unit array-name -lock on -keycode manual-attached-keycode
Are you sure you want to lock the option? (y/n [n]): y
The option is locked.
.%
```

3. Execute the auopt command to confirm that TCMD is uninstalled. For example:

```
% auopt -unit array-name -refer
Option Name      Type      Term      Status
  Reconfigure Memory Status
TC-EXTENDED      Permanent ---      Enable
      N/A
%
```

Uninstalling TCMD is now complete

Enabling and disabling

TCMD can be disabled or enabled.

Prerequisites for disabling

- All of TCE or TrueCopy pairs must be deleted (the status of all volumes must be **Simplex**).
- The remote path must be deleted.
- All the remote port CHAP secret settings must be deleted.

To enable / disable TCMD

1. From the command prompt, register the array in which the status of the feature is to be changed, and then connect to the array.
2. Execute the auopt command to change TCMD status (enable or disable).

The following is an example of changing the status from enable to disable. If you want to change the status from disable to enable, enter enable after the -st option.

```
% auopt -unit array-name -option TC-EXTENDED -st disable
Are you sure you want to disable the option? (y/n [n]): y
The option has been set successfully.

%
```

3. Execute the auopt command to confirm that the status has been changed. For example:

```
% auopt -unit array-name -refer
Option Name      Type      Term      Status
  Reconfigure Memory Status
TC-EXTENDED      Permanent ---      Enable
      N/A
TC-DISTRIBUTED   Permanent ---      Disable
      N/A

%
```

Enabling or disabling TCMD is now complete.

Setting the Distributed Mode

To set the remote paths between one array and two or more arrays using TCMD, set the Distributed mode to Hub for one array.

Before setting the Distributed mode, perform the following:

- Decide the configuration of the array that uses TCMD in advance, and check the array in which Distributed mode is set to Hub and the array in which Distributed mode remains as Edge.
- If you install the arrays in TCMD, the Distributed mode is set to Edge in the initial status.

Changing the Distributed mode to Hub from Edge

Prerequisites

- All the remote paths settings must be deleted.
- All the remote port CHAP secret settings must be deleted.

To change the distributed mode to Hub from Edge

1. From the command prompt, register the array in which you want to set to the Hub array, and then connect to the array.
2. Execute the aurmtpath command to set the Distributed mode. For example:

```
% aurmtpath -unit array-name -set -distributedmode hub
Are you sure you want to set the remote path information? (y/n [n]): y
The remote path information has been set successfully.
%
%
```

3. Execute the aurmtpath command to confirm whether the Distributed mode has been set. For example:

```
% aurmtpath -unit array-name -refer
Initiator Information
Local Information
  Array ID          : 93000026
  Distributed Mode  : Hub

Path Information
  Interface Type    : ---
  Remote Array ID   : ---
  Remote Path Name  : ---
  Bandwidth [0.1 Mbps] : ---
  iSCSI CHAP Secret : ---

Path      Status      Local      Remote      Remote Port      TCP Port No. of
0         Undefined    ---        ---        IP Address       Remote Port
1         Undefined    ---        ---        ---              ---
%
```

Changing the Distributed mode to Edge from Hub is now complete.

Changing the Distributed Mode to Edge from Hub

Prerequisites

- All the remote paths settings must be deleted.
- All the remote port CHAP secret settings must be deleted.

To change the distributed mode to Edge from Hub

1. Execute the aurmtpath command to set the Distributed mode.

```
% aurmtpath -unit array-name -set -distributedmode edge
Are you sure you want to set the remote path information? (y/n [n]): y
The remote path information has been set successfully.
%
```

2. Execute the aurmtpath command to confirm whether the Distributed mode has been set.

```
% aurmtpath -unit array-name -refer
Initiator Information
Local Information
  Array ID          : 93000026
  Distributed Mode  : Edge

Path Information
Interface Type      : ---
Remote Array ID     : ---
Remote Path Name    : ---
Bandwidth [0.1 Mbps] : ---
iSCSI CHAP Secret   : ---

Path   Status      Local   Remote   Remote Port   TCP Port No. of
      0   Undefined   ---     ---     IP Address    Remote Port
      1   Undefined   ---     ---     ---           ---
%
```

Changing the Distributed mode to Hub from Edge is now complete.

Setting the remote port CHAP secret

For iSCSI array model, the remote path can be use the CHAP secret. Set the CHAP secret mode to the remote array to be the connection destination of the remote path. If you set the CHAP secret to the remote array, you can prevent creation of the remote path from the array that the same character string is not set as the CHAP secret.

To set the remote port CHAP secret :

1. From the command prompt, register the array in which you want to set the remote port CHAP secret, and then connect to the array
2. Execute the `aurmtpath` command with the `-set` option and perform the CHAP secret of the remote port. The input example and the result are shown below.

Example:

```
% aurmtpath -unit array-name -set -target -local 91200027 -secret
Are you sure you want to set the remote path information? (y/n [n]): y
Please input Path 0 Secret.
Path 0 Secret:
Re-enter Path 0 Secret:
Please input Path 1 Secret.
Path 1 Secret:
Re-enter Path 1 Secret:
The remote path information has been set successfully.
%
```

The setting of the remote port CHAP secret is completed.



NOTE: If setting the remote port CHAP in the array, the remote path whose CHAP secret is set to automatic input cannot be connected for the array. When setting the remote port CHAP secret while using the remote path whose CHAP secret is set to automatic input, see [Adding the Edge array in the configuration of the set TCMD on page 24-3](#) and recreate the remote path.

Setting the remote path

Data is transferred from the local to the remote array over the remote path. The settings of the remote path procedure are different with iSCSI and Fibre channel.

Prerequisites

- Both local and remote disk arrays must be connected to the network for the remote path.
- The remote disk array ID will be required. This is shown on the main disk array screen.
- Network bandwidth will be required.
- For the iSCSI array model, you can identify the IP address for the remote path in the IPv4 or IPv6 format. Be sure to use the same format when specifying the port IP addresses for the remote path for the local array and the remote array.
- If the interface between the arrays is iSCSI, you need to set the remote paths from the controller 0 to the other controller 0 and from the controller 1 to the other controller 1.

To set up the remote path for the Fibre Channel array

1. From the command prompt, register the array in which you want to set the remote path, and then connect to the array.
2. To refer the remote array ID, use the `auunitinfo` command. The example is shown below. The remote array ID is displayed to the Array ID field (remote array ID=91100026). You must obtain array IDs for the number of Edge arrays

Example:

