

**Hitachi TagmaStore®
Adaptable Modular Storage
and Workgroup Modular Storage
Copy-on-Write Snapshot Software
User's Guide**

© 2007 Hitachi Data Systems, ALL RIGHTS RESERVED

Notice: 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 Data Systems Corporation (hereinafter referred to as “Hitachi Data Systems”).

Hitachi Data Systems reserves the right to make changes to this document at any time without notice and assumes no responsibility for its use. 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.

This document contains the most current information available at the time of publication. When new and/or revised information becomes available, this entire document is updated and distributed to all registered users.

Trademarks

Hitachi Data Systems is a registered trademark and service mark of Hitachi, Ltd., and the Hitachi Data Systems design mark is a trademark and service mark of Hitachi, Ltd.

TagmaStore, ShadowImage, and TrueCopy are trademarks of Hitachi Data Systems Corporation.

Microsoft, Windows, Windows NT, and Windows Server are registered trademarks or trademarks of Microsoft Corporation.

UNIX is a registered trademark of The Open Group in the United States and other countries.

All other brand or product names are or may be trademarks or service marks of and are used to identify products or services of their respective owners.

Notice of Export Controls

Export of technical data contained in this document may require an export license from the United States government and/or the government of Japan. Please contact the Hitachi Data Systems Legal Department for any export compliance questions.

Document Revision Level

Revision	Date	Description
MK-95DF708-00	June 2005	Initial Release
MK-95DF708-01	August 2005	Revision 1, supersedes and replaces MK-95DF708-00
MK-95DF708-02	December 2005	Revision 2, supersedes and replaces MK-95DF708-01
MK-95DF708-03	February 2006	Revision 3, supersedes and replaces MK-95DF708-02
MK-95DF708-04	March 2006	Revision 4, supersedes and replaces MK-95DF708-03
MK-95DF708-05	March 2006	Revision 4, supersedes and replaces MK-95DF708-04
MK-95DF708-06	May 2006	Revision 6, supersedes and replaces MK-95DF708-05
MK-95DF708-07	June 2006	Revision 7, supersedes and replaces MK-95DF708-06
MK-95DF708-08	July 2006	Revision 8, supersedes and replaces MK-95DF708-07
MK-95DF708-09	October 2006	Revision 9, supersedes and replaces MK-95DF708-08
MK-95DF708-10	January 2007	Revision 10, supersedes and replaces MK-95DF708-09
MK-95DF708-11	May 2007	Revision 11, supersedes and replaces MK-95DF708-10
MK-95DF708-12	June 2007	Revision 12, supersedes and replaces MK-95DF708-11

Changes in this Revision

- Changed product version number to version 7.2A and higher
- Added Chapter 1 Overview of Hitachi Local Replication Solutions
- Added Figure 1.1 Data Protection Options
- Added Chapter 2 Overview of Copy-on-Write Snapshot Software
- Added Chapter 3 Copy-on-Write SnapShot Requirements
- Added reference to Volume Migration product
- Added recommendations for setting differential management LUs in Table 3.1 and Table 7.1
- Added support and restrictions for concurrent use of Volume Migration in Table 3.1 and Table 7.1
- Added information on P-VOL/POOL drive types in Table 3.1 and Table 7.1

Preface

This document describes and provides instructions for performing Copy-on-Write Snapshot operations on the Hitachi TagmaStore® Adaptable Modular Storage and Workgroup Modular Storage array subsystem.

This user's guide assumes the following:

- The user has a background in data processing and understands RAID storage subsystems and their basic functions.
- The user is familiar with the Hitachi TagmaStore Adaptable Modular Storage and Workgroup Modular Storage and Workgroup Modular Storage array subsystem.
- The user has read and understands the *Hitachi TagmaStore Adaptable Modular Storage and Workgroup Modular Storage Command Control Interface (CCI) User and Reference Guide (MK-95DF701)*.
- The user is familiar with the Windows® 98, Windows NT®, and/or the Windows® 2000 operating system.

Notice:

- The use of Copy-on-Write Snapshot software and all other Hitachi Data Systems products is governed by the terms of your license agreement(s) with Hitachi Data Systems.
- Throughout this manual, the term “Snapshot” refers to the Copy-on-Write Snapshot Software program.
- Throughout this manual, the term “Navigator” refers to the Storage Navigator Modular program.
- Throughout this manual, the term “TrueCopy” refers to the TrueCopy Synchronous Remote Replication Software program.
- Throughout this manual, the term “ShadowImage” refers to the ShadowImage In-System Replication Software program.
- Throughout this manual, the term “Volume Migration” refers to the Modular Volume Migration.
- The term “TagmaStore” refers to the Hitachi TagmaStore Adaptable Modular Storage and Workgroup Modular Storage subsystem, unless otherwise noted. Please refer to the *Hitachi TagmaStore Workgroup Modular Storage 100 User and Reference Guide (MK-95DF738)*, or *Hitachi TagmaStore Adaptable Modular Storage 200 User and Reference Guide (MK-95DF713)*, or the *Hitachi TagmaStore Adaptable Modular Storage 500 User and Reference Guide (MK-95DF714)* for further information on your Hitachi TagmaStore Adaptable Modular Storage and Workgroup Modular Storage disk array subsystem.
- The term DF700 refers to the Hitachi TagmaStore Adaptable Modular Storage and Workgroup Modular Storage subsystem.
- For further information, please contact your Hitachi Data Systems account team or visit the Hitachi Data Systems worldwide web site at <http://www.hds.com>.

Software Version

This document revision applies to Hitachi TagmaStore Adaptable Modular Storage and Workgroup Modular Storage Products version 7.2/A and higher.

Release Notes

The Release Notes for this product (located on the installation CD) contain requirements and/or restrictions that may not be fully described in this document. The Release Notes may also contain updates and/or corrections to this document. Make sure to read the Release Notes before installation and use of the product.

Convention for Storage Capacity Values

Storage capacity values for logical units (LUs) are calculated based on the following values:

- 1 KB (kilobyte) = 1,024 bytes
- 1 MB (megabyte) = 1,024² bytes
- 1 GB (gigabyte) = 1,024³ bytes
- 1 TB (terabyte) = 1,024⁴ bytes

Storage capacity values for hard disk drives (HDDs) are calculated based on the following values:

- 1 KB (kilobyte) = 1,000 bytes
- 1 MB (megabyte) = 1,000² bytes
- 1 GB (gigabyte) = 1,000³ bytes
- 1 TB (terabyte) = 1,000⁴ bytes

Referenced Documents

- *Hitachi TagmaStore Adaptable Modular Storage and Workgroup Modular Storage Navigator Modular Graphical User Interface (GUI) User's Guide (MK-95DF711)*
- *Hitachi TagmaStore Adaptable Modular Storage and Workgroup Modular Storage Navigator Modular Command Line Interface (CLI) (MK-95DF712)*
- *Hitachi TagmaStore Adaptable Modular Storage and Workgroup Modular Storage Command Control Interface (CCI) User and Reference Guide (MK-95DF701)*
- *Hitachi TagmaStore Adaptable Modular Storage and Workgroup Modular Storage ShadowImage In-System Replication Software User's Guide (MK-95DF709)*
- *Hitachi TagmaStore Adaptable Modular Storage and Workgroup Modular Storage TrueCopy Synchronous Remote Replication Software User's Guide (MK-95DF710)*

Comments

Please send us your comments on this document. Make sure to include the document title, number, and revision. Please refer to specific section(s) and paragraph(s) whenever possible.

- E-mail: doc.comments@hds.com
- Fax: 858-695-1186
- Mail:
Technical Writing, M/S 35-10
Hitachi Data Systems
10277 Scripps Ranch Blvd.
San Diego, CA 92131

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

Contents

Chapter 1 Overview of Hitachi Local Replication Solutions

1.1	Replication & Disaster Preparedness Business Objectives	2
1.1.1	Disaster Scales	3
1.1.2	Disaster Preparedness Measures	4
1.1.3	RPO and RTO	5
1.1.4	Cost	6
1.2	Replication Concepts	7
1.2.1	Backup and Replication	8
1.2.2	Local Data Replication Options	9
1.2.3	Snapshots	10
1.2.4	Copy-on-Write Technology.....	11
1.2.5	Integrating Local and Remote Replication	12
1.3	Hitachi's Local Replication Product Portfolio.....	13
1.3.1	Copy-on-Write SnapShot Software.....	13
1.3.2	Other Complementary Copy Solutions	14
1.3.3	Hitachi TrueCopy Synchronous Software	14
1.4	Copy-on-Write SnapShot Product Differentiation	15
1.4.1	Product Market Differentiation	15
1.4.2	Differentiation from ShadowImage	17

Chapter 2 Overview of Copy-on-Write Snapshot Software

2.1	Typical Snapshot System Environment	20
2.2	Snapshot Components	21
2.2.1	Snapshot Volume Pairs (P-VOLs and V-VOLs).....	22
2.2.2	Snapshot Data Pools (POOLS)	22
2.2.3	Consistency Group (CTG).....	22
2.3	SnapShot Functional Overview	23
2.3.1	SnapShot Operations.....	26

Chapter 3 Copy-on-Write SnapShot Requirements

3.1	Copy-on-Write SnapShot Requirements	30
3.2	Supported Platforms	33
3.3	Copy-on-Write SnapShot Management Software.....	34
3.3.1	Command Control Interface Software (CCI).....	34
3.3.2	Storage Navigator Modular Software (SNM)	34
3.3.3	Optional Software	35

Chapter 4 Setting up Snapshot

4.1	Supported Snapshot Capacity.....	38
4.1.1	Maximum Supported Capacity of P-VOL and POOL for Each Cache Memory Capacity	38
4.1.2	Maximum Supported Capacity of Concurrent Use of Other Copy Functions ..	42
4.2	Installing and Uninstalling Snapshot	44
4.2.1	Installing Snapshot	44
4.2.2	Uninstalling Snapshot.....	51
4.2.3	Enabling and Disabling Snapshot.....	52

4.3	Determining POOL Capacity	55
4.4	Creating the Configuration Definition File	58
4.5	Setting the Environment Variable.....	61
4.6	Configuring Snapshot.....	62
4.6.1	Setting the POOL and V-VOL.....	62
4.6.2	Path Connection	63
4.6.3	Locating P-VOLs and POOLS.....	64
4.6.4	P-VOLs and POOLS in a RAID Configuration.....	66
4.6.5	Monitoring POOL Capacity	67
4.6.6	Monitoring Transition to PSUE Status	69
4.7	Recommended Pair Assignment	70
4.7.1	Windows 2000/2003 and Dynamic Disk	71
4.7.2	Applying Snapshot to an Existing System	73
4.8	Configuration Restrictions	74

Chapter 5 Performing SnapShot Operations

5.1	Snapshot Operations	78
5.2	Setting the Command Device.....	79
5.3	Setting the Differential Management LU	84
5.4	Setting the POOL	86
5.5	Setting P-VOL and V-VOL.....	89
5.6	Setting Mapping Information.....	93
5.6.1	Selecting Mapping Mode.....	93
5.6.2	Setting Mapping Information	95
5.6.3	Paircreate Operation	99
5.6.4	Updating V-VOL.....	99
5.6.5	Restoration	100
5.6.6	Pair Failures.....	102
5.6.7	Pair Splitting	102
5.7	Differential Data.....	103
5.8	Snapshot Pair Status	104
5.9	Concurrent Use of LUN Expansion	106
5.10	Cascade Configurations.....	108
5.10.1	Cascading TrueCopy with Snapshot.....	108
5.10.2	Cascading with a Snapshot P-VOL.....	109
5.11	Cascade Restrictions	111
5.11.1	Cascade Restrictions with Snapshot	111
5.11.2	Cascade Restrictions on TrueCopy with Snapshot.....	112
5.11.3	P-VOL cascade restrictions	113
5.11.4	V-VOL Cascade Restrictions.....	115
5.11.5	Cascade Restrictions with a Snapshot POOL	116
5.11.6	Cascade Restrictions with Snapshot and ShadowImage	116
5.12	Cascade Connection of TCE with SnapShot.....	117
5.12.1	Cascade Restrictions with P-VOL of TCE.....	119
5.13	Notes and Restrictions	121
5.14	Creating Snapshot Pairs	122
5.15	Updating the V-VOL	124
5.16	Restoring a V-VOL to the P-VOL.....	125
5.17	Releasing Snapshot Pairs	127
5.18	Confirming Pairs Status.....	128

Chapter 6	Troubleshooting	
6.1	Hardware Failure	130
6.2	Data Pool Capacity Shortage.....	131
6.3	Backup Operation for Quick Recovery.....	132
6.4	Online Backup Operation Using an Inexpensive Configuration	133
6.5	Restoring Backup Data.....	134
6.5.1	Backup Operation for Quick Recovery	135
6.5.2	Restoring Backup Data from a Tape Device	136
6.6	Previously Designed Procedure for Recovering Data during Malfunction	138
Chapter 7	Specifications and Function Comparison between Snapshot and ShadowImage	
7.1	Differences between Snapshot and ShadowImage	139
7.1.1	External Specifications	139
7.1.2	Redundancy	148
7.1.3	Snapshot and ShadowImage Functions.....	151
Appendix A	Operations Using CLI	
A.1	Installing Snapshot	153
A.2	Uninstalling Snapshot	157
A.3	Enabling or Disabling Snapshot	159
A.4	Setting Command Device.....	161
A.5	Setting Differential Management LU	163
A.6	Setting the POOL	164
A.7	Setting P-VOL and V-VOL.....	166
A.8	Setting Mapping Information.....	167
A.9	Confirming Status of Pairs	168
Appendix B	Installing SnapShot when Cache Partition Manager is Being Used.....	169
Acronyms and Abbreviations	171
Glossary	173

