

Hitachi AMS 2000 Family Copy-on-Write SnapShot User Guide

FASTFIND LINKS

- [Document organization](#)
- [Release notes and readme](#)
- [Getting help](#)
- [Table of Contents](#)

Copyright © 2010 Hitachi Ltd., Hitachi Data Systems Corporation, 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 Hitachi Data Systems' applicable agreement(s). The use of Hitachi Data Systems products is governed by the terms of your agreement(s) 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 in the United States and other countries.

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



Table of Contents

Preface	vii
Intended audience	viii
Product version	viii
Release notes and readme	viii
Document revision level	viii
Changes in this revision	viii
Document organization	ix
Document conventions	x
Convention for storage capacity values	xi
Accessing product documentation	xi
Getting help	xvii
Comments	xvii
1 SnapShot overview	1-1
Copy-on-Write SnapShot software	1-2
Hardware and software configuration	1-2
How SnapShot works	1-3
Volume pairs — P-VOLs and V-VOLs	1-3
Data pools	1-4
Consistency group (CTG)	1-6
Differential Management LUs (DMLU)	1-6
LU Ownership of P-VOLs and Data Pools	1-6
Interfaces for performing SnapShot operations	1-6
Cascade connection of SnapShot with Simple DR	1-8
Cascade connection of SnapShot with Simple DR P-VOL	1-9
Cascade connection of SnapShot with Simple DR S-VOL	1-10
V-VOLs number of SnapShot	1-10
Cascade connection of SnapShot with TrueCopy	1-11
Cascade restrictions with SnapShot P-VOL	1-11
LU shared with P-VOL on SnapShot and P-VOL on TrueCopy	1-12
One LU used for P-VOL on SnapShot and S-VOL on TrueCopy	1-12

V-VOLs number of SnapShot	1-13
Cascade restrictions with SnapShot V-VOL	1-13
Configuration restrictions on the Cascade of TrueCopy with SnapShot	1-15
Cascade restrictions with SnapShot Data Pool	1-15
2 Planning and design.	2-1
The plan and design workflow	2-2
Assessing business needs	2-2
Establishing how often a copy is made (copy frequency)	2-2
Determine how often a snapshot should be taken	2-3
Selecting a reasonable time between Snapshots	2-3
Establishing how long a copy is held (copy lifespan).	2-3
Lifespan based on backup requirements	2-4
Lifespan based on business uses	2-4
Establishing the number of V-VOLs that Are required	2-4
Establishing data pool size	2-5
Measuring workload data, sizing the data pool.	2-6
Rule-of-thumb calculation	2-9
Requirements and recommendations for SnapShot Logical Units	2-10
RAID configuration for LUs assigned to SnapShot	2-10
Operating system considerations	2-13
Identifying P-VOL and V-VOL LUs on Windows	2-13
LU mapping	2-13
Cluster and path switching software	2-14
Microsoft Cluster Server (MSCS).	2-14
AIX	2-14
Veritas Volume Manager (VxVM)	2-14
Windows 2000	2-14
Windows Server 2003/2008 and SnapShot configuration	2-15
Linux and LVM configuration	2-16
Tru64 UNIX and SnapShot configuration	2-16
Concurrent use of Cache Partition Manager	2-16
Concurrent use of Dynamic Provisioning	2-16
Windows Server 2003 and 2008/Windows 2000 and Dynamic Disk	2-17
Cascading SnapShot with TrueCopy.	2-20
Cascading SnapShot with TrueCopy Extended Distance	2-21
Maximum supported capacity	2-22
SnapShot and TCE capacity	2-23
SnapShot, TCE, ShadowImage concurrent capacity	2-25
Cache limitations on Data and Data Pool volumes	2-27
3 System requirements	3-1
System requirements	3-2
Displaying the hardware revision number	3-3

Supported platforms.	3-4
4 Installing and enabling SnapShot	4-1
Important prerequisite information	4-2
Installing, uninstalling SnapShot	4-2
Enabling, disabling SnapShot.	4-4
5 Configuring SnapShot	5-1
Configuration workflow.	5-2
Set up the data pool.	5-2
Set Up the Virtual Volume (V-VOL) (manual method)	5-3
Set up the Differential Management LU (DMLU)	5-4
Set up the command device	5-5
Setting the LU ownership	5-6
6 Using SnapShot	6-1
SnapShot replication workflow.	6-2
Back up your volume — creating a pair	6-2
Using the Backup Wizard	6-2
Using the Create Pair procedure	6-3
Update the V-VOL	6-5
Restore the P-VOL from the V-VOL	6-6
Use the V-VOL for tape backup, testing, reports, etc.	6-7
Tape backup recommendations	6-8
Restoring data from a tape backup	6-9
Edit pairs, data pool	6-10
Delete pairs, V-VOLs, data pools, DMLU	6-10
7 Monitoring and maintenance	7-1
Monitoring SnapShot	7-2
Monitoring pair status	7-2
Monitoring data pool usage.	7-3
Expanding data pool capacity	7-3
Other methods for lowering data pool load	7-3
8 Troubleshooting	8-1
Pair failure.	8-2
Recovering from pair failure due to POOL FULL	8-2
Recovering from pair failure due to a hardware failure	8-3
Data pool capacity exceeded.	8-3
A Specifications	A-1

SnapShot specifications	A-2
B Operations using CLI.	B-1
Installing and uninstalling SnapShot	B-2
Important prerequisite information	B-2
Uninstalling SnapShot	B-4
Enabling and disabling SnapShot.	B-5
Operations for SnapShot configuration.	B-6
Setting the DMLU	B-6
Setting the POOL	B-7
Setting the V-VOL.	B-9
Setting the LU ownership	B-9
Performing SnapShot operations	B-10
Creating SnapShot pairs	B-10
Updating SnapShot Logical Unit	B-11
Restoring V-VOL to P-VOL	B-11
Releasing SnapShot pairs	B-12
Changing pair information	B-13
Creating pairs that belong to a group	B-13
Sample back up script for Windows	B-14
C Operations using CCI	C-1
Setting up CCI.	C-2
Setting the command device	C-2
Setting mapping information	C-3
Defining the configuration definition file	C-4
Setting the environment variable	C-6
Performing SnapShot operations	C-8
Confirming pair status.	C-8
Paircreate operation	C-9
Pair creation using a consistency group.	C-10
Updating the V-VOL	C-10
Restoring a V-VOL to the P-VOL.	C-11
Releasing SnapShot pairs	C-11
Pair and group name differences in CCI and Navigator 2	C-12
D Using SnapShot with Cache Partition Manager	D-1
SnapShot with Cache Partition Manager.	D-2

Glossary

Index



Preface

This document provides instructions on assessing your snapshot requirements, designing an implementation to meet those requirements, and implementing and operating Copy-on-Write Snapshot software using the Storage Navigator 2 graphical user interface.

This preface includes the following information:

- [Intended audience](#)
- [Product version](#)
- [Release notes and readme](#)
- [Document revision level](#)
- [Changes in this revision](#)
- [Document organization](#)
- [Document conventions](#)
- [Convention for storage capacity values](#)
- [Accessing product documentation](#)
- [Getting help](#)

Intended audience

This document is intended for system administrators, Hitachi Data Systems representatives, and Authorized Service Providers who install, configure, and operate Hitachi Adaptable Modular System (AMS) 2000 family storage systems.

Product version

This document applies to Hitachi AMS 2000 Family firmware version 0890/A or later.

Release notes and readme

Read the release notes and readme file before installing and using this product. They may contain requirements or restrictions that are not fully described in this document and/or updates or corrections to this document.

Document revision level

This section provides a history of the revision changes to this document.

Revision	Date	Description
MK-97DF8124-01	October 2008	Initial release.
MK-97DF8124-02	December 2008	This release supersedes release 01.
MK-97DF8124-03	March 2009	This release supersedes release 02.
MK-97DF8124-04	June 2009	This release supersedes release 03.
MK-97DF8124-05	August 2009	This release supersedes release 04.
MK-97DF8124-06	November 2009	This release supersedes release 05.
MK-97DF8124-07	January 2010	This release supersedes release 06.
MK-97DF8124-08	April 2010	This release supersedes release 07.

Changes in this revision

The following information has been added for this release:

- Added new section [LU Ownership of P-VOLs and Data Pools on page 1-6](#).
- Updated [Table 3-1 on page 3-2](#).
- Updated [Table 3-2 on page 3-4](#) to reflect Windows Server 2008 information.
- Added new section [Setting the LU ownership on page 5-6](#).
- Added new section [Setting the LU ownership on page B-9](#).




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, SnapShot overview	Provides descriptions of SnapShot components and how they work together.
Chapter 2, Planning and design	Guides you in planning a configuration that meets organization needs and the workload requirements of your host application.
Chapter 3, System requirements	Provides SnapShot requirements.
Chapter 4, Installing and enabling SnapShot	Provides instructions for installing SnapShot.
Chapter 5, Configuring SnapShot	Provides configuration information.
Chapter 6, Using SnapShot	Provides information and procedures for using SnapShot.
Chapter 7, Monitoring and maintenance	Provides information and instructions for monitoring and maintaining the SnapShot system.
Chapter 8, Troubleshooting	Provides information for correcting system problems.
Appendix A, Specifications	Provides SnapShot specifications.
Appendix B, Operations using CLI	Provides Navigator 2 Command Line Interface instructions for configuring and using SnapShot.
Appendix C, Operations using CCI	Provides detailed Command Control Interface instructions for configuring and using SnapShot.
Appendix D, Using SnapShot with Cache Partition Manager	Provides information for using SnapShot with Cache Partition Manager.
Glossary	Provides definitions for terms and acronyms found in this document.
Index	Provides locations to specific information in this document.

Document conventions

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.

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 the user or system. Example: copy source-file target-file. Note: Angled brackets (< >) are also used to indicate variables.
screen/code	Indicates text that is displayed on screen or entered by the user. Example: # pairdisplay -g oradb
< > angled brackets	Indicates a variable, which is a placeholder for actual text provided by the user or system. Example: # pairdisplay -g <group> Note: 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: [a b]

Convention for storage capacity values

Physical storage capacity values (e.g., 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 (e.g., 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 AMS 2000 Family user documentation is available on the Hitachi Data Systems Portal: <https://portal.hds.com>. Please check this site for the most current documentation, including important updates that may have been made after the release of the product.

This documentation set consists of the following documents.


Release notes

- Adaptable Modular Storage System Release Notes
- Storage Navigator Modular 2 Release Notes



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

Installation and getting started

The following documents provide instructions for installing an AMS 2000 Family storage system. They include rack information, safety information, site-preparation instructions, getting-started guides for experienced users, and host connectivity information. The symbol  identifies documents that contain initial configuration information about Hitachi AMS 2000 Family storage systems.

 **AMS2100/2300 Getting Started Guide**, MK-98DF8152

Provides quick-start instructions for getting an AMS 2100 or AMS 2300 storage system up and running as quickly as possible.

 **AMS2500 Getting Started Guide**, MK-97DF8032

Provides quick-start instructions for getting an AMS 2500 storage system up and running as quickly as possible.

AMS 2000 Family Site Preparation Guide, MK-98DF8149

Contains site planning and pre-installation information for AMS 2000 Family storage systems, expansion units, and high-density expansion units. This document also covers safety precautions, rack information, and product specifications.

AMS 2000 Family Fibre Channel Host Installation Guide,
MK-08DF8189

Describes how to prepare Hitachi AMS 2000 Family Fibre Channel storage systems for use with host servers running supported operating systems.

AMS 2000 Family iSCSI Host Installation Guide, MK-08DF8188

Describes how to prepare Hitachi AMS 2000 Family iSCSI storage systems for use with host servers running supported operating systems.

Storage and replication features

The following documents describe how to use Storage Navigator Modular 2 (Navigator 2) to perform storage and replication activities.

Storage Navigator 2 Advanced Settings User's Guide, MK-97DF8039

Contains advanced information about launching and using Navigator 2 in various operating systems, IP addresses and port numbers, server certificates and private keys, boot and restore options, outputting configuration information to a file, and collecting diagnostic information.

Storage Navigator Modular 2 User's Guide, MK-99DF8208

Describes how to use Navigator 2 to configure and manage storage on an AMS 2000 Family storage system.

AMS 2000 Family Dynamic Provisioning Configuration Guide, MK-09DF8201

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

Storage Navigator 2 Storage Features Reference Guide for AMS, MK-97DF8148

Contains concepts, preparation, and specifications for Account Authentication, Audit Logging, Cache Partition Manager, Cache Residency Manager, Data Retention Utility, LUN Manager, Performance Monitor, SNMP Agent, and Modular Volume Migration.

AMS 2000 Family Copy-on-write SnapShot User Guide, MK-97DF8124 — this document

Describes how to create point-in-time copies of data volumes in AMS 2100, AMS 2300, and AMS 2500 storage systems, without impacting host service and performance levels. Snapshot copies are fully read/write compatible with other hosts and can be used for rapid data restores, application testing and development, data mining and warehousing, and nondisruptive backup and maintenance procedures.

AMS 2000 Family ShadowImage In-system Replication User Guide, MK-97DF8129

Describes how to perform high-speed nondisruptive local mirroring to create a copy of mission-critical data in AMS 2100, AMS 2300, and AMS 2500 storage systems. ShadowImage keeps data RAID-protected and fully recoverable, without affecting service or performance levels. Replicated data volumes can be split from host applications and used for system backups, application testing, and data mining applications while business continues to operate at full capacity.

AMS 2000 Family TrueCopy Remote Replication User Guide, MK-97DF8052

Describes how to create and maintain multiple duplicate copies of user data across multiple AMS 2000 Family storage systems to enhance your disaster recovery strategy.

AMS 2000 Family TrueCopy Extended Distance User Guide, MK-97DF8054

Describes how to perform bi-directional remote data protection that copies data over any distance without interrupting applications, and provides failover and recovery capabilities.


AMS 2000 Data Retention Utility User's Guide, MK-97DF8019

Describes how to lock disk volumes as read-only for a certain period of time to ensure authorized-only access and facilitate immutable, tamper-proof record retention for storage-compliant environments. After data is written, it can be retrieved and read only by authorized applications or users, and cannot be changed or deleted during the specified retention period.

Storage Navigator Modular 2 online help

Provides topic and context-sensitive help information accessed through the Navigator 2 software.

Hardware maintenance and operation

The following documents describe how to operate, maintain, and administer an AMS 2000 Family storage system. They also provide a wide range of technical information and specifications for the AMS 2000 Family storage systems. The symbol  identifies documents that contain initial configuration information about Hitachi AMS 2000 Family storage systems.

AMS 2100/2300 Storage System Hardware Guide, MK-97DF8010

Provides detailed information about installing, configuring, and maintaining an AMS 2100/2300 storage system.

AMS 2500 Storage System Hardware Guide, MK-97DF8007

Provides detailed information about installing, configuring, and maintaining an AMS 2500 storage system.

AMS 2000 Family Storage System Reference Guide, MK-97DF8008

Contains specifications and technical information about power cables, system parameters, interfaces, logical blocks, RAID levels and configurations, and regulatory information about AMS 2100, AMS 2300, and AMS 2500 storage systems. This document also contains remote adapter specifications and regulatory information.

AMS 2000 Family Storage System Service and Upgrade Guide, MK-97DF8009

Provides information about servicing and upgrading AMS 2100, AMS 2300, and AMS 2500 storage systems.

AMS 2000 Family Power Savings User Guide, MK-97DF8045

Describes how to spin down volumes in selected RAID groups when they are not being accessed by business applications to decrease energy consumption and significantly reduce the cost of storing and delivering information.

Command and Control (CCI)

The following documents describe how to install the Hitachi AMS 2000 Family Command Control Interface (CCI) and use it to perform TrueCopy and ShadowImage operations.

AMS 2000 Family Command Control Interface (CCI) Installation Guide, MK-97DF8122

Describes how to install CCI software on open-system hosts.

AMS 2000 Family Command Control Interface (CCI) Reference Guide, MK-97DF8121

Contains reference, troubleshooting, and maintenance information related to CCI operations on AMS 2100, AMS 2300, and AMS 2500 storage systems.

AMS 2000 Family Command Control Interface (CCI) User's Guide, MK-97DF8123

Describes how to use CCI to perform TrueCopy and ShadowImage operations on AMS 2100, AMS 2300, and AMS 2500 storage systems.

Command Line Interface (CLI)

The following documents describe how to use Hitachi Storage Navigator Modular 2 to perform management and replication activities from a command line.

Storage Navigator Modular 2 Command Line Interface (CLI) Unified Reference Guide, MK-97DF8089

Describes how to interact with all Navigator 2 bundled and optional software modules by typing commands at a command line.

Storage Navigator 2 Command Line Interface Replication Reference Guide for AMS, MK-97DF8153

Describes how to interact with Navigator 2 to perform replication activities by typing commands at a command line.

Dynamic Replicator documentation

The following documents describe how to install, configure, and use Hitachi Dynamic Replicator to provide AMS Family storage systems with continuous data protection, remote replication, and application failover in a single, easy-to-deploy and manage platform.

Dynamic Replicator - Scout Release Notes, RN-99DF8211

Dynamic Replicator - Scout Host Administration Guide, MK-98DF8212

Dynamic Replicator - Scout Installation and Configuration Guide,
MK-98DF8213

Dynamic Replicator - Scout Quick Start Guide, MK-98DF8214

Dynamic Replicator - Scout Host Troubleshooting Guide,
MK-98DF8215

Dynamic Replicator DR-Scout ICAT Utility Guide, MK-98DF8216

Dynamic Replicator - Scout RX Server Deployment Guide,
MK-98DF8217

Dynamic Replicator VX Solution for Oracle (Solaris), MK-98DF8218

Dynamic Replicator - Scout Solution for SharePoint 2007,
MK-98DF8219

Dynamic Replicator - Scout Solution for MySQL (Windows),
MK-98DF8220

**Protecting Citrix XenServer Using Hitachi Dynamic Replicator -
Scout**, MK-98DF8221

Dynamic Replicator Quick Install/Upgrade Guide, MK-98DF8222

Dynamic Replicator - Scout Protecting MS SQL Server, MK-98DF8223

Dynamic Replicator - Scout - Protecting Microsoft Exchange Server,
MK-98DF8224

Dynamic Replicator - Scout File Server Solution, MK-98DF8225

Dynamic Replicator - Scout ESX - Protecting ESX Server (RCLI),
MK-99DF8226

Getting help

If you need to contact the Hitachi Data Systems support center, please provide as much information about the problem as possible, including:

- The circumstances surrounding the error or failure.
- The exact content of any messages displayed on the host system(s).
- The exact content of any messages displayed on Storage Navigator Modular 2.
- The Storage Navigator Modular 2 configuration information. This information is used by service personnel for troubleshooting purposes.

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 Hitachi Data Systems 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, number, and revision, and refer to specific section(s) and paragraph(s) whenever possible.

Thank you! (All comments become the property of Hitachi Data Systems.)

SnapShot overview

Copy-on-Write SnapShot creates virtual copies of data volumes within the Hitachi Adaptable Modular Storage (AMS) 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.

This guide provides instructions for planning and designing, configuring and testing, and using and monitoring SnapShot. In this chapter, see the following:

- ❑ [Copy-on-Write SnapShot software](#)
- ❑ [Hardware and software configuration](#)
- ❑ [How SnapShot works](#)
- ❑ [Interfaces for performing SnapShot operations](#)
- ❑ [Cascade connection of SnapShot with Simple DR](#)
- ❑ [Cascade connection of SnapShot with TrueCopy](#)



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 AMS 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 an AMS 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.

SnapShot volumes include primary data volumes (P-VOLs) belonging to the same consistency group, secondary volumes referred to as virtual volumes (V-VOLs), data pool, a differential management logical unit (DMLU), and command device. These elements are explained in this chapter.

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 1-1](#) shows the SnapShot configuration.

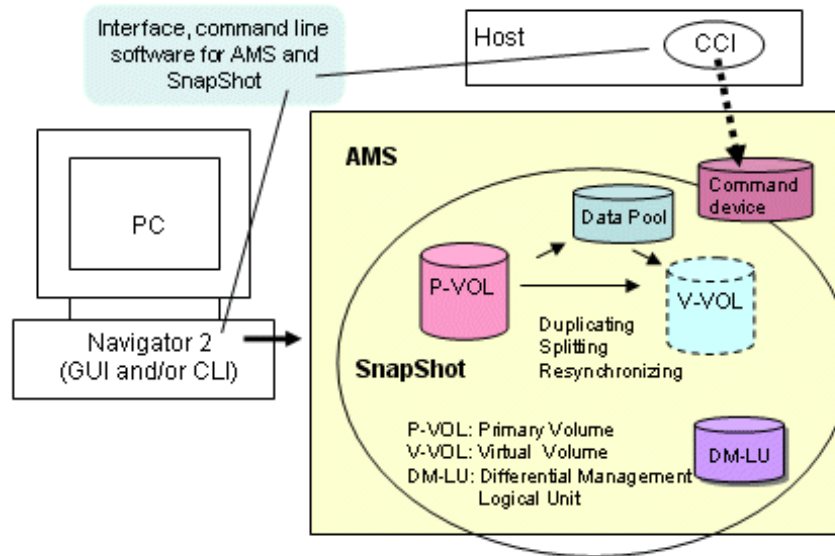


Figure 1-1: SnapShot functional component

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 32 virtual volumes (V-VOLs), which contain virtual copies of the P-VOL. The P-VOL and its V-VOL(s) are located in the same array.

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 data pool. V-VOL pointers in cache memory are updated to reference the original data's new location in the data pool. [Figure 1-2](#) illustrates SnapShot volumes and data pool interaction.

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 data pool.

V-VOL's are set up with LUs 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 array logic. Also, V-VOL pointers to data in the data pool and P-VOL actually reside in cache memory. Because of this, part of your array's cache is reserved for SnapShot when it is enabled. (See [Maximum supported capacity on page 2-22](#) and [Appendix D, Using SnapShot with Cache Partition Manager.](#))

Data pools

The data pool holds data from the P-VOL that is being replaced.