```
% auunitinfo -unit remote-array-name
Array Unit Type      : HUS110
H/W Rev.            : 0100
Construction        : Dual
Serial Number       : 91100026
Array ID            : 91100026
Firmware Revision(CTL0) : 0917/A-W
Firmware Revision(CTL1) : 0917/A-W
CTL0
:
:
%
```

- Execute the `aurmtpath` command to set the remote path. The example shows that the array ID of the remote-side array is 91100026, path 0 is the 0A port of the local-side array and 0A port of the remote-side array, path 1 is the 1A port of the local-side array and the 1A port of the remote-side array.

Example:

```
% aurmtpath -unit local-array-name -set -remote 91100026 -band
auto
                -path0 0A 0A -path1 1A 1A -remotename Array_91100026
Are you sure you want to set the remote path information? (y/n
[n]): y
The remote path information has been set successfully.
%
```

- You must set the remote paths for the number of Edge arrays. Execute the `aurmtpath` command to confirm whether the remote path has been set. For example:

Example:

```
% aurmtpath -unit local-array-name -refer
Initiator Information
  Local Information
    Array ID       : 93000026
    Distributed Mode : Hub

  Path Information
    Interface Type   : FC
    Remote Array ID  : 91100026
    Remote Path Name : Array_91100026
    Bandwidth [0.1 Mbps] : Over 10000
    iSCSI CHAP Secret : N/A

Port No. of Path Status Local Remote IP Address Remote Port TCP
0 Normal 0A 0A N/A N/A
1 Normal 1A 1A N/A N/A

  Path Information
    Interface Type   : FC
    Remote Array ID  : 91100027
    Remote Path Name : Array_91100027
    Bandwidth [0.1 Mbps] : Over 10000
    iSCSI CHAP Secret : N/A

:
:
%
```

Creation of the remote path is now complete. You can start the copy operations.

To set up the remote path for the iSCSI array

1. From the command prompt, register the array in which you want to set the remote path, and then connect to the array.
2. The following is an example of referencing the remote path status where remote path information is not yet specified.

Example:

```
% aurmtpath -unit array-name -refer
Initiator Information
  Local Information
    Array ID       : 93000026
    Distributed Mode : Hub

  Path Information
    Interface Type      : ---
    Remote Array ID     : ---
    Remote Path Name    : ---
    Bandwidth [0.1 Mbps] : ---
    iSCSI CHAP Secret   : ---

    Path  Status      Local      Remote Port      TCP Port No. of
    0     Undefined   ---      ---      IP Address      Remote Port
    1     Undefined   ---      ---      ---              ---

  Target Information
    Local Array ID      :
%
```

3. Execute the aurmtpath command to set the remote path.

Example:v

```
% aurmtpath -unit array-name -set -initiator -remote 91200027 -
secret disable
    -path0 0B -path0_addr 192.168.1.201 -band 100
    -path1 1B -path1_addr 192.168.1.209
Are you sure you want to set the remote path information? (y/n
[n]): y
The remote path information has been set successfully.
%
```

4. Execute the aurmtpath command to confirm whether the remote path has been set.

Example:

```
% aurmtpath -unit array-name -refer
Initiator Information
  Local Information
    Array ID       : 93000026
    Distributed Mode : Hub

  Path Information
    Interface Type   : iSCSI
    Remote Array ID  : 91200027
    Remote Path Name : N/A
    Bandwidth [0.1 Mbps] : 100
    iSCSI CHAP Secret : Disable

    Path  Status      Local  Remote  Remote Port  TCP Port No. of
    0  Normal        0B    N/A    192.168.0.201  3260
    1  Normal        1B    N/A    192.168.0.209  3260

  Target Information
    Local Array ID   : 93000026

%
```

Creation of the remote path is now complete. You can start the copy operations.

Deleting the remote path

When the remote path becomes unnecessary, delete the remote path.

Prerequisites

- To delete the remote path, change all the TrueCopy pairs or all the TCE pairs in the array to the Simplex or Split status.
- Do not do a pair operation for a TrueCopy/TCE pair when the remote path for the pair is not set up, because then the pair operation may not complete correctly.



NOTE: When performing the planned shutdown of the remote array, the remote path should not necessarily be deleted. Change all the TrueCopy pairs or all the TCE pairs in the array to the Split status, and then perform the planned shutdown of the remote array. After restarting the array, perform the pair resynchronization. However, when the Warning notice to the Failure Monitoring Department at the time of the remote path blockade or the notice by the SNMP Agent Support Function or the E-mail Alert Function is not desired, delete the remote path, and then turn off the power of the remote array.

To delete the remote path

1. From the command prompt, register the array in which you want to delete the remote path, and then connect to the array.
2. Execute the `aurmtpath` command to delete the remote path. For example:

```
% aurmtpath -unit array-name -rm -remote 91100027
Are you sure you want to delete the remote path information? (y/n
[n]): y
The remote path information has been deleted successfully.
%
```

You delete the remote paths for the number of Edge arrays if necessary.

Deletion of the remote path is now complete.



Glossary

This glossary provides definitions for replication terms as well as terms related to the technology that supports your Hitachi modular array. Click the letter of the glossary section to display the related page.

A

array

A set of hard disks mounted in a single enclosure and grouped logically together to function as one contiguous storage space.

asynchronous

Asynchronous data communications operate between a computer and various devices. Data transfers occur intermittently rather than in a steady stream. Asynchronous replication does not depend on acknowledging the remote write, but it does write to a local log file. Synchronous replication depends on receiving an acknowledgement code (ACK) from the remote system and the remote system also keeps a log file.

B

background copy

A physical copy of all tracks from the source volume to the target volume.

bps

Bits per second, the standard measure of data transmission speeds.

A	B	C	D	E	F	G	H	I	J	K	L	M	N	O	P	Q	R	S	T	U	V	W	X	Y	Z

C

cache

A temporary, high-speed storage mechanism. It is a reserved section of main memory or an independent high-speed storage device. Two types of caching are found in computers: memory caching and disk caching. Memory caches are built into the architecture of microprocessors and often computers have external cache memory. Disk caching works like memory caching; however, it uses slower, conventional main memory that on some devices is called a memory buffer.

capacity

The amount of information (usually expressed in megabytes) that can be stored on a disk drive. It is the measure of the potential contents of a device; the volume it can contain or hold. In communications, capacity refers to the maximum possible data transfer rate of a communications channel under ideal conditions.

cascading

Cascading is connecting different types of replication program pairs, like ShadowImage with Snapshot, or ShadowImage with TrueCopy. It is possible to connect a local replication program pair with a local replication program pair and a local replication program pair with a remote replication program pair. Cascading different types of replication program pairs allows you to utilize the characteristics of both replication programs at the same time.

CCI

See command control interface.

CLI

See command line interface.

cluster

A group of disk sectors. The operating system assigns a unique number to each cluster and then keeps track of files according to which clusters they use.

cluster capacity

The total amount of disk space in a cluster, excluding the space required for system overhead and the operating system. Cluster capacity is the amount of space available for all archive data, including original file data, metadata, and redundant data.

A	B	C	D	E	F	G	H	I	J	K	L	M	N	O	P	Q	R	S	T	U	V	W	X	Y	Z

command control interface (CCI)

Hitachi's Command Control Interface software provides command line control of Hitachi array and software operations through the use of commands issued from a system host. Hitachi's CCI also provides a scripting function for defining multiple operations.

command devices

Dedicated logical volumes that are used only by management software such as CCI, to interface with the arrays. Command devices are not used by ordinary applications. Command devices can be shared between several hosts.

command line interface (CLI)

A method of interacting with an operating system or software using a command line interpreter. With Hitachi's Storage Navigator Modular Command Line Interface, CLI is used to interact with and manage Hitachi storage and replication systems.

concurrency of S-VOL

Occurs when an S-VOL is synchronized by simultaneously updating an S-VOL with P-VOL data AND data cached in the primary host memory. Discrepancies in S-VOL data may occur if data is cached in the primary host memory between two write operations. This data, which is not available on the P-VOL, is not reflected on to the S-VOL. To ensure concurrency of the S-VOL, cached data is written onto the P-VOL before subsequent remote copy operations take place.

concurrent copy

A management solution that creates data dumps, or copies, while other applications are updating that data. This allows end-user processing to continue. Concurrent copy allows you to update the data in the files being copied, however, the copy or dump of the data it secures does not contain any of the intervening updates.

configuration definition file

The configuration definition file describes the system configuration for making CCI operational in a TrueCopy Extended Distance Software environment. The configuration definition file is a text file created and/or edited using any standard text editor, and can be defined from the PC where the CCI software is installed. The configuration definition file describes configuration of new TrueCopy Extended Distance pairs on the primary or remote array.

consistency group (CTG)

A group of two or more logical units in a file system or a logical volume. When a file system or a logical volume which stores application data, is configured from two or more logical units, these multiple logical units are managed as a consistency group (CTG) and treated as a single

A	B	C	D	E	F	G	H	I	J	K	L	M	N	O	P	Q	R	S	T	U	V	W	X	Y	Z

entity. A set of volume pairs can also be managed and operated as a consistency group.

consistency of S-VOL

A state in which a reliable copy of S-VOL data from a previous update cycle is available at all times on the remote array. A consistent copy of S-VOL data is internally pre-determined during each update cycle and maintained in the remote data pool. When remote takeover operations are performed, this reliable copy is restored to the S-VOL, eliminating any data discrepancies. Data consistency at the remote site enables quicker restart of operations upon disaster recovery.

CRC

Cyclical Redundancy Checking. A scheme for checking the correctness of data that has been transmitted or stored and retrieved. A CRC consists of a fixed number of bits computed as a function of the data to be protected, and appended to the data. When the data is read or received, the function is recomputed, and the result is compared to that appended to the data.

CTG

See Consistency Group.

cycle time

A user specified time interval used to execute recurring data updates for remote copying. Cycle time updates are set for each array and are calculated based on the number of consistency groups CTG.

cycle update