List of Figures

Figure 1.1	Data Protection Options.....	6
Figure 1.2	SnapShot ShadowImage Product Differentiators	17
Figure 2.1	Snapshot Components	21
Figure 2.2	SnapShot Operation Overview.....	23
Figure 2.3	Shared Data Pool	24
Figure 2.4	Instant Restore using SnapShot	25
Figure 2.5	Restoration from Checkpoints.....	25
Figure 2.6	Before Modification of Data Block on the P-VOL	26
Figure 2.7	After Writing to the Data Block on the P-VOL	27
Figure 2.8	Another Snapshot Image Created	27
Figure 2.9	Write Command to the same Data Block Again.....	28
Figure 2.10	Restore Is Possible from Any Snapshot Image (V-VOLI).....	28
Figure 4.1	Array System Viewer Window (Logical Status Page)	46
Figure 4.2	Install/Unlock Options Dialog Box	46
Figure 4.3	Options Selection Dialog	47
Figure 4.4	Resulting Dialog Box	48
Figure 4.5	Array System Viewer Window (Logical Status Page Option: Enable)	50
Figure 4.6	Array System Viewer Window (Logical Status Page Option: Disable)	54
Figure 4.7	Horcm0.conf Example	59
Figure 4.8	Horcm1.conf Example	60
Figure 5.1	Array System Viewer Window (Command Device Page: Before Setting)	80
Figure 5.2	Command Devices Settings Dialog Box (Before Setting)	81
Figure 5.3	Command Devices Settings Dialog (After Setting)	82
Figure 5.4	Array System Viewer Window (Command Device Page: After Setting)	83
Figure 5.5	Array System Viewer Window (Differential Management LU Page: Before Setting)	84
Figure 5.6	Select Logical Unit Dialog Box.....	84
Figure 5.7	Array System Viewer Window (Differential Management LU Page: After Setting)	85
Figure 5.8	Setting the POOL	86
Figure 5.9	Add Logical Unit Dialog.....	87
Figure 5.10	Added POOL Updates	87
Figure 5.11	Property Dialog Box	88
Figure 5.12	Setting the P-VOL and SnapShot Image (V-VOL).....	89
Figure 5.13	Add Pair Dialog Box	90
Figure 5.14	Adding the P-VOL and SnapShot Image (V-VOL)	91
Figure 5.15	Create Snapshot Image Dialog.....	91
Figure 5.16	Added Snapshot Image Updates.....	92
Figure 5.17	Deleting the SnapShot Image (V-VOL).....	92
Figure 5.18	Array System Viewer Window (Mapping Mode)	93
Figure 5.19	Mapping Mode Dialog Box	94
Figure 5.20	Array System Viewer Window (Setting Mapping Information).....	95
Figure 5.21	Mapping Property Window (Before Setting).....	96
Figure 5.22	Mapping Property Window (After Setting).....	97

Figure 5.23	Creating a Snapshot Pair	99
Figure 5.24	Operation Example when Snapshot Instruction is Issued to the Other V-VOL During the Restoration	101
Figure 5.25	Differential Data	103
Figure 5.26	Snapshot Pair Status Transitions	104
Figure 5.27	Example: Restriction Placed on Composing Unified LU	106
Figure 5.28	Example: Restriction Placed on Uniting LUs with Different Drive Types	107
Figure 5.29	Cascade Connection of TrueCopy with Snapshot	108
Figure 5.30	Cascade Structure: Snapshot P-VOL	109
Figure 5.31	Cascade Structure: Snapshot V-VOL	110
Figure 5.32	Restrictions in Cascade Connection-1	111
Figure 5.33	Configuration Restrictions on Cascade of TrueCopy with Snapshot	112
Figure 5.34	Restrictions in Cascade Connection-2	116
Figure 5.35	Cascade Connection of TCE with SnapShot	117
Figure 5.36	Paircreate -split Command Example	122
Figure 5.37	Paircreate -m grp Command Example	122
Figure 5.38	Pairresync Command Example	124
Figure 5.39	Pairresync -restore command example	126
Figure 5.40	Pairsplit Command Example	126
Figure 5.41	Pairsplit -S Command Example	127
Figure 6.1	Ordinary Quick Recovery Operation	132
Figure 6.2	Ordinary Operation	133
Figure 7.1	Concept	143
Figure 7.2	P-VOL Failures	148
Figure 7.3	POOL (S-VOL) Failures	149
Figure 7.4	POOL (S-VOL) Failures during Restore Operation	150

List of Tables

Table 1.1	Examples of Disaster Scales	3
Table 1.2	SnapShot-ShadowImage Differentiators	17
Table 3.1	Snapshot Requirements	30
Table 3.2	External Specifications	33
Table 4.1	Formula for Calculating Maximum Supported Capacity Value for P-VOL/POOL (WMS100)	38
Table 4.2	Formula for Calculating Maximum Supported Capacity Value for P-VOL/POOL (AMS200).....	38
Table 4.3	Formula for Calculating the Maximum Supported Capacity Value for P-VOL/POOL (AMS500).....	39
Table 4.4	Formula for Calculating the Maximum Supported Capacity Value for P-VOL/POOL (AMS1000)	39
Table 4.5	Supported Capacity Value of the P-VOL/POOL (When Cache Memory is 1 GB/CTL: WMS100).....	39
Table 4.6	Supported Capacity Value for P-VOL/POOL (Cache Memory=1 GB/CTL, AMS200).....	39
Table 4.7	Supported Capacity Value for P-VOL/POOL (Cache Memory=2 GB/CTL, AMS200).....	40
Table 4.8	Supported Capacity Value for P-VOL/POOL (Cache Memory=2GB/CTL, AMS500).....	40
Table 4.9	Supported Capacity Value for P-VOL/POOL (Cache Memory=3 GB/CTL, AMS500).....	40
Table 4.10	Supported Capacity Value for P-VOL/POOL (Cache Memory=4 GB/CTL, AMS500).....	40
Table 4.11	Supported Capacity Value for P-VOL/POOL (Cache Memory = 2GB/CTL, AMS1000)	40
Table 4.12	Supported Capacity Value for P-VOL/POOL (Cache Memory = 4 GB/CTL, AMS1000)	41
Table 4.13	Supported Capacity Value for P-VOL/POOL (Cache Memory = 6 GB/CTL, AMS1000)	41
Table 4.14	Supported Capacity Value for P-VOL/POOL (Cache Memory = 8 GB/CTL, AMS1000)	41
Table 4.15	Single Maximum Supported Capacity of SnapShot (TB)	42
Table 4.16	Recommended Value of the POOL Capacity (When the P-VOL Capacity is 1 TB).....	56
Table 4.17	P-VOL and POOL RAID Configuration	66
Table 4.18	Operation Performed when the Threshold Value is Exceeded	67
Table 4.19	Operation Performed when Pair Status Changed to PSUE	69
Table 5.1	SnapShot Functions	78
Table 5.2	SnapShot Pair Status	105
Table 5.3	Read/Write Instruction to a P-VOL on the Local Side	113
Table 5.4	Read/Write Instruction to P-VOL on the Remote Side	114
Table 5.5	V-VOL Read/Write Instruction on the Local Side (TrueCopy)	115
Table 5.6	A Read/Write Instruction to a P-VOL of SnapShot on the Local Side (TCE)....	119

Table 5.7	A Read/Write Instruction to a P-VOL of SnapShot on the Remote Side (TCE) .	120
Table 5.8	Time Required to Examine Differential.....	125
Table 5.9	Icons for Pair Status	128
Table 7.1	External Specifications	139
Table 7.2	Characteristics of FC Drive and SATA Drive	142
Table 7.3	Single Maximum Supported Capacity of SnapShot Function (TB)	144
Table 7.4	Formula for Calculating Maximum Supported Capacity Value for P-VOL/POOL (AMS200).....	145
Table 7.5	Formula for Calculating Maximum Supported Capacity Value for P-VOL/POOL (AMS200).....	145
Table 7.6	Formula for Calculating the Maximum Supported Capacity Value for P-VOL/POOL (AMS500).....	145
Table 7.7	Formula for Calculating the Maximum Supported Capacity Value for P-VOL/POOL (AMS1000)	145
Table 7.8	Snapshot and ShadowImage Functions.....	151

Chapter 1 Overview of Hitachi Local Replication Solutions

This chapter discusses local replication concepts; provides an overview of Hitachi's local replication solutions and an introduction to Hitachi's Copy-on-Write SnapShot solution. This chapter includes the following sections:

- Replication & Disaster Preparedness Business Objectives (section 1.1)
- Replication Concepts (section 1.2)
- Hitachi's Local Replication Product Portfolio (section 1.3)
- Copy-on-Write SnapShot Product Differentiation (section 1.4)

1.1 Replication & Disaster Preparedness Business Objectives

Today, businesses are increasingly dependent on the availability of critical data resources. Most enterprises require access to business data 24/7 and cannot tolerate a downtime of more than a few hours, without serious impact to the bottom line. Beyond the loss of revenue, there are adverse headlines and the potential impact on company valuation to consider, along with lower employee productivity caused by sporadic outages.

Governments have also established mandates, requiring corporations to comply with business data protection regulations. For example, United States regulations on data protection now apply to health care (HIPAA), financial services (SEC 17a4), corporate accountability (Sarbanes-Oxley Act), life sciences (21 CFR Part 11), and government (DoD 5015.2-STD). Beyond government regulators, investors and even insurers insist that businesses should establish feasible disaster recovery and business continuity plans, to protect critical information. In addition, organizations constantly face increasing internal and external threats to the stability of information systems.

In today's extremely competitive business environment, organizations cannot afford to ignore the impact of all the above factors. Since business continuity and high availability have become the top concerns of businesses globally, it has become important for them to address the area of disaster preparedness and proactively plan for successful disaster recovery. Data protection and preparedness measures are often undertaken by analyzing the following business objectives:

1. Disaster scale: What are the assumptions for the scale of a disaster or failure?
2. Recovery Point Objective: Up to what point in time, should data be recovered in the event of a disaster?
3. Recovery Time Objective: How soon business operations should be resumed after the disaster?
4. Cost: How much can be budgeted for preparedness measures?
5. Backup window: How long can a backup process take to complete without impacting critical application processing?
6. Consistency groups: The number of multiple applications that need to be restored to the same point in time.

The overriding goals of any business continuity plan must be to survive a disaster and resume operations as quickly as possible. The best recovery plan to achieve these goals depends on how an organization chooses to prioritize the above business objectives. The following sections discuss each of the above business objectives in more detail.

1.1.1 Disaster Scales

Before determining necessary preparedness measures, assumptions must be made about the scale of the disaster. Disaster scales describe the scope of a disaster, in terms of damage to information systems and the overall business environment. These assumptions often determine how information systems need to be configured.

Table 1.1 shows examples of disaster scales. For example, when a file is accidentally erased, the damage is limited only to the specific file. On the other hand, in the event of a disaster such as an earthquake, the damage can easily extend to an entire city or region.

Table 1.1 Examples of Disaster Scales

Type of Disaster	Example	Disaster Scale
Data destruction	File erased by mistake	File
Drive failure	Dual failure on a drive on which RAID is being reconfigured	Logical Unit
Subsystem failure	Hardware failure in controller, etc.	Subsystem
Site disaster	Fire in building, etc	Site (within some hundreds of meters)
Natural disaster	Earthquake, hurricane, etc.	Extensive (hundreds of kilometers or more)

1.1.2 Disaster Preparedness Measures

Businesses are continuously faced with risks associated with catastrophes resulting from natural calamities, human operational errors, planned or unplanned downtime, security and terrorist incidents, content or application third-party dependencies, increased dependency on heterogeneous technologies etc. To eliminate these risks, it is necessary for organizations to undertake basic disaster preparedness measures. There are two fundamental types of disaster preparedness measures:

- Data protection measures
- Disaster recovery measures

Data protection measures: These measures are steps taken to guard critical data against physical destruction (such as when a dual failure occurs on a drive), logical destruction (such as when a file is accidentally erased), site failures or natural disasters. One way to protect data is to periodically create data backups on physically independent secondary storage devices. If a primary data storage device is physically destroyed, most of the data can be retrieved from the latest backup data in the independent secondary storage device.

Another way to protect data is to create snapshots of primary data on the local storage device, at specific time intervals. In the case of logical destruction in the storage device, data can be restored from the point-in-time snapshots. Backup data can also be created at a remote location as a protection against site failures or natural disasters. This type of backup is called remote replication.

Disaster recovery measures: These measures are contingency plans that can be implemented in the event of a disaster. These measures help achieve a high RTO and may involve creating a clustering configuration to ensure work can be resumed using the copied data, immediately after a primary site failure is detected. Recovery measures may also involve establishing backup sites far from the location of a possible disaster.

Once the above business objectives are addressed and risks established, businesses can design resilient information systems, to ensure quick recovery time of critical application processing. Assessment based on the discussed business objectives also helps determine which applications require the most protection and the business impact of unexpected downtime.

1.1.3 RPO and RTO

Recovery Point Objective (RPO) and the Recovery Time Objective (RTO) are parameters that define how quickly an organization needs to recover to survive a disaster and how much data loss can be tolerated.

Recovery Point Objective (RPO) is defined by the Storage Networking Industry Association (SNIA) as follows:

- The maximum desired time period prior to a failure or disaster during which changes to data may be lost as a consequence of recovery. Data changes preceding the failure or disaster by at least this time period are preserved by recovery. Zero is a valid value and is equivalent to a “zero data loss” requirement.

In other words, RPO is a target value that specifies the point in time starting from which data needs to be recovered, in the event of a failure or disaster. It answers the question: How much data can your business afford to lose? For example, the RPO value for temporary data generated during a calculation is considered low because such data can be generated again if the calculation can be made again. On the other hand, if the transaction data of an Internet shopping site is lost, the image of the business suffers and, consequently, the RPO value is considered high.

Recovery Time Objective (RTO) is defined by the Storage Networking Industry Association (SNIA) as follows:

- The maximum desired time period required to bring one or more applications and associated data back to a correct operational state.

In other words, RTO defines the time frame within which specific business operations or data must be restored to avoid any business disruption. It answers the question: how long can a business afford to be down? RTO values for data are determined by how indispensable the data is for continuous site operation. For example, although auditing data is important for business operation and the RPO value is high, the RTO value of this data is considered low if the data is used only for auditing. On the other hand, the RTO value of the transaction data mentioned in the RPO example above is considered high because such data is indispensable to business continuity.

Both RTO and RPO influence the data replication and recovery options an organization needs to adopt. The Recovery Time Objective combined with the Recovery Point Objective of business data, can be mapped to a range of technical approaches, associated costs, and degrees of data protection.

1.1.4 Cost

Cost is the monetary requirement for building and operating a system which achieves its disaster scale, RPO, and RTO targets. Cost and the RPO and RTO values exist in a trade-off relationship. For example, cost increases as the loss of data in preparation for an extensive disaster is minimized and the target for a quick recovery is achieved. Recovery-time objective combined with a recovery point objective based on the value of particular data, map to a range of technical approaches, costs and degrees of data protection. The following figure illustrates the data protection choices based on cost versus value of data.




Figure 1.1 Data Protection Options

When determining cost it is also essential to analyze the cost associated with downtime. This cost can be estimated by determining the effects of downtime on factors such as employee productivity, business reputation, repair expenses, revenue and financial performance. Business processes and services must also be classified based on mission critical functions (customer-facing, revenue-producing, back-office etc) when determining business continuity and disaster recovery strategies.