After a snapshot is taken, the V-VOL maintains pointers to P-VOL data. If changes occur, before the updated block is written to the P-VOL, the data that is being replaced is first copied to the data pool. The V-VOL pointer to this block is updated to the new address in the data pool. Thus, the V-VOL maintains the point-in-time image of the P-VOL, until the next snapshot is taken.

A data pool can be shared by multiple SnapShot pairs.

The data pool's function in the SnapShot process is illustrated in [Figure 1-2](#).

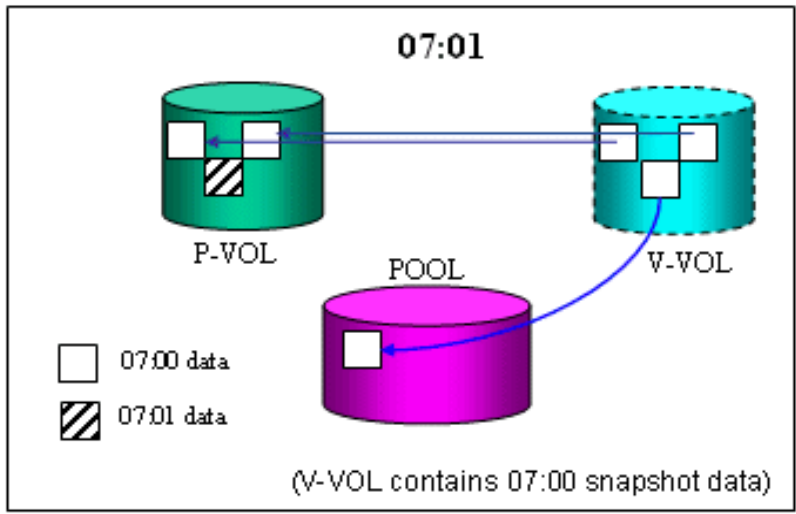
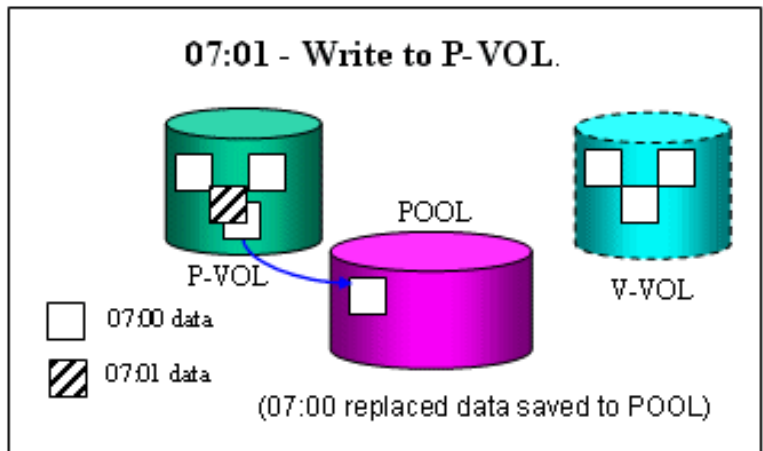
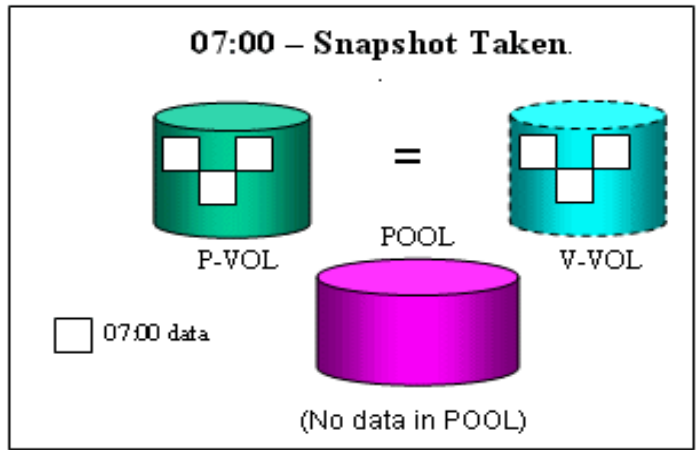


Figure 1-2: V-VOL Maintains Snapshot Data

Consistency group (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 “consistency group” (CTG), all primary logical volumes are treated as a single entity.

Managing SnapShot primary volumes as a consistency 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.

Differential Management LUs (DMLU)

The DMLU is an exclusive volume used for storing SnapShot information when the array system is powered down. The DMLU is treated the same as other volumes in the storage system, but is hidden from a host. The DMLU requires user setup. See [Set up the Differential Management LU \(DMLU\) on page 5-4](#) for details.

LU Ownership of P-VOLs and Data Pools

The load balancing function is not applied to the LUs specified as a SnapShot pair. Since the ownership of the LUs specified as a SnapShot pair is the same as the ownership of the LUs specified as a data pool, perform the setting so that the ownership of LUs specified as a data pool is balanced in advance. Refer to [Setting the LU ownership on page 5-6](#) or, for CLI, [Setting the LU ownership on page B-9](#) for more information.

Interfaces for performing SnapShot operations

SnapShot can be operated using of 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 arrays may want to continue using CLI, though the GUI is an option. The same recommendation applies to CCI users.

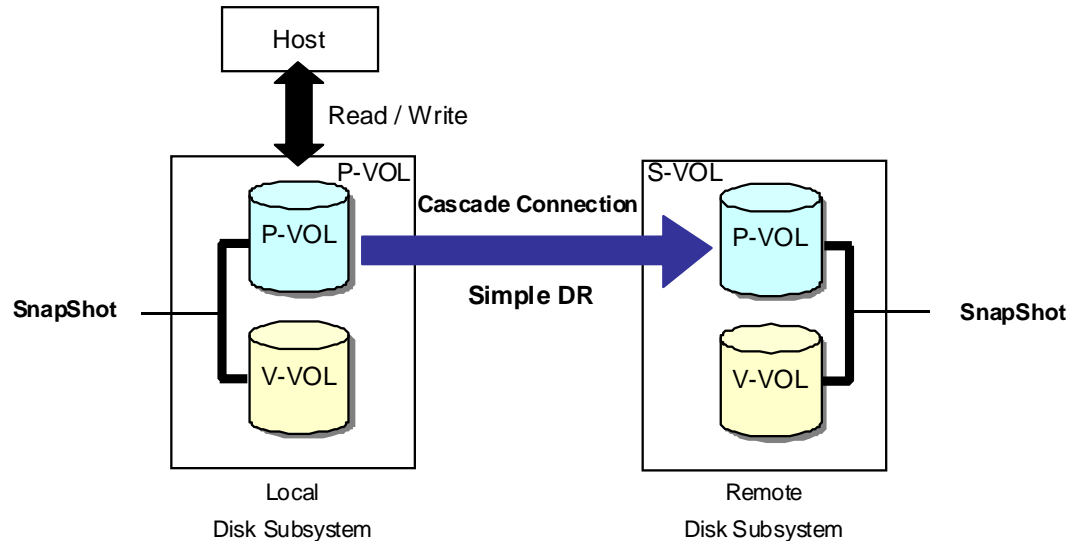


NOTE: Hitachi Replication Manager is used to manage and integrate Copy-on-Write. It provides a GUI topology view of the SnapShot system, with monitoring, scheduling, and alert functions. For more information on purchasing Replication Manager, visit the Hitachi Data Systems website.

Cascade connection of SnapShot with Simple DR

Volumes of SnapShot P-VOL can be cascaded with those of Simple DR, as shown in Figure 1-3. Cascading of SnapShot with Simple DR lowers performance, however.

- Cascade with a P-VOL of SnapShot



- Cascade with a V-VOL of SnapShot

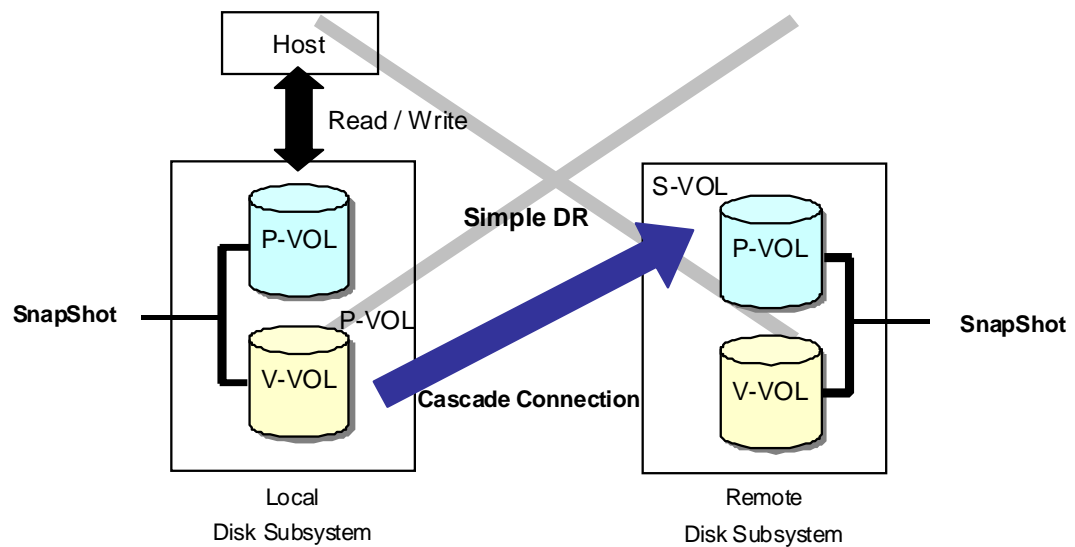


Figure 1-3: Cascade Connection of SnapShot with Simple DR

The Simple DR pair can be cascaded only with the SnapShot P-VOL with the following restrictions placed on the cascade connection:

- Restoration of the SnapShot pair cascaded with the Simple DR P-VOL can be done only when the status of the Simple DR pair is **Simplex**, **Split**, or **Pool Full**.

- When restoration of the SnapShot pair cascaded with the Simple DR S-VOL, it is required to make the status of the Simple DR pair **Simplex** or **Split**. The restoration can be done in the **Takeover** status, but it cannot be done when the status is **Busy** in which the S-VOL is being restored using the pool data.
- When the Simple DR S-VOL is in the **Busy** status in which it is being restored using the pool data, the Read/Write instruction cannot be issued to the SnapShot V-VOL cascaded with the Simple DR S-VOL.

Cascade connection of SnapShot with Simple DR P-VOL

LU Shared with P-VOL on SnapShot and P-VOL on Simple DR.

Table 1-1 shows whether a read/write from/to a P-VOL of SnapShot on the local side is possible when a P-VOL of SnapShot and a P-VOL of Simple DR are the same LU.

Table 1-1: A Read/Write Instruction to a SnapShot P-VOL on the Local Side (Simple DR)

Simple DR P-VOL	SnapShot P-VOL				
	Paired	Reverse Synchronizing	Split	Failure	Failure (Restore)
Paired	+, R/W	–	+, R/W	+, R/W	–
Synchronizing	+, R/W	–	+, R/W	+, R/W	–
Split	+, R/W	+, R/W	+, R/W	+, R/W	Δ , R/W
Pool Full	+, R/W	+, R/W	+, R/W	+, R/W	Δ , R/W
Failure	+, R/W	Δ , R/W	+, R/W	Δ , R/W	Δ , R/W

+ indicates a possible case, – indicates an impossible 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 possible.
 R: Read by a host is possible but write is impossible.
 W: Write by a host is possible but read is impossible.
~~R/W~~: Read/Write by a host is impossible.



NOTE: Failure in this table excludes a condition in which access of an LU is not allowed (for example, LU blockage).

Cascade connection of SnapShot with Simple DR S-VOL

One LU used for P-VOL on SnapShot and S-VOL on Simple DR

Table 1-2 shows whether a read/write from/to a P-VOL of SnapShot on the remote side is possible when a P-VOL of SnapShot and an S-VOL of Simple DR are the same LU.

Table 1-2: A Read/Write Instruction to a SnapShot P-VOL on the Remote Side (Simple DR)

Simple DR S-VOL		SnapShot P-VOL				
		Paired	Reverse Synchronizing	Split	Failure	Failure (Restore)
Paired		+, R	–	+, R	+, R	–
Synchronizing		+, R	–	+, R	+, R	–
Split	R/W	+, R/W	+, R/W	+, R/W	+, R/W	Δ , R/W
	R	+, R	–	+, R	+, R	–
Inconsistent		Δ , R/W	–	Δ , R/W	Δ , R/W	–
Takeover		+, R/W	+, R/W	+, R/W	+, R/W	Δ , R/W
Busy		Δ , R/W	–	+, R/W	+, R/W	–
Pool Full		+, R	–	+, R	Δ , R	–

+ indicates a possible case, – indicates an impossible 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 possible.

R: Read by a host is possible but write is impossible.

W: Write by a host is possible but read is impossible.

$\overline{R/W}$: Read/Write by a host is impossible.



NOTE: Failure in this table excludes a condition in which access of an LU is not allowed (for example, LU blockage).

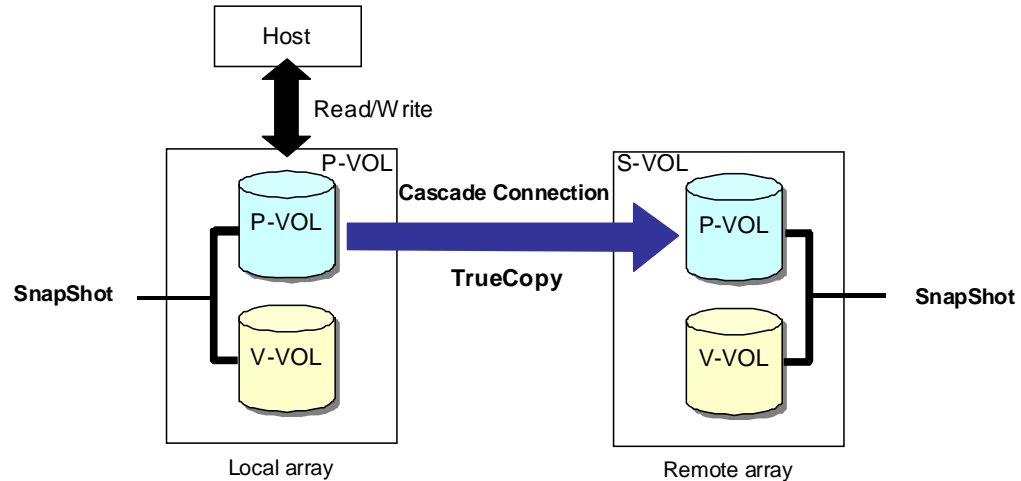
V-VOLs number of SnapShot

V-VOLs of up to 32 generations can be made even in the case where the P-VOL of SnapShot is cascaded with the P-VOL and S-VOL of Simple DR in the same way as in the case where no cascade connection is made.

Cascade connection of SnapShot with TrueCopy

Volumes of SnapShot can be cascaded with those of TrueCopy as shown in Figure 1-4. Because the cascade of SnapShot with TrueCopy lowers the performance, only use it when necessary. Incidentally, SnapShot cannot be cascaded with ShadowImage.

■ Cascade with a P-VOL of SnapShot



■ Cascade with a V-VOL of SnapShot

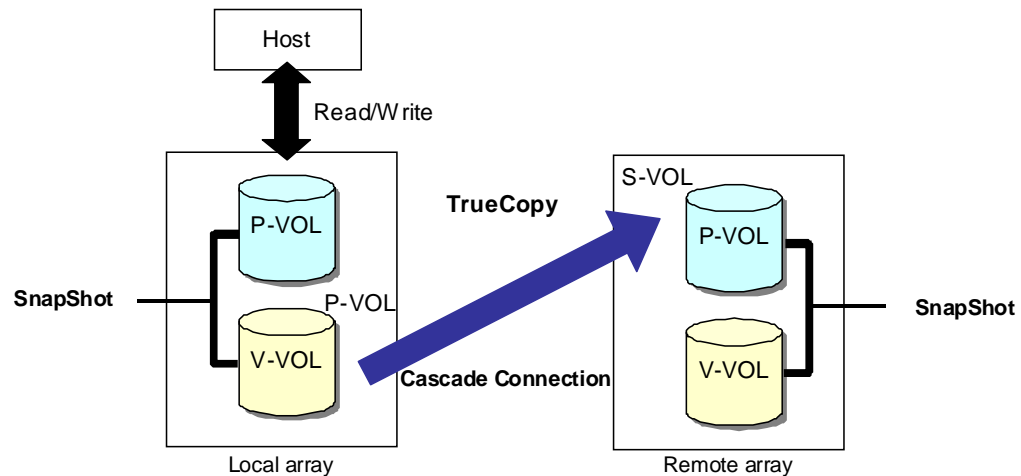


Figure 1-4: Cascade Connection of SnapShot with TrueCopy

Cascade restrictions with 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 LUs for P-VOL that are cascaded using TrueCopy on the local side and the remote side cannot be assured of equality.

LU shared with P-VOL on SnapShot and P-VOL on TrueCopy

Table 1-3 shows whether a read/write to/from 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 LU.

Table 1-3: A Read/Write Instruction to a SnapShot P-VOL on the Local Side (TrueCopy)

TrueCopy P-VOL		SnapShot P-VOL				
		Paired	Synchronizing (Restore)	Split	Failure	Failure (Restore)
Paired		+, R/W	–	+, R/W	+, R/W	–
Synchronizing		+, R/W	–	+, R/W	+, R/W	–
Split	R/W	+, R/W	+, R/W	+, R/W	+, R/W	Δ , R/W
	R	+, R	–	+, R	+, R	–
Failure	R/W	+, R/W	Δ , R/W	+, R/W	Δ , R/W	Δ , R/W
	R	+, R	–	+, R	Δ , R	–
	R/W	+, R/W	–	+, R/W	Δ , R/W	–

+ indicates a possible case, – indicates an impossible 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 possible.

R: Read by a host is possible but write is impossible.

W: Write by a host is possible but read is impossible.

~~R/W~~: Read/Write by a host is impossible.



NOTE: Failure in this table excludes a condition in which access of an LU is not allowed (for example, LU blockage).

One LU used for P-VOL on SnapShot and S-VOL on TrueCopy

Table 1-4 shows whether a read/write from/to a P-VOL of SnapShot on the remote side is possible or not in the case where a P-VOL of SnapShot and an S-VOL of TrueCopy are the same LU.

Table 1-4: A Read/Write Instruction to a SnapShot P-VOL on the Remote Side (TrueCopy)

TrueCopy S-VOL		SnapShot P-VOL				
		Paired	Synchronizing (Restore)	Split	Failure	Failure (Restore)
Paired		+, R	–	+, R	+, R	–
Synchronizing		+, R	–	+, R	+, R	–
Split	R/W	+, R/W	+, R/W	+, R/W	+, R/W	Δ, R/W
	R	+, R	–	+, R	+, R	–
Failure		+, R	–	+, R	+, R	–

+ indicates a possible case, – indicates an impossible 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 possible.
 R: Read by a host is possible but write is impossible.
 W: Write by a host is possible but read is impossible.
 R/W: Read/Write by a host is impossible.



NOTE: Failure in this table excludes a condition in which access of an LU is not allowed (for example, LU blockage).

V-VOLs number of SnapShot

V-VOLs of up to 32 generations can be made even in the case where the SnapShot P-VOL is cascaded with the P-VOL and S-VOL of TrueCopy in the same way as in the case where no cascade connection is made.

Cascade restrictions with SnapShot V-VOL

The following explains the transition of statuses of TrueCopy and SnapShot pairs

About cascading of an LU of TrueCopy with a SnapShot V-VOL, it is supported only when the SnapShot V-VOL and a TrueCopy P-VOL are the same LU. Besides, 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 1-5 shows whether a read/write to/from 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 LU.

Table 1-5: A Read/Write Instruction to a SnapShot V-VOL on the Local Side (TrueCopy)

TrueCopy P-VOL		SnapShot V-VOL				
		Paired	Synchronizing (Restore)	Split	Failure	Failure (Restore)
Paired		–	–	+, R	–	–
Synchronizing		–	–	+, R	–	–
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

+ indicates a possible case, – indicates an impossible 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 possible.

R: Read by a host is possible but write is impossible.

W: Write by a host is possible but read is impossible.

~~R/W~~: Read/Write by a host is impossible.



NOTE: Failure in this table excludes a condition in which access of an LU is not allowed (for example, LU blockage).

Configuration restrictions on the Cascade of TrueCopy with SnapShot

Figure 1-5 shows an example of a configuration in which restrictions are placed on the cascade of TrueCopy with SnapShot.