Involves periodically transferring differential data updates from the P-VOL to the S-VOL. TrueCopy Extended Distance Software remote replication processes are implemented as recurring cycle update operations executed in specific time periods (cycles).

D

data pool

One or more disk volumes designated to temporarily store un-transferred differential data (in the local array or snapshots of backup data in the remote array). The saved snapshots are useful for accurate data restoration (of the P-VOL) and faster remote takeover processing (using the S-VOL).

data volume

A volume that stores database information. Other files, such as index files and data dictionaries, store administrative information (metadata).

A	B	C	D	E	F	G	H	I	J	K	L	M	N	O	P	Q	R	S	T	U	V	W	X	Y	Z

differential data control

The process of continuously monitoring the differences between the data on two volumes and determining when to synchronize them.

differential data copy

The process of copying the updated data from the primary volume to the secondary volume. The data is updated from the differential data control status (the pair volume is under the suspended status) to the primary volume.

Differential Management Logical Unit (DMLU)

The volumes are used to manage differential data in a array. In a TrueCopy Extended Distance system, there may be up to two DM logical units configured per array. For Copy-on-Write and ShadowImage, the DMLU is an exclusive volume used for storing data when the array system is powered down.

differential-data

The original data blocks replaced by writes to the primary volume. In Copy-on-Write, differential data is stored in the data pool to preserve the copy made of the P-VOL to the time of the snapshot.

disaster recovery

A set of procedures to recover critical application data and processing after a disaster or other failure. Disaster recovery processes include failover and failback procedures.

disk array

An enterprise storage system containing multiple disk drives. Also referred to as “disk array device” or “disk storage system.”

DMLU

See Differential Management-Logical Unit.

DP Pool

Dynamic Provisioning Pool.

dual copy

The process of simultaneously updating a P-VOL and S-VOL while using a single write operation.

duplex

The transmission of data in either one or two directions. Duplex modes are full-duplex and half-duplex. Full-duplex is the simultaneous transmission of data in two direction. For example, a telephone is a full-duplex device, because both parties can talk at once. In contrast, a

A	B	C	D	E	F	G	H	I	J	K	L	M	N	O	P	Q	R	S	T	U	V	W	X	Y	Z

walkie-talkie is a half-duplex device because only one party can transmit at a time.

E

entire copy

Copies all data in the primary volume to the secondary volume to make sure that both volumes are identical.

extent

A contiguous area of storage in a computer file system that is reserved for writing or storing a file.

F

failover

The automatic substitution of a functionally equivalent system component for a failed one. The term failover is most often applied to intelligent controllers connected to the same storage devices and host computers. If one of the controllers fails, failover occurs, and the survivor takes over its I/O load.

fallback

Refers to the process of restarting business operations at a local site using the P-VOL. It takes place after the arrays have been recovered.

Fault tolerance

A system with the ability to continue operating, possibly at a reduced level, rather than failing completely, when some part of the system fails.

FC

See Fibre Channel.

Fibre Channel

A gigabit-speed network technology primarily used for storage networking.

firmware

Software embedded into a storage device. It may also be referred to as Microcode.

FMD

Flash module drive.

A	B	C	D	E	F	G	H	I	J	K	L	M	N	O	P	Q	R	S	T	U	V	W	X	Y	Z

full duplex

The concurrent transmission and the reception of data on a single link.

G**Gbps**

Gigabit(s) per second.

granularity of differential data

Refers to the size or amount of data transferred to the S-VOL during an update cycle. Since only the differential data in the P-VOL is transferred to the S-VOL, the size of data sent to S-VOL is often the same as that of data written to the P-VOL. The amount of differential data that can be managed per write command is limited by the difference between the number of incoming host write operations (inflow) and outgoing data transfers (outflow).

GUI

Graphical user interface.

H**HA**

High availability.

HLUN

A unique host logical unit. The logical host LU within the storage system that is tied to the actual physical LU on the storage system. Each H-LUN on all nodes in the cluster must point to the same physical LU.

I**I/O**

Input/output.

initial copy

An initial copy operation involves copying all data in the primary volume to the secondary volume prior to any update processing. Initial copy is performed when a volume pair is created.

initiator ports

A port-type used for main control unit port of Fibre Remote Copy function.

A	B	C	D	E	F	G	H	I	J	K	L	M	N	O	P	Q	R	S	T	U	V	W	X	Y	Z

IOPS

I/O per second.

iSCSI

Internet-Small Computer Systems Interface. A TCP/IP protocol for carrying SCSI commands over IP networks.

iSNS

Internet-Small Computer Systems Interface. A TCP/IP protocol for carrying SCSI commands over IP networks.

L

LAN

Local Area Network. A computer network that spans a relatively small area, such as a single building or group of buildings.

load

In UNIX computing, the system load is a measure of the amount of work that a computer system is doing.

logical

Describes a user's view of the way data or systems are organized. The opposite of logical is physical, which refers to the real organization of a system. A logical description of a file is that it is a quantity of data collected together in one place. The file appears this way to users. Physically, the elements of the file could live in segments across a disk.

logical unit

See logical unit number.

logical unit number (LUN)

An address for an individual disk drive, and by extension, the disk device itself. Used in the SCSI protocol as a way to differentiate individual disk drives within a common SCSI target device, like a disk array. LUNs are normally not entire disk drives but virtual partitions (or volumes) of a RAID set.

LU

Logical unit.

LUN

See logical unit number.

A	B	C	D	E	F	G	H	I	J	K	L	M	N	O	P	Q	R	S	T	U	V	W	X	Y	Z

LUN Manager

This storage feature is operated through Storage Navigator Modular 2 software and manages access paths among host and logical units for each port in your array.

M

metadata

In sophisticated data systems, the metadata -- the contextual information surrounding the data -- will also be very sophisticated, capable of answering many questions that help understand the data.

microcode