1.2 Replication Concepts

An essential component of business continuity is the ability to quickly replicate and restore data. Replication provides the first line of protection and involves processes that create and maintain copies of data volume in the storage system. Replicated data can be leveraged for a variety of needs, well beyond those of data recovery, addressing such business challenges as backups, change/control, maintenance, for analytical purposes, etc. More importantly, using data replicas on a cyclical basis protects data from corruption, and allows for reverting to an uncorrupted version of data when needed.

Data replication increases data availability in the following ways:

- Allowing non-disruptive backup of current production data with no impact to the production application
- Speeding data restoration by replacing slow and labor-intensive tape-based restores with continuously available online backups
- Enabling frequent, non-disruptive disaster recovery testing with an online copy of current and accurate production data

The following sections discuss data protection concepts, local replication options and how local and remote replication solutions can be combined to provide near continuous data availability.

1.2.1 Backup and Replication

Traditional tape backup processes can be labor-intensive and span several hours. In a world of just-in-time delivery, global supply chains, around-the-clock customer demand, and competition, businesses cannot afford to gamble with the threat of planned or unplanned downtime. These tasks must be performed without impacting production data or causing business downtime.

Traditional tape backup also poses the following disaster recovery challenges:

- Inability to meet some requirements with really large databases and file systems
- Inability to restore multiple systems to the same point in time
- Difficult to back up remote offices
- Involve risks associated with physical media leaving business facility

In addition, most organizations deal with the following two conflicting demands:

- Necessity to backup increasing amounts of data
- Necessity to meet the business requirement for continuous data availability

The above challenges have led data protection and disaster recovery planners to switch to “replication” for protecting data and restoring critical applications. Addressing the need for continuous uptime and data availability, storage hardware manufacturers have responded with exceptionally fast server-less split mirror backup solutions. Software-based solutions using copy-on-write snapshot technology are also now available to address the same issues. These solutions are slower than hardware based backup and restore solutions but are faster than conventional backup procedures. Some backup implementations combine both software and hardware methods to provide more versatile solutions.

1.2.2 Local Data Replication Options

Local data and replication involves creating and maintaining replicas of primary data volumes on local storage subsystems. There are primarily two techniques of local data replication:

- Mirroring or Full volume replication
- Differential data replication

Full Volume Replication: Full volume replication uses split mirror technology which mirrors or clones a set of data volumes or drives to an existing alternate mirrored set of drives, often using a SAN. The mirrored data volumes can be rapidly synchronized and split from the original data volumes after replication. This replication technique creates multiple full volume point-in-time copies (snapshots) of critical information, on local storage subsystems or secondary storage mediums.

Differential Data Replication: Differential data replication uses copy-on-write snapshot technology to capture point-in-time images (snapshots) of differential data in a data volume. Differential data can be recently updated data (in a primary data volume) resulting from host I/O operations or residual data (in a secondary data volume) resulting from a previous copy operation. These point-in-time images are temporarily stored in a group of one or more designated disk volumes (data pools), and made available to replication functions (during periodic copy cycles) or to restoration functions (during recovery processes). This replication technique allows for accurate restoration of original data volumes or faster takeover processing using the mirrored volumes.

Your choice of data replication method should depend on the varying data protection requirements of your different applications and your organization's data retention needs. When evaluating and understanding data protection techniques and alternatives, it is necessary to take the following factors into consideration:

- Applicable threats to your applications and data
- RTO and RPO requirements associated with your applications
- Hardware and software technologies already in place
- Usability of data protection solution's interface
- Budget, timeframe, tolerance to disruption and risk aversion

1.2.3 Snapshots

In a storage environment, a snapshot is a point-in-time copy of data generated for backup or replication purposes. Snapshots are images representing the state of a storage volume at a particular point in time. Snapshots can be either partial copies of changed data (copy-on-write snapshots) or full volume copies (split-mirror snapshots, clones or mirrored volumes).

Copy-on-Write Snapshots: Copy-on-write snapshots are instantaneous point-in-time copies of data that has changed on the original storage volumes. Copy-on-write snapshots are virtual images (views) of the original data volume, generated by referencing both the original data volume and the changed data blocks in the data pool. Each virtual snapshot comprises of references to data that exists partly in the data pool and partly in the original/primary volume.

It is required to maintain multiple generations of snapshots to effectively restore or recover data. Generations of copy-on-write snapshots can be stored in designated physical disk volumes (pools). A significant savings of storage space is realized when using copy-on-write snapshots instead of full cloning methods, since the amount of storage space required for retaining each snapshot is substantially smaller than the storage space required for storing full volume copies. Since snapshots are available instantly they are useful for speedy data restoration or recovery after data loss due to file system corruption, application malfunction or erroneous write operations.

Split Mirror Snapshots: Split-mirror snapshots are volume captures of all data in the original storage volume. All data blocks in the original data volume are sequentially copied onto a secondary volume. The original and secondary data volumes are synchronized so that all updates to the original data volume are continually mirrored to the secondary data volume. Full volume copies can be split from the original data volumes. When they are split, they are essentially mirror images of each other at that point in time. The secondary volume can be used for offline testing or analytical purposes since there is no common data sharing with the original volume. Since there are no dependencies between the original and secondary volumes, these replicas allow for more simplified data recovery and archival processes.

It generally takes longer to create full volume captures and each split mirror snapshot requires more storage space than copy-on-write snapshots. Since they are physical copies of data volumes, they require the same amount of storage space as the original data volume.

1.2.4 Copy-on-Write Technology

Copy-on-Write SnapShot technology uses software utilities to continually monitor write operations to a storage subsystem, and copy differential data (resulting from writes to the original volumes) onto a designated storage volume (data pool). Virtual snapshots (views) are captured at particular instants in time when a snapshot creation instruction is processed. These snapshots are maintained in virtual volumes (V-VOLs) and comprise of pointer references to physical data in the original volume and the data pool.

This technique derives its name (copy-on-write) from the fact that there is a copy operation associated with every write operation received by the storage medium. This copy operation is performed first, even before the original data volume is updated with new or changed data. When a host issues a write operation to the original data volume, the changed data block or new data block is first copied to the data pool. Following this copy operation, pointers associated with existing (previous) snapshots in the V-VOLs, are updated to reference the new or changed data block in the data pool. The write operation is then completed by updating or creating the data block on the original data volume.

The copy-on-write replication technique replaces older virtual snapshot images on the virtual volumes when new snapshots are generated, maintaining only predetermined generations of point-in-time snapshots. This technique involves processes that run in the background without demanding server resources and impacting application performance. Copy-on-write technology is usually used in combination with other hardware replication options to implement backup to secondary storage volumes.

Snapshot technologies can be used for the following purposes:

- To simplify and speed up restoration after data loss, data corruption or failure to backup critical data
- To make copies of data for test purposes including software development, regression testing etc
- To make copies of data for application processing such as data warehousing, reporting and data mining
- To make copies for non-disruptive backups and data migration

1.2.5 Integrating Local and Remote Replication

Local replication functions (such as snapshot or image backups) can be combined with remote data replication functionality to provide extremely resilient and powerful replication solutions. The ability to integrate these functions differentiates local data replication technology from traditional backups. The static point-in-time frozen replicas (snapshots) of data can be transferred to the secondary storage mediums during different types of remote backup operations. The multiple snapshots created prior to the onset of a disaster, are useful for data recovery, particularly from rolling disasters. Frozen virtual copies created at frequent checkpoints, allow fast data restoration and rapid recovery of critical data.

Note: Copy-on-Write Snapshot and ShadowImage software can also be active simultaneously in the same subsystem.

1.3 Hitachi's Local Replication Product Portfolio

Hitachi's replication solutions enhance operational efficiency and resilience of your information systems. Using the available solutions you can implement customized business continuity solutions, addressing key business problems relating to disaster recovery, data protection and availability. Different solutions are available for the large-size TagmaStore Universal Storage Platform (USP/NSC) subsystems and the mid-range TagmaStore AMS/WMS subsystems.

Small and medium-sized businesses, facing the same business continuity challenges as their enterprise-level counterparts, can keep data assets protected and available with the local replication solutions available from Hitachi Data Systems for mid-range TagmaStore AMS/WMS subsystems. The portfolio of midrange local replication solutions includes the Copy-on-Write SnapShot solution (hereafter referred to as SnapShot) and the ShadowImage In-system Replication solution, either of which can be leveraged depending on an organization's data protection objectives. The following sections discuss the Copy-on-Write SnapShot local replication solution for mid-range TagmaStore AMS/WMS subsystems and provide an overview of the other Hitachi copy solutions.

1.3.1 Copy-on-Write SnapShot Software

Hitachi's Copy-on-Write Snapshot software uses a high-speed, non-disruptive, logical snapshot technology to rapidly create up to 15 point-in-time snapshot copies of any data volume within AMS/WMS subsystems. It is a local replication solution that creates instantaneous virtual backup copies without impacting host service or performance levels. These snapshots significantly reduce backup time and accelerate recovery processes and are suitable for immediate use in decision support, software testing and development, data backup, or rapid recovery operations.

Copy-on-Write SnapShot reduces critical application downtime. This replication solution 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 substantially smaller than the original data volume. As a result, Copy-on-Write Snapshot software can provide significant savings over full cloning methods.

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 Copy-on-Write Snapshot V-VOL
- Periodic point-in-time disk copies for rapid restores in the event of a corrupted data volume

When integrated with advanced disaster recovery techniques, it provides the ability to create high-availability solutions for quick restoration of production environments, in the event of a disaster.

1.3.2 Other Complementary Copy Solutions

The following section provides a brief overview of some of Hitachi's other local and remote data replication technologies.

1.3.2.1 Hitachi TrueCopy™ Software

Hitachi TrueCopy™ software is the industry's first storage hardware-based, server-less, asynchronous, remote copy solution that replicates data over any distance. TrueCopy remote data replication technology can be used in conjunction with local replication functions such as Copy-on-Write SnapShot software or ShadowImage software, to provide versatile data protection and disaster recovery solutions.

1.3.2.2 Hitachi TrueCopy Extended Distance Software (TCE)

Hitachi has integrated its TrueCopy remote replication technology with Copy-on-write Snapshot operations, to provide TrueCopy Extended Distance (TCE), a continuous, non-disruptive, host-independent asynchronous remote replication and disaster recovery solution for midrange Hitachi TagmaStore® Adaptable Modular Storage models AMS500 and AMS1000. TrueCopy Extended Distance enables open-system users to perform remote copy operations between TagmaStore storage sub-systems in geographically distant locations, using fibre-channel interface connections.

1.3.3 Hitachi TrueCopy Synchronous Software

For distances within the same metropolitan area, TrueCopy Synchronous Remote Replication software (TrueCopy) provides a no data loss, rapid restart solution. TrueCopy provides the ability to copy data over short and medium distances, by utilizing local clones and data movement capabilities of TrueCopy™ technology. Combined with Universal Replicator and the Universal Storage Platform, TrueCopy allows for advanced three data center configurations for optimal data protection.

1.3.3.1 Hitachi ShadowImage In-system Replication Software

Hitachi's ShadowImage is an in-system replication software that uses a high-speed, undistruptive, local mirroring technology to rapidly create multiple full-volume copies or clones of mission-critical information within storage subsystems, while keeping data RAID-protected and fully recoverable. ShadowImage can be combined with TrueCopy to provide a versatile remote replication solution with maximum flexibility in data backup and restoration.

1.4 Copy-on-Write SnapShot Product Differentiation

Hitachi's Copy-on-write SnapShot is a data protection solution for the following demand drivers and customer requirements:

- Increase storage utilization and ROI
- Meet objectives for availability and performance
- Develop new applications faster
- Provide reliable access to, and the sharing and safeguarding of critical business information
- Protect more data with less effort
- Execute tape backups while critical applications remain productive and online
- Provide current test data for application development without disrupting production
- Cost effectively maintain frequent point-in-time copies on disk for faster recovery
- Provide nondisruptive copies of online databases for decision support

The following sections provide competitive benefits and differentiate Copy-on-Write SnapShot from Hitachi ShadowImage In-system Replication solution.

1.4.1 Product Market Differentiation

Hitachi's Copy-on-Write SnapShot offers the following competitive business continuity benefits as well as productivity and process improvements:

Ensures Business Continuity and Data Availability

- Provides non-disruptive, volume "snapshots" which use less space than full copies or "clones"
- Allows frequent, cost-effective, point-in-time copies (up to 15 virtual copies of a primary volume)
- Shortens restart and recovery times with the consistency group function, which provides multivolume, point-in-time copies for applications and databases that share or span multiple volumes
- Dramatically reduces recovery from data corruption or human error through immediate restores from a disk-resident, point-in-time data copy
- Eliminates the backup window and enables immediate read/write access to time-critical information using virtual copies

Improves Productivity and Processes

- Reduces testing and deployment time and increases the accuracy of application development by providing always-available copies of current production data
- Increases competitive advantage by facilitating the sharing of and immediate access to time-critical information for decision support, populating data warehouses, performing analysis or other data mining operations
- Enables normal backup operations on a copy of up-to-date production data while critical applications run unaffected
- Improves operational efficiency by allowing multiple processes to run in parallel with access to the same information
- Improves backup efficiency and eliminates costly downtime due to tape backups
- Offers shorter development cycles which translates into faster time-to-market for new business processes
- Accelerates application testing and development and provides fast access to data for decision support

Reduces Operational and Capital Costs

- Allows disk space to be used efficiently through volume snapshots rather than full-copy clones, which helps maximize disk utilization and improved ROI
- Allows business to remain online during data center activities, eliminating the need for 24/7 resources to perform these tasks
- Enables IT to refocus resources onto revenue-driven projects
- Enables fast restore and rapid recovery with frequent point-in-time checkpoints

Note on product pricing: Copy-on-Write Snapshot software product pricing is on an array basis for midrange AMS/WMS systems. For RAID, pricing is tiered based on the total of source and Pool Groups (target volumes) of usable storage capacity.

1.4.2 Differentiation from ShadowImage

Hitachi ShadowImage™ copies differ from Copy-on-Write snapshots in that they are physically separate copies of the data and no portion is shared between primary and secondary volumes, therefore performance and protection issues are resolved.