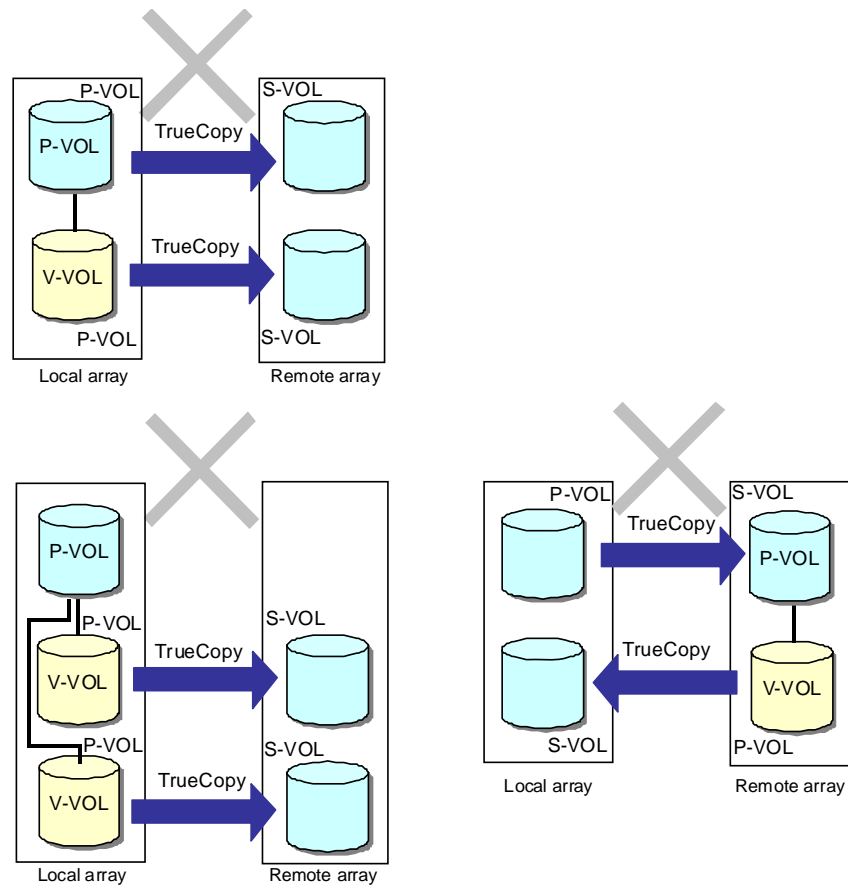


Figure 1-5: Configuration restrictions on the Cascade of TrueCopy with SnapShot

Cascade restrictions with SnapShot Data Pool

Neither TrueCopy/Simple DR pair nor ShadowImage pair can be created using a data pool.

Planning and design

A snapshot ensures that volumes with bad or missing data can be restored. With Copy-on-Write 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.

- ❑ [The plan and design workflow](#)
- ❑ [Assessing business needs](#)
- ❑ [Establishing data pool size](#)
- ❑ [Requirements and recommendations for SnapShot Logical Units](#)
- ❑ [Operating system considerations](#)
- ❑ [Cascading SnapShot with TrueCopy](#)
- ❑ [Cascading SnapShot with TrueCopy Extended Distance](#)
- ❑ [Maximum supported capacity](#)

The plan and design workflow

The SnapShot planning effort consists of determining the number of V-VOLs required by your organization, the V-VOL(s)' lifespan, that is, how long they must be held before being updated again, the frequency that snapshots are taken, and the size of the data 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.
- Size the data pool. (For a description of the data pool, see [Data pools on page 1-4](#).)

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 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.

Establishing how often a copy is made (copy frequency)

The frequency that copies need to be made is determined by how much data can be lost in a disaster before business is significantly impacted.

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 data pool size. The more time that elapses between snapshots, the more data accumulates in the data pool. Copy frequency may need to be modified to reduce the data 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 data pool. Longer periods between backups require more space in the data 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 that Are required

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 2-1](#).

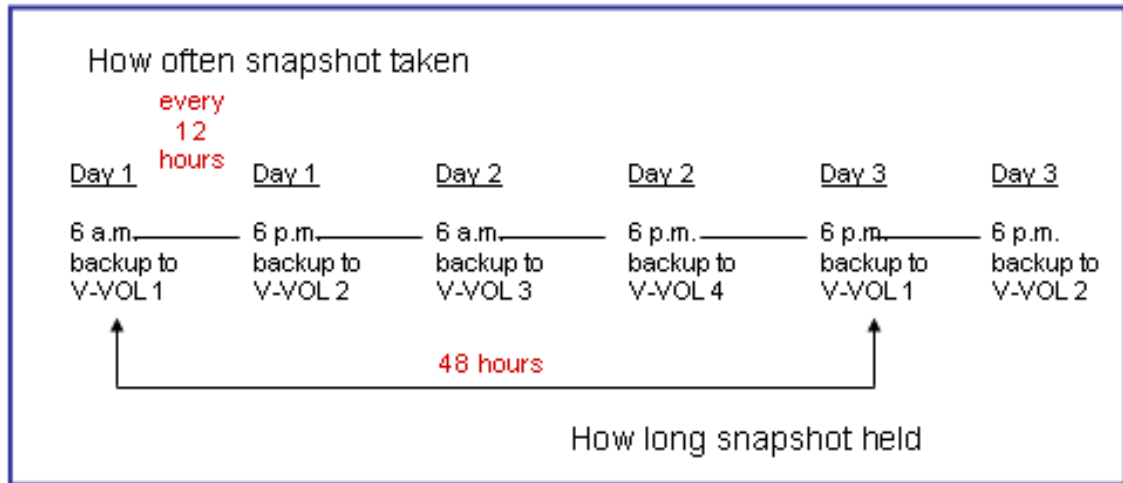


Figure 2-1: V-VOL Frequency, Lifespan

Establishing data pool size

The data pool holds data from the P-VOL that is being replaced. By holding this original data in the data pool, the mirror image of the P-VOL at the time of the snapshot is maintained.

You determine the size of the data pool that your system requires by:

- Measuring the amount of write workload that passes from the host application to the P-VOL. Write workload is the megabytes per second that are written to the primary volume over a specific time.
- Calculating the amount of data that would accumulate during the lifespan of your V-VOL.
- Multiplying the amount of data that accumulates by the number of V-VOLs.

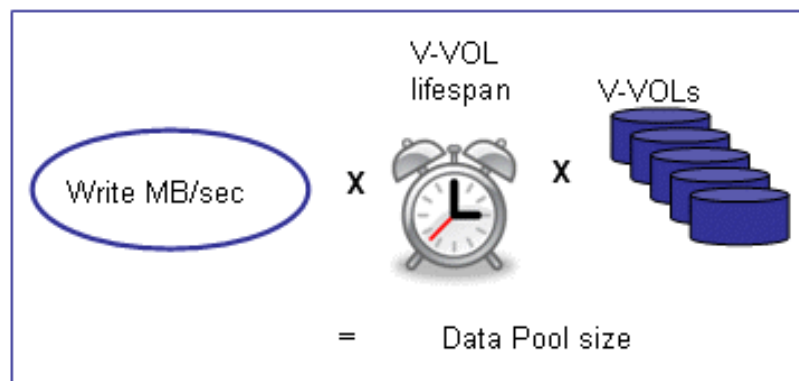


Figure 2-2: Write workload and data pool size

Measuring workload data, sizing the data pool

To set up SnapShot, you must measure the amount of data that changes in your production system. The amount of data written to the primary volume indicates how large the data pool must be.

Workload data is collected using performance monitoring software on your operating system—preferably during the busiest time of month, quarter, and year. The goal is to collect data that shows your system's actual workloads during high peaks and spikes, when more is changing and the demands on the system are greatest.

To collect workload data and size the data pool

1. Using your operating system's performance monitoring software, collect the following:
 - Disk-write bytes/second for every physical volume that will be replicated.
 - Collect this data at 10 minute intervals.
 - Collect this data over a 4-6 week period that includes high peaks and spikes, and when the demands on the system are greatest.
2. At the end of the period, convert the data to MB/second and import into a spreadsheet tool. [Figure 2-3](#) shows collected raw data, in megabytes per second in 10 minute segments.

	A	B
1	Time	Raw Data
2	0:00	16.34
3	0:10	28.19
4	0:20	15.86
5	0:30	20.58
6	0:40	4.57
7	0:50	11.36
8	1:00	20.04
9	1:10	35.69
10	1:20	26.44
11	1:30	23.97
12	1:40	20.69

Figure 2-3: Raw Data Example in MB/sec

3. Using the copy frequency established earlier, calculate averages over the collection period. Most spreadsheet tools have an average function. For example:

If copy frequency is 1 hour, then calculate 60 minute rolling averages using the values in 6 10-minute intervals.

If copy frequency is 4 hours, then calculate 240 minute rolling averages using the values in 24 10-minute intervals.

Figure 2-4 illustrates 60-minute rolling averages.

	A	B	C
1	Time	Raw Data	60 Min Rolling Avg
2	0:00	16.34	
3	0:10	28.19	
4	0:20	15.86	16.15
5	0:30	20.58	16.77
6	0:40	4.57	18.02
7	0:50	11.36	19.78
8	1:00	20.04	20.35
9	1:10	35.69	23.03
10	1:20	26.44	24.28
11	1:30	23.97	25.11
12	1:40	20.69	23.41

Figure 2-4: Rolling Averages Calculated Using V-VOL Frequency

Example rolling-average procedure using Microsoft Excel:

- a. In cell C4, type =Average(b2:b7).
- b. Press Enter.
This instructs the tool to calculate the average value in cells B2 through B7 and populates C4 with that data.
- c. Copy the value in C4.
- d. Highlight cells C5 to the last C cell in the last row of workload data in the spreadsheet.
- e. Right-click the highlighted cells and select the paste option.
Excel maintains the logic and increments the formula values initially entered in C4. It then calculates all the point in time averages and populates the C cells.

Figure 2-5 illustrates rolling averages graphed over raw measurement data averages.

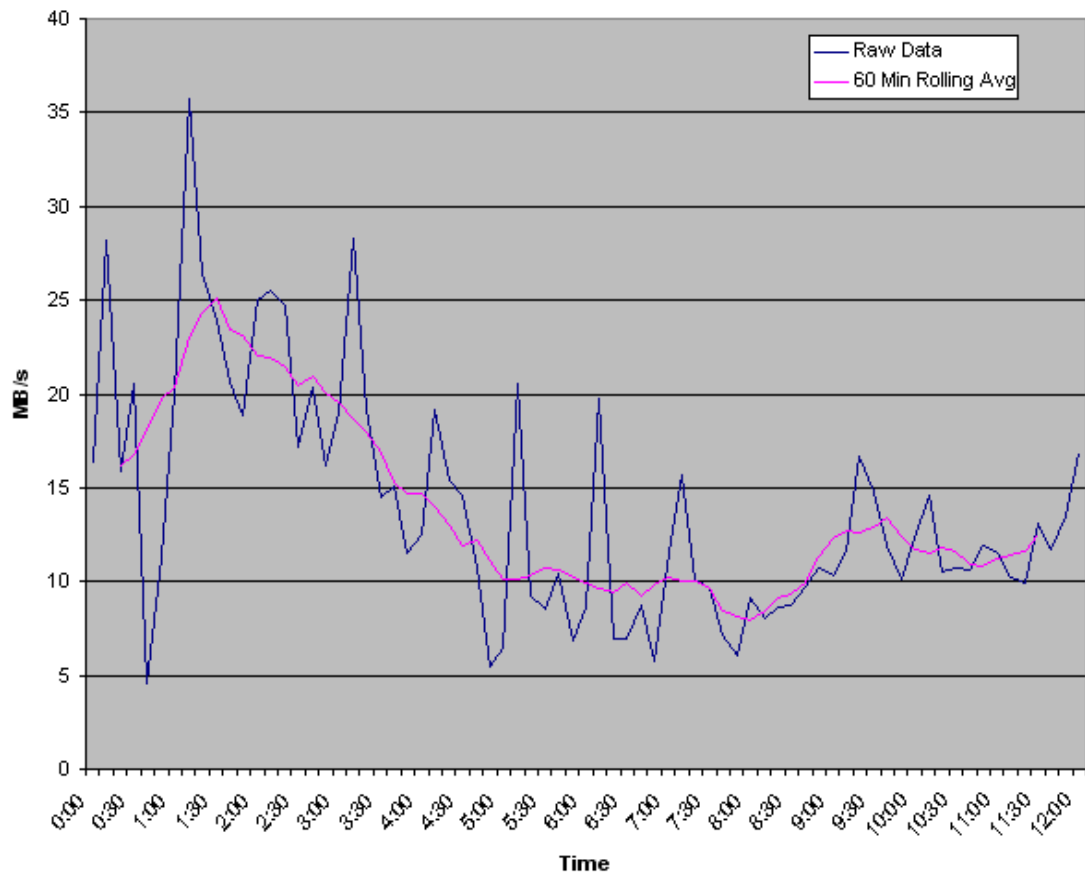


Figure 2-5: Rolling Averages Graphed Over Raw Averages

4. Locate the maximum rolling average (RA) value in the C column. Using this peak value and the following formula, calculate the *cumulative peak data change* over the lifespan of a copy (V-VOL):

$$(\text{RA peak MB/sec}) \times (\text{V-VOL lifespan seconds}) = (\text{Cumulative data over V-VOL lifespan})$$

For example, if the RA peak is 25 MB/sec, and the V-VOL lifespan is 3600 seconds (1 hour), then:

$$25\text{MB/sec} \times 3600 \text{ seconds} = 90,000 \text{ MB}$$

Cumulative data over a V-VOL's 1-hour lifespan is 90,000 MB.

5. Calculate the *base data pool* size for your primary/virtual volumes by multiplying the MB size of one V-VOL in Step 4 by the number of V-VOLs, which was established earlier. For example:

$$90,000 \text{ MB} \times 4 \text{ V-VOLS} = 360,000 \text{ MB}$$

This is the base data pool size for a SnapShot system in which the copy frequency is 1 hour, the copy lifespan is 4 hours, and the number of copies (V-VOLs) is 4.

6. It is highly recommended that a safety factor of 20%, be calculated. Do so using the following formula:

$$(\text{Base data pool size}) \times 1.2. \text{ For example:}$$

$$360,000 \text{ MB} \times 1.2 = 432,000 \text{ MB}$$

- It is also advisable to factor in annual increases in data transactions. Do this by multiplying the base pool size by the percentage of expected annual growth. For example:

$$432,000 \text{ MB} \times 1.2 \text{ (20 percent growth rate for per year)} \\ = 518,400 \text{ MB}$$

This is the size of the data pool with growth factored for the first year.

- Repeat this step for each year the solution will be in place. For example:

$$518,400 \text{ MB} \times 1.2 \text{ (20 percent growth rate for second year)} \\ = 622,080 \text{ MB}$$

This is the size of the data pool with growth factored for the second year.

Rule-of-thumb calculation

When write-workload has not been measured, Hitachi suggests the change rates shown in [Table 2-1](#).

Table 2-1: Workload Rates when No Measurement

Snapshot lifespan	Suggested write workload change rate
1-4 hours	10%
4-8 hours	15%
8-12 hours	20%
12-24 hours	25%

To calculate data pool size using rule-of-thumb change rate

- Use the following formula:

$$\text{Data Pool size} = (\text{P-VOL} \times \% \text{ of changed data} \times 2.5 \text{ safety rate}) \times \text{number of V-VOLs}$$

For example, if the P-VOL = 1 TB and one snapshot is taken per 24 hours, then:

$$1 \text{ TB} \times 25\% \text{ change rate} \times 25\% \text{ safety rate of} = 625 \text{ GB}$$

- Multiply the data pool size by the number of V-VOLs. Thus:

$$4 \text{ V-VOLs} \times 625 \text{ GB} = 2500 \text{ GB (2.5 TB)}$$

Data pool key points

- The data pool must be on the same controller as the P-VOL and V-VOL(s).
- Data pool capacity should be at least 20 GB.

- Up to 64 volumes can be assigned to a data pool.
- When a volume is assigned to a data pool, it is no longer recognized by a host.

Requirements and recommendations for SnapShot Logical Units

Please review the following key rules and recommendations regarding P-VOLS, V-VOLS, and data pools. See [Appendix A, Specifications](#), for general specifications required for SnapShot.

- Primary and secondary volumes must be set up prior to making SnapShot copies.
- The P-VOL and S-VOL must be assigned to the same controller.
- Assign four or more disks to SnapShot LUs 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.

RAID configuration for LUs assigned to SnapShot

Please observe the following regarding RAID levels when setting up SnapShot pair volumes, data pools, command devices, and Differential Management LUs.

- More than one pair may exist on the same RAID group on the 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 LUs in the RAID group
- Performance is best when P-VOL and data pool are assigned to a RAID group consisting of SAS or SSD drives. Performance decreases when located on SATA drives.
- Multiple data pool volumes can exist on a single array, but the RAID group for each data pool volume should be restricted only to that POOL, otherwise performance is impacted.
- Locate a P-VOL and associated data pool in different ECC groups within a RAID group, as shown in [Figure 2-6](#). When they are in the same ECC group, performance decreases and the chance of failure increase.

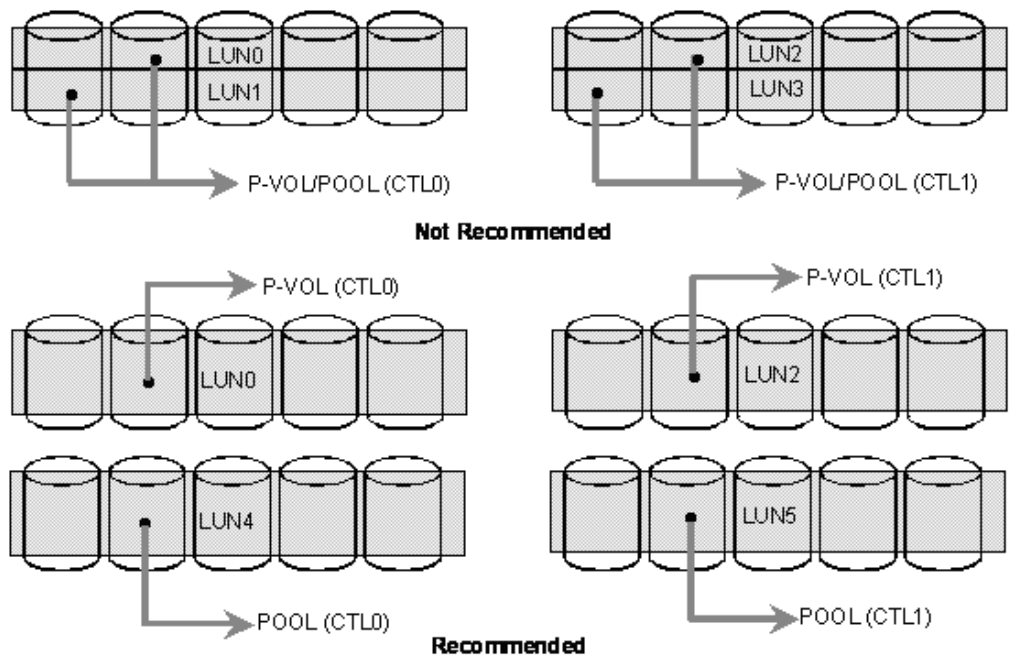


Figure 2-6: Locating P-VOL, Data Pool in Separate ECC Groups

- The P-VOL and data pool require a RAID level with redundancy. RAID 0 is not supported.
- The P-VOL and data pool may exist on different RAID levels and/or number of drives in RAID. Performance improves, however, when the RAID levels and number of drives are the same. [Table 2-2](#) shows RAID configurations for the P-VOL and data pool.

Table 2-2: P-VOL and Data Pool RAID Configuration

P-VOL	Data Pool	Amount of User Data	Total Amount of SnapShot	Share of User Data
RAID 1+0 (N = 1 to 8)	RAID 1+0 (N = 1 to 8)	1	4	1/4
RAID 1+0 (N = 1 to 8)	RAID 5 (N = 4) (see Note 1)	1	2+1.25 = 3.25	1/3.25
RAID 5 (N = 4) (see Note 1)	RAID 1+0 (N = 1 to 8)	1	1.25+2 = 3.25	1/3.25
RAID 5 (When N = 4) (see Note 1)	RAID 5 (When N = 4) (see Note 1)	1	1.25+1.25 = 2.5	1/2.5
RAID 5 (When N = 8) (see Note 1)	RAID 5 (When N = 8) (see Note 1)	1	1.125+1.125 = 2.25	1/2.25
RAID 6 (When N = 4) (see Note 2)	RAID 6 (When N = 4) (see Note 2)	1	1.5+1.5 = 3	1/3
RAID 6 (When N = 8) (see Note 2)	RAID 6 (When N = 8) (see Note 2)	1	1.25+1.25 = 2.5	1/2.5
RAID 6 (When N = 4) (see Note 2)	RAID 5 (When N = 4) (see Note 1)	1	1.5+1.25 = 2.75	1/2.75

Note 1: Capacity = $(1+1/N)$ where N = Number of data drives in RAID

Note 2: Capacity = $(1+2/N)$ where N = Number of data drives in RAID

- RAID 5 (4D+1)/RAID 5 (4D+1) is the recommended configuration because 4D+1P is a recommended configuration for performance. It is also a balanced ratio of redundancy and user data related to RAID 5.
- When two or more **command devices** are located in the same AMS array, assign them to respective RAID groups. This safeguards the SnapShot system in the event of system malfunction, in which case both command devices would become unavailable.
- When two **Differential Management LUs** are located in the same AMS array, assign them to respective RAID groups. This safeguards the SnapShot system in the event of system malfunction, in which case both DM-LUs would become unavailable.

Operating system considerations

The following sections provide recommendations and restrictions for SnapShot volumes.

Identifying P-VOL and V-VOL LUs on Windows

In the Navigator 2 GUI, the P-VOL and S-VOL are identified by their LU number. In Windows, LUs are identified by their HLUN. To understand the mapping of a LUN on Windows, proceed as follows: These instructions provide procedures for the iSCSI and fibre channel interfaces.

1. Identify the HLUN of your Windows disk.
 - a. From the Windows Server 2003 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-LUN Mapping for the iSCSI interface as follows. (If using fibre channel, skip to Step 3.)
 - a. In the Navigator 2 GUI, select the desired array.
 - b. In the array tree, click the **Group** icon, then click **Host Groups**.
 - c. Click the host group that the volume is mapped to.
 - d. On the screen for the host group, click the **Edit Host Group** button. The volumes mapped to the Host Group display. You can confirm the LUN is mapped to the H-LUN.
3. Identify HLUN-to-LUN Mapping for the Fibre Channel interface, as follows:
 - a. In the Navigator 2 GUI, select the desired array.
 - b. In the array tree that displays, click the **Groups** icon, then click the **Host Groups** icon in the Groups tree.
 - c. On the Host Groups screen, select a Host group.
 - d. On the host group screen, select the **Logical Units** tab. Find the identified H-LUN. The LUN displays in the next column.
 - e. If the HLUN is not present on a host group target screen, on the Host Groups screen, select another Host group and repeat Step d.