The lowest-level instructions directly controlling a microprocessor. Microcode is generally hardwired and cannot be modified. It is also referred to as firmware embedded in a storage array.

Microsoft Cluster Server

Microsoft Cluster Server is a clustering technology that supports clustering of two NT servers to provide a single fault-tolerant server.

mount

To mount a device or a system means to make a storage device available to a host or platform.

mount point

The location in your system where you mount your file systems or devices. For a volume that is attached to an empty folder on an NTFS file system volume, the empty folder is a mount point. In some systems a mount point is simply a directory.

P

pair

Refers to two logical volumes that are associated with each other for data management purposes (e.g., replication, migration). A pair is usually composed of a primary or source volume and a secondary or target volume as defined by the user.

pair splitting

The operation that splits a pair. When a pair is "Paired," all data written to the primary volume is also copied to the secondary volume. When the pair is "Split," the primary volume continues being updated, but data in the secondary volume remains as it was at the time of the split, until the pair is re-synchronized.

A	B	C	D	E	F	G	H	I	J	K	L	M	N	O	P	Q	R	S	T	U	V	W	X	Y	Z

pair status

Internal status assigned to a volume pair before or after pair operations. Pair status transitions occur when pair operations are performed or as a result of failures. Pair statuses are used to monitor copy operations and detect system failures.

paired volume

Two volumes that are paired in a disk array.

parity

The technique of checking whether data has been lost or corrupted when it's transferred from one place to another, such as between storage units or between computers. It is an error detection scheme that uses an extra checking bit, called the parity bit, to allow the receiver to verify that the data is error free. Parity data in a RAID array is data stored on member disks that can be used for regenerating any user data that becomes inaccessible.

parity groups

RAID groups can contain single or multiple parity groups where the parity group acts as a partition of that container.

peer-to-peer remote copy (PPRC)

A hardware-based solution for mirroring logical volumes from a primary site (the application site) onto the volumes of a secondary site (the recovery site).

point-in-time logical copy

A logical copy or snapshot of a volume at a point in time. This enables a backup or mirroring application to run concurrently with the system.

pool volume

Used to store backup versions of files, archive copies of files, and files migrated from other storage.

primary or local site

The host computer where the primary volume of a remote copy pair (primary and secondary volume) resides. The term "primary site" is also used for host failover operations. In that case, the primary site is the host computer where the production applications are running, and the secondary site is where the backup applications run when the applications on the primary site fail, or where the primary site itself fails.

A	B	C	D	E	F	G	H	I	J	K	L	M	N	O	P	Q	R	S	T	U	V	W	X	Y	Z

primary volume (P-VOL)

The storage volume in a volume pair. It is used as the source of a copy operation. In copy operations a copy source volume is called the P-VOL while the copy destination volume is called "S-VOL" (secondary volume).

P-VOL

See primary volume.

Q

quiesce

Used to describe pausing or altering the state of running processes on a computer, particularly those that might modify information stored on disk during a backup, in order to guarantee a consistent and usable backup. This generally requires flushing any outstanding writes.

R

RAID

Redundant Array of Independent Disks. A disk array in which part of the physical storage capacity is used to store redundant information about user data stored on the remainder of the storage capacity. The redundant information enables regeneration of user data in the event that one of the array's member disks or the access path to it fails.

Recovery Point Objective (RPO)

After a recovery operation, the RPO is the maximum desired time period, prior to a disaster, in which changes to data may be lost. This measure determines up to what point in time data should be recovered. Data changes preceding the disaster are preserved by recovery.

Recovery Time Objective (RTO)

The maximum desired time period allowed to bring one or more applications, and associated data back to a correct operational state. It defines the time frame within which specific business operations or data must be restored to avoid any business disruption.

remote or target site

Maintains mirrored data from the primary site.

remote path

A route connecting identical ports on the local array and the remote array. Two remote paths must be set up for each array (one path for each of the two controllers built in the array).

A	B	C	D	E	F	G	H	I	J	K	L	M	N	O	P	Q	R	S	T	U	V	W	X	Y	Z

remote volume

In TrueCopy operations, the remote volume (R-VOL) is a volume located in a different array from the primary host array.

resynchronization

Refers to the data copy operations performed between two volumes in a pair to bring the volumes back into synchronization. The volumes in a pair are synchronized when the data on the primary and secondary volumes is identical.

RPO

See Recovery Point Objective.

RTO

See Recovery Time Objective.

S**SAS**

Serial Attached SCSI. An evolution of parallel SCSI into a point-to-point serial peripheral interface in which controllers are linked directly to disk drives. SAS delivers improved performance over traditional SCSI because SAS enables up to 128 devices of different sizes and types to be connected simultaneously.

secondary volume (S VOL)

A replica of the primary volume (P-VOL) at the time of a backup and is kept on a standby array. Recurring differential data updates are performed to keep the data in the S-VOL consistent with data in the P-VOL.

SMPL

Simplex.

snapshot

A term used to denote a copy of the data and data-file organization on a node in a disk file system. A snapshot is a replica of the data as it existed at a particular point in time.

SNM2

See Storage Navigator Modular 2.

A	B	C	D	E	F	G	H	I	J	K	L	M	N	O	P	Q	R	S	T	U	V	W	X	Y	Z

SSD

Solid State Disk (drive). A data storage device that uses solid-state memory to store persistent data. An SSD emulates a hard disk drive interface, thus easily replacing it in most applications.

Storage Navigator Modular 2

A multi-featured scalable storage management application that is used to configure and manage the storage functions of Hitachi arrays. Also referred to as "Navigator 2."

suspended status