Table 1.2 SnapShot-ShadowImage Differentiators

	ShadowImage Clone	Copy-on-Write Virtual Volume
Features	Superior protection because it is a complete copy Immediate restore if P-VOL becomes corrupted	Provides a very quick copy because it is a virtual volume
Time to Create	From minutes to hours, depending on the size of the P-VOL	Instantaneous
Disk Space Used	Same amount as P-VOL	Size will vary depending upon rate of data change but will be far less than the P-VOL
Data Recovery Time After P-VOL is corrupted	Instantaneous	Two stage restore: 1. V-VOL verify, data not available to host 2. Copy back, data available

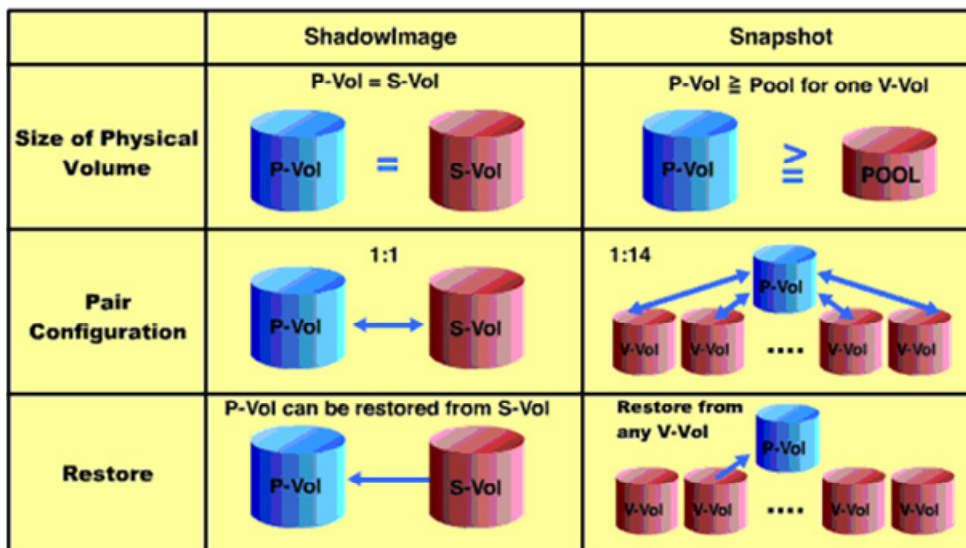


Figure 1.2 SnapShot ShadowImage Product Differentiators

Note on product availability: Copy-on-Write SnapShot and ShadowImage solutions can be either purchased separately or as a competitively priced in-system replication bundle.

Chapter 2 Overview of Copy-on-Write Snapshot Software

This chapter presents an overview of Copy-on-Write SnapShot software and discusses the functional and operational details of the product. This chapter includes the following sections:

- Typical Snapshot System Environment (section 2.1)
- Snapshot Components (section 2.2)
- SnapShop Functional Overview (section 2.3)

2.1 Typical Snapshot System Environment

A typical Copy-on-Write SnapShot hardware configuration includes a TagmaStore AMS subsystem (local subsystem or remote subsystem), a host connected to the TagmaStore subsystem and a management server. The host is connected to the TagmaStore AMS subsystem via fibre channel connections. The management server is connected to the TagmaStore AMS subsystems via a management LAN.

The logical configuration of the TagmaStore AMS subsystem includes a command device, a differential management logical unit (DM-LU), primary data volumes (P-VOLs) belonging to the same consistency group (CTG), virtual volumes (V-VOLs) and a logical unit for the data pool. Copy-on-Write Snapshot creates a volume pair from a primary volume (P-VOL), which contains the original data, and a Snapshot Image (V-VOL), which contains the snapshot data. Copy-on-Write Snapshot uses the V-VOL as the secondary volume (S-VOL) of the volume pair. Since each P-VOL is paired with its V-VOL independently, each volume can be maintained as an independent copy set.

In addition to the above components, the Copy-on-Write SnapShot system architecture also includes the Storage Navigator Modular (SNM) software and the Command Control Interface (CCI) software. Storage Navigator Modular is installed on the management server and is used to configure the Copy-on-Write SnapShot system environment. Command Control Interface software is installed on the host and manages Copy-on-Write SnapShot volume pair operations. Snapshot operations can be performed from a UNIX and/or PC server host using CCI.

2.2 Snapshot Components

Snapshot operations involve the primary volumes (P-VOLs), Snapshot Images (V-VOL), POOL in the TagmaStore subsystem, the Navigator, and the Command Control Interface (CCI). Figure 2.1 shows a typical Snapshot configuration. The Snapshot system components include:

- Snapshot volume pairs (P-VOLs and V-VOLs) (see section 2.2.1)
- Snapshot Data Pools (POOLS) (see section 2.2.2)
- Consistency Group (CTG) (see section 2.2.3)

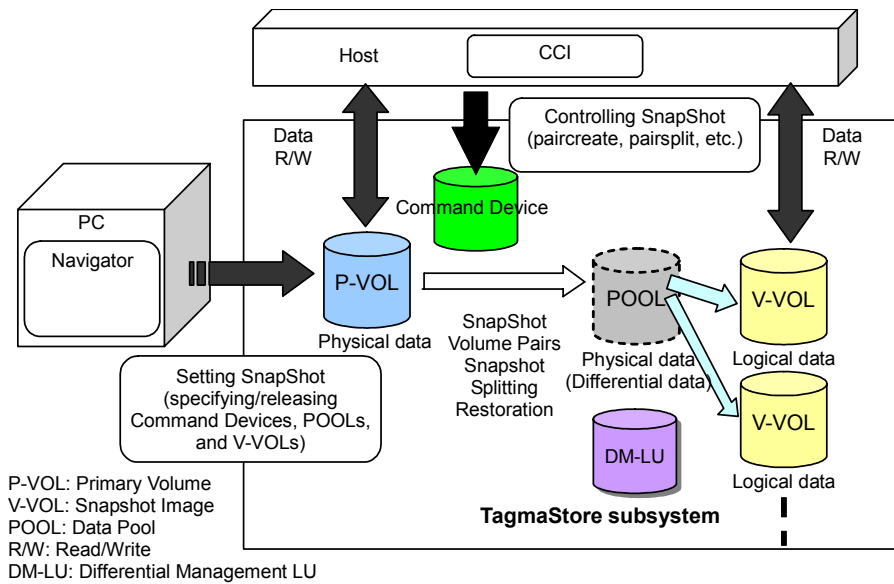


Figure 2.1 Snapshot Components

Note: Snapshot requires a dual controller configuration.

2.2.1 Snapshot Volume Pairs (P-VOLs and V-VOLs)

The TagmaStore AMS subsystem manages both the original and copied data that exist at the time of Snapshot instruction. WMS100/AMS200 supports up to 254 Snapshot pairs (AMS500: 1,022 pairs, AMS1000: 2,046 pairs). Up to 254 (1,022/2,046) P-VOLs can be created.

Each Snapshot pair consists of one primary volume (P-VOL) and up to 15(14: the micro program earlier than version 0750/A) (V-VOLs), which are located in the same TagmaStore AMS subsystem. The P-VOLs are the primary volumes, which contain the original data. The V-VOLs are duplicated volumes which contain the data that exists from the moment the Snapshot instruction is executed.

2.2.2 Snapshot Data Pools (POOLs)

The Snapshot Image (V-VOL) contains physical data from the primary volume and differential data stored in the POOL. (This differs from the ShadowImage function where all the data is retained in the secondary volume only; it is actually a pseudo volume with no capacity.) One POOL can be set for each controller. The POOL can handle two or more primary volumes and the differential data of two or more V-VOLs, sharing them for the each controller.

2.2.3 Consistency Group (CTG)

There is a case where a file system that stores the application data or a logical volume on an OS is configured with two or more logical units. In this case, it is required to assure that the data of those logical units are of the same time. In order to assure that the data of two or more Snapshot Images are of the same time and to manage them as a group, the TagmaStore AMS subsystem provides the consistency group (CTG). The TagmaStore AMS subsystem assures that data of the Snapshot Images in the group are of the same time by splitting the pairs in the group using the CTG. When the CTG is not used, the timing to split each Snapshot Image varies because the pairs in the group are split in order. Therefore, if an I/O instruction is received from a host in the middle of the splitting of the pairs in the group, data of the different time is stored in each Snapshot Image.

To use the CTG when splitting a pair, it is required to create the pair adding the `-m grp` option at the time of the pair creation.

2.3 SnapShot Functional Overview

This section provides a functional overview of SnapShot operations.

Virtual snapshots (views) are captured at particular instants in time when a snapshot creation instruction is processed. These snapshots are maintained in virtual volumes (V-VOLs) which comprise of pointer references to physical data. Each virtual volume (V-VOL) maintains a view of the primary volume (P-VOL) at a particular point in time and is a composite of data residing in the P-VOL and the data pool. Since the V-VOL does not copy all data it can be created or deleted nearly instantly. Each V-VOL uses only a fraction of the P-VOL storage space.




Figure 2.2 SnapShot Operation Overview

The SnapShot data pool is shared by multiple P-VOLs and V-VOLs allowing for maximum storage utilization and improved ROI.

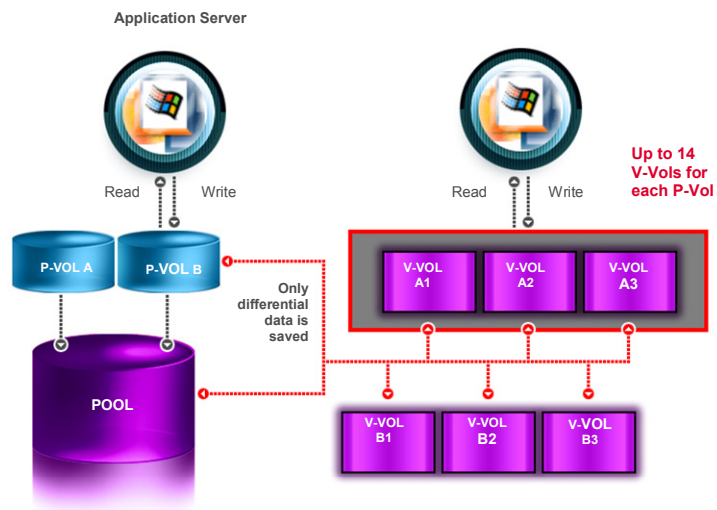


Figure 2.3 Shared Data Pool

SnapShot enables data restoration from any point-in-time image (V-VOL). Data is copied in the background and restored data is available to the host after the V-VOL certification is completed. The V-VOL appears as a full volume copy to any secondary host.

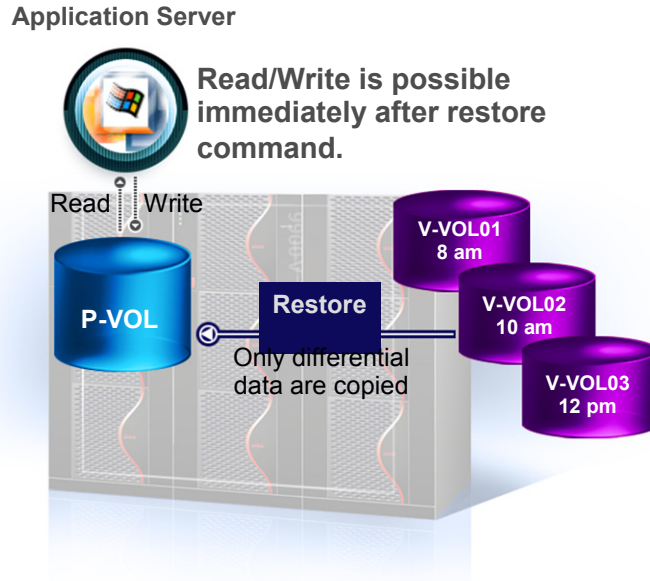


Figure 2.4 Instant Restore using SnapShot

Rapid data restoration is made possible using periodic checkpoint copies stored on the disk. Data can be restored from any previous checkpoint depending on the number of generations of snapshots retained in the subsystem. Figure 2.5 shows how transactions can be reapplied from multiple checkpoints before the instant at which data corruption occurs.

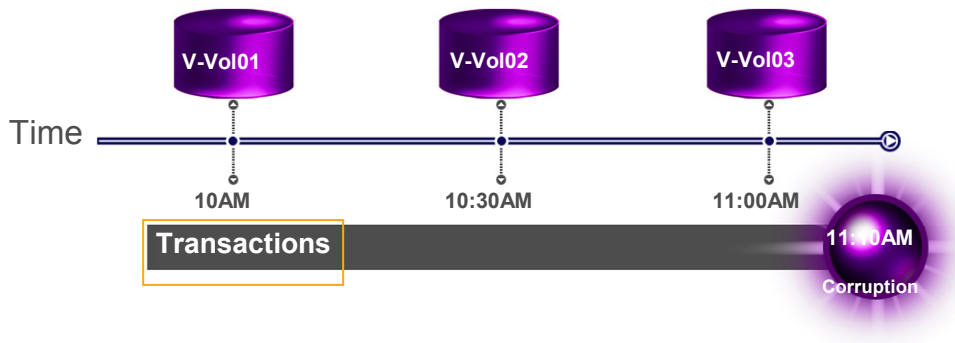


Figure 2.5 Restoration from Checkpoints

2.3.1 SnapShot Operations

This section explains the internal operations used to implement the SnapShot functionality using a typical SnapShot usage scenario. Figure 2.6 illustrates a scenario in which two snapshots of the P-VOL data have been created. When access to any data block in a snapshot is requested through the V-VOL, it is made available from the P-VOL. This situation lasts as long as the corresponding block on the P-VOL is not altered.




Figure 2.6 Before Modification of Data Block on the P-VOL

When there is a write operation from the host to a data block, before the actual write to the P-VOL is executed, that data block is first saved in the data pool. The data pool and snapshot images in the V-VOL are linked by addresses (pointer references) managed in the cache. After the original block is copied to the pool, the set of pointers that represent the V-VOL are updated to point to the copy of the original data block in the pool instead of the original data block in the P-VOL. Any request for the original block through a V-VOL is now physically taken from POOL, until the next snapshot is created. From the host's perspective, the V-VOL (snapshot image) seems to be unchanged.

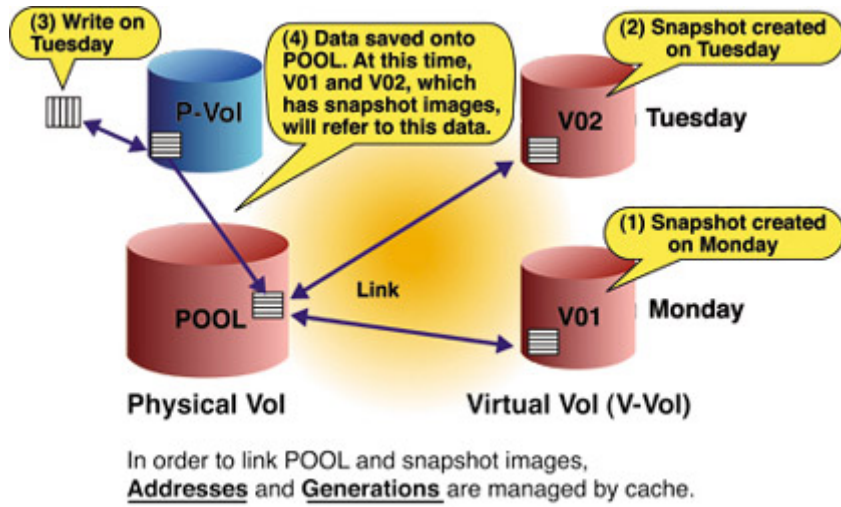


Figure 2.7 After Writing to the Data Block on the P-VOL

The write operation to the P-VOL is completed by physically updating the data block on the P-VOL. When another snapshot of the P-VOL is generated, the corresponding V-VOL pointers are updated to reference the original P-VOL until another write operation is executed. When a particular data block is requested through the V-VOL, data will physically be read from the data pool or from the P-VOL depending on which snapshot is requested. Figure 2.8 illustrates a sample scenario.