LU mapping

To perform pair operation using CCI, the P-VOL and V-VOL must be configured for the port specified in the configuration definition file. Doing this makes them recognized by a host. Otherwise they cannot be paired. Use the LUN Manager if you do not want them recognized by a host

Cluster and path switching software

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

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.
- Shutdown MSCS before executing the CCI `sync` command.
- The command device cannot be shared between different hosts in the cluster.
- Assign the exclusive command device to each host.

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.

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.

Windows 2000

- A P-VOL and S-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 Adaptable Modular Storage Command Control Interface (CCI) Reference Guide*.

Windows Server 2003/2008 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 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 Simple Modular Storage Command Control Interface (CCI) Reference 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 Simple Modular Storage Command Control Interface (CCI) Reference Guide* for details of the restrictions of Windows Server 2008 when using the `mount/umount` command.
- In Windows Server™ 2008, set only the P-VOL to be recognized by the host and let another host recognize the V-VOL. When you have created two or more V-VOLs for one P-VOL, do not make the same host recognize two or more V-VOLs (shared P-VOLs) at the same time.
- Using CCI to operate pairs. When describing a command device in the configuration definition file, specify it as `Volume{GUID}`. (For more detail, see the *Hitachi Simple Modular Storage Command Control Interface (CCI) Reference Guide*.)
- 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 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.
- When mounting volumes, use either the CCI `mount` command or Windows `mountvol /P` command. If using Windows' `mountvol`,

the command specification should be understood and sufficient testing performed before including it in your operation. When using the CCI `mount` command, you must use the CCI `umount` to unmount. For more information on CCI commands, see the *Hitachi Adaptable Modular Storage Command Control Interface (CCI) Reference Guide*.

- When using CCI, if 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 recover. Restart CCI if Windows cannot access the command device, even if CCI is able to recognize Windows.
- When using CCI, specify a command device in the configuration definition file as `Volume{GUID}`. For more information, see the *Hitachi Adaptable Modular Storage Command Control Interface (CCI) Reference Guide*.

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.

Concurrent use of Cache Partition Manager

When SnapShot is used together with Cache Partition Manager, please refer to the *Hitachi Storage Navigator Modular 2 Storage Features Reference Guide*.

Concurrent use of Dynamic Provisioning

Observe the following items:

- Dynamic Provisioning and SnapShot cannot be unlocked at the same time in an AMS2500 where a cache memory of 2 GB/CTL is installed. To unlock Dynamic Provisioning and SnapShot at the same time, add cache memory.
- The data pool used by SnapShot and TCE cannot be used as a Dynamic Provisioning DP pool. Moreover, the DP pool used by Dynamic Provisioning cannot be used as data pools of SnapShot and TCE.
- The DP-VOLs created by Dynamic Provisioning cannot be set for a DMLU and a command device.



CAUTION! The DP-VOLs created by Dynamic Provisioning cannot be set for a SnapShot P-VOL. Moreover, the DP-VOLs cannot be added to the data pool used by SnapShot and TCE.

Windows Server 2003 and 2008/Windows 2000 and Dynamic Disk

Observe the following when using Windows Server 2003 dynamic disk:

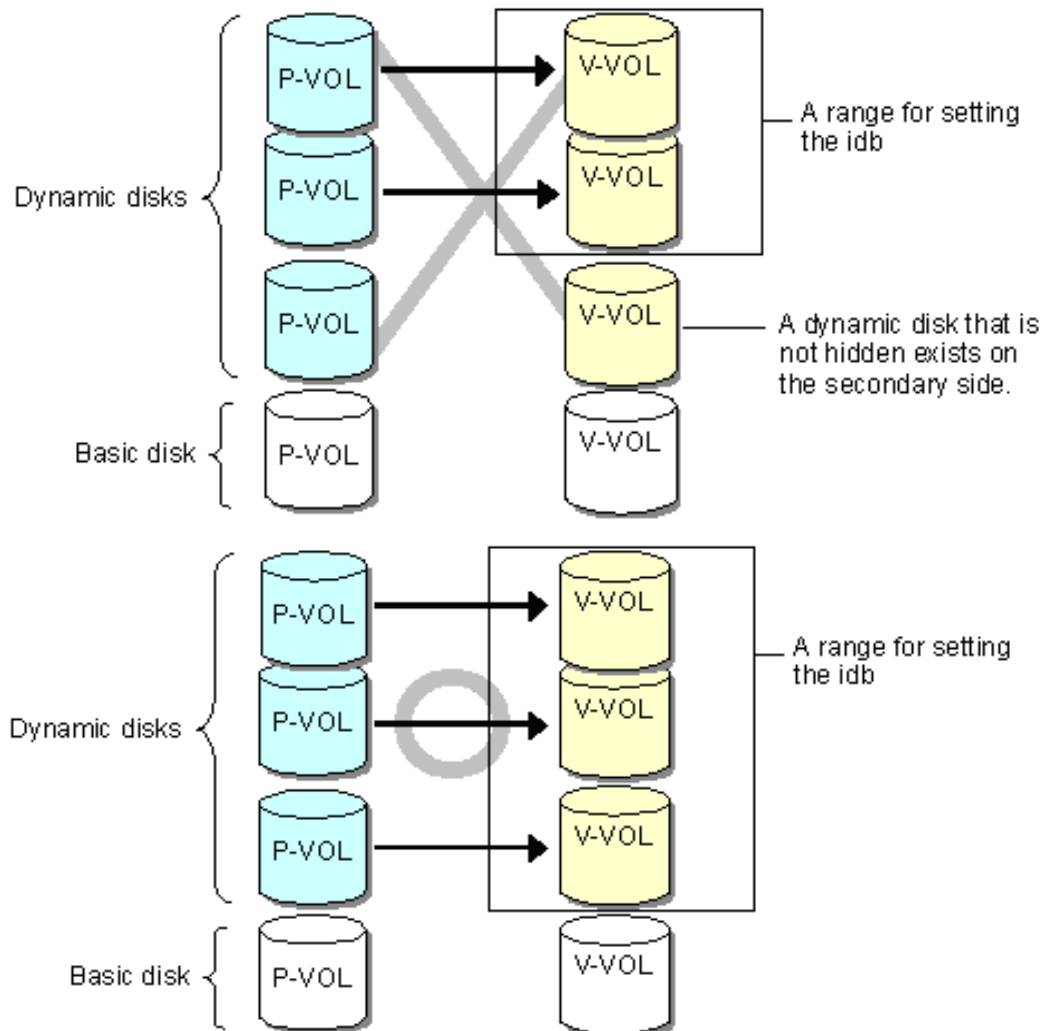
- A P-VOL and S-VOL *can* be made into a dynamic disk on Windows Server 2003 but cannot be made on Windows 2000 and Windows Server 2008.
- You cannot make a P-VOL and a V-VOL into a dynamic disk; however you can use a P-VOL and a V-VOL as a dynamic disk.
- When using a V-VOL with a secondary host, ensure that the pair status is Split.
- 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.
- An LU in which two or more dynamic disk volumes co-exist cannot be copied.
- On the secondary host side, use a dynamic disk function for V-VOLs only.

When copying, hide all the dynamic disks that exist on the primary side using the `raidvchkset -vg idb` command. No restriction is placed on the primary side. Hide all the dynamic disk volumes to be restored on the primary side at the time of restoration.

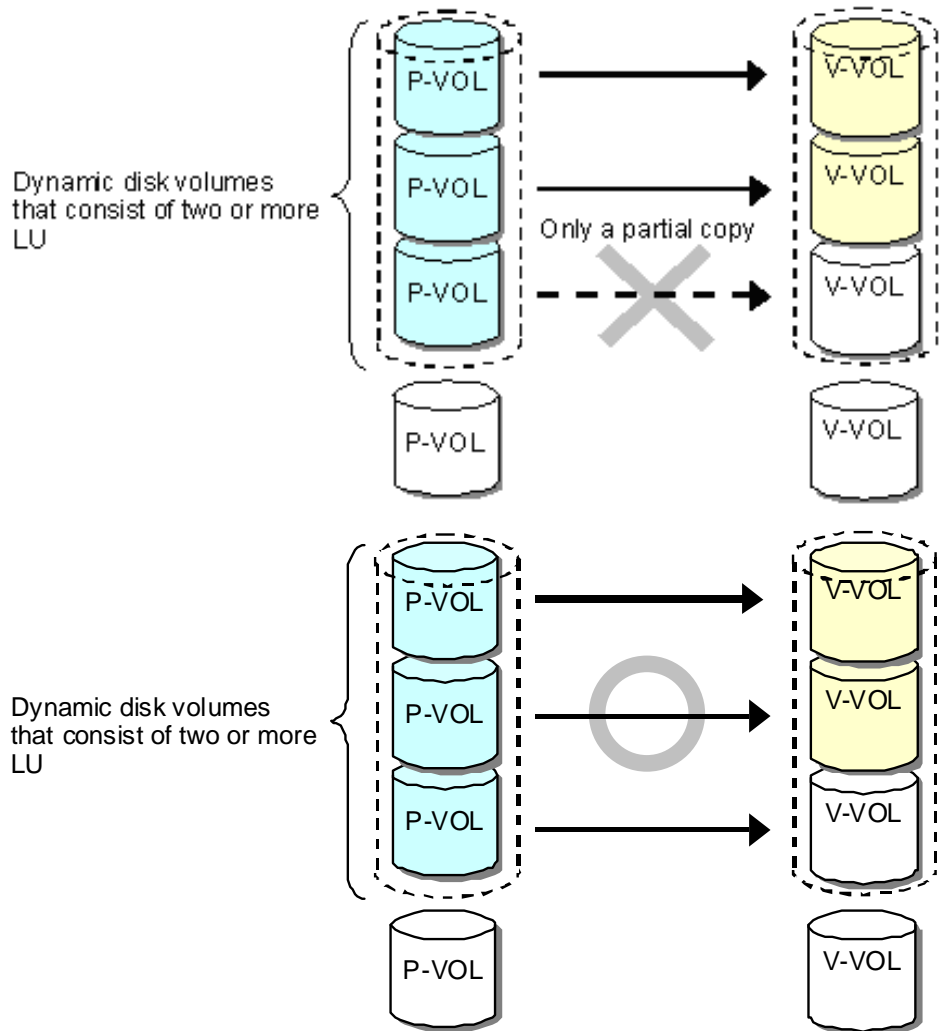


NOTE: A dynamic disk is a Windows Server 2008/Windows Server 2003/Windows 2000 feature.

If any one of the dynamic disks is left unhidden, a **Missing** drive occurs. When this occurs, delete it manually using the `diskpart delete` command.



- Copy dynamic disk volumes that consist of two or more LUs only after hiding all LUs from a host. When the copy is completed, you can have them recognized by a host.

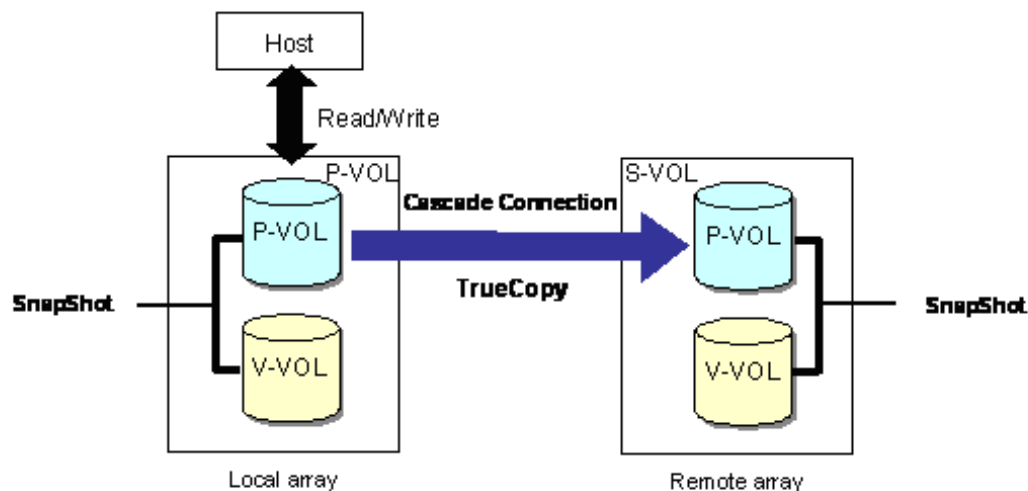


- A dynamic disk cannot be used with a cluster (MSCS, VCS, etc.) nor with VxVM and HDLM.

Cascading SnapShot with TrueCopy

Both the SnapShot P-VOL and V-VOL can be cascaded (shared) with the TrueCopy P-VOL or S-VOL, as shown in Figure 2-7.

- Cascade with a P-VOL of SnapShot



- Cascade with a V-VOL of SnapShot

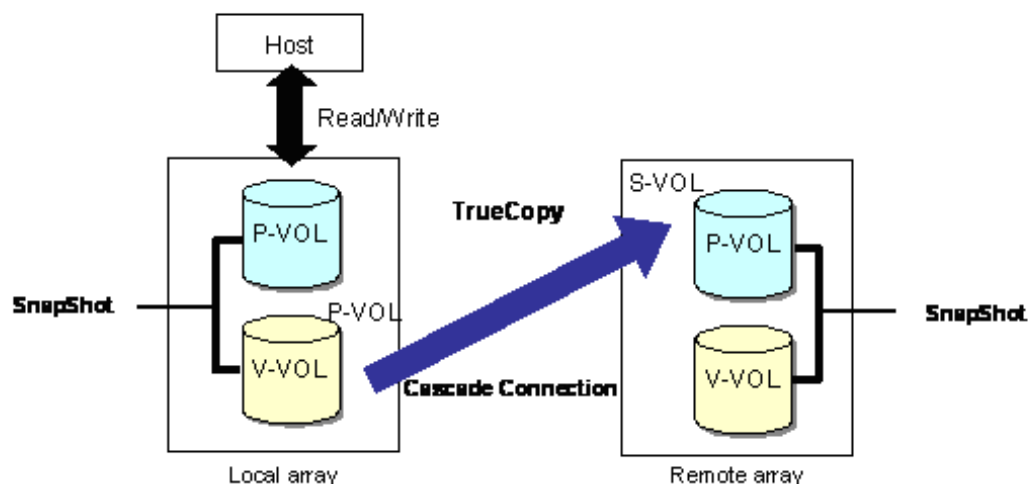


Figure 2-7: Basic SnapShot Cascade Examples with TrueCopy

Read/write and pair operations are limited when a SnapShot pair is cascaded with TrueCopy. Please see the appendix on cascading in the *Hitachi Adaptable Modular Storage TrueCopy Remote Replication User Guide* (MK-DF978052) for full information.

Cascading SnapShot with TrueCopy Extended Distance

The SnapShot P-VOL can be cascaded (shared) with both the primary and secondary TrueCopy Extended Distance (TCE) volumes, as shown in Figure 2-8. A V-VOL cannot be cascaded.

- Cascade with a P-VOL of SnapShot

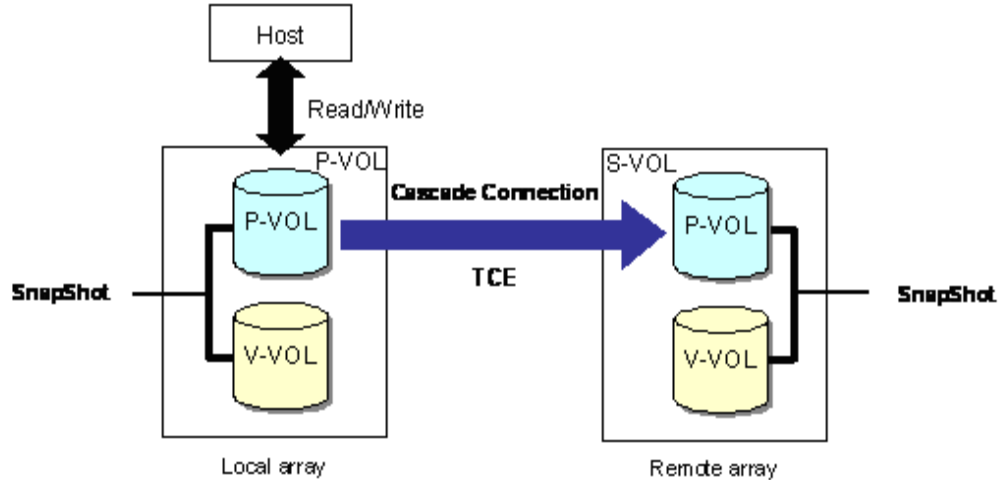


Figure 2-8: Basic SnapShot Cascade Examples with TrueCopy Extended

Read/write and pair operations are limited when a SnapShot pair is cascaded with TCE. Please see the appendix on cascading in the *Hitachi Adaptable Modular Storage TrueCopy Extended Distance User Guide* (MK-DF978054) for full information.

Maximum supported capacity

- When SnapShot and ShadowImage are used on the same array

When ShadowImage is used in addition to SnapShot

The capacity you can assign to SnapShot replication volumes per controller is limited, for the following reasons:

- The maximum capacity supported by a SnapShot pair depends on the ratio of P-VOL to data pool and cache memory.
- When using other copy systems and SnapShot together, the maximum supported capacity of the P-VOL may be restricted further.

In addition to this, capacity is managed by the AMS array in blocks of 15.75 KB for data volumes and 3.2 KB for data pools. For example, when a P-VOL block's actual size is 16 KB, the array manages it as two blocks of 15.75 KB, or 31.5 KB. Data pool capacity is managed in the same way but at 3.2 KB per block.



NOTE: In a dual-controller array, the calculations must be performed for both controllers.

This section provides formulas for calculating your existing or planned SnapShot volume capacity and comparing it to the maximum supported capacity for your particular controller and its cache memory size.

SnapShot capacity must be calculated for both of the following:

1. The ratio of SnapShot and TCE (if used) capacity to data pool capacity. Capacity is calculated using the following volumes:
 - SnapShot P-VOLs
 - TCE P-VOLs and S-VOLs (if used)
 - All data pools
2. Concurrent use of TCE and ShadowImage. If SnapShot is used concurrently also, it is included in this calculation. Capacity is calculated using the following volumes, if used:
 - SnapShot P-VOLs
 - TCE P-VOLs and S-VOLs
 - TrueCopy P-VOLs and S-VOLs
 - ShadowImage S-VOLs



NOTE: A portion of cache memory is assigned to SnapShot for internal operations. For more information, please see [Appendix D, Using SnapShot with Cache Partition Manager](#).

SnapShot and TCE capacity

When TCE is used, SnapShot and TCE share common data pool resources. Therefore, you calculate their managed capacities together. If TCE is not used, you only calculate SnapShot's capacity.

All formulas, tables, graphs and examples pertain to one controller. On dual controller arrays, you must perform calculations for both controllers.

Managed capacity is calculated here, per controller, using the following formula:

$$\begin{aligned} & \text{Size all SnapShot P-VOLs} + \\ & \text{Size of all TCE P-VOLs (if used)} + \\ & \text{Size of all TCE S-VOLs (if used)} + \\ & / 5 \\ & + \text{size of all data pool volumes} \\ & < \text{Maximum Supported Capacity} \end{aligned}$$

Maximum supported capacity is shown in [Table 2-3](#).

Table 2-3: Maximum Supported Capacities, Cache Size per Controller

Cache Memory per Controller	(SnapShot P-VOLs, TCE P-VOLs, S-VOLs, Data Pools)		
	AMS2100	AMS2300	AMS2500
2 GB per CTL	1.4 TB	Not supported	Not supported
4 GB per CTL	6.2 TB	6.2 TB	Not supported
8 GB per CTL	Not supported	12.0 TB	12.0 TB
16 GB per CTL	Not supported	Not supported	24.0 TB

Example:

To calculate capacity when only SnapShot is used

In this example, the array and cache memory per controller is AMS 2300/4 GB.

1. List the size of each P-VOL in the Storage system. For example:

P-VOL 1 = 100 GB

P-VOL 2 = 50 GB

2. Calculate managed P-VOL capacity, using the formula:

ROUNDUP (P-VOL capacity / 15.75) * 15.75

For example:

P-VOL1: ROUNDUP (100 / 15.75) = 7

7 * 15.75 = 110.25 GB, the managed P-VOL Capacity

P-VOL2: ROUNDUP (50 / 15.75) = 4

4 * 15.75 = 63 GB, the managed P-VOL Capacity

3. List all data pools and their sizes. For example:

Data pool 1 = 100 GB

Data pool 2 = 60 GB

4. Calculate managed data pool capacity, using the formula:

ROUNDUP (data pool capacity / 3.2) * 3.2

For example:

Data pool 1: ROUNDUP (100 / 3.2 = 32)

32 * 3.2 = 102.4 GB, the managed data pool capacity

Data pool 2: ROUNDUP (60 / 3.2 = 19)

19 * 3.2 = 60.8 GB, managed data pool capacity

5. Calculate maximum capacity using the following equation:

**(Total P-VOL capacity) / 5 + (Total Data Pool capacity) <
= 800 GB**

For example:

Total PVOL size = 173.25 GB

Total data pool size = 163.2 GB

Thus:

173.25 GB / 5 = 34.65 GB

34.65 GB + 163.2 GB = 197.85 GB

In this example, the SnapShot maximum capacity is 197.85 GB, well below the maximum supported capacity shown in [Table 2-3](#).

To calculate capacity when SnapShot and TCE are used

1. Add the total managed capacities for SnapShot P-VOLs and TCE P-VOLs and S-VOLs. For example:

Total SnapShot P-VOL managed capacity = 173 GB

Total TCE P-VOL, S-VOL managed capacity = 221 GB

173 GB + 221 GB = 394 GB

2. Divide the sum by 5. For example:

394 GB / 5 = 77 GB

3. Add the quotient to total data pool capacity. The following uses the example data pool capacity from above:

77 GB + 61 = 138 GB

Managed capacity is well below maximum supported capacity.