Occurs when the update operation is suspended while maintaining the pair status. During suspended status, the differential data control for the updated data is performed in the primary volume.

S-VOL

See secondary volume.

S-VOL determination

Independent of update operations, S-VOL determination replicates the S-VOL on the remote array. This process occurs at the end of each update cycle and a pre-determined copy of S-VOL data, consistent with P-VOL data, is maintained on the remote site at all times.

T

target copy

A file, device, or any type of location to which data is moved or copied.

TCMD

TrueCopy Modular Distributed

TrueCopy

Refers to the TrueCopy remote replication.

V

virtual volume (V-VOL)

In Copy-on-Write, a secondary volume in which a view of the primary volume (P-VOL) is maintained as it existed at the time of the last snapshot. The V-VOL contains no data but is composed of pointers to data in the P-VOL and the data pool. The V-VOL appears as a full volume copy to any secondary host.

A	B	C	D	E	F	G	H	I	J	K	L	M	N	O	P	Q	R	S	T	U	V	W	X	Y	Z

volume (VOL)

A disk array object that most closely resembles a physical disk from the operating environment's viewpoint. The basic unit of storage as seen from the host.

volume copy

Copies all data from the P-VOL to the S-VOL.

volume pair

Formed by pairing two logical data volumes. It typically consists of one primary volume (P-VOL) on the local array and one secondary volume (S-VOL) on the remote arrays.

VLAN

Virtual Local Area Network

V-VOL

See virtual volume.

V-VOLTL

Virtual Volume Tape Library.

W**WDM**

Wavelength Division Multiplexing

WOC

WAN Optimization Controller

WMS

Workgroup Modular Storage.

write order guarantee

Ensures that data is updated in an S-VOL, in the same order that it is updated in the P-VOL, particularly when there are multiple write operations in one update cycle. This feature is critical to maintain data consistency in the remote S-VOL and is implemented by inserting sequence numbers in each update record. Update records are then sorted in the cache within the remote system, to assure write sequencing.

write workload

The amount of data written to a volume over a specified period of time.

A	B	C	D	E	F	G	H	I	J	K	L	M	N	O	P	Q	R	S	T	U	V	W	X	Y	Z

Index

Symbols

27-90

A

adding a group name 15-11, 20-10
AMS, version 8-2
array problems, recovering pairs after 21-26
arrays, supported combinations 14-3
arrays, swapping I/O to maintain 20-18
assessing business needs 9-3
assigning pairs to a consistency group 5-7, 10-9, 15-8, 20-4

B

backing up the S-VOL 20-12
backup requirements 4-3
backup script, CLI C-20
backup script, using CLI B-20
backup, protecting from read/write access 5-8
bandwidth
 calculating 19-7
 changing 21-14
 measuring workload for 19-4
bandwidth, calculating 14-28
basic operations 20-2
behavior when data pool over D-42
best practices for data paths 14-52
best practices for remote path 19-32
block size, checking 19-37
business uses of S-VOLs 4-3

C

Cache Partition Manager, initializing for TCE
 installation D-43
Cache Partition Manager, using with SnapShot B-40
Cascade Connection of SnapShot with
 TrueCopy 27-21
cascading
 overview 27-29

with another TrueCopy system 27-30, 27-35, 27-59
with ShadowImage 27-30
with SnapShot 27-59

CCI

change command device D-27
create pairs D-33
define config def file D-28
description 2-22, 7-18, 12-6, 17-14
monitor pair status D-32
release command device D-27
release pairs D-35
resync pairs D-34
set command device D-26
set environment variable D-30
split pairs D-33
suspend pairs D-34
version 8-2

CCI, using to

change a command device C-25
confirm pair status A-29, B-30
create pairs A-30, B-30, C-32
define config def file C-26
define the config def file A-23
release a command device C-25
release pairs A-33, B-34, C-35
restore the P-VOL B-33
resynchronize pairs A-32, C-34
set environment variable A-26, B-27, C-29
set LU mapping A-22, B-23, C-25, D-28
set the command device A-21, C-24
split pairs A-32, C-34

changing a command device using CCI D-27

channel extenders 14-40

checking pair status 5-2, 6-3

CLI

back up S-VOL 20-14
create pairs D-19
description 2-21, 7-18, 12-6, 17-14
display pair status D-18
enable, disable TCE D-9, E-9
install TCE D-8, E-6

- resynchronize pairs [D-20](#)
- set the remote path [D-14](#)
- split pairs [D-20](#)
- swap pairs [D-20](#)
- uninstall TCE [D-10](#)
- CLI, using to
 - change pair info [C-19](#)
 - change pair information [B-18](#)
 - check pair status [A-12](#)
 - create a pair [A-12](#), [C-16](#)
 - create multiple pairs in a group [C-17](#)
 - create pairs [B-14](#)
 - define DMLU [C-8](#)
 - delete a pair [C-18](#)
 - delete the remote path [C-13](#)
 - display pair status [C-15](#)
 - edit pair information [A-16](#)
 - enable and disable SnapShot [B-10](#)
 - enable, disable ShadowImage [A-8](#)
 - install [B-7](#)
 - install ShadowImage [A-7](#)
 - install, enable, disable TrueCopy [C-6](#)
 - release a pair [A-16](#)
 - release DMLU [C-8](#)
 - release pairs [B-17](#)
 - restore the P-VOL [A-15](#), [B-16](#)
 - resync a pair [A-15](#)
 - resynchronize a pair [C-17](#)
 - set up the DMLU [A-9](#)
 - split a pair [A-14](#), [C-17](#)
 - swap a pair [C-18](#)
 - uninstall ShadowImage [A-8](#)
 - update the V-VOL [B-15](#), [B-16](#)
- collecting write-workload data [19-4](#)
- Command Control Interface, see CCI.
- Command Control Interface. See CCI
- command device
 - changing [D-27](#)
 - recommendation for LUs [4-7](#)
 - releasing [A-22](#), [B-23](#), [C-25](#), [D-27](#)
 - set up using GUI [9-35](#)
 - setting up [A-21](#)
 - setup [D-26](#)
- Command Line Interface. See CLI
- configuration definition file [C-26](#)
- configuration definition file, defining [D-28](#)
- Configuration Restrictions on the Cascade of TrueCopy with SnapShot [27-25](#)
- configuration workflow [9-31](#)
- configuring ShadowImage [4-22](#)
- consistency group
 - checking status with CLI [D-23](#)
 - creating in GUI [15-8](#)
 - creating, assigning pairs to [20-4](#)
 - description [12-6](#), [17-11](#)
 - specifications [C-3](#)
 - using CCI for operations [D-37](#)
- Consistency Groups
 - creating and assigning pairs to using GUI [10-9](#)
 - creating pairs for using CLI [B-18](#)
 - creating, assigning pairs to [5-7](#)
 - description [2-18](#), [7-11](#)
 - number allowed [A-3](#)
 - copy pace [5-5](#), [15-4](#)
 - Copy Pace, changing [21-16](#)
 - Copy Pace, specifying [20-5](#)
 - create pair [10-6](#)
 - create pair procedure [20-3](#)
 - creating a pair [5-2](#), [15-3](#)
 - creating the V-VOL [10-6](#)
 - CTG. See consistency group
 - cycle time, monitoring, changing in GUI [21-15](#)

D

- dark fibre [D-45](#)
- data fence level [15-6](#)
- data path
 - defining [14-24](#)
 - description [12-4](#)
 - failure and data recovery [15-20](#)