Figure 2.8 Another Snapshot Image Created

When there is another write to the same data block, the pointer references of the previous V-VOL (previously linked to the P-VOL) are again updated to point to the copy of data block created in the data pool. Figure 2.9 illustrates this scenario.

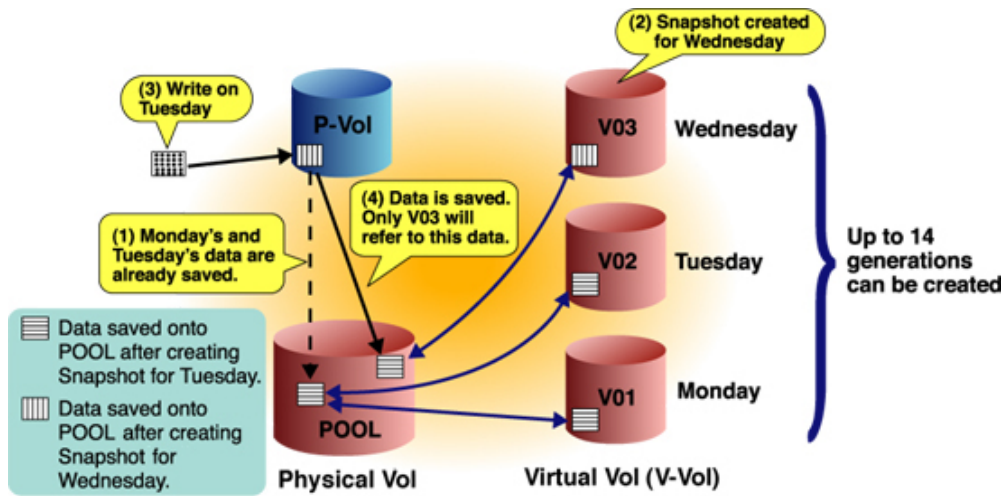


Figure 2.9 Write Command to the same Data Block Again

When P-VOL data restoration is required, any of the previous generations of snapshots in the V-VOLs can be used. Only the differential data in the corresponding V-VOL is copied onto the P-VOL during the restore operation. The restoration operation is performed in the background and is completely invisible to the host.



Figure 2.10 Restore Is Possible from Any Snapshot Image (V-VOLI)

Chapter 3 Copy-on-Write SnapShot Requirements

This chapter describes Copy-on-Write SnapShot operational requirements and provides an overview of Copy-on-Write SnapShot management software. This chapter includes the following sections:

- Copy-on-Write Snapshot Requirements (section 3.1)
- Supported Platforms (section 3.2)
- Copy-on-Write Snapshot Management Software (section 3.3)

3.1 Copy-on-Write SnapShot Requirements

Snapshot operations provide subsystem-internal copies of UNIX®/PC server volumes (LUs) on the TagmaStore AMS subsystem. Table 3.1 describes Snapshot operational requirements.

Table 3.1 Snapshot Requirements

Parameter	Requirement
User interface	The following software is necessary: CCI: used for issuing commands to create and split Snapshot pairs Navigator: used for specifying and releasing Command Devices, Data Pools, and P-VOLs
Cache memory	WMS100: 1 GB/CTL AMS200: 1, 2 GB/CTL AMS500: 2, 3, 4 GB/CTL AMS1000: 2, 4, 6, 8 GB/CTL
Number of controllers	Use the dual controller configuration.
Max number of Command Devices	Up to 2 per disk subsystem can be set.
Max number of Differential Management LUs	Up to 2 per disk subsystem can be set. The Differential Management LU size must be greater than or equal to 5 GB. We recommended that you set two Differential Management LUs according to following conditions. <ul style="list-style-type: none"> ▪ To be created in different RAID group ▪ To be allocated in different controllers
Unit of pair management	LUs are the target of Snapshot pairs and are managed per logical unit.
Max number of pairs	Up to 254 pairs (254 P-VOLs and 254 V-VOLs) per subsystem can be created (AMS500: 1,022 pairs, AMS1000: 2,046 pairs).
Pair structure	Up to 15(14: the micro program earlier than version 0750/A) copies (V-VOLs) per P-VOL can be created.
Supported RAID level	RAID 5 (2D+1P to 15D+1P), RAID 1 (1D+1D), RAID 1+0 (2D+2D to 8D+8D), and RAID 6 (2D+2P to 28D+2P)
Combination of RAID levels	All combinations are supported. The number of data disks does not have to be the same.
LU size	Make the LU size of the P-VOL the same as or larger than that of the POOL. (The "POOL," in this case, means a single POOL capacity used by a single V-VOL.) The Command Device LU size must be greater than or equal to 33 MB.
POOL to be used	The POOL that is common to TCE is used. The subsystem must be restarted in order to secure the POOL resources. When the POOL is already used by TCE, a restart is not needed.
P-VOL/POOL drive types	Both LUs with FC drives and SATA drives can be assigned to any P-VOL and a POOL. However, an LU with FC drives assigned to a P-VOL and a POOL is recommended. When creating a pair with the LUs configured in the SATA drive, the use conditions of the SATA drive may be different. (See Table 7.2)
Max supported capacity of P-VOL and POOL	The supported capacity of Snapshot has restrictions. For details, see section 4.1.
Allocation of P-VOL and V-VOL	The P-VOL and V-VOL can only be allocated in the same controller.
Consistency Group (CTG) number	Max 128/subsystem

Parameter	Requirement
	WMS100/AMS200: Max 254 pairs/CTG AMS500: Max 1,022 pairs/CTG AMS1000: Max 2,046 pairs/CTG
Management of differential	Enabled
Number of POOLs that can be set	One per controller (POOL number "0" is for the Controller 0 and POOL number "1," for the Controller 1).
Access to the POOL from a host	The POOL is not recognizable from the host.
LU number for POOL that can be set	Up to 64 per disk subsystem can be set.
Expansion of POOL capacity	Possible. The capacity is expanded through an addition of LU[s] to the POOL. The extension of a POOL can be made while a pair is formed. However, an LU with an FC drive and an LU with a SATA drive cannot coexist in a POOL.
Reduction of POOL capacity	Possible only when all the V-VOLs have been deleted.
Unification of the LU for the POOL	Not possible.
Formatting/LU unification regarding the V-VOL	Not possible.
Number of V-VOLs that can be set	AMS200/WMS100: Up to 512 including normal LU(s) AMS500: Up to 2,048 including normal LU(s) AMS1000: Up to 4,096 including normal LU(s)
Deletion of the V-VOL	Possible only when the pair status is SMPL.
Initial copy during execution of paircreate	Not required.
Resynchronization (from the P-VOL to the V-VOL)	Not required
Resynchronization (from the V-VOL to the P-VOL)	Possible.
Pair splitting	Possible (The V-VOL is made invalid when a pair is split.)
Splitting	Split always.
Concurrent use of LUN Expansion	Available. For details, see section 5.9.
Concurrent use of Data Retention Utility	When the S-VOL Disable is set for an LU, a pair formation using the LU as an S-VOL (POOL) is suppressed. A setting of the S-VOL Disable of a volume that has already become an S-VOL (V-VOL or POOL) is not suppressed only when the pair status is PSUS. If the S-VOL Disable is set for a P-VOL, restoration of SnapShot is suppressed.
Concurrent use of Cache Residency Manager	Available. Cache partition information is initialized when SnapShot is installed and Cache Partition Manager is being used.
Concurrent use of SNMP Agent	Available.
Concurrent use of Password Protection	Available.
Concurrent use of TrueCopy	TrueCopy can be cascaded with Snapshot. For details, see section 5.10.

Parameter	Requirement
Concurrent use of TCE	TCE can be cascaded with SnapShot P-VOL. For details, see section 5.12.
Concurrent use of Volume Migration	Available. However, a P-VOL, an S-VOL, and reserved LU of Volume Migration cannot be specified as the Snapshot image (V-VOL) and creating SnapShot.
Concurrent use of LUN Manager	Available.
Concurrent use of ShadowImage	Snapshot and ShadowImage can be use together, but Snapshot cannot be cascaded with ShadowImage. The number of CTG at the time of using SnapShot and ShadowImage together is limited to the maximum of 128 combining that of SnapShot and ShadowImage.
Concurrent use of Cache Partition Manager	Possible. They must be installed in the order of Snapshot and Cache Partition Manager. When Cache Partition Manager has been installed, install SnapShot function after setting up a Cache Partition Manager function invalid. Set up a Cache Partition Manager function after installing SnapShot function.
Interchange between the P-VOL and V-VOL	Not possible.
Unification of LUs in a pair	Not possible.
Pairing with the unified LU	Available only P-VOL (Number of unified LUs (16 or less) and a unified LU capacity of 1 GB or larger is required.)
Formatting LUs in a pair	Not possible.
Microprogram changed on line	Possible.
Potential effect caused by installing Snapshot function	Reboot is required.
Potential effect caused by a detachment of one of the controllers	Detaching one controller has no effect on V-VOL data.
Action to be taken when the usable POOL capacity threshold value is exceeded	Return a warning to CCI. (Each user may set threshold value.)
Action to be taken when the limit of usable POOL capacity is exceeded	Status of all V-VOLs becomes PSUE when POOL capacity usage reaches 100%.
Potential effect caused by a P-VOL failure	V-VOL data also exists in the P-VOL; therefore, P-VOL failure causes V-VOL failure.
Reduction of memory	The memory cannot be reduced when the ShadowImage, Snapshot, or TrueCopy, or TCE function is validated. Make the reduction after invalidating the function.

3.2 Supported Platforms

The following table lists the platforms on which SnapShot is supported.

Table 3.2 External Specifications

No.	Platforms	Operating System Version
1	SUN	Solaris 8 (SPARC)
		Solaris 9 (SPARC)
		Solaris 10 (SPARC)
		Solaris 10 (x86)
		Solaris 10 (x64)
2	PC Server (Microsoft)	Windows 2000
		Windows Server 2003 (IA32)
		Windows Server 2003 (x64)
		Windows Server 2003 (IA64)
3	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
4	IBM	AIX 5.1
		AIX 5.2
		AIX 5.3
5	Red Hat	Red Hat Linux AS2.1 (IA32)
		Red Hat Linux AS3.0 (IA32)
		Red Hat Linux AS4.0 (IA32)
		Red Hat Linux AS3.0 (EM64T)
		Red Hat Linux AS4.0 (EM64T)
		Red Hat Linux AS3.0 (IA64)
		Red Hat Linux AS4.0 (IA64)
6	SGI	IRIX 6.5.x

3.3 Copy-on-Write SnapShot Management Software

3.3.1 Command Control Interface Software (CCI)

CCI software is used to display Copy-on-Write SnapShot volume information, create and manage Copy-on-Write SnapShot pairs, and issue commands for replication operations. CCI software resides on the UNIX®/ MS Windows® management server and interfaces with the TagmaStore AMS subsystems through dedicated logical volumes. CCI commands can be issued from the UNIX command line or using a script file. For more information on CCI, see *Hitachi TagmaStore™ Adaptable Modular Storage and Workgroup Modular Storage Command Control Interface (CCI) User and Reference Guide (MK-95DF701)*.

3.3.2 Storage Navigator Modular Software (SNM)

Storage Navigator Modular software GUI displays detailed Copy-on-Write SnapShot information and is used for configuring the Copy-on-Write SnapShot environment. SNM software communicates directly with the TagmaStore AMS subsystems via a local area network (LAN). SNM software supports both a GUI and CLI user interface to perform configuration, monitoring, maintenance and system reduction.

Copy-on-Write SnapShot configuration tasks include setting up data pool, cycle time, logical units and path for remote Copy-on-Write SnapShot operations, enabling or releasing a command device etc. SNM interfaces display important failure information for subsystems and maintenance information including the amount of differential data, the time difference between the P-VOL and the S-VOL, and the used capacity of a data pool. Navigator can be used to increase the pool capacity by adding more LUs to the pool.

In the event of a system failure or disaster at the primary site, the SNM software simplifies and expedites disaster recovery procedures. System Reduction tasks such as deleting command devices and data pools can also be performed using SNM software.

For more information on SNM software, see *Hitachi TagmaStore Adaptable Modular Storage and Workgroup Modular Storage Storage Navigator Modular Graphical User Interface (GUI) User's Guide (MK-95DF711)* or *Hitachi TagmaStore Adaptable Modular Storage and Workgroup Modular Storage Storage Navigator Modular Command Line Interface (CLI) User's Guide (MK-95DF712)*.

3.3.3 Optional Software

The following optional management software is also used in the Copy-on-Write SnapShot system environment:

- 3.3.3.1 HiCommand Device Manager
- 3.3.3.2 HiCommand Tuning Manager

3.3.3.1 HiCommand Device Manager

HiCommand Device Manager is a storage management product which enables users to consolidate storage operations and manage capacity in systems containing multiple storage subsystems. Targeted for users managing multiple storage subsystems in open or shared environments, Device Manager quickly discovers the key configuration attributes of storage subsystems and allows storage administrators to begin proactively managing complex and heterogeneous storage environments using its easy-to-use, browser-based GUI. Device Manager's built-in report facility compiles and presents key information in preformatted reports (HTML) and as comma-separated values for export.

Device Manager provides link-and-launch integration with other software such as HiCommand Tuning Manager, HiCommand Protection Manager and Storage Navigator Modular (for Web) software. For more information on HiCommand Device Manager, see *Hitachi HiCommand Device Manager Web Client User's Guide* (MK-91HC001) or *Hitachi HiCommand Device Manager Command Line Interface (CLI) User's Guide* (MK-91HC007).

3.3.3.2 HiCommand Tuning Manager

HiCommand Tuning Manager is a storage resource management product which enables centralized management of a storage area network environment. This software monitors and analyzes the performance information and capacity information of the host, fibre channel switches, and storage devices. Changes in host I/O performance and capacity over time are displayed graphically and useful for planning future system expansion and reduction. For more information on HiCommand Tuning Manager, see *Hitachi HiCommand Tuning Manager User's Guide* (MK-92HC022).