[Table 2-5 on page 2-27](#) through [Table 2-8 on page 2-28](#) show how closely capacity between data volumes and data pool volumes must be managed. These tables are provided for your information. Also, [Figure 2-9 on page 2-28](#) shows a graph of how the ratio of data volume to data pool volume relates to maximum supported capacity.

SnapShot, TCE, ShadowImage concurrent capacity

If ShadowImage is used on the same controller as SnapShot, managed capacity for concurrent use must also be calculated and compared to maximum supported capacity. If TCE or TrueCopy is used also, it is included in concurrent-use calculations.

Managed concurrent-use capacity is calculated using the following formula:

$$\begin{aligned} &\text{SnapShot: Maximum supported capacity value of P-VOL (TB)} \\ &= \text{Maximum SnapShot single capacity} \\ &- (\text{Total ShadowImage S-VOL capacity} / 17) \\ &- (\text{Total TrueCopy P-VOL and S-VOL capacity} / 17) \\ &- (\text{Total TCE P-VOL and S-VOL capacity} \times 3) \end{aligned}$$

SnapShot maximum single capacity is shown in [Table 2-4](#).

Table 2-4: SnapShot Maximum Single Capacity, per Controller

Equipment Type	Mounted Memory Capacity	Single Maximum Supported Capacity (TB)
AMS2100	1 GB per CTL	Not supported
	2 GB per CTL	46
	4 GB per CTL	56
AMS2300	1 GB per CTL	Not supported
	2 GB per CTL	42
	4 GB per CTL	116
	8 GB per CTL	233
AMS2500	2 GB per CTL	30
	4 GB per CTL	116
	6 GB per CTL	163
	8 GB per CTL	210
	10 GB per CTL	280
	12 GB per CTL	350
	16 GB per CTL	420

Example

In the following example, array and cache memory capacity are AMS2100 and 2 GB.

SnapShot: Maximum supported capacity value of P-VOL (TB)
= Maximum SnapShot single capacity

- (Total ShadowImage S-VOL capacity / 17)
- (Total TrueCopy P-VOL and S-VOL capacity / 17)
- (Total TCE P-VOL and S-VOL capacity x 3)

1. SnapShot Maximum single capacity = 46 TB
2. Divide the total ShadowImage S-VOL capacity by 17. For example:

Total SI S-VOL = 4 TB (4000 GB)

$4000 \text{ GB} / 17 = 235.3 \text{ GB}$

3. Subtract the quotient from the TCE maximum single capacity. For example:

$46 \text{ TB (46000 GB)} - 235 \text{ GB} = 45765 \text{ GB}$

4. Divide the total TrueCopy P-VOL and S-VOL capacity by 17. For example:

Total TrueCopy P-VOL and S-VOL capacity = 20 TB (20000 GB)

$20000 \text{ GB} / 17 = 1176 \text{ GB}$

5. Subtract this quotient from the remaining SnapShot maximum single capacity. For example:

$45765 \text{ GB} - 1176 \text{ GB} = 44589 \text{ GB}$

6. Multiply total TCE P-VOL and S-VOL capacity by 3. For example:

Total TCE P-VOL and S-VOL capacity = 600 GB

$600 \text{ GB} \times 3 = 1800 \text{ GB}$

7. Subtract this product from the remaining SnapShot maximum single capacity. For example:

$44589 \text{ GB} - 1800 \text{ GB} = 42789 \text{ GB}$, the capacity left for SnapShot P-VOLs on the controller.



NOTE: When SnapShot is enabled, a portion of cache memory is assigned to it. Please review [Appendix D, SnapShot with Cache Partition Manager](#) for more information.

If your system's maximum capacity exceeds the maximum allowed capacity, you can do one or more of the following:

- Change the P-VOL size
- Reduce the number of P-VOLs
- Change the data pool size
- Reduce the number of V-VOLs
- Reduce the lifespan of the V-VOL

- Reduce ShadowImage P-VOL/S-VOL sizes

Cache limitations on Data and Data Pool volumes

This section provides comparisons in capacity between the data volumes and the data pool volumes under the limitations of the AMS controllers' cache memory. The values in the tables and graph in this section are calculated from the formulas and maximum supported capacity in [SnapShot and TCE capacity on page 2-23](#).



NOTE: "Data volumes" in this section consist of SnapShot P-VOLs and TCE P-VOLs and S-VOLs (if used).

Table 2-5: P-VOL to Data Pool Capacity Ratio on AMS 2100 when Cache Memory is 2 GB per CTL

Ratio All P-VOL Capacity to All Data Pool Capacity	All P-VOL Capacity to All Data Pool Capacity (TB)
1:0.5	2.0 : 1.0
1:1	1.1 : 1.1
1:3	0.4 : 1.2

Table 2-6: P-VOL to Data Pool Capacity Ratio on AMS 2100/2300 when Cache Memory is 4 GB per CTL

Ratio of All P-VOL Capacity to All Data Pool Capacity	All P-VOL Capacity to All Data Pool Capacity (TB)
AMS 2100/2300	AMS 2100/2300
1:0.5	8.8 : 4.4
1:1	5.1 : 5.1
1:3	1.9 : 5.7

Table 2-7: P-VOL to Data Pool Capacity Ratio on AMS 2300/2500 when Cache Memory is 8 GB per CTL

Ratio of All P-VOL Capacity to All Data Pool Capacity	All P-VOL Capacity to All Data Pool Capacity (TB)
1:0.5	17.1 : 8.5
1:1	10.0 : 10.0
1:3	3.7 : 11.1

Table 2-8: P-VOL to Data Pool Capacity Ratio on AMS 2500 when Cache Memory is 16 GB per CTL

Ratio of All P-VOL Capacity to All Data Pool Capacity	All P-VOL Capacity to All Data Pool Capacity (TB)
1:0.5	34.2 : 17.1
1:1	20.0 : 20.0
1:3	7.5 : 22.5

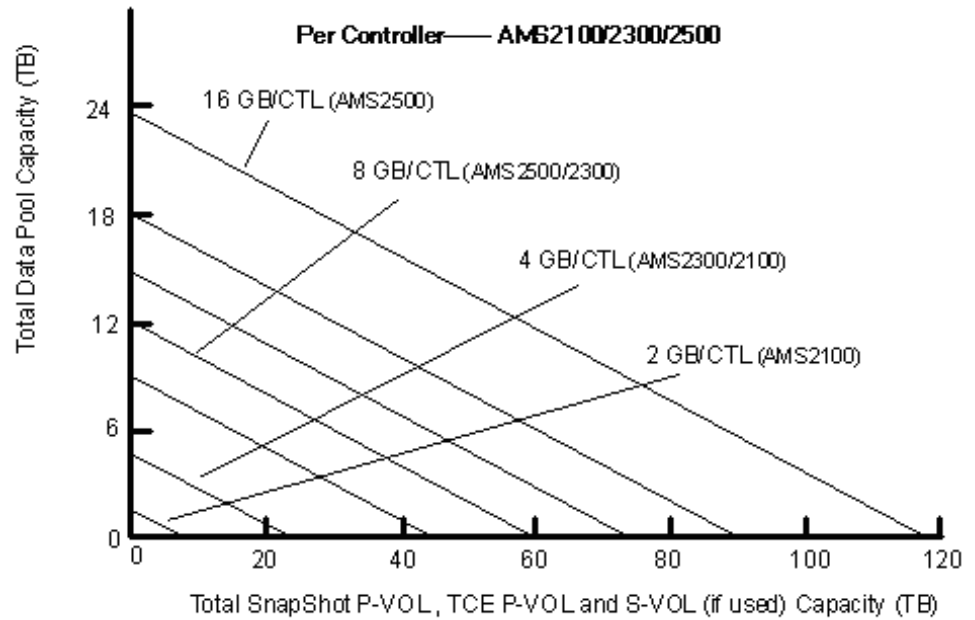


Figure 2-9: Relation of Data Volume, Data Pool Volume Capacities to Cache Size — per Controller

System requirements

This chapter describes minimum requirements.

- ❑ [System requirements](#)
- ❑ [Supported platforms](#)

System requirements

Table 3-1 shows the minimum requirements for SnapShot. See [Appendix A, Specifications](#) for additional information.

Table 3-1: Storage System Requirements

Item	Requirements
Firmware	<p>Version 0832/B or later is required for an AMS 2100 or 2300 array where the hardware Rev. is 0100.</p> <p>Version 0840/A or later is required for an AMS 2500 array where the hardware Rev. is 0100.</p> <p>Version 0890/A or later is required for an AMS 2100, 2300, and 2500 array where the hardware Rev. is 0200.</p>
Storage Navigator Modular 2	<p>Version 3.21 or later is required for the management PC for an AMS 2100 or 2300 array where the hardware Rev. is 0100.</p> <p>Version 4.00 or later is required for the management PC for an AMS 2500 array where the hardware Rev. is 0100.</p> <p>Version 9.00 or later is required for the management PC for an AMS2100, 2300, and 2500 where the hardware Rev. is 0200.</p>
CCI	Version 01-21-03/06 or later is required for the host when CCI is used for the SnapShot operation.
Number of controllers	Two. The primary volume and data pool must be defined under the same controller.
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.
Differential Management LUs	Maximum of 2. DMLU size must be equal to or greater than 10 GB. Two DM-LUs are recommended, each in a different RAID group.
Data pool	<p>Maximum of 64.</p> <ul style="list-style-type: none"> - One per controller required; two per controller highly recommended. - One or more pairs can be assigned to a data pool.
LU size	V-VOL size must equal P-VOL size.

Displaying the hardware revision number

The hardware revision (Rev.) number can be displayed when an individual array is selected from the Arrays list using Navigator 2, version 9.00 or later.

The screenshot shows the Hitachi AMS2300 array management interface. The left sidebar contains a tree view with the following items: Components, Groups, Replication, Settings, Security, Performance, and Alerts & Events. The main content area displays the details for the array **AMS2300_85000026**. The **Summary** section contains a table with the following data:

Summary			
Status	Ready	Capacity of All LU	1.0TB
Type	AMS2300	Raw Capacity of All Drives	103.5TB
H/W Rev.	0100	IPv4 Address	Controller 0 172.16.11.230 Controller 1 172.16.11.231
Serial No.	85000026	IPv6 Address	Controller 0 2001::200:87ff:fec6:46e7 Controller 1 2001::200:87ff:fec6:46e9
Array ID	85000026		
Firmware	0890/A-M		

The **H/W Rev.** value of 0100 is circled in red. Below the summary table is the **Common Array Tasks** section, which lists several tasks with icons and descriptions:

- Initial Setup**: Configure several items on the array to make it ready to use.
- Install License**: Install Licenses for optional storage features.
- Create Logical Unit and Mapping**: Create logical unit and mapping easily.
- Update Firmware**: Update the firmware in the array from a local file or from the support website.
- Backup Volume**: Copy the selected volume to prevent data loss.
- Check for Errors**: View the Alerts & Events screen and show the latest status of the array.
- Look at All Arrays**: Log out of this array and return to the main arrays list. You can then choose another array to manage.

Supported platforms

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

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)
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 and enabling SnapShot

SnapShot must be installed on AMS using a license key. It can also be disabled or uninstalled. This chapter provides instructions for performing these tasks.

- ❑ [Important prerequisite information](#)
- ❑ [Installing, uninstalling SnapShot](#)
- ❑ [Enabling, disabling SnapShot](#)

Important prerequisite information

If TrueCopy or TrueCopy Extended Distance (TCE) are installed, and you are installing/uninstalling or enabling/disabling SnapShot on the *remote* array, the operation causes the following to occur:

- The data paths for TrueCopy or TCE become blocked. The path is recovered from the blockade automatically after the array is restarted.
- TrueCopy or TCE pairs in Paired or Synchronizing status are changed to Failure pair status when the array is restarted for the installation. If the array is not restarted, there is not effect on TCE pair status.
- Hitachi recommends changing TrueCopy or TCE pair status to Split before installing SnapShot on the remote array.

Installing, uninstalling SnapShot

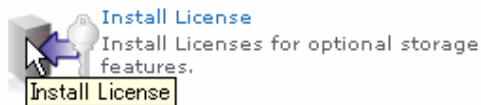
A key code or key file is required to install or uninstall. If you do not have the key file or code, you can obtain it from the download page on the HDS Support Portal, <http://support.hds.com>.

- Installation instructions are provided here for Navigator 2 GUI.
- For CLI instructions, see (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.

To install SnapShot

1. In the Navigator 2 GUI, click the check box for the array where you want to install SnapShot.
2. Click the **Show & Configure Array** button.
3. Under Common Array Tasks, click the **Install License** icon.



The Install License screen displays as shown in [Figure 4-1](#).

License Property

Enter the information for the license to be installed.

*Install with:

Key File:

Input the Key File Name.

Key Code:

Input the Key Code.

* Required field



Figure 4-1: Install License Screen

4. Select the **Key File** or **Key Code** radio button, then enter the file name or key code. You may browse for the Key File.
5. Click **OK**, then click **Confirm** on the confirmation screen.
6. A reboot array message display. The array cannot access the host until the reboot is completed and the system restarts. Make sure that the host has stopped accessing data before beginning the restart process. When ready, click the check box, then click **Reboot Array**. This may take a few minutes.
7. Click **Close**.

To uninstall SnapShot

SnapShot pairs must be released and their status returned to Simplex before uninstalling. The key code is required.

1. In the Navigator 2 GUI, click the check box for the array where you want to uninstall SnapShot.
2. Click the **Show & Configure Array** button.
3. In the tree view, click **Settings**, then select the **Licenses** icon.
4. On the Licenses screen, select **SNAPSHOT** in the Licenses list and click the **De-install License** button.
5. On the De-Install License screen, enter the code in the **Key Code** box, and then click **OK**. On the confirmation screen, click **Confirm**.
6. A reboot array message display. Make sure that the host has stopped accessing data before beginning the restart process. When ready, click the check box, then click **Reboot Array**. This may take a few minutes.
7. Click **Close**.

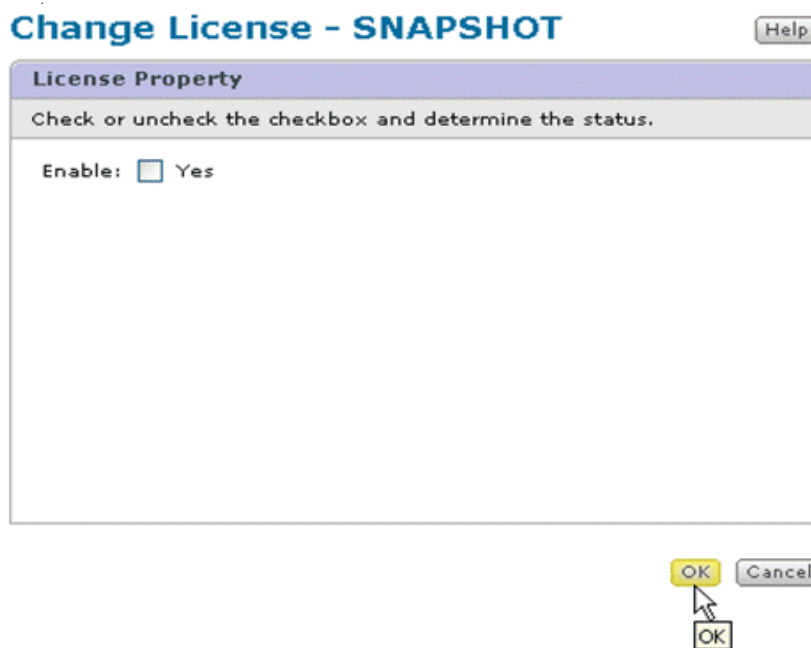
Enabling, disabling SnapShot

To enable/disable SnapShot

(For instructions using CLI, see [Enabling and disabling SnapShot on page B-5.](#))

1. In the Navigator 2 GUI, select the array where you want to enable SnapShot and click the **Show & Configure Array** button.
2. In the tree view, click **Settings**, then click **Licenses**.
3. Select **SnapShot** in the Licenses list, then click the **Change Status** button at the bottom of the page.

The Change License screen appears.



4. To enable, check the **Enable: Yes** box.
To disable, clear the **Enable: Yes** box.



NOTE: When disabling, SnapShot pairs must be deleted.

5. Click **OK**.
6. In the confirmation screen that appears for the enabling and disabling actions, click **OK**. The Reboot Array message appears confirming the operation is complete. Before beginning the restart process, make sure that the host has stopped accessing data.

Reboot Array - SA800_81000026



7. When ready, click the **Yes, I have read...** check box at the bottom of the message, then click **Reboot Array**. The reboot time displays. If the Storage system does not respond after 15 minutes, check its condition.
8. When rebooting has finished, a message appears stating that the restart is successful. Click **Close**.

Configuring SnapShot

This chapter describes the steps for setting up SnapShot.

- ❑ [Configuration workflow](#)
- ❑ [Set up the data pool](#)
- ❑ [Set Up the Virtual Volume \(V-VOL\) \(manual method\)](#)
- ❑ [Set up the Differential Management LU \(DMLU\)](#)
- ❑ [Set up the command device](#)
- ❑ [Setting the LU ownership](#)

Configuration workflow

Setup for SnapShot consists of assigning logical units for the following:

- Data pools
- V-VOL(s)
- DM-LUs
- command device (if using CCI)

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

Please refer to the following for requirements and recommendations:

- [Requirements and recommendations for SnapShot Logical Units on page 2-10](#)
- [Operating system considerations on page 2-13](#)
- [Chapter 3, System requirements](#)
- [Appendix A, Specifications](#)

Set up the data pool

The data pool stores differential data after the snapshot is created. For more information on the data pool see [Data pools on page 1-4](#).

Prerequisites

- LUs for a data pool must be on the same controller as the P-VOL LUs.
- Up to 64 LUs can be assigned to a data pool.
- Hitachi recommends a minimum of 20 GB for data pool size.
- To review the data pool sizing procedure, see [Establishing data pool size on page 2-5](#).
- When Cache Partition Manager is used with SnapShot, the segment size of the LU belonging to a data pool must be the default size (16 KB) or less. See *Hitachi Storage Navigator Modular 2 Storage Features Reference Guide* for more information.

To create and assign volumes for data pools

(For instructions using CLI, see [Setting the POOL on page B-7](#).)

1. In Navigator 2 GUI, select the desired array, then click the **Show & Configure Array** button.
2. From the Replication tree, select the **Local Replication** icon, then select the **Setup** icon. The Setup screen displays.
3. Select **Data Pools**. View screen instructions by clicking the Help button.



NOTE: The default **Threshold** value is 70%. When usage reaches the Threshold plus 1 percent, both data pool and pair status change to “Threshold over”, and the Storage system issues a warning. Data pool status returns to “Normal” when usage decreases 5 percent below the Threshold value.

If capacity reaches 100 percent, the pair fails and all data in the V-VOL is lost.

Set Up the Virtual Volume (V-VOL) (manual method)

You must set up the V-VOL manually when using the Create Pair procedure. Proceed with the instructions below.

If you will use the Backup Wizard in the GUI to create an initial pair, the V-VOL is created and set up automatically during the pair creation operation. However, the prerequisites apply to both methods.

Prerequisites

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

To assign volumes as V-VOLs

(For instructions using CLI, see [Setting the V-VOL on page B-9.](#))

1. In Navigator 2 GUI, select the desired array, then click the **Show & Configure Array** button.
2. From the Replication tree, select the **Local Replication** icon, then select the **Setup** icon. The Setup screen displays.
3. Select **SnapShot Logical Units**. The SnapShot Logical Units page displays.
4. Click **Create SnapShot LU**. The Create SnapShot Logical Unit window appears.
5. Enter the **LUN** to be used for the V-VOL. You can use any unused LUN that matches the P-VOL in size. The lowest available LU number is the default.
6. Enter the V-VOL size in the **Capacity** field. The Capacity range is 1 MB - 128 TB.
7. Click **OK**.

Set up the Differential Management LU (DMLU)

The DMLU is an exclusive volume used for storing SnapShot information when the array is powered down. You must set up the DMLU before using SnapShot.

Prerequisites

- The DMLU must be set at a minimum size of 10 GB.
- Though only one DMLU is needed, two are recommended, with the second used as a backup.
- For RAID considerations, see the bullet on DMLU in [RAID configuration for LUs assigned to SnapShot on page 2-10](#).
- Also see [Appendix A, Specifications](#).

To set up a DMLU

(For instructions using CLI, see [Setting the DMLU on page B-6](#).)

1. In the navigation tree, select **Settings**, then **DMLU**. The Differential Management Logical Units screen displays.
2. Click **Add DMLU**. The Add DMLU screen displays, as shown in [Figure 5-1](#).