- Data path, planning [19-14](#)
- data path. See remote path
- data paths
 - best practices [14-52](#)
 - channel extenders [14-40](#)
 - designing [14-34](#)
 - preventing blockage [14-52](#)
 - supported configurations [14-34](#)
- data pools
 - description [17-9](#)
 - editing [10-15](#)
 - expanding [11-7](#), [21-10](#)
 - measuring workload for [19-4](#)
 - specifications [B-3](#)
- data recovery, versus performance [14-31](#)
- Data Retention Utility [C-3](#)
- data, measuring write-workload [19-4](#)
- definitions, pair status [6-3](#)
- deleting
 - remote path [21-20](#)
 - volume pair [21-19](#)
- deleting a pair [5-12](#), [15-12](#), [20-11](#)
- deleting the remote path [15-13](#), [19-60](#), [24-15](#), [C-13](#), [D-17](#), [E-18](#)
- design workflow [4-2](#)
- designating a command device [A-21](#)
- designing the SnapShot system [9-2](#)
- designing the system [19-2](#)
- Differential Management Logical Unit. See DMLU
- Differential Management-Logical Unit. See DMLU
- direct connection [14-35](#)
- disabling ShadowImage [3-6](#)
- disabling SnapShot [8-1](#)
- disaster recovery process [20-21](#)
- DMLU
 - defining [14-20](#)

- description 12-5, 17-14
- recommendation for LUs 4-7
- setup 4-25
- setup, CLI A-9

drive types supported C-3

dynamic disk with Windows 2000 Server 14-11

dynamic disk with Windows Server 2000 19-40

Dynamic Provisioning 4-12, 9-26, 14-14, 19-45

E

- editing data pool information 10-15
- editing pair information 5-13, 10-15, 15-11, 20-10
- enabling ShadowImage 3-6
- enabling SnapShot 8-1
- enabling, disabling TCE 18-5, 23-7
- enabling, with CLI D-9, E-9
- enabling/disabling 13-4
- environment variable D-30
- error codes, failure during resync 21-30
- Event Log, using 21-32
- expanding data pool size 11-7, 21-10
- extenders D-45

F

- failback procedure 20-22
- fence level 15-6
- fibre channel extenders 14-40
- Fibre Channel remote path requirements and configurations 19-14
- Fibre Channel, port transfer-rate 19-21
- fibre channel, port transfer-rate 14-41
- frequency, snapshot 9-4

G

- graphic, SnapShot hardware and software 7-2
- Group Name field 15-8
- Group Name, adding 5-8, 10-9, 20-4
- group name, adding 15-11, 20-10
- GUI, description 2-21, 7-18, 12-6, 17-14
- GUI, using to
 - assign pairs to a Consistency Group 5-7, 15-8
 - assign pairs to a consistency group 20-4
 - check pair status 5-2
 - create a pair 5-2, 15-3
 - define DMLU 14-20
 - define remote path 14-24
 - delete a pair 5-12, 10-14, 15-12, 20-11, 21-19
 - delete a remote path 21-20
 - delete a V-VOL 10-14
 - delete remote path 15-13, 19-60, 24-15, C-13, D-17, E-18
 - edit a pair 5-13
 - edit pair information 15-11, 20-10
 - edit pairs 10-15

- enable, disable ShadowImage 3-6
- install 8-4
- install ShadowImage 3-4
- install, enable/disable TrueCopy 13-4
- monitor pair status 6-3, 16-4, 21-4
- restore the P-VOL 5-13, 10-13
- resync a pair 5-10
- resynchronize a pair 15-10, 20-8
- set up remote path 19-58
- set up the command device 9-35
- set up the DMLU 4-25
- set up the V-VOL 9-34
- split a pair 5-9, 15-9, 20-6
- swap a pair 15-11, 20-9
- uninstall 8-6
- uninstall ShadowImage 3-7
- update the V-VOL 10-11

H

- horctakeover 20-21
- host group, connecting to HP server 14-8, 19-38
- host recognition of P-VOL, S-VOL 14-7, 19-38
- host server failure, recovering the data 15-21
- host time-out recommendation 14-7, 19-38
- how long to hold snapshots 9-5
- how long to keep S-VOL 4-3
- how often to copy P-VOL 4-2
- how often to take snapshots 9-4

I

- I/O performance, versus data recovery 14-31
- I/O Switching Mode
 - description A-35
 - enabling using GUI A-38
 - setup with CLI A-11
 - specifications A-36
- initial copy 20-2
- installation 3-4, 8-4, 13-4
- installing SnapShot 8-1
- installing TCE with CLI D-8, E-6
- installing TCE with GUI 18-3, 23-3
- interfaces for ShadowImage 2-21
- interfaces for SnapShot 7-18
- interfaces for TCE 17-14
- interfaces for TrueCopy 12-6
- iSCSI remote path requirements and configurations 19-22

K

- key code, key file 13-3

L

- LAN requirements 14-33, 19-13
- license A-5
- lifespan, snapshot 9-5
- lifespan, S-VOLs 4-3
- logical units, pair recommendations 14-4

logical units, recommendations [19-36](#)
LUN expansion [14-5](#)

M

maintaining local array, swapping I/O [20-18](#)
maintaining the SnapShot system [11-2](#)
MC/Service Guard [14-8](#), [19-38](#)
measuring write-workload [14-28](#), [19-4](#)
memory, reducing [C-4](#)
monitoring
 pair status [21-4](#)
 remote path [16-9](#), [21-14](#)
monitoring data pool usage [11-2](#)
monitoring pair status [11-2](#)
monitoring ShadowImage [6-1](#)
moving data procedure [20-19](#)

N

never fence level [15-6](#)
number of copies to make [4-4](#)
number of V-VOLs, establishing [9-6](#)

O

operating systems, restrictions with [14-7](#), [19-38](#)
operations [20-2](#)
overview [7-1](#)

P

Pace field [5-5](#), [15-4](#)
Pair Name field, differences on local, remote array [20-4](#)
pair names and group names, Nav2 differences from CCI [D-42](#)
pair operation
 restrictions [27-5](#)
pair operations using CCI [C-30](#)
pair-monitoring script, CLI [C-21](#)
pairs
 assigning to a consistency group [5-7](#), [15-8](#), [20-4](#)
 creating [5-2](#), [15-3](#), [19-36](#)
 deleting [5-12](#), [15-12](#), [20-11](#), [21-19](#)
 description [17-6](#)
 displaying status with CLI [D-18](#)
 editing [5-13](#)
 monitoring status with GUI [21-4](#)
 monitoring with CCI [D-32](#)
 number allowed [A-2](#)
 recommendations [19-36](#)
 recommendations for volumes [4-6](#)
 resynchronizing [15-10](#), [20-8](#)
 resyncing [5-10](#)
 splitting [5-9](#), [15-9](#), [20-6](#)
 status definitions [21-5](#)
 status definitions and checking [6-3](#)
 status monitoring, definitions [16-4](#)
 swapping [15-11](#), [20-9](#)

pairs resyncing [5-10](#)
pairs, assigning to a consistency group [10-9](#)
path failure, recovering the data [15-20](#)
performance info for multiple paths [14-41](#)
planning
 LUN expansion [14-5](#)
 remote path [14-34](#), [19-14](#)
 TCE volumes [19-36](#)
 workflow [14-3](#)
planning a ShadowImage system [4-2](#)
Planning the remote path [19-14](#)
planning the SnapShot system [9-2](#)
planning workflow [4-2](#)
platforms supported [3-3](#)
platforms, supported [8-3](#)
port transfer-rate [14-41](#), [19-21](#)
Power Saving [C-4](#)
prerequisites for pair creation [19-36](#)
primary volume [2-2](#)
production site failure, recovering the data [15-22](#)
P-VOL
 and S-VOL setup [4-22](#)
 and S-VOL, definition [2-4](#)
 restoring [5-13](#)
P-VOL and S-VOL, definitions [12-2](#)
P-VOLs and V-VOLs [7-4](#)

R

RAID grouping for volume pairs [A-2](#)
RAID groups and volume pairs [19-37](#)
RAID level for volume pairs [A-2](#)
RAID levels for SnapShot volumes [9-16](#)
RAID levels supported [C-3](#)