Chapter 4 Setting up Snapshot

This chapter contains the following information.

- Supported Snapshot Capacity (section 4.1)
- Installation and Uninstalling Snapshot (section 4.2)
- Determining POOL Capacity (section 4.3)
- Creating the Configuration Definition File(section 4.4)
- Setting the Environment Variables (section 4.5)
- Configuring Snapshot (section 4.6)
- Recommended Pair Assignment (section 4.7)
- Configuration Restrictions (section 4.8)

4.1 Supported Snapshot Capacity

The Snapshot function restricts P-VOL/POOL capacity. The supported maximum capacity value varies, depending on the capacity ratio of P-VOL to POOL and cache memory. When using other copy system functions and SnapShot together, the maximum supported capacity of the P-VOL may be restricted further. Therefore, the supported capacity of P-VOL needs to meet the following two conditions.

- Must be less than or equal to the maximum supported capacity calculated by the capacity ratio with POOL (see section 4.1.1)
- Must be less than or equal to the maximum supported capacity at the time of the combined use with other copy system functions (see section 4.1.2)

Furthermore, the POOL used by SnapShot is in common with that used by TCE, it is required to construct a system in which a sufficient POOL capacity is secured when the SnapShot function is used together with TCE.

4.1.1 Maximum Supported Capacity of P-VOL and POOL for Each Cache Memory Capacity

Table 4.1 - Table 4.4 show the maximum supported capacities of the P-VOL and the POOL for each cache memory capacity, and the formula for calculating the previous capacities.

Table 4.1 Formula for Calculating Maximum Supported Capacity Value for P-VOL/POOL (WMS100)

Capacity of the Cache Memory Installed	Capacity Spared for the Differential Data
512 MB/CTL	Not supported.
1 GB/CTL	Total P-VOL capacity ÷ 4 + Total POOL capacity < 1.0 TB

Table 4.2 Formula for Calculating Maximum Supported Capacity Value for P-VOL/POOL (AMS200)

Capacity of the Cache Memory Installed	Capacity Spared for the Differential Data
1 GB/CTL	Total P-VOL capacity ÷ 4 + Total POOL capacity < 1.0 TB
2 GB/CTL	Total P-VOL capacity ÷ 4 + Total POOL capacity < 3.7 TB

Table 4.3 Formula for Calculating the Maximum Supported Capacity Value for P-VOL/POOL (AMS500)

Capacity of the Cache Memory Installed	Capacity Spared for the Differential Data
1 GB/CTL	Not supported.
2 GB/CTL	Total P-VOL of SnapShot and P-VOL (S-VOL) of TCE capacity ÷ 4 + Total POOL capacity < 1.8 TB
3 GB/CTL	Total P-VOL of SnapShot and P-VOL (S-VOL) of TCE capacity ÷ 4 + Total POOL capacity < 3.7 TB
4 GB/CTL	Total P-VOL of SnapShot and P-VOL (S-VOL) of TCE capacity ÷ 4 + Total POOL capacity < 7.0 TB

Table 4.4 Formula for Calculating the Maximum Supported Capacity Value for P-VOL/POOL (AMS1000)

Capacity of the Cache Memory Installed	Capacity Spared for the Differential Data (shared by SnapShot and TCE)
2 GB/CTL	Total P-VOL of SnapShot and P-VOL (S-VOL) of TCE capacity ÷ 4 + Total POOL capacity < 1.8 TB
4 GB/CTL	Total P-VOL of SnapShot and P-VOL (S-VOL) of TCE capacity ÷ 4 + Total POOL capacity < 5.5 TB
6 GB/CTL	Total P-VOL of SnapShot and P-VOL (S-VOL) of TCE capacity ÷ 4 + Total POOL capacity < 11.0 TB
8 GB/CTL	Total P-VOL of SnapShot and P-VOL (S-VOL) of TCE capacity ÷ 4 + Total POOL capacity < 15.0 TB

Table 4.5 - Table 4.14 shows the maximum supported capacity per capacity ratio calculated from the formula of Table 4.1 to Table 4.3.

Table 4.5 Supported Capacity Value of the P-VOL/POOL (When Cache Memory is 1 GB/CTL: WMS100)

Total capacity of all the P-VOLs: Total capacity of all the POOLs	Supported total capacity of all the P-VOLs (TB)	Supported total capacity of all the POOLs (TB)
1:0.5	1.3	0.6
1:1	0.7	0.7
1:3	0.3	0.9

Table 4.6 Supported Capacity Value for P-VOL/POOL (Cache Memory=1 GB/CTL, AMS200)

Total capacity of all the P-VOLs: Total capacity of all the POOLs	Supported total capacity of all the P-VOLs (TB)	Supported total capacity of all the POOLs (TB)
1:0.5	1.3	0.6
1:1	0.7	0.7
1:3	0.3	0.9

Table 4.7 Supported Capacity Value for P-VOL/POOL (Cache Memory=2 GB/CTL, AMS200)

Total capacity of all the P-VOLs: Total capacity of all the POOLS	Supported total capacity of all the P-VOLs (TB)	Supported total capacity of all the POOLS (TB)
1:0.5	4.9	2.4
1:1	2.9	2.9
1:3	1.1	3.3

Table 4.8 Supported Capacity Value for P-VOL/POOL (Cache Memory=2GB/CTL, AMS500)

Total capacity of all the P-VOLs: Total capacity of all the POOLS	Supported total capacity of all the P-VOLs (TB)	Supported total capacity of all the POOLS (TB)
1:0.5	2.4	1.2
1:1	1.4	1.4
1:3	0.5	1.5

Table 4.9 Supported Capacity Value for P-VOL/POOL (Cache Memory=3 GB/CTL, AMS500)

Total capacity of all the P-VOLs: Total capacity of all the POOLS	Supported total capacity of all the P-VOLs (TB)	Supported total capacity of all the POOLS (TB)
1:0.5	4.9	2.4
1:1	2.9	2.9
1:3	1.1	3.3

Table 4.10 Supported Capacity Value for P-VOL/POOL (Cache Memory=4 GB/CTL, AMS500)

Total capacity of all the P-VOLs: Total capacity of all the POOLS	Supported total capacity of all the P-VOLs (TB)	Supported total capacity of all the POOLS (TB)
1:0.5	9.3	4.6
1:1	5.6	5.6
1:3	2.1	6.3

Table 4.11 Supported Capacity Value for P-VOL/POOL (Cache Memory = 2GB/CTL, AMS1000)

Total capacity of all the P-VOLs: Total capacity of all the POOLS	Supported total capacity of all the P-VOLs (TB)	Supported total capacity of all the POOLS (TB)
1:0.5	2.4	1.2
1:1	1.4	1.4
1:3	0.5	1.5

Table 4.12 Supported Capacity Value for P-VOL/POOL (Cache Memory = 4 GB/CTL, AMS1000)

Total capacity of all the P-VOLs: Total capacity of all the POOLS	Supported total capacity of all the P-VOLs (TB)	Supported total capacity of all the POOLS (TB)
1:0.5	7.3	3.6
1:1	4.4	4.4
1:3	1.6	4.8

Table 4.13 Supported Capacity Value for P-VOL/POOL (Cache Memory = 6 GB/CTL, AMS1000)

Total capacity of all the P-VOLs: Total capacity of all the POOLS	Supported total capacity of all the P-VOLs (TB)	Supported total capacity of all the POOLS (TB)
1:0.5	7.3	3.6
1:1	4.4	4.4
1:3	1.6	4.8

Table 4.14 Supported Capacity Value for P-VOL/POOL (Cache Memory = 8 GB/CTL, AMS1000)

Total capacity of all the P-VOLs: Total capacity of all the POOLS	Supported total capacity of all the P-VOLs (TB)	Supported total capacity of all the POOLS (TB)
1:0.5	20.0	10.0
1:1	12.0	12.0
1:3	4.6	13.8

Notes:

- The capacity of each P-VOL is managed in units of 15.75 GBs. When P-VOL capacity is 17 GB, it is regarded as using a capacity of 31.5 GB. When there are two P-VOLs; each has a capacity of 17 GB. They use a total capacity of 63 GB (= 31.5 GB×2), though the actual capacity is 34 GB (= 17 GB×2).
- The capacity of each LU, which has been registered to POOL, is managed in of 4 GB units. When the LU capacity is 5 GB, the LU is regarded as using an 8 GB capacity. When there are two LUs registered to a POOL, each has a 5 GB capacity. They use a total capacity of 16 GB (= 8 GB×2), though the actual capacity is 10 GB (= 5 GB×2).

4.1.2 Maximum Supported Capacity of Concurrent Use of Other Copy Functions

The maximum capacity supported by SnapShot can be calculated from the following formula. The single maximum capacity of SnapShot is shown in Table 4.15. The capacity of all the P-VOLs must be less than or equal to the maximum supported, and it must meet the requirements shown in Table 4.1 to Table 4.4.

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

- = Maximum SnapShot single capacity
- (2 × Total S-VOL capacity of ShadowImage ÷ 17)
- (Total P-VOL and S-VOL capacity of TrueCopy ÷ 17)
- (Total P-VOL and S-VOL capacity of TCE × 3)

Table 4.15 Single Maximum Supported Capacity of SnapShot (TB)

Equipment Type	Mounted Memory Capacity	Single Maximum Supported Capacity (TB)
WMS100	512 MB/CTL	—
	1 GB/CTL	11
AMS200	1 GB/CTL	7
	2 GB/CTL	22
AMS500	1 GB/CTL	—
	2 GB/CTL	27
	3 GB/CTL	33
	4 GB/CTL	48
AMS1000	2 GB/CTL	47
	4 GB/CTL	69
	6 GB/CTL	84
	8 GB/CTL	100

4.1.2.1 Snapshot Setup Notes

- The Snapshot function uses RAID 5, RAID 1, RAID 1+0, or RAID 6. If a usable RAID level does not exist in the disk array subsystem, it must be prepared using Navigator.
- It is necessary to determine which logical unit to specify as the P-VOL and to specify the controller for the P-VOL prior to starting operations.
- It is necessary to designate the number of duplicated volumes, V-VOLs, that are to be paired with the P-VOLs.
- Capacity of the logical unit to be designated as the POOL must be determined prior to starting operations. POOL capacity is determined by:
 - The total capacity of the P-VOLs
 - The number of V-VOLs
 - The interval of the snapshot instructions (a period for holding the V-VOL)
 - The amount of data updated during the interval
- Approximately 10 percent of the data is updated per day. When one V-VOL is made per P-VOL of 1 TB and one Snapshot instruction is issued a day, the recommended POOL capacity is about 2.5 times as large as the general updated data amount a day, which is 10 percent, or about 250 GB depending on the change in the amount of updated data and the operation or locality of access.
- The recommended POOL capacity is about 1.2 TB (five times the capacity in the case where only one V-VOL is made).

When five V-VOLs are made per 1 TB P-VOL and one Snapshot instruction is issued to one of the five V-VOLs a day (each V-VOL holds data for a five day period).
- The amount of data actually accumulated in the POOL varies depending on the application, amount of processed data, and time zone of operation. When using Snapshot, check the POOL capacity operation and watch the capacity actually used in operation.
- This manual does not contain procedures for creating a RAID group and logical unit and that for specifying a controller that controls the logical unit. For the information above, refer to the *Hitachi TagmaStore Adaptable Modular Storage and Workgroup Modular Storage Navigator Modular Graphical User Interface (GUI) User's Guide (MK-95DF711)*.

4.2 Installing and Uninstalling Snapshot

Since Snapshot is an option, it cannot usually be selected when first using TagmaStore. To make Snapshot available, you must install and unlock it.

Snapshot can be installed from the Navigator. This section describes installation/uninstallation procedures performed by using the Navigator via the GUI. For details on the GUI version of Navigator, refer to the *Hitachi TagmaStore Adaptable Modular Storage and Workgroup Modular Storage, Storage Navigator Modular Graphical User Interface (GUI) User's Guide (MK-95DF711)*.

For procedures performed by using the Command Line Interface (CLI) of Navigator, see Appendix B. (Procedures for setting a Command Device, Data Pool, and Snapshot Image are explained also.)

Note: When installing, uninstalling, enabling, or disabling Snapshot where the disk array subsystem is used on a TrueCopy or TCE remote site, the following occurs when the disk array subsystem is restarted:

- Both paths of TrueCopy or TCE are blocked. When a path is blocked, a TRAP (notification to the SNMP Agent Support Function) occurs. Should this occur, notify all affected departments immediately. A blocked path of TrueCopy or TCE recovers automatically when the disk array subsystem is restarted.
- The pair status of TrueCopy or TCE is PAIR or COPY changes to PSUE.

When you restart the disk array subsystem, install, uninstall, enable, or disable Snapshot **after** changing the pair status of TrueCopy or TCE to PSUS.

Notes:

- The installing, uninstalling, enabling, and disabling functions for the Snapshot feature are set for each disk array subsystem.
- Before installing, uninstalling, enabling, or disabling Snapshot, verify that the disk array subsystem is operating in a normal state. If a failure such as a controller blockade occurred, Snapshot cannot be installed, uninstalled, enable, or disable.
- If you install, uninstall, enable, or disable Snapshot on disk array subsystem connected to a NAS, you must also stop the clusters between NAS units. When restarting the subsystem, you must restart the clusters.

4.2.1 Installing Snapshot

To install Snapshot, the key code or key file is required. The following describes the installation procedure.

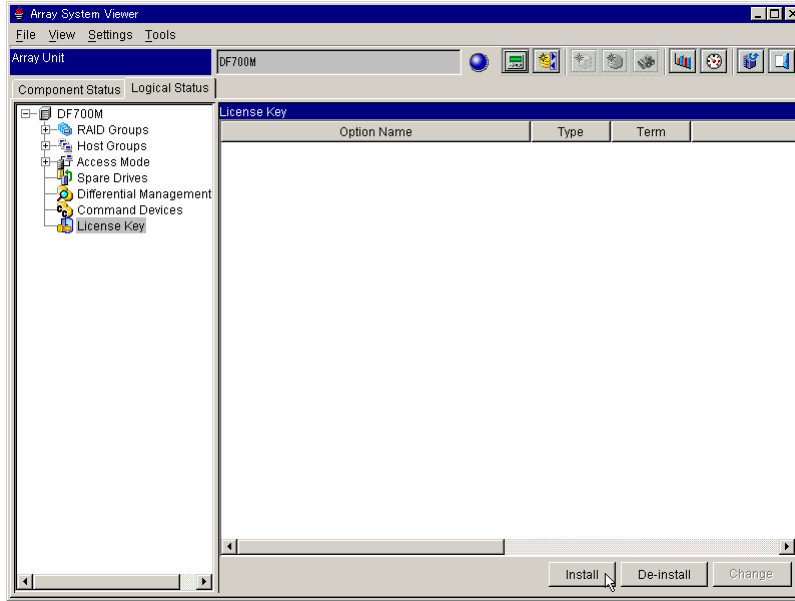
To install Snapshot:

1. Start Navigator and change to **Management Mode** (administrator mode).
2. Register the subsystem (array unit) in which Snapshot is to be installed and connect to the subsystem.