Add DMLU

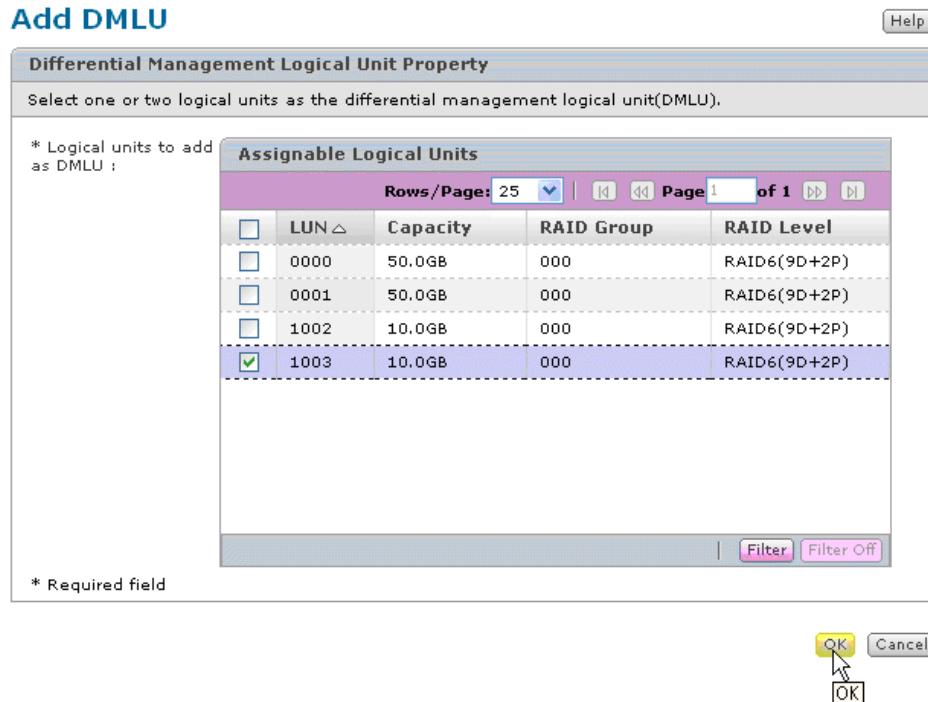


Figure 5-1: Add DMLU Screen

3. Select the **LUNs** that you want to assign as DM-LUs, and then click **OK**. A confirmation message displays.
4. Select the **Yes, I have read...** check box, then click **Confirm**. When a success message displays, click **Close**.

Set up the command device

CCI can be used in place of the Navigator 2 GUI and/or CLI to operate SnapShot. CCI interfaces with AMS through the command device, which is a dedicated logical volume.

Prerequisites

- Setup of the command device is required only if using CCI.
- LUs 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 AMS array.
- Please review [RAID configuration for LUs assigned to SnapShot on page 2-10](#).
- Also see [Appendix A, Specifications](#).
- Refer to *Hitachi Adaptable Modular Storage Command Control Interface (CCI) User's 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 C-2](#)

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 Logical Units box, click the check box for the **LUN** 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 LUN listed in the Command Device column.
5. Click **OK**.

Setting the LU ownership

The load balancing function is not applied to the LUs specified as a SnapShot pair. Since the ownership of the LUs specified as a SnapShot pair is the same as the ownership of the LUs specified as a data pool, perform the setting so that the ownership of LUs specified as a data pool is balanced in advance.

Perform the following steps to set the LU ownership:

1. Select LU Ownership icon in the Tuning Parameter tree view of the Performance tree view.

The screenshot shows the 'LU Ownership' configuration window for array AMS2500_87000026. The window title is 'LU Ownership' and the breadcrumb is 'AMS2500_87000026 > Performance > Tuning Parameter > LU Ownership'. On the left is a tree view with 'LU Ownership' selected. The main area contains a table with columns: LUN, Controller, Core, RAID Group, DP Pool, Cache Partition, and Type. The table has 6 rows, with the last row (LUN 2001, Controller 1, Core Y) selected. Below the table are buttons for 'Change LU Ownership', 'Filter', and 'Filter Off'. A tooltip 'Change Logical Unit Ownership' is visible over the 'Change LU Ownership' button.

<input type="checkbox"/>	LUN	Controller	Core	RAID Group	DP Pool	Cache Partition	Type
<input type="checkbox"/>	0000	0	X	000	N/A	00	SAS
<input type="checkbox"/>	0001	1	X	000	N/A	00	SAS
<input type="checkbox"/>	1000	0	X	000	N/A	00	SAS
<input type="checkbox"/>	1001	1	X	000	N/A	00	SAS
<input type="checkbox"/>	2000	0	X	000	N/A	00	SAS
<input checked="" type="checkbox"/>	2001	1	Y	000	N/A	00	SAS

The Change Logical Unit Ownership screen appears.

2. Select Controller 0 or Controller 1 and X Core or Y Core, and then click OK.

Change Logical Unit Ownership - 2001

The dialog box is titled 'Logical Unit Ownership' and contains the instruction 'Specify the logical unit ownership setting.' It has two sections: 'Controller' with radio buttons for 'Controller 0' (selected) and 'Controller 1'; and 'Core' with radio buttons for 'X Core' (selected) and 'Y Core'. At the bottom right are 'OK' and 'Cancel' buttons. A tooltip 'OK' is visible over the 'OK' button.

The Core item is not displayed: AMS2100/2300

A message appears. Click **Close**.

Using SnapShot

This chapter describes SnapShot copy operations.

- ❑ [SnapShot replication workflow](#)
- ❑ [Back up your volume — creating a pair](#)
- ❑ [Update the V-VOL](#)
- ❑ [Restore the P-VOL from the V-VOL](#)
- ❑ [Use the V-VOL for tape backup, testing, reports, etc.](#)
- ❑ [Edit pairs, data pool](#)
- ❑ [Delete pairs, V-VOLs, data pools, DMLU](#)

SnapShot replication workflow

The SnapShot workflow includes the following:

- Copy a volume
- Update the V-VOL
- Restore the P-VOL from its V-VOL
- Delete a pair, V-VOL, or data pool
- Edit a pair or data pool

The following sections describe these processes.

Back up your volume — creating a pair

(For instructions using CLI, please see [Creating SnapShot pairs on page B-10.](#))

The Navigator 2 GUI allows you to create a pair in two ways:

- Using the backup wizard, described below. This is the simplest and quickest method.
- Using the create pair procedure, which allows more customization. This is described on [Page 6-3](#). You can set copy pace, assign the pair to a group (and create a group), and automatically split the pair after it is created.

When you copy a volume, 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.

Prerequisites

- Make sure the primary volume is set up on the array. See [SnapShot specifications on page A-2](#) for primary volume specifications.
- Create a data pool by assigning a volume. See [Set up the data pool on page 5-2](#).

Using the Backup Wizard

To create a pair using the backup wizard

1. In Navigator 2 GUI, select the desired array, then click the **Show & Configure Array** button.
2. On the array page under Common Array Tasks, click the **Backup Volume** link. The Backup Your Volume Wizard window opens.
3. On the Introduction screen, click **Next**. The Select Primary Volume screen displays.
4. Existing volumes in the array are listed in the **Primary Volume** box. Click the button next to the **LUN** that you want to back up, then click **Next**. The Prepare Secondary Volume screen displays.

5. The system will create a secondary volume (SnapShot virtual volume) in the array, with the same capacity as the selected P-VOL. In the **LUN** box, use the default value (if present), or the enter an available logical unit number. Then click **Next**. The Set Pair Parameters screen displays.
6. Use the default **Pair Name**, or enter a new name.
7. From the **Pool Number** drop-down box, select a data pool for the pair, then click **Next**. The Confirm screen displays.
8. Click **Confirm**, then click **Finish**. The backup pair is created.

Using the Create Pair procedure

Copying a volume using 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

In addition to those prerequisites shown in the previous section, you must also set up the V-VOL when using the Create Pair method. See [Set Up the Virtual Volume \(V-VOL\) \(manual method\) on page 5-3](#)

To create a pair using the create pair procedure

1. In Navigator 2 GUI, select the desired array, then click the **Show & Configure Array** button.
2. From the Replication tree, select the **Local Replication** icon. The **Pairs** screen displays.
3. Select the **Create Pair** button. The Create Pair screen displays.
4. In the Copy Type area, click the **SnapShot radio** button. There may be a brief delay while the screen refreshes.
5. In the **Pair Name** box, enter a name for the pair.
6. In the **Primary** and **Secondary Volume** fields, select the primary and secondary volumes that you want to pair. For SnapShot, the secondary volume list contains the virtual volumes that have been previously set up. The size of the V-VOL must be equal to the P-VOL size.



NOTE: The LUN may be different from the volume's H-LUN. Refer to [Identifying P-VOL and V-VOL LUs on Windows on page 2-13](#) to map LUN to H-LUN.

7. From the **Data Pool Number** drop-down list, select the data pool, previously set up, that you want to assign to the pair.
8. Click the **Advanced** tab.
9. 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.
10. In the **Group Assignment** area, you have the option of assigning the new pair to a consistency group. See [Consistency group \(CTG\) on page 1-6](#) 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:

- 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**.
-

11. 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.
12. Click **OK**, then click **Close** on the confirmation screen that appears. The pair has been created.

Update 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 [Updating SnapShot Logical Unit on page B-11.](#))

1. In Navigator 2 GUI, select the desired array, then click the **Show & Configure 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. When the Resync is completed, click the **Split Pair** button.

Split Pair - SS_LU0000_LU1000 Help

Local Pair Split Option

Select the split pair operation option.

Option :

Attach description to identify the pair upon split :

31 characters or less (alphanumeric characters, special symbols "%", "*", "+", "-", ".", "/", "=", "@", "_", ":", "[", or "]")

OK Cancel

OK

This operation is completed quickly. When finished, the V-VOL is updated.



NOTE: Differential data is deleted from the data pool when a V-VOL is updated. Time required for deletion of 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.

Restore 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 AMS 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 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.



NOTE: Splitting a pair while status is Reverse Synchronizing causes other V-VOLs paired with the P-VOL to be placed in Failure status. Also, the P-VOL data being restored cannot be used logically.

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 AMS2100/2300, and eight per controller for AMS2500 at one time. 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 on page B-11](#).)

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 on page B-11](#).
5. Click **Restore Pair**. View subsequent screen instructions by clicking the Help button.

Use the V-VOL for tape backup, testing, reports, etc.

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 [Update the V-VOL on page 6-5](#).
 - For instructions using CLI, see the resync pair instruction in [Updating SnapShot Logical Unit on page B-11](#).



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 [Update the V-VOL on page 6-5](#).
 - For instructions using CLI, please see the split pair instruction in [Updating SnapShot Logical Unit on page B-11](#).
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.
- Data 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 data pool should be as large or larger than this. It is recommended that the data pool be sized to 1.5 times the P-VOL capacity as a safety precaution.

Figure 6-1 illustrates this example scenario.

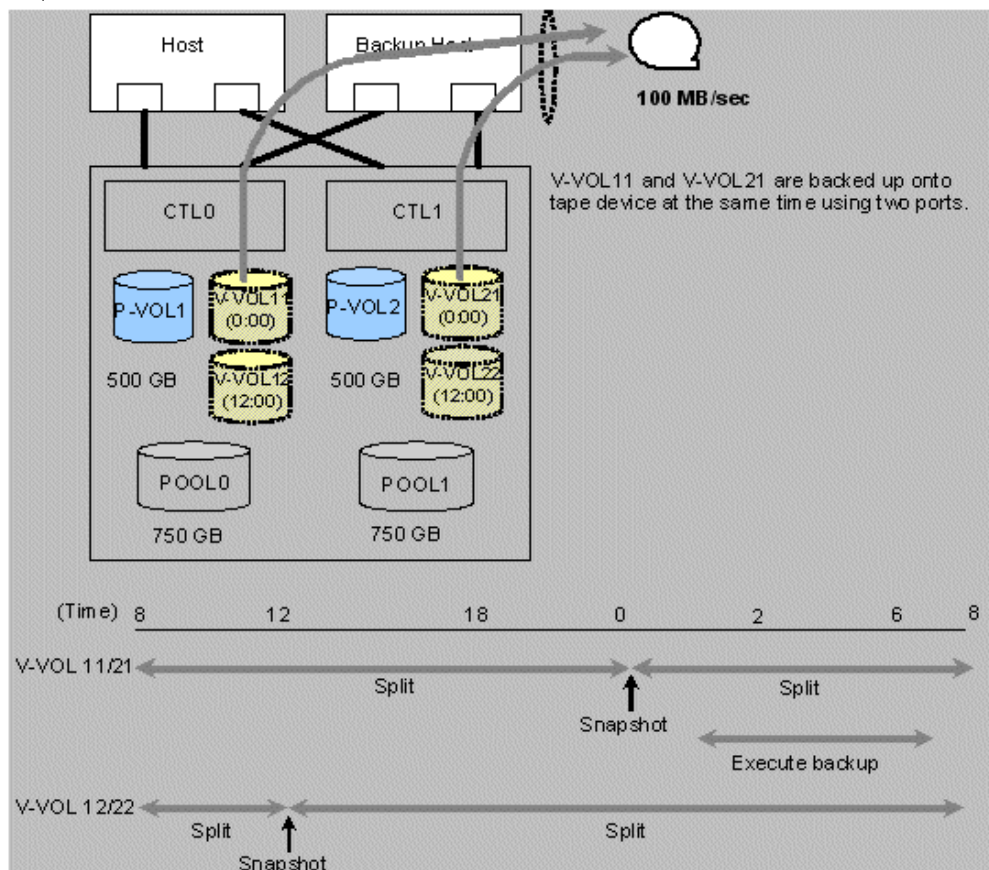


Figure 6-1: Backing-up to Tape

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.
 - Data 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 data pool and impact to performance).
 - Use this method when the V-VOL is in Failure status (V-VOL data is corrupt), or when capacity of the data pool is less than 1.5-times the capacity of the P-VOL.

Edit pairs, data pool

You can edit certain information concerning a pair and a data pool.

- For pairs, you can change the name, group name, and copy pace.
- For data pools, you can increase the size, edit capacity threshold, and add LUs (volumes).

To edit pairs

(For instructions using Navigator 2 CLI, see [Changing pair information on page B-13.](#))

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. View screen instructions for specific information by clicking the Help button.

To edit a data pool

1. In the Storage Navigator 2 GUI, select the **Local Replication** icon in the Replication tree view.
2. In the GUI, select **Setup**, then select **Data Pool**.
3. Select the LUN (data pool) that you want to edit.
4. Make changes to editable fields as needed, then click **OK**.

Delete pairs, V-VOLs, data pools, DMLU

You can delete a pair, V-VOL, data pool, and DMLU to free space or 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.
- Data pool: When a data pool is deleted, the V-VOLs must also be deleted.
- DMLU: When only one DMLU exists, it cannot be deleted. If two DM-LUs exist, one can be deleted.

To delete a pair

(For instructions using the Storage Navigator 2 CLI, see [Releasing SnapShot pairs on page B-12.](#))

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**.

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 Logical Units** icon in the tree view.
3. In the Logical Units for Snapshot list, select the V-VOL that you want to delete.
4. Click **Delete LU for SnapShot**. A message appears.
5. Click **Close**. The V-VOL is deleted.

To delete a data pool

1. Select the **Data Pools** icon in the tree view.
2. Select a data pool you want to delete in the **Data Pool** list.
3. Click **Delete Data Pool**.
4. A message appears. Click **Close**.

To delete a DMLU

1. In the Replication tree view, select **Setup** and then **DMLU**. The Differential Management Logical Units list appears.
1. Select the **LUN** you want to remove.
2. Click the **Remove DMLU** button. A success message displays.
3. Click **Close**.

Monitoring and maintenance

It is important that a data pool's capacity is sufficient to handle the differential data sent to it from the P-VOLs associated with it. If a data 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](#)
- ❑ [Expanding data pool capacity](#)

Monitoring SnapShot

The SnapShot data pool must have sufficient capacity to handle the write workload demands placed on it. You can check that the data pool is large enough to handle workload by monitoring pair status and data pool usage.

Monitoring pair status

To monitor pair status

(If using CLI, see the *Storage Navigator Modular 2 Command Line Interface (CLI) User's Guide*.)

1. In Navigator 2 GUI, select the desired array, then click the **Show & Configure Array** button.
2. From the Replication tree, select the **Local Replication** icon. The **Pairs** screen displays.
3. Locate the pair whose status you want to review in the **Pair** list, then review the **Status** column. Click the **Refresh Information** button to make sure the data is current.

[Table 7-1](#) shows Navigator2 GUI status and descriptions. For CCI status (which differ), see [Confirming pair status on page C-8](#).

Table 7-1: Pair Statuses

Pair Status	Description	P-VOL	V-VOL
Simplex	Status when the volume has not been paired or when the pair has been deleted. The volume has no SnapShot association with another volume.	Read/Write	No Read/Write
Paired	Initial or update copy is completed.	Read/Write	No Read/Write
Synchronizing	Initial or update copy is in progress.	Read/Write	No Read/Write
Reverse Synchronizing	P-VOL restoration from V-VOL is in progress. Status returns to Paired when restoration is completed. When a failure occurs, or the pair is split during restoration, V-VOL status is correlated with the P-VOL status. Any status other than Simplex being restored becomes Failure	Read/Write	No Read/Write
Split	The snapshot is executed and P-VOL data is updated in the V-VOL. In this state, access from a host to a P-VOL is lowered.	Read/Write	Read/Write. (Read/write not accepted during P-VOL restore.)
Threshold Over	Threshold value for the data pool is exceeded; warning status.	P-VOL is being restored	
Failure	Copying is suspended due to a hardware failure or data pool overflow. All V-VOLs become invalid.	Read/Write	Read/Write. (Read/write not accepted during P-VOL restore.)

Monitoring data pool usage

To monitor data pool usage level

1. In Navigator 2 GUI, select the desired array, then click the **Show & Configure Array** button.
2. From the Replication tree, select the **Local Replication** icon, then select the **Setup** icon. The Setup screen displays.
3. Select **Data Pools**. The Data Pools screen displays.
4. Locate the desired data pool and review the **% Used** column. This shows the percentage of the data pool that is being used. Click the **Refresh Information** button to make sure the data is current.

If usage reaches the Threshold level or is close to it on a regular basis, the data pool should be expanded, and/or the lifespan and number of V-VOLs reduced.



NOTE: Threshold is set by the user during data pool setup. It is a percentage of the data pool that, when reached, indicates that maximum capacity is close to being reached. The default Threshold level is 70%.

Expanding data pool capacity

When monitoring indicates that the data pool is in danger of filling, you can add new volumes to expand its size.

To expand data pool capacity

The array allows a maximum of 128 volumes for data pools. One data pool may consist of up to 64 volumes.

1. Split the pair associated with a data pool that is to be expanded. For instructions, see [Update the V-VOL on page 6-5](#).
2. Add a volume or volumes to the data pool.

Other methods for lowering data pool load

When a data pool is in danger of being full, the following actions can be taken as alternatives or in addition to expanding the data pool:

- Delete one or more V-VOLs. With fewer V-VOLs, less data accumulates in the data pool.
- Reduce V-VOL lifespan. By holding snapshots for a shorter length of time, less data accumulates, which relieves the load on the data pool.
- A re-evaluation of your SnapShot system's design may show that not enough data pool space was originally allocated. See [Planning and design on page 2-1](#) for more information.

Troubleshooting

Two types of problem can be experienced with a SnapShot system: pair failure and data pool capacity exceeded. This chapter discusses the causes and provides solutions for these problems.

- ❑ [Pair failure](#)
- ❑ [Data pool capacity exceeded](#)

Pair failure

Pair failure is caused for one of the following reasons:

- A hardware failure occurred that affects pair or data pool volumes.
- A data pool's capacity was exceeded.

To determine the cause of pair failure

1. Check the status of the data pool whose associated pairs' status is changed to Failure. See [Monitoring data pool usage on page 7-3](#) for instructions.
2. If the status of the data pool is POOL FULL, the pair failure is due to exceeded data pool capacity.
3. If the status of the data pool is other than POOL FULL, the pair failure is due to hardware failure.

The procedure for restoring the pair differs according to the cause.

Recovering from pair failure due to POOL FULL

To recover pairs when status is POOL FULL

1. Delete all the pairs that are using the full data pool.
2. Review SnapShot system configuration. See [Chapter 2, Planning and design](#) for detailed information on the following:
 - Measuring write workload and sizing of the data pool. This addresses the amount of data that accumulates in the data pool and provide calculations for determining the correct size of the data pool.
 - Assessing business requirements to establish the lifespan and number of V-VOLs
3. Do one or more of the following:
 - Increase the size of the data pool. See [Expanding data pool capacity on page 7-3](#).
 - Reduce the number of V-VOLs.
 - Reduce the lifespan of V-VOLs.s
4. Re-create the pairs.

Recovering from pair failure due to a hardware failure

To recover the SnapShot system after a hardware failure

1. Review the information log to see what the hardware failure is.
2. Restore the Storage system. See Navigator 2 program Help for details.
3. When the systems is restored, delete the pair. See [Delete pairs, V-VOLs, data pools, DMLU on page 6-10](#) for more information.
4. Re-create the pair.

Data pool capacity exceeded

When your data pool capacity is exceeded, the size of the data pool must be increased. Most likely, the POOL FULL condition has caused pair failures as well.



A

Specifications

This appendix provides external specifications for SnapShot.

- ❑ [SnapShot specifications](#)

SnapShot specifications

Table A-1 lists external specifications for SnapShot.