recovering after array problems [21-26](#)
recovering data
 data path failure [15-20](#)
 host server failure [15-21](#)
 production site failure [15-22](#)
recovering from failure during resync [21-30](#)
release a command device [C-25](#)
release a command device, using CCI [D-27](#)
releasing a command device [A-22](#), [B-23](#)
remote array restriction, Sync Cache Ex Mode [14-4](#)
remote array, shutdown, TCE tasks [21-20](#)
Remote path
 planning [19-14](#)
remote path
 best practices [19-32](#)
 defining [14-24](#)
 deleting [15-13](#), [19-60](#), [21-20](#), [24-15](#), [C-13](#), [D-17](#), [E-18](#)
 description [12-4](#), [19-14](#)
 guidelines [14-34](#)
 monitoring [16-9](#), [21-14](#)
 planning [14-34](#), [19-14](#)
 preventing blockage [19-32](#)
 requirements [19-14](#)

Index-4

- setup with CLI [D-14](#)
- setup with GUI [14-25, 19-58, 24-12, 25-13, 25-14, E-11, E-12, E-14](#)
- supported configurations [19-14](#)
- Replication Manager [2-22, 4-30, 12-7](#)
- reports, using the V-VOL for [10-16](#)
- Requirements
 - bandwidth, for WANs [19-7](#)
 - LAN [14-33, 19-13](#)
- requirements [3-2, 18-2, 23-2](#)
- SnapShot system [8-2](#)
- response time for pairsplit [D-40](#)
- restoring the P-VOL [5-13, 10-13](#)
- restrictions on cascading TCE with SnapShot [27-28](#)
- resync a pair [10-11](#)
- resynchronization error codes [21-30](#)
- resynchronization errors, correcting [21-30](#)
- resynchronizing a pair [15-10, 20-8](#)
- resyncing a pair [5-10](#)
- RPO, checking [21-17](#)
- RPO, update cycle [19-3](#)

S

- scripts
 - CLI backup [C-20](#)
 - CLI pair-monitoring [C-21](#)
- scripts for backups (CLI) [20-12, D-25](#)
- secondary volume [2-2](#)
- setting port transfer-rate [14-41, 19-21](#)
- ShadowImage
 - configuring [4-22](#)
 - enable, disable [3-6](#)
 - environment [2-2](#)
 - how it works [2-4](#)
 - installing [3-4](#)
 - interface [2-21](#)
 - maintaining [6-1](#)
 - plan and design [4-2](#)
 - specifications [A-2](#)
 - uninstalling [3-7](#)
 - using [5-1](#)
 - workflow [5-2, 10-2](#)
- ShadowImage, cascading with [27-30](#)
- SnapShot
 - behaviors vs TCE's [D-42](#)
 - enabling, disabling [8-1](#)
 - how it works [7-3](#)
 - installing [8-4](#)
 - installing, uninstalling [8-1](#)
 - interface [7-18](#)
 - interfaces [12-7](#)
 - maintaining [11-2](#)
 - overview [7-1](#)
 - planning [9-2](#)
 - restoring the P-VOL operation [10-13](#)
 - uninstalling [8-6](#)
 - using with Cache Partition Manager [B-40](#)
- SnapShot* versus *snapshot* [7-1](#)

- SnapShot, cascading with [27-59](#)
- snapshots
 - how long to keep [9-5](#)
 - how often to make [9-4](#)
- specifications [A-2, B-3, C-2, D-2, E-2, E-4](#)
- split pair procedure [20-6](#)
- splitting a pair [10-11](#)
- splitting the pair [5-9, 15-9](#)
- status definitions [6-3](#)
- statuses, pair [21-5](#)
- Storage Navigator Modular 2
 - description [12-7](#)
 - version [8-2](#)
- supported data path configurations [14-34](#)
- supported platforms [3-3, 8-3](#)
- supported remote path configurations [19-14](#)
- S-VOL
 - description [2-4, 12-2](#)
 - frequency, lifespan, number of [4-2](#)
 - number allowed [A-2](#)
 - specifying as backup only [5-8](#)
 - updating [5-10, 15-10](#)
 - using [5-15](#)
- S-VOL, backing up [20-12](#)
- S-VOL, updating [20-8](#)
- swapping pairs [15-11, 20-9](#)
- switch connection [14-36](#)
- Synchronize Cache Execution Mode [14-4](#)
- system requirements [3-2](#)

T

- takeover [20-21](#)
- tape backups [10-16](#)
- TCE
 - backing up the S-VOL [20-12](#)
 - behaviors vs SnapShot's [D-42](#)
 - calculating bandwidth [19-7](#)
 - changing bandwidth [21-14](#)
 - create pair procedure [20-3](#)
 - data pool
 - description [17-9](#)
 - environment [17-6](#)
 - how it works [17-2](#)
 - interface [17-14](#)
 - monitoring pair status [21-4](#)
 - operations [20-2](#)
 - operations before firmware updating [21-20](#)
 - pair recommendations [19-36](#)
 - procedure for moving data [20-19](#)
 - remote path configurations [14-34, 19-14](#)
 - requirements [18-2, 23-2](#)
 - setting up the remote path [19-58](#)
 - setup [19-51](#)
 - SnapShot cascade restrictions [27-90](#)
 - splitting a pair [20-6](#)
 - typical environment [17-5](#)
- TCMD
 - aggregation backup [25-1](#)
 - CLI operations [E-1](#)

- configuration [25-1](#)
- installation [23-3](#)
- overview [22-2](#)
- planning and design [24-1](#)
- setting distributed mode [25-13](#)
- setup procedures [24-1](#)
- system requirements [23-2](#)
- system specifications [E-1](#)
- troubleshooting [26-2](#)
- testing, using the V-VOL for [10-16](#)
- troubleshooting [16-10](#)
- TrueCopy
 - defining the remote path (GUI) [14-24](#)
 - how it works [12-2](#)
 - installing, enabling, disabling [13-4](#)
 - interface [12-6](#)
 - operations overview [12-7](#)
 - pair status monitoring, definitions [16-5](#)
 - troubleshooting pair failure [16-10](#)
 - troubleshooting path blockage [16-10](#)
 - typical environment [12-3](#)
 - using unified LUs [14-5](#)

U

- unified LUs, in TrueCopy volumes [14-5](#)
- uninstalling [13-5](#)
- uninstalling ShadowImage [3-7](#)
- uninstalling SnapShot [8-1](#), [8-6](#)
- uninstalling with CLI [D-10](#)
- uninstalling with GUI [18-6](#), [23-5](#)
- update cycle [17-2](#), [17-10](#), [19-3](#)
 - specifying cycle time [21-15](#)
- updating firmware, TCE tasks [21-20](#)
- updating the S-VOL [5-10](#), [15-10](#), [20-8](#)
- using the S-VOL [5-15](#)

V

- version
 - AMS [8-2](#)
 - CCI [8-2](#)
 - Navigator 2 [8-2](#)
- Volume Migration [C-3](#)
- volume pair description [17-6](#)
- volume pairs
 - creating [10-6](#)
 - description [2-4](#), [7-4](#)
 - editing [10-15](#)
 - monitoring status [11-2](#)
 - RAID levels and grouping [A-2](#)
 - setup recommendations [4-6](#)
- volume pairs, recommendations [19-36](#), [19-37](#)
- volumes
 - setup recommendations [9-14](#)
- V-VOLs
 - creating [10-6](#)
 - description [7-4](#)
 - establishing number of [9-6](#)
 - procedure for secondary uses [10-16](#)
 - updating [10-11](#)

W

- WAN
 - bandwidth requirements [19-7](#)
 - configurations supported [14-44](#), [19-24](#)
 - general requirements [14-33](#), [19-13](#)
 - types supported [14-33](#), [19-13](#)
- WDM [D-45](#)
- Windows 2000 Server, restrictions [14-11](#)
- Windows Server 2000 restrictions [19-40](#)
- Windows Server 2003 restrictions [19-40](#)
- WOCs, configurations supported [19-27](#)
- write order [17-10](#)
- write-workload [19-4](#)
- write-workload, measuring [14-28](#)

Hitachi Data Systems

Corporate Headquarters

2845 Lafayette Street
Santa Clara, California 95050-2639
U.S.A.
www.hds.com

Regional Contact Information

Americas

+1 408 970 1000
info@hds.com

Europe, Middle East, and Africa

+44 (0)1753 618000
info.emea@hds.com

Asia Pacific

+852 3189 7900
hds.marketing.apac@hds.com



MK-91DF8274-18