The Array System Viewer window is displayed and shows the connected subsystem.

3. Click the **Logical Status** tab.
4. Click the **License Key** icon.

Storage Navigator Modular: Version 5.00 or higher



Storage Navigator Modular: Earlier than Version 5.00

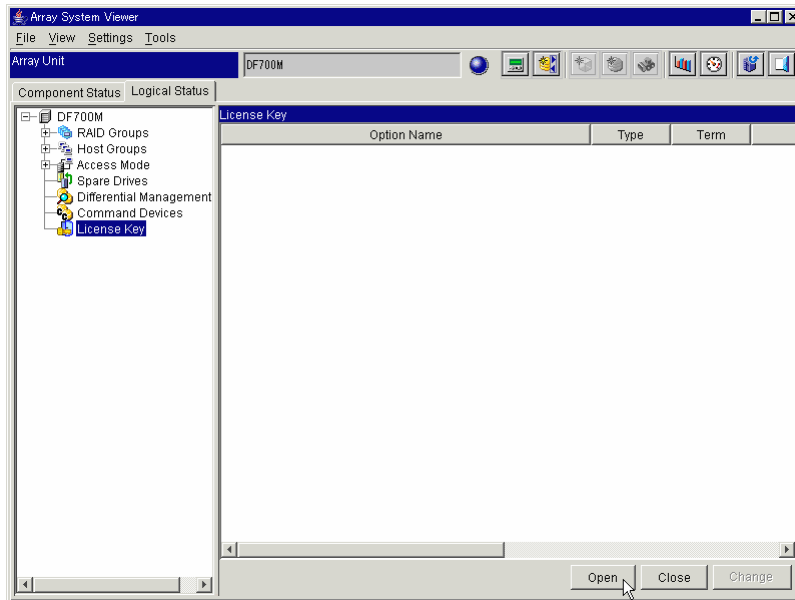



Figure 4.1 Array System Viewer Window (Logical Status Page)

5. Click **Install**. (Storage Navigator Modular: Version 5.00 or higher)
The **Install Options** dialog box is displayed.
Click **Open**. (Storage Navigator Modular: Earlier than Version 5.00)
The **Unlock Options** dialog box is displayed.
Storage Navigator Modular: Version 5.00 or higher



Storage Navigator Modular: Earlier than Version 5.00


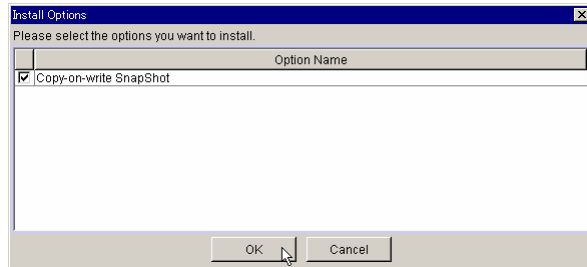


Figure 4.2 Install/Unlock Options Dialog Box

6. When you install the option using the key code, click the **Key Code** radio button. Set up the key code. When you install the options using the key file, click the **Key File** radio button and set up the path for the key file. Click **OK**.
Browse is used to set the path to a key file correctly.

- When you install the options using the key file, the options selection dialog displays. make sure the **Option Name** and click **OK**.

Storage Navigator Modular: Version 5.00 or higher



Storage Navigator Modular: Earlier than Version 5.00

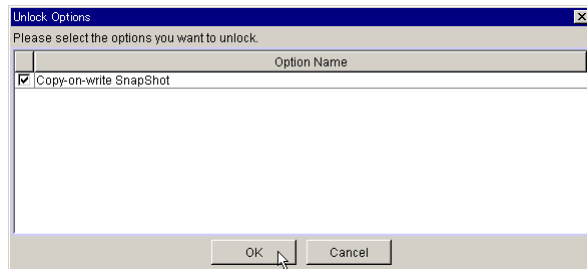



Figure 4.3 Options Selection Dialog

- A message appears, asking you to confirm that you want to install Snapshot. Click **OK**.

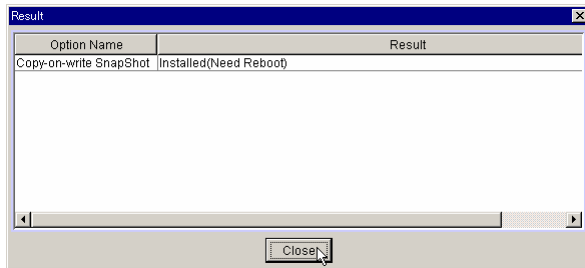
The following message is displayed when Storage Navigator Modular version 3.00 or higher is used with Cache Partition Manager enabled. Since SnapShot uses the data pool, the information on the cache partition is reset to the default settings. This occurs because a part of the cache memory is used for the internal resource management.

If you would like to override the configuration information, click **OK**. When you do not want to lose the configuration information on the partition immediately, click **Cancel**. You must then install the SnapShot function again after checking the partition information.



- When you install options using the key file, the following dialog box displays. Click **Close**.

Storage Navigator Modular: Version 5.00 or higher



Storage Navigator Modular: Earlier than Version 5.00

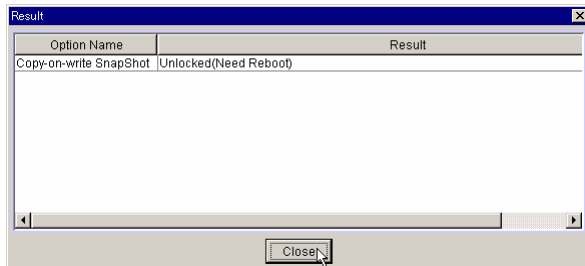
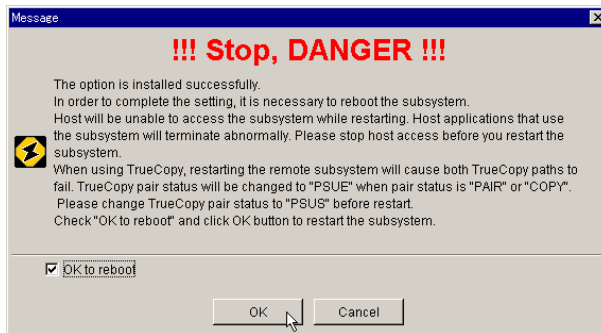


Figure 4.4 Resulting Dialog Box

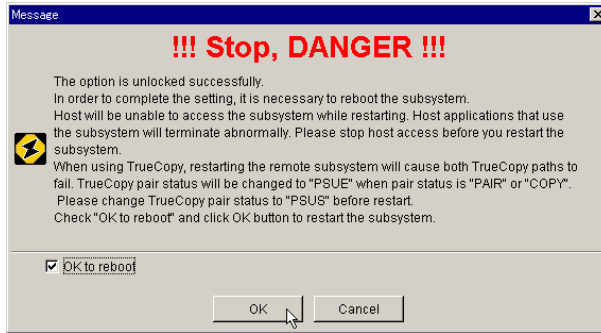
10. A message appears, confirming that this optional feature is installed. It also requests confirmation to restart the subsystem. Click **OK**.

When SnapShot is used together with TCE, a restart is not required if the subsystem has already been restarted to use TCE.


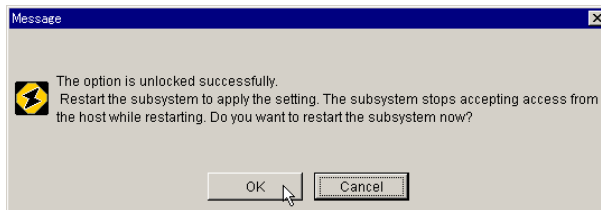
Storage Navigator Modular: Version 5.00 or higher





Storage Navigator Modular: Version 4.03 or higher



Storage Navigator Modular: Earlier than Version 4.03 (case by key code)



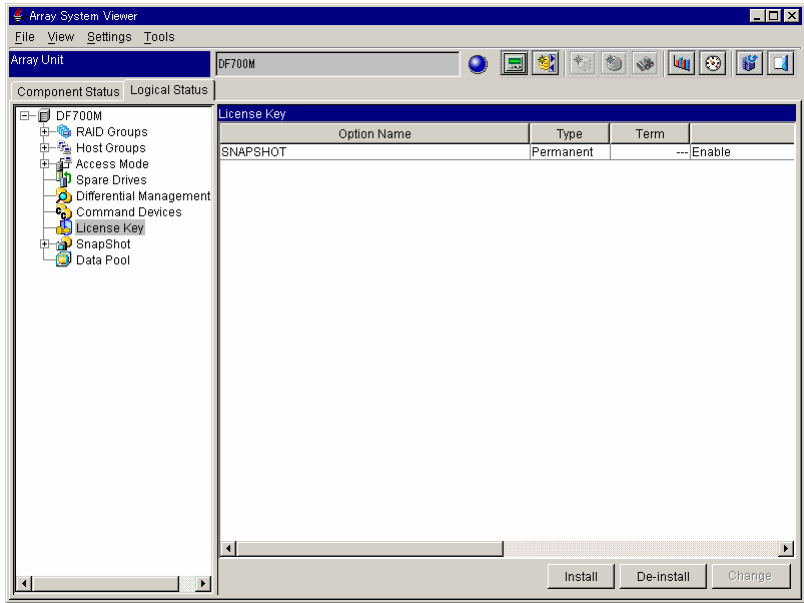
Storage Navigator Modular: Earlier than Version 4.03 (case by key file)



Note: To validate the installing of this optional feature, restart the array unit. The previous setting stays valid until the subsystem is restarted. The array unit cannot access the host until the restart is completed. Therefore, be certain the host has stopped accessing data before beginning the restart process.

When a restart is not in process, this set-up optional feature will be updated and displayed.

Storage Navigator Modular: Version 5.00 or higher



Storage Navigator Modular: Earlier than Version 5.00

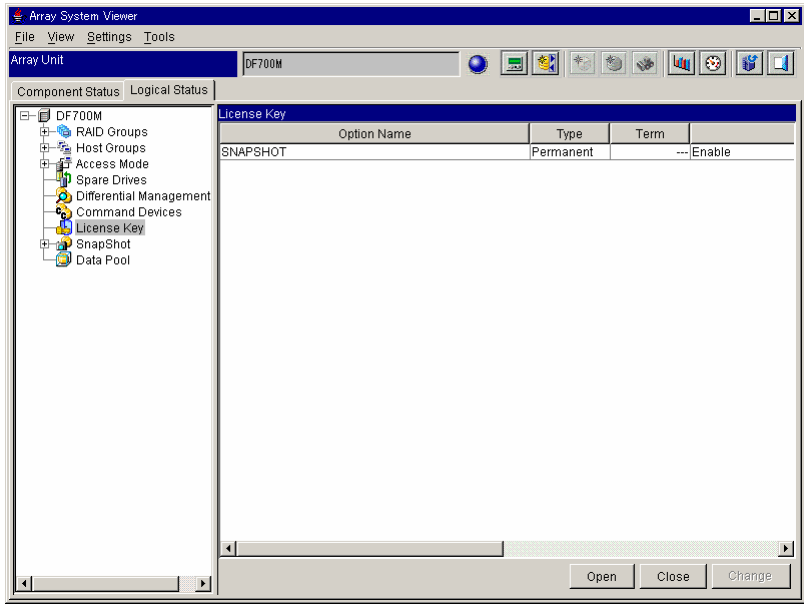


Figure 4.5 Array System Viewer Window (Logical Status Page Option: Enable)

When you choose to restart the array unit, the time the restart began is displayed. Restarting takes approximately four to 15 minutes.

Note: It may take up to 15 minutes for an array unit to respond, depending on the configuration of the array unit.

11. Click **OK**. A message appears, stating that restart was successful.

4.2.2 Uninstalling Snapshot

The key code is required to uninstall Snapshot. After successful uninstall, Snapshot cannot be used again until it is installed using the original key code or key file.

Note: The following conditions must be met in order to uninstall Snapshot.

- All Snapshot pairs must be released (the statuses of all LUs are SMPL).
- All the Data POOLS must be deleted.
- All the V-VOL must be deleted.

To uninstall Snapshot:

1. Start Navigator and change the operation mode to **Management Mode** (administrator mode).
2. Register the subsystem in which the Snapshot is to be uninstalled, then connect to the subsystem.

The **Array System Viewer** window (see Figure 4.1) opens and displays the connected subsystem.

3. Click **Logical Status** tab.
4. Click **License Key** icon (see Figure 4.5).
5. Click **De-install**. (Storage Navigator Modular: Version 5.00 or higher)
The **De-install Options** dialog box is displayed.
Click **Close**. (Storage Navigator Modular: Earlier than Version 5.00)
The **Lock Options** dialog box is displayed.
6. Enter a **key code** in the text box and click **OK**.
7. A message displays, asking you to confirm that you want to uninstall Snapshot. Click **OK**.
When SnapShot is used together with TCE, a restart is not required if TCE continues to be used.
8. A message displays, confirming that this optional feature is uninstalled. It also requests confirmation to restart the subsystem. Click **OK**.

Note: To validate uninstalling this optional feature, restart the array unit. The previous setting stays valid until the subsystem is restarted. The array unit cannot access the host until the restart is completed. Therefore, be certain the host has stopped accessing data before beginning the restart process.

When a restart is not in process, this set-up optional feature will be updated and displayed (see Figure 4.1).

Note: It may take up to 15 minutes for an array unit to respond, depending on the configuration of the array unit.

When you choose to restart the array unit, the time the restart began is displayed. Restarting takes approximately four to 15 minutes.

9. Click **OK**. A message displays, stating that the restart was successful.
The **Array System Viewer** window closes.
Snapshot is now uninstalled.

4.2.3 Enabling and Disabling Snapshot

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

Note: The following conditions must be satisfied in order to disable Snapshot.

- All Snapshot pairs must be released (the status of all LUs are SMPL).
- All Data Pools must be deleted.
- All Snapshot Images (V-VOL) must be deleted.