Table A-1: General Specifications

Item	Specification
AMS model	AMS2100/AMS2300/AMS2500 (For dual configuration only.)
Host interface	Fibre channel or iSCSI.
Number of pairs	<ul style="list-style-type: none"> AMS2300/AMS2500: 2,046 maximum AMS2100: 1,022 maximum Note: When a P-VOL is paired with 32 S-VOLs, the number of pairs is 32.
Cache Memory	AMS2500: 8, 16 GB per controller AMS2300: 4, 8 GB per controller AMS2100: 2, 4 GB per controller
Command devices	Required for CCI. <ul style="list-style-type: none"> Maximum: 128 per array Volume size: 33MB or greater
Differential Management LU (DMLU)	<ul style="list-style-type: none"> Minimum size: 10 GB One DMLU is required, two (maximum) are recommended. When using two, each must be: <ul style="list-style-type: none"> - In different RAID groups than the other - Under different controller than the other
Unit of pair management	Volumes are the target of SnapShot pairs, and are managed per logical unit.
Number of V-VOLs per P-VOL	1:32
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 between LUs for the P-VOL and POOL. The number of data disks may be different.
LU size	LUs for the V-VOL must be equal in size to the P-VOL.
Drive types for the P-VOL and data pool	If the drive types are supported by the array, they can be set for the P-VOL and data pool. SAS and SSD drives are recommended. A DP-VOL CANNOT be a P-VOL.
Consistency Group (CTG) number	<ul style="list-style-type: none"> Up to 256 consistency groups allowed per array AMS2300/AMS2500: 2,046 pairs per CTG (maximum) AMS2100: 1,022 pairs per CTG (maximum)

Table A-1: General Specifications (Continued)

Item	Specification
Data pools	<ul style="list-style-type: none"> • Maximum of 128 LUs per array can be used for data pools • A data pool must be on the same controller as the P-VOL(s) • Unified LUs can be assigned as a data pool only when array firmware version is 0852/A or higher • Maximum of 64 data pools per array (data pool number is 0 to 63) • Up to 64 LUs can be set for one data pool
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 Logical Units on page 2-10 .
Access to data pool LUs from a host	Data pool volumes are not recognizable from the host
Expansion of data pool capacity	Data pool capacity can be expanded by adding LUs
Reduction of data pool capacity	Possible only when all the pairs that use the data pool have been deleted
Unifying, growing and shrinking of an LU assigned to a data pool	No
Formatting, deleting, growing, shrinking of an LU 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 LU	When the array firmware version is less than 0852/A, the capacity of each LU 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	No. The P-VOL and data pool are out of the target for load balancing.
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.
Concurrent use with ShadowImage	SnapShot and ShadowImage can be used at the same time on the same array. However, SnapShot volumes cannot be paired with ShadowImage volumes. If SnapShot is used concurrently with ShadowImage, CTGs limited to 256.
SnapShot use with expanded volumes	Yes

Table A-1: General Specifications (Continued)

Item	Specification
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 Extended Distance on page 2-21 for more information.
Concurrent use of Dynamic Provisioning	Yes. See Concurrent use of Dynamic Provisioning on page 2-16 .
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 LU cannot be specified as a SnapShot P-VOL.
Concurrent use of SNMP Agent Support Function	Yes. When pair status is changed to Failure or Threshold Over, a trap is reported.
Concurrent use of Cache Residency Manager	Yes, however the LU specified for Cache Residency (LU 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 LU segment size must be the default size (16 KB) or less. See Appendix D, Using SnapShot with Cache Partition Manager .
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: <ul style="list-style-type: none"> • When S-VOL Disable is set for an LU, the LU cannot be used in a SnapShot pair. • When S-VOL Disable is set for an LU 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.
Concurrent use of Power Saving	Yes. However, when a P-VOL is included in a RAID group in which Power Saving 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.
Requirement for SnapShot installation	Reboot is required to acquire pool resources.
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 appendix describes Storage Navigator 2 Command Line Interface (CLI) procedures for SnapShot enabling, configuration and copy operations.

- ❑ [Installing and uninstalling SnapShot](#)
- ❑ [Enabling and disabling SnapShot](#)
- ❑ [Operations for SnapShot configuration](#)
- ❑ [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 *Navigator 2 Command Line Interface (CLI) Reference Guide for Replication*.

Installing and 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 (advanced users only).

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

Important prerequisite information

If TrueCopy or TrueCopy Extended Distance (TCE) are installed, and you are installing/uninstalling or enabling/disabling SnapShot on the *remote* array of one of those systems, the operation causes the following to occur:

- The data paths for TrueCopy or TCE become blocked. When a path is blocked, a TRAP occurs, that is, a notification to the SNMP Agent Support Function. The path is recovered from the blockade automatically after the array is restarted.
- TrueCopy or TCE pairs in Paired or Synchronizing status are changed to Failure pair status.
- Hitachi recommends changing TrueCopy or TCE pair status to Split before installing SnapShot on the remote array.

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 lock the option? (y/n[n]): y
The option is unlocked.
In order to complete the setting, it is necessary to reboot
the subsystem.
Host will be unable to access the subsystem while restarting.
Host applications that use the subsystem will terminate
abnormally. Please stop host access before you restart the
subsystem.
Also, if you are logging in, the login status will be canceled
when restarting begins.
When using Remote Replication, restarting the remote subsystem
will cause both Remote Replication paths to fail. Remote
Replication pair status will be changed to "Failure(PSUE)"
when pair status is "Paired(PAIR)" or "Synchronizing(COPY)".
Please change Remote Replication pair status to "Split(PSUS)"
before restart.
Do you agree with restarting? (y/n [n]): y
Are you sure you want to execute?
(y/n [n]): y
Now restarting the subsystem. Start Time hh:mm:ss Time
Required 4 - 15min.
The subsystem restarted successfully.
%

```

It may take time for the array to respond, depending on the condition of the array. If it does not respond after 15 minutes, check the condition of the array.

3. Execute the `auopt` command to confirm whether SnapShot has been installed. An example is shown below. For example:

```

% auopt -unit array-name -refer
Option Name                Type      Term      Status
SNAPSHOT                   Permanent ---      Enable
%

```

SnapShot is installed and the status is "Enable". SnapShot installation is complete.

Uninstalling SnapShot

The key code is required to uninstall SnapShot.

Prerequisites

- SnapShot pairs must be released and their status returned to Simplex
- Data pools must be deleted
- V-VOLs must be deleted

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.  
In order to complete the setting, it is necessary to reboot  
the subsystem.  
Host will be unable to access the subsystem while restarting.  
Host applications that use the subsystem will terminate  
abnormally. Please stop host access before you restart the  
subsystem.  
Also, if you are logging in, the login status will be canceled  
when restarting begins.  
When using Remote Replication, restarting the remote subsystem  
will cause both Remote Replication paths to fail. Remote  
Replication pair status will be changed to "Failure(PSUE)"  
when pair status is "Paired(PAIR)" or "Synchronizing(COPY)".  
Please change Remote Replication pair status to "Split(PSUS)"  
before restart.  
Do you agree with restarting? (y/n [n]): y  
Are you sure you want to execute?  
(y/n [n]): y  
Are you sure you want to execute? (y/n [n]): y  
Now restarting the subsystem. Start Time hh:mm:ss Time  
Required 4 - 15min.  
The subsystem restarted successfully.  
%
```

It may take time for the array to respond, depending on the condition of the array. If it does not respond after 15 minutes, check the condition of the array.

3. Execute the `auopt` command to confirm whether SnapShot has been uninstalled. An example is shown below. For example:

```
% auopt -unit array-name -refer  
DMEC002015: No information displayed.  
%
```

Enabling and disabling SnapShot

SnapShot is bundled with the Adaptable array. You must enable it before using.

The following describes the enabling/disabling procedure.



NOTE: The following conditions must be satisfied in order to disable SnapShot:

- All SnapShot pairs must be released; that is, the status of all LUs are SMPL
- All Data Pools must be deleted, and
- All SnapShot Images (V-VOL) must be deleted

1. 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.
In order to complete the setting, it is necessary to reboot
the array.
Host will be unable to access the array while restarting. Host
applications that use the array will terminate abnormally.
Please stop host access before you restart the array.
Also, if you are logging in, the login status will be canceled
when restarting begins.
When using Remote Replication, restarting the remote array
will cause both R
emote Replication paths to fail.
Remote Replication pair status will be changed to
"Failure(PSUE)" when pair status is "Paired(PAIR)" or
"Synchronizing(COPY)". Please change Remote Replication
pair status to "Split(PSUS)" before restart.
Do you agree with restarting? (y/n [n]): y
Are you sure you want to execute?
(y/n [n]): y
Now restarting the array. Start Time hh:mm:ss Time Required 4
- 15min.
The array restarted successfully.
%
```

It may take time for the array to respond, depending on the condition of the array. If it does not respond after 15 minutes, check the condition of the array.

3. Execute `auopt` to confirm whether the status has been changed. An example is shown below.

```

% auopt -unit array-name -refer
Option Name                Type      Term      Status
SNAPSHOT                   Permanent ---      Disable
%

```

SnapShot Enable/Disable is complete.

Operations for SnapShot configuration

Setting the DMLU

The DMLU (Differential Management Logical Unit) is an exclusive logical unit for storing the differential data while the volume is being copied. 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).

Prerequisites

- DMLU setup is required.
- The DMLU must be set with a minimum size of 10 GB.
- One DMLU is required, though two are recommended. The second one is used as a backup.

To designate DM-LUs:

1. From the command prompt, register and connect to the array to which you want to create the DMLU.
2. Execute the `audmlu` command to create a DMLU.

This command first displays LUs that can be assigned as DMLUs and later creates a DMLU.

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. A DMLU can be released only if two DM-LUs exist. If only one DMLU exists, it cannot be released.

To release a DMLU, specify the `-rm` and `-lu` options in the `audmlu` command.

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.
%
```

Setting the POOL

Up to 64 data pools can be designated for each array, by assigning a logical unit that has been created and formatted. Up to 64 logical units can be assigned to each data pool. The accurate capacity of a data pool cannot be determined immediately after an LU has been assigned. Data pool capacity can only be confirmed approximately 3 minutes per 100 GB.

The following restrictions apply to LUs assigned to a data pool:

- Logical units once assigned to a data pool are no longer recognized by a host.
- Because data will be lost when excess over the limited value of the POOL capacity occurs, 20 GB at least is recommended as a standard POOL capacity. Incidentally, when the POOL capacity being used exceeds the threshold value (default value: usage rate of 70%), the pair in the Split status is changed to the Pool full status.
- An LU with a SAS/SSD drive and an LU with a SATA drive cannot coexist in a data pool.
- When using SnapShot with Cache Partition Manager, the segment size of the LU belonging to a data pool must be the default size (16 KB) or less.

The following is the procedure for creating a POOL for storing differential data for use by SnapShot.

To designate data Pool(s) (POOL(s)):

1. From the command prompt, register the array to which you want to create the Data Pool, then connect to the array.
2. Execute the `aupool` command create a Data Pool.

First, display the LUs to be assigned to a Data Pool, and then create a Data Pool.

The following is the example of specifying LU 100 for Data Pool 0.

```

% aupool -unit array-name -availablelist -poolno 0
Data Pool      : 0
Available Logical Units
  LUN  Capacity RAID Group RAID Level  Type Status
  100   30.0 GB           0  6( 9D+2P) SAS Normal
  200   35.0 GB           0  6( 9D+2P) SAS Normal
%
% aupool -unit array-name -add -poolno 0 -lu 100
Are you sure you want to add the logical unit(s) to the data
pool 0?
(y/n[n]): y
The logical unit has been successfully added.
%

```

3. Execute the `aupool` command to verify that the Data Pool has been created. Refer to the following example.

```

% aupool -unit array-name -refer -poolno 0
Data Pool      : 0
Data Pool Usage Rate: 6% (2.0/30.0 GB)
Threshold      : 70%
Usage Status   : Normal
  LUN  Capacity      RAID Group RAID Level  Type Status
  100   30.0 GB           0  6( 9D+2P) SAS Normal
%

```

4. When deleting the logical unit set as the Data Pool, it is necessary to delete all SnapShot images (V-VOLs). To delete an existing Data Pool, refer to the following example.

```

% aupool -unit array-name -rm -poolno 0
Are you sure you want to delete all logical units from the
data pool 0?
(y/n[n]): y
The logical units have been successfully deleted.
%

```

5. To change an existing threshold value for a Data Pool, refer to the following example.

```

% aupool -unit array-name -cng -poolno 0 -thres 70
Are you sure you want to change the threshold of usage rate
in the data pool?
(y/n[n]): y
The threshold of the data pool usage rate has been
successfully changed.
%

```

Setting the V-VOL

To create a SnapShot pair you must first set up a V-VOL.

If a specification for the logical unit assigned to a V-VOL is omitted when setting the V-VOL, Navigator 2 assigns the smallest undefined number to the logical unit.

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 LU ownership

The load balancing function is not applied to the LUs specified as a SnapShot pair. Since the ownership of the LUs specified as a SnapShot pair is the same as the ownership of the LUs specified as a data pool, perform the setting so that the ownership of LUs specified as a data pool is balanced in advance.

Perform the following steps to set LU ownership by CLI:

1. From the command prompt, register the array to which you want to set the LU ownership, and then connect to the array.
2. Execute the `autuningluown` command to confirm an LU ownership.

```
% autuningluown -unit array-name -refer
LU   CTL   Core  RAID Group  DP Pool  Cache Partition  Type
  0   0     X      0         N/A      0                SAS
  1   1     X      0         N/A      0                SAS
1000 0     X      0         N/A      0                SAS
1001 1     X      0         N/A      0                SAS
2000 0     X      0         N/A      0                SAS
2001 1     Y      0         N/A      0                SAS
%
```

The Core shows N/A: AMS2100/2300

- Execute the **autuningluown** command to change the LU 2001 ownership.

```
% autuningluown -unit array-name -set -lu 2001 -ctl0 -coreX
Are you sure you want to set the LU ownership? (y/n [n]): y
The LU ownership has been set successfully.
%
```

- Execute the **autuningluown** command to confirm an LU ownership.

```
% autuningluown -unit array-name -refer
  LU  CTL  Core  RAID Group  DP Pool  Cache Partition  Type
   0   0    X      0          N/A          0 SAS
   1   1    X      0          N/A          0 SAS
 1000  0    X      0          N/A          0 SAS
 1001  1    X      0          N/A          0 SAS
 2000  0    X      0          N/A          0 SAS
 2001  0    X      0          N/A          0 SAS
%
```

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

To create SnapShot pairs:

- From the command prompt, register the array to which you want to create the SnapShot pair, then connect to the array.
- Execute the `aureplicationlocal` command create a pair.

First, display the LUs 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 RAID Level  Type Status
   100   30.0 GB      0  6( 9D+2P)  SAS Normal
   200   35.0 GB      0  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.
%
```

- 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.

Updating SnapShot Logical Unit

To update the V-VOL:

1. From the command prompt, register the array to which you want to update the SnapShot pair, then connect to the array.
2. Execute the `aureplicationlocal` command update the pair.

Change the Split status of the SnapShot pair to Paired status using `-resync` option. Then, change the status to Split using `-split` option. Refer to the following example.

```

% aureplicationlocal -unit array-name -ss -resync -pvol 200 -
svol 1001
Are you sure you want to re-synchronize pair?
(y/n[n]): y
The pair has been re-synchronized successfully.
%
% aureplicationlocal -unit array-name -ss -split -pvol 200 -
svol 1001
Are you sure you want to split pair?
(y/n[n]): y
The pair has been split successfully.
%

```

3. Execute `aureplicationlocal` to update 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
%

```

The V-VOL was updated.

Restoring V-VOL to P-VOL

To restore the V-VOL to the P-VOL:

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.

Releasing SnapShot pairs

To release the SnapShot pair and change the status to Simplex:

1. From the command prompt, register the array to which you want to release the SnapShot pair, then connect to the array.
2. Execute the `aureplicationlocal` command to release 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 release the pair. Refer to the following example.

```
% aureplicationlocal -unit array-name -ss -refer
DMEC002015: No information is displayed.
%
```

The SnapShot pair is released.

Changing pair information

You can change the pair name, 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.
%
```

The SnapShot pair information is changed.

Creating pairs that belong to a group

To create multiple SnapShot 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. 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.
%
```

2. Add the name to the group if necessary using command 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.

```
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-xxxxxxxxxxxxxxxx
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 2000, Windows Server 2003, or Windows Server 2008 is used, the CCI mount command must be used when mounting/unmounting 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 chapter 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 Command Device is a dedicated logical volume on the array that functions as the interface to CCI software. 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 array.

The following sub-sections describe necessary set up procedures for CCI for SnapShot.

Setting the command device

The command device(s) and LU mapping setting is used Navigator 2.

To set up a command device(s)

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 LUs 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 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 1 200
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 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 array downtime, designate two or more command devices. For details on alternate Command Device function, refer to the *Hitachi Adaptable Modular Storage Command Control Interface (CCI) User's Guide*.

4. The following example releases a command device:

```

% 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 an already set command device, release the command device, then change the LU number. The following example specifies LU 201 for command device 1.

```

% aucmddev -unit 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 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 LU Mapping. The following is an example of setting LU 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      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.

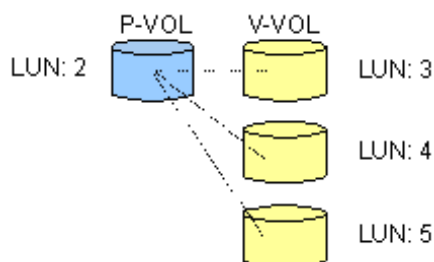
A sample configuration definition file, `HORCM_CONF`, is included with the CCI software. It should be used as the basis for creating your configuration definition file(s). The system administrator should copy the sample file, set the necessary parameters in the copied file, and place the copied file in the proper directory. For details on configuration definition file, refer to the *Hitachi Adaptable Modular Storage Command Control Interface (CCI) User's Guide*.

The configuration definition file can be automatically created using the `mkconf` command tool. For details on the `mkconf` command, refer to the *Hitachi Adaptable Modular Storage Command Control Interface (CCI) Reference 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 the following figure.



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 array. See the *Hitachi Adaptable Modular Storage Command Control Interface (CCI) User's Guide* for more information.

- In the **HORCM_CMD** section, specify the physical drive (command device) on the array. For example:

```

HORCM_MON
#ip_address      service      poll(10ms)    timeout(10ms)
xxxxxxxxxxx      5000         12000         3000

HORCM_CMD
#dev_name        dev_name      dev_name
\\.\CMD-8100012385000123-200-CL1-A

HORCM_LDEV
#dev_group       dev_name      Serial#       CU:LDEV(LDEV#)  MU#
VG01             oradb1        8100012385000123  02              0
VG01             oradb2        8100012385000123  02              1
VG01             oradb3        8100012385000123  02              2

HORCM_INST
#dev_group       ip_address    service
VG01             xxxxxxxxxxxx  5001

```

- Set the necessary parameters in the **HORCM_LDEV** section, then in the **HORCM_INST** section.
- Save the configuration definition file.
- Repeat Steps 3 to 7 for the horcm1.conf file. Example:

```

HORCM_MON
#ip_address      service      poll(10ms)    timeout(10ms)
xxxxxxxxxxx      5001         12000         3000

HORCM_CMD
#dev_name        dev_name      dev_name
\\.\CMD-8100012385000123-200-CL1-A

HORCM_LDEV
#dev_group       dev_name      Serial#       CU:LDEV(LDEV#)  MU#
VG01             oradb1        8100012385000123  03              0
VG01             oradb2        8100012385000123  04              0
VG01             oradb3        8100012385000123  05              0

HORCM_INST
#dev_group       ip_address    service
VG01             xxxxxxxxxxxx  5000

```

- Enter the following example lines in the command prompt to verify the connection between CCI and the array:

```

C:\>cd HORCM\etc

C:\HORCM\etc>echo hd1-7 | .\inraid
Harddisk 1 -> [ST] CL1-A Ser =8100012385000123 LDEV = 200 [HITACHI ] [DF600F-
CM      ]
Harddisk 2 -> [ST] CL1-A Ser =8100012385000123 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 =8100012385000123 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 =8100012385000123 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 =8100012385000123 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 =8100012385000123 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 =8100012385000123 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 Adaptable Modular Storage Command Control Interface (CCI) User's 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
```

3. Execute the horcmstart script, and then execute the pairdisplay 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>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 )85000123    2.SMPL  ----,-----  ---- -
VG01    oradb1(R)   (CL1-A , 1, 3-0 )85000123    3.SMPL  ----,-----  ---- -
VG01    oradb2(L)   (CL1-A , 1, 2-1 )85000123    2.SMPL  ----,-----  ---- -
VG01    oradb2(R)   (CL1-A , 1, 4-0 )85000123    4.SMPL  ----,-----  ---- -
VG01    oradb3(L)   (CL1-A , 1, 2-2 )85000123    2.SMPL  ----,-----  ---- -
VG01    oradb3(R)   (CL1-A , 1, 5-0 )85000123    5.SMPL  ----,-----  ---- -

```

Performing SnapShot operations

Pair operation using CCI are shown in Figure D-1.

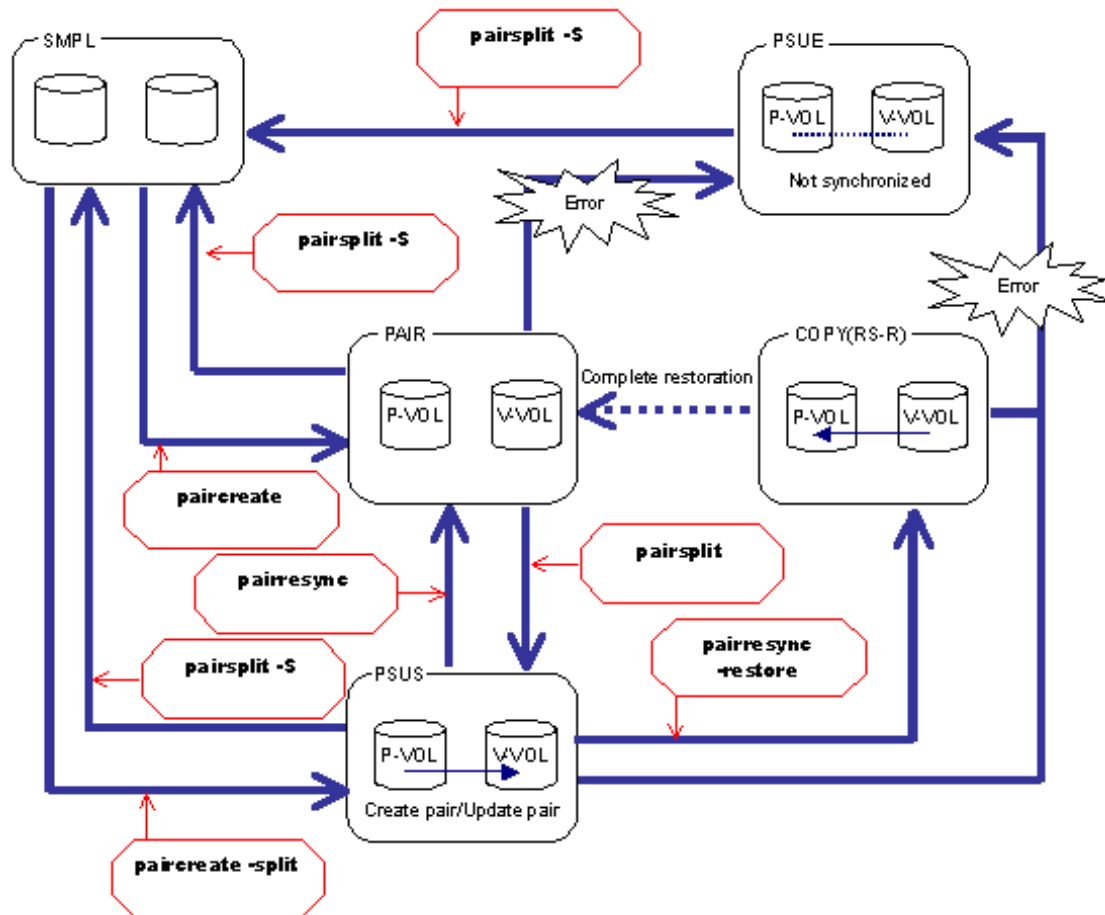


Figure C-1: Snapshot Pair Status for CCI

Confirming pair status

Table C-1 shows the related CCI and Navigator 2 GUI pair status.

Table C-1: 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 SnapShot instruction is retained in the V-VOL.

Table C-1: CCI/Navigator 2 GUI Pair Status

CCI	Navigator 2	Description
PFUS	Threshold Over	Status in which the used rate of data pool reaches the threshold of data pool.
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 `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, 2-0 )85000123      2.P-VOL PSUS,-----  ---- -
VG01   oradb1(R)   (CL1-A , 1, 3-0 )85000123      3.S-VOL SSUS,-----  ---- -
```

The pair status is displayed. For details on the `pairdisplay` command and its options, refer to *Hitachi AMS 2000 Family Command Control Interface (CCI) Reference Guide*.

Paircreate operation

To create SnapShot pairs

In the examples below, the group name in the configuration definition file is VG01.

1. Execute `pairdisplay` 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 )8100012385000123  2.P-VOL PSUS,-----  ---- -
VG01   oradb1(R)   (CL1-A , 1, 3-0 )8100012385000123  3.S-VOL SSUS,-----  ---- -
VG01   oradb2(L)   (CL1-A , 1, 2-1 )8100012385000123  2.SMPL  ----,-----  ---- -
VG01   oradb2(R)   (CL1-A , 1, 4-0 )8100012385000123  4.SMPL  ----,-----  ---- -
VG01   oradb3(L)   (CL1-A , 1, 2-2 )8100012385000123  2.SMPL  ----,-----  ---- -
VG01   oradb3(R)   (CL1-A , 1, 5-0 )8100012385000123  5.SMPL  ----,-----  ---- -
```

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 1-6](#).

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 )8100012385000123 2.P-VOL PSUS,----- ---- -
VG01   oradb1(R)   (CL1-A , 1, 3-0 )8100012385000123 3.S-VOL SSUS,----- ---- -
VG01   oradb2(L)   (CL1-A , 1, 2-1 )8100012385000123 2.P-VOL PSUS,----- ---- -
VG01   oradb2(R)   (CL1-A , 1, 4-0 )8100012385000123 4.S-VOL SSUS,----- ---- -
VG01   oradb3(L)   (CL1-A , 1, 2-2 )8100012385000123 2.P-VOL PSUS,----- ---- -
VG01   oradb3(R)   (CL1-A , 1, 5-0 )8100012385000123 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.

Updating the V-VOL

To update the V-VOL

In the examples below, the group name in the configuration definition file is VG01.

1. Change the PSUS status of the SnapShot pair to PAIR status using `pairresync`; then, change the status to PSUS using `pairsplit`. For example:

```
C:\HORCM\etc>pairresync -g VG01 -d oradb1
C:\HORCM\etc>pairevtwait -g VG01 -s pair -t 300 10
pairevtwait : Wait status done.
C:\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 )8100012385000123 2.P-VOL PSUS,----- ---- -
VG01   oradb1(R)   (CL1-A , 1, 3-0 )8100012385000123 3.S-VOL SSUS,----- ---- -
VG01   oradb2(L)   (CL1-A , 1, 2-1 )8100012385000123 2.SMPL ----,----- ---- -
VG01   oradb2(R)   (CL1-A , 1, 4-0 )8100012385000123 4.SMPL ----,----- ---- -
VG01   oradb3(L)   (CL1-A , 1, 2-2 )8100012385000123 2.SMPL ----,----- ---- -
VG01   oradb3(R)   (CL1-A , 1, 5-0 )8100012385000123 5.SMPL ----,----- ---- -
```

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 `pairdisplay` to display 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 )8100012385000123 2.P-VOL RCCOPY,----- ---- -
VG01   oradb1(R)   (CL1-A , 1, 3-0 )8100012385000123 3.S-VOL RCCOPY,----- ---- -
VG01   oradb2(L)   (CL1-A , 1, 2-1 )8100012385000123 2.SMPL ----,----- ---- -
VG01   oradb2(R)   (CL1-A , 1, 4-0 )8100012385000123 4.SMPL ----,----- ---- -
VG01   oradb3(L)   (CL1-A , 1, 2-2 )8100012385000123 2.SMPL ----,----- ---- -
VG01   oradb3(R)   (CL1-A , 1, 5-0 )8100012385000123 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.

Releasing SnapShot pairs

To release 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 `pairdisplay` command to verify that the status of the SnapShot pair is PSUS or PSUE. 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 )8100012385000123 2.P-VOL PSUS,----- ---- -
VG01   oradb1(R)   (CL1-A , 1, 3-0 )8100012385000123 3.S-VOL SSUS,----- ---- -
VG01   oradb2(L)   (CL1-A , 1, 2-1 )8100012385000123 2.SMPL ----,----- ---- -
VG01   oradb2(R)   (CL1-A , 1, 4-0 )8100012385000123 4.SMPL ----,----- ---- -
VG01   oradb3(L)   (CL1-A , 1, 2-2 )8100012385000123 2.SMPL ----,----- ---- -
VG01   oradb3(R)   (CL1-A , 1, 5-0 )8100012385000123 5.SMPL ----,----- ---- -
```

- Execute the `pairsplit -S` command to release the SnapShot pair. For example:

```
C:\HORCM\etc>pairsplit -S -g VG01 -d oradb1
```

- 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-M) ,Seq#,LDEV#.P/S,Status, Seq#,P-LDEV# M
VG01   oradb1(L)   (CL1-A , 1, 2-0 )8100012385000123 2.SMPL ----,----- ---- -
VG01   oradb1(R)   (CL1-A , 1, 3-0 )8100012385000123 3.SMPL ----,----- ---- -
VG01   oradb2(L)   (CL1-A , 1, 2-1 )8100012385000123 2.SMPL ----,----- ---- -
VG01   oradb2(R)   (CL1-A , 1, 4-0 )8100012385000123 4.SMPL ----,----- ---- -
VG01   oradb3(L)   (CL1-A , 1, 2-2 )8100012385000123 2.SMPL ----,----- ---- -
VG01   oradb3(R)   (CL1-A , 1, 5-0 )8100012385000123 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 Adaptable Modular Storage Command Control Interface (CCI) Reference Guide*.

Using SnapShot with Cache Partition Manager

This chapter provides special instructions for using SnapShot with Cache Partition Manager.

- [SnapShot with Cache Partition Manager](#)

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 D-1 shows partitions before SnapShot is installed; Figure D-2 shows them with SnapShot.

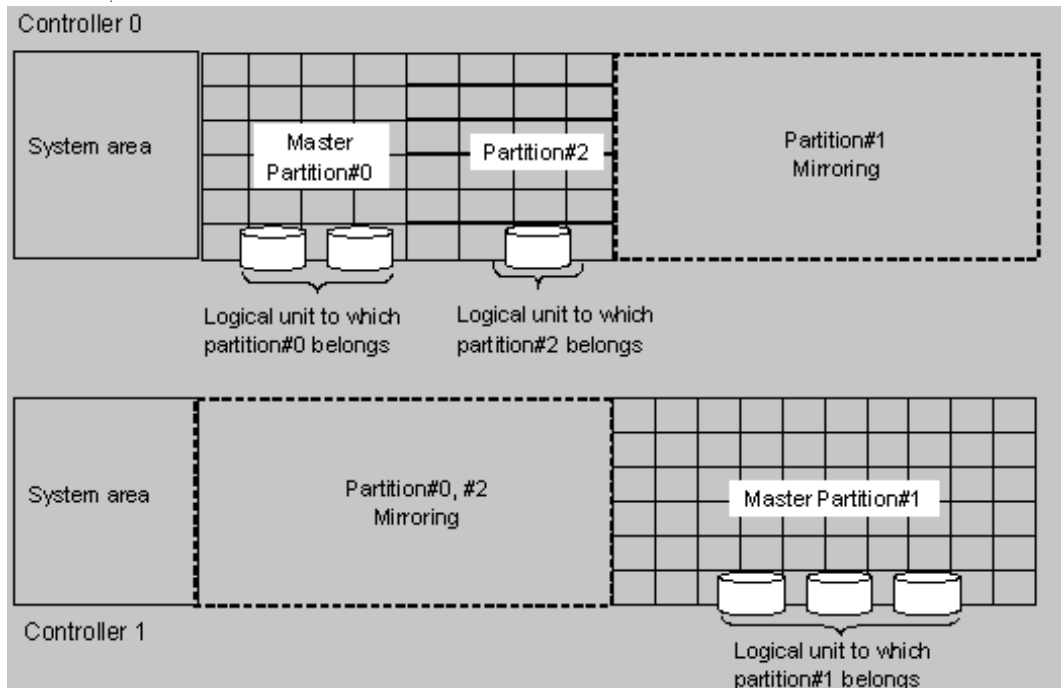


Figure D-1: Cache Partitions with Cache Partition Manager

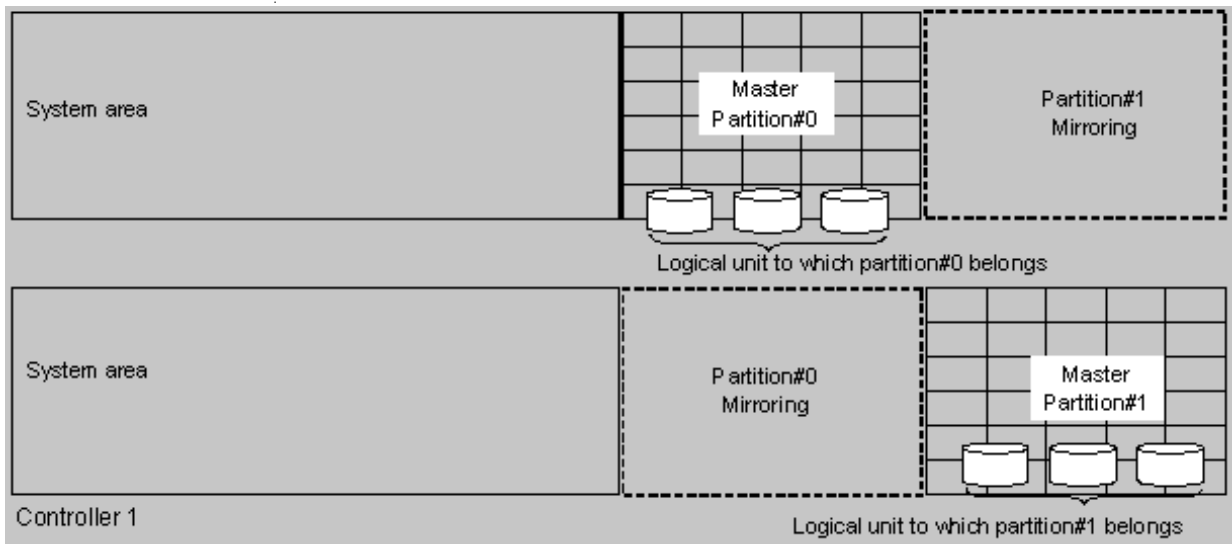


Figure D-2: Cache Partitions when SnapShot with Cache Partition Manager



Glossary

This glossary provides definitions for replication terms as well as terms related to the technology that supports your Hitachi AMS 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 for 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 gigabytes) 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.

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.

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 storage systems. Command devices are not used by ordinary applications. Command devices can be shared between several hosts.

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 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 storage system.

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 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 storage system. 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.

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
---	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---

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 storage system 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 storage system or snapshots of backup data in the remote storage system). 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).

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.

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 Management Logical Unit (DMLU)

The volumes used to manage differential data in a storage system. In a TrueCopy Extended Distance system, there may be up to two DM logical units configured per storage system. 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.

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 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.

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
---	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---

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 storage systems 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.

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

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
---	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---

number of incoming host write operations (inflow) and outgoing data transfers (outflow).

GUI

Graphical user interface.

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.

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

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
---	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---

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.

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 subsystem.

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.

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
---	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---

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.

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).

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
---	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---

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.

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.

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.

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 or target site

Maintains mirrored data from the primary site.

remote path

A route connecting identical ports on the local storage system and the remote storage system. Two remote paths must be set up for each storage system (one path for each of the two controllers built in the storage system).

remote volume stem

In TrueCopy operations, the remote volume (R-VOL) is a volume located in a different subsystem from the primary host subsystem.

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.

SATA

Serial ATA is a computer bus technology primarily designed for the transfer of data to and from hard disks and optical drives. SATA is the evolution of the legacy Advanced Technology Attachment (ATA) interface from a parallel bus to serial connection architecture.

secondary volume (S VOL)

A replica of the primary volume (P-VOL) at the time of a backup and is kept on a standby storage system. Recurring differential data updates are performed to keep the data in the S-VOL consistent with data in the P-VOL.

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
---	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---

ShadowImage

Software that uses local mirroring technology to create a copy of any volume in the Simple Modular Storage (SMS) array.

Simple DR

Simple DR refers to the Simple Data Recovery.

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.

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 storage system. 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.

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
---	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---

T

target copy

A file, device, or any type of location to which data is moved or copied.

TrueCopy

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.

volume

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 storage system and one secondary volume (S-VOL) on the remote storage systems.

V-VOL

See virtual volume.

V-VOLTL

Virtual Volume Tape Library.

W

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

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
---	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---

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

A

- AMS, version [3-2](#)
- assessing business needs [2-2](#)
- assigning pairs to a consistency group [6-4](#)

B

- backup script, using CLI [B-14](#)

C

- Cache Partition Manager, using with SnapShot [D-2](#)
- capacity, supported maximum [2-22](#)
- Cascade Connection of SnapShot with TrueCopy [1-11](#)
- Cascade Restrictions with SnapShot Data Pool [1-15](#)
- Cascade Restrictions with SnapShot P-VOL [1-11](#)
- Cascade Restrictions with SnapShot V-VOL [1-13](#)
- CCI
 - description [1-6](#)
 - version [3-2](#)
- CCI, using to
 - confirm pair status [C-8](#)
 - create pairs [C-9](#)
 - define config def file [C-4](#)
 - release pairs [C-11](#)
 - restore the P-VOL [C-11](#)
 - set environment variable [C-6](#)
 - set LU mapping [C-3](#)
 - set up the command device [C-2](#)
 - update the V-VOL [C-10](#)
- CLI
 - description [1-6](#)
- CLI, using to
 - change pair information [B-13](#)
 - create pairs [B-10](#)
 - enable and disable SnapShot [B-5](#)
 - install [B-2](#)
 - release pairs [B-12](#)
 - restore the P-VOL [B-11](#)

- set up the pool [B-7](#)
- set up the V-VOL [B-9](#)
- uninstall [B-4](#)
- update the V-VOL [B-11](#)
- Command Control Interface, see CCI.
- command device
 - RAID configuration [2-12](#)
 - releasing [C-2](#)
 - set up using CCI [C-2](#)
 - set up using GUI [5-5](#)
- Command Line Interface. See CLI
- Configuration Restrictions on the Cascade of TrueCopy with SnapShot [1-15](#)
- configuration workflow [5-2](#)
- Consistency Groups
 - creating and assigning pairs to using GUI [6-4](#)
 - creating pairs for using CLI [B-13](#)
 - description [1-6](#)
- create pair [6-3](#)
- creating the V-VOL [6-3](#)

D

- data pools
 - creating [5-2](#)
 - editing [6-10](#)
 - expanding [7-3](#)
 - general [1-4](#)
 - monitoring usage [7-3](#)
 - RAID configuration [2-11](#)
 - requirements [2-9](#)
 - sizing [2-5](#)
 - specifications [A-3](#)
 - Threshold field [5-3](#)
- delete
 - data pool [6-10](#)
 - DMLU [6-10](#)
 - volume pair [6-10](#)
 - V-VOL [6-10](#)
- designing the SnapShot system [2-1](#)
- Differential Management LUs. See DMLU
- disabling SnapShot [4-1](#)

DMLU

- delete [6-10](#)
- description [1-6](#)
- RAID configuration [2-12](#)

E

- editing data pool information [6-10](#)
- editing pair information [6-10](#)
- enabling SnapShot [4-1](#)
- expanding data pool size [7-3](#)

F

- frequency, snapshot [2-2](#)

G

- graphic, SnapShot hardware and software [1-2](#)
- Group Name, adding [6-4](#)
- GUI, description [1-6](#)
- GUI, using to
 - create a pair [6-2](#)
 - delete a DMLU [6-11](#)
 - delete a pair [6-10](#)
 - delete a V-VOL [6-11](#)
 - edit a data pool [6-10](#)
 - edit pairs [6-10](#)
 - enable, disable SnapShot [4-4](#)
 - install [4-2](#)
 - monitor data pool usage [7-3](#)
 - monitor pair status [7-2](#)
 - restore the P-VOL [6-6](#)
 - set up data pools [5-2](#)
 - set up the command device [5-5](#)
 - set up the DMLU [5-4](#)
 - set up the V-VOL [5-3](#)
 - uninstall [4-3](#)
 - update the V-VOL [6-5](#)

H

- hardware failure, recovering [8-2](#)
- how long to hold snapshots [2-3](#)
- how often to take snapshots [2-2](#)

I

- installation [4-2](#)
- installing SnapShot [4-1](#)
- interfaces for SnapShot [1-6](#)

L

- lifespan, snapshot [2-3](#)

M

- maintaining the SnapShot system [7-2](#)
- measuring write workload [2-6](#)
- monitoring data pool usage [7-2](#), [7-3](#)
- monitoring pair status [7-2](#)

N

- number of V-VOLs, establishing [2-4](#)

O

- overview [1-1](#)

P

- pair status
 - definitions [7-2](#)
 - monitoring [7-2](#)
- pairs, assigning to a consistency group [6-4](#)
- planning the SnapShot system [2-1](#)
- platforms, supported [3-4](#)
- Pool Full, recovery from [8-2](#)
- P-VOLs and V-VOLs [1-3](#)

R

- RAID levels for SnapShot volumes [2-10](#)
- recovering from Pool Full, hardware failure [8-2](#)
- releasing (deleting) a pair [6-10](#)
- releasing a command device [C-2](#)
- Replication Manager [1-7](#)
- reports, using the V-VOL for [6-7](#)
- requirements
 - data pool [2-9](#)
 - SnapShot system [3-2](#)
- restoring the P-VOL [6-6](#)
- resync a pair [6-5](#)
- rolling average [2-7](#)

S

- Setting up a command device [C-2](#)
- sizing the data pool [2-5](#)
- SnapShot
 - data pools [1-4](#)
 - data, graphic [1-5](#)
 - enabling, disabling [4-1](#)
 - how it works [1-3](#)
 - installing [4-2](#)
 - installing, uninstalling [4-1](#)
 - interface [1-6](#)
 - interfaces [1-8](#)
 - maintaining [7-2](#)
 - overview [1-1](#)
 - planning [2-1](#)
 - restoring the P-VOL operation [6-6](#)
 - specifications [A-1](#)
 - uninstalling [4-3](#)
 - using with Cache Partition Manager [D-2](#)
- SnapShot* versus *snapshot* [1-1](#)
- snapshots
 - how long to keep [2-3](#)
 - how often to make [2-2](#)
- specifications [A-3](#)
- splitting a pair [6-5](#)
- status definitions [7-2](#)
- Storage Navigator Modular [2](#)

Index-2

- description [1-8](#)
- version [3-2](#)
- supported capacity calculation [2-22](#)
- supported platforms [3-4](#)

T

- tape backups [6-7](#)
- testing, using the V-VOL for [6-7](#)
- threshold reached, consequences [5-3](#)

U

- uninstalling SnapShot [4-1](#), [4-3](#)

V

- version
 - AMS [3-2](#)
 - CCI [3-2](#)
 - Navigator 2 [3-2](#)
- volume pairs
 - creating [6-3](#)
 - description [1-3](#)
 - editing [6-10](#)
 - monitoring status [7-2](#)
 - RAID configuration [2-11](#)
 - statuses [7-2](#)
- volumes
 - assigning to data pools [5-2](#)
 - setup recommendations [2-10](#)
- V-VOLs
 - creating [6-3](#)
 - deleting [6-10](#)
 - description [1-3](#)
 - establishing number of [2-4](#)
 - procedure for secondary uses [6-7](#)
 - updating [6-5](#)

W

- write workload
 - measuring [2-6](#)
 - rule of thumb change rates [2-9](#)

Hitachi Data Systems

Corporate Headquarters
750 Central Expressway
Santa Clara, California 95050-2627
U.S.A.
Phone: 1 408 970 1000
www.hds.com
info@hds.com

Asia Pacific and Americas

750 Central Expressway
Santa Clara, California 95050-2627
U.S.A.
Phone: 1 408 970 1000
info@hds.com

Europe Headquarters

Sefton Park
Stoke Poges
Buckinghamshire SL2 4HD
United Kingdom
Phone: + 44 (0)1753 618000
info.eu@hds.com



MK-97DF8124-08