The following describes the enabling/disabling procedure:

1. Start Navigator and change the operation mode to **Management Mode** (administrator mode).
2. Register the subsystem from which the status of Snapshot is to be changed and connect to the subsystem.
The **Array System Viewer** window (see Figure 4.1) opens displaying the connected subsystem.
3. Click **Logical Status** tab.
4. Click **License Key** icon (see Figure 4.5).
5. From **Option Name**, select **SNAPSHOT**, then click **Change**.
6. A message displays, asking you to confirm that you want to change the status (enable or disable). Click **OK**.
7. A message displays, asking you to confirm that this optional feature is set. It also requests confirmation to restart the subsystem. Click **OK**.


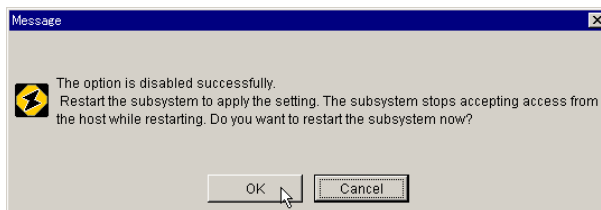
The restart for validation: The restart is not required at the time of the unlocking or when the subsystem is already restarted to use TCE.

The restart for invalidation: Restart the subsystem unless the restart is not to be done at the time of the locking or TCE is still to be used successively when SnapShot is used together with TCE.

Storage Navigator Modular: Version 4.03 or higher



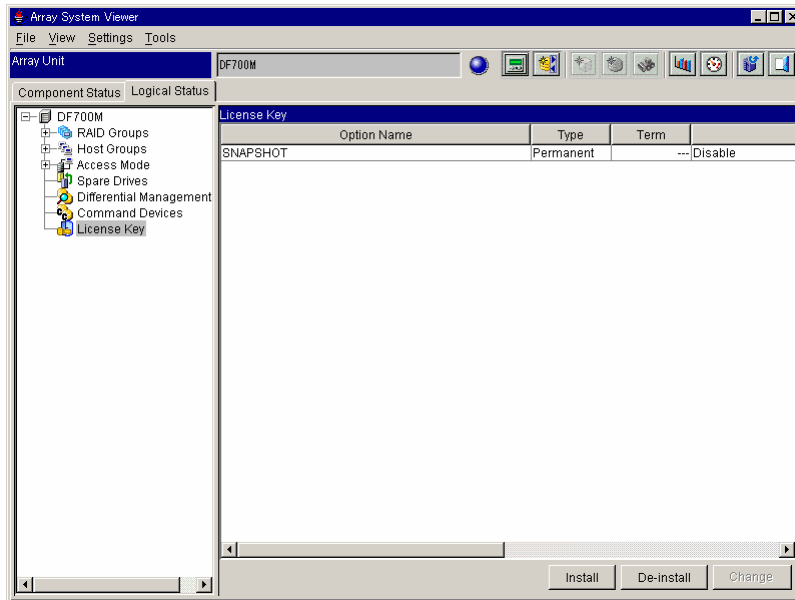
Storage Navigator Modular: Earlier than Version 4.03



Note: This setting is not active until the system is restarted. The subsystem cannot access the host until the restart is completed. Therefore, be certain the host has stopped accessing data before beginning the restart process.

When a restart is not in process, this set-up optional feature will be updated and displayed.

Storage Navigator Modular: Version 5.00 or higher



Storage Navigator Modular: Earlier than Version 5.00

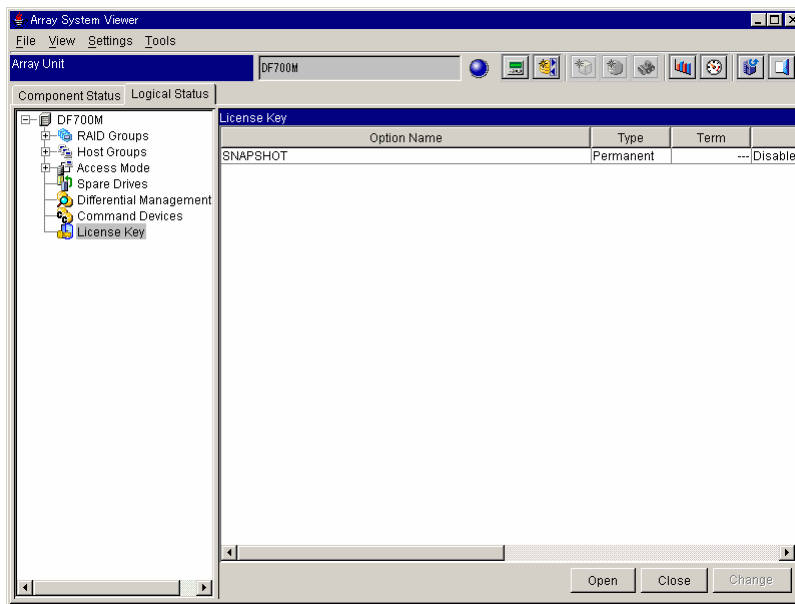


Figure 4.6 Array System Viewer Window (Logical Status Page Option: Disable)

When you choose to restart the array unit, the time the restart began is displayed. Restart takes approximately four to 15 minutes.

Note: It may take up to 15 minutes for an array unit to respond, depending on the configuration of the array unit.

8. A message appears, stating that restart was successful. Click **OK**.
The Array System Viewer window closes.

4.3 Determining POOL Capacity

- The factors that determine the POOL capacity include:
 - A total capacity of P-VOLs
 - A number of generations (a number of V-VOLs)
 - An interval of the snapshot instructions (a period for holding the V-VOL)
 - An amount of data updated during the interval and the spare capacity (safety rate)

The method for calculating the POOL capacity is:

POOL Capacity = P-VOL capacity x (Amount of renewed data x safety rate) x a number of generations.

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

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

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

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

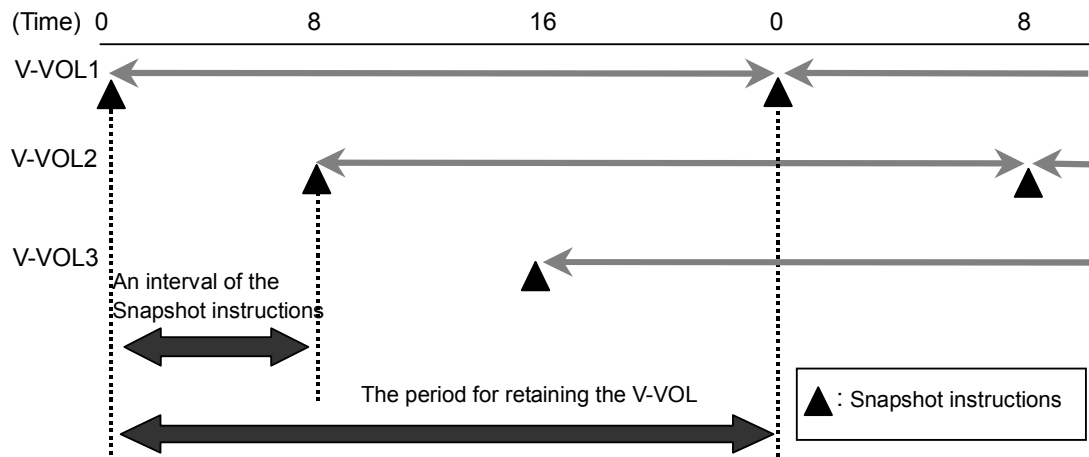
The recommended value of the POOL capacity when a capacity of one P-VOL is 1 TB is shown in the following table. Multiply the value in the following table by a value of a capacity per V-VOL in TB.

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

Table 4.16 Recommended Value of the POOL Capacity (When the P-VOL Capacity is 1 TB)

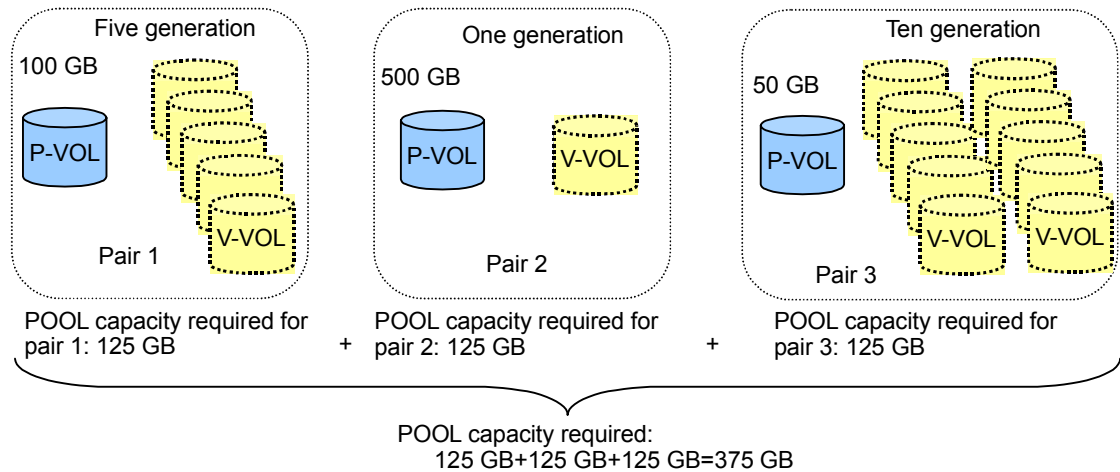
An Interval of Snapshot Instructions (see <i>Note 1</i>)	V-VOL Number (n)					
	1	2	3	4	5	6 to 14
From one to four hours	0.10 TB	0.20 TB	0.30 TB	0.40 TB	0.50 TB	0.10 x n TB
From four to eight hours	0.15 TB	0.30 TB	0.45 TB	0.60 TB	0.75 TB	0.15 x n TB
From eight to 12 hours	0.20 TB	0.40 TB	0.60 TB	0.80 TB	1.00 TB	0.20 x n TB
From 12 to 24 hours	0.25 TB	0.50 TB	0.75 TB	1.00 TB	1.25 TB	0.25 x n TB
Longer than 24 hours	Individual consultation (see <i>Note 2</i>)					

Note 1: An interval of Snapshot instructions means a time between Snapshot instructions issued to the designated P-VOL. When there is only one V-VOL, the interval of the Snapshot instructions is as long as the period for retaining the V-VOL. When there are two or more V-VOLs, the interval of the Snapshot instructions multiplied by the number of the V-VOLs is the period for retaining the one V-VOL.



Note 2: Construct a system in which the interval of Snapshot instructions is less than one day. It becomes difficult, depending on system environment to estimate the amount of data accumulated in the data pool when the interval of the Snapshot instructions is long.

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



Up to 64 LUs can be set for a POOL and a capacity of the POOL can be expanded through addition of the LU(s) in the online status (while the Snapshot pair is formed). Therefore, at the time of the initial system configuration, prepare a spare LU to be used when the POOL capacity becomes insufficient. (This is recommended.) Additional care should be taken that all the V-VOLs using the POOL are dissolved in order to return the LU assigned to the POOL to an ordinary LU.

When returning the backup data from a tape device to the V-VOL, a free POOL with a capacity larger than the P-VOL capacity is required. More than 1.5 times the P-VOL capacity is recommended.

Note: An LU with an FC drive and an LU with a SATA drive cannot coexist in the POOL.

4.4 Creating the Configuration Definition File


This section includes an example that shows how to create pairs of three V-VOLs from one P-VOL.

The configuration definition file describes the system configuration in order to make CCI operational. 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. This sample configuration definition file (HORCM_CONF) is included with the CCI software and this file 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 TagmaStore Adaptable Modular Storage and Workgroup Modular Storage Command Control Interface (CCI) User and Reference Guide (MK-95DF701)*.

The configuration definition file can be automatically created using the `mkconf` command tool. However, the parameters must be set manually (see step 4 below).

The following describes an example for manually defining the configuration definition file, when the system is configured with two instances within the same server (Windows® 2000).

The P-VOL and V-VOLs are conceptually diagrammed in the following figure.



1. On the host where CCI is installed, verify that the CCI is not running. If the CCI software is still running, shut down the CCI software using `horcmshutdown`.
2. In the command prompt, make two copies of the sample file (`horcm.conf`).

Example:

```
c:\HORCM\etc> copy \HORCM\etc\horcm.conf \WINNT\horcm0.conf
c:\HORCM\etc> copy \HORCM\etc\horcm.conf \WINNT\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). Refer to the *Hitachi TagmaStore Adaptable Modular Storage and Workgroup Modular Storage Command Control Interface (CCI) User and Reference Guide (MK-95DF701)* for details on calculating the poll (10ms) value. Specifying the value incorrectly may cause resource contention in the internal process; the process is temporarily suspended and pauses subsystem internal processing. For more details on configuration parameters, refer to the *Hitachi TagmaStore Adaptable Modular Storage and Workgroup Modular Storage Command Control Interface (CCI) User and Reference Guide (MK-95DF701)*.

5. In the HORCM_CMD section, specify the subsystem's physical drive (Command Device).

```

horcm0.conf - Notepad
File Edit Search Help
HORCM_MON
#ip_address      service      poll(10ms)    timeout(10ms)
XXXXXXXXX        5000         12000         3000

HORCM_CMD
#dev_name        dev_name      dev_name
\\.\PHYSICALDRIVE2

HORCM_DEV
#dev_group       dev_name      port#         TargetID      LU#          MU#
UG01             oradb1        CL1-A         1             2            0
UG01             oradb2        CL1-A         1             2            1
UG01             oradb3        CL1-A         1             2            2

HORCM_INST
#dev_group       ip_address    service
UG01             XXXXXXXXX    5001

```

Figure 4.7 Horcm0.conf Example

6. Save the configuration definition file and use `horcmstart` to start the CCI software.
7. Execute `raidscan` and write down the target ID displayed in the execution result.
8. Shut down the CCI software and open the configuration definition file again.
9. In the HORCM_DEV section, set the necessary parameters. For the target ID, set the ID of the raidscan result you wrote down. Also, the item MU# must be added after the LU.
10. In the HORCM_INST section, set the necessary parameters and save (overwrite) the file.

11. Repeat step 3 through 10 for the horcm1.conf file (see Figure 4.8).

```

horcm1.conf - Notepad
File Edit Search Help
HORCH_MON
#ip_address      service      poll(10ms)    timeout(10ms)
XXXXXXXXX        5001         12000         3000

HORCH_CMD
#dev_name        dev_name      dev_name
\\.\PHYSICALDRIVE2

HORCH_DEU
#dev_group       dev_name      port#          TargetID       LU#           MU#
UG01             oradb1        CL1-A         1              3             0
UG01             oradb2        CL1-A         1              4             0
UG01             oradb3        CL1-A         1              5             0

HORCH_INST
#dev_group       ip_address    service
UG01             XXXXXXXXX    5000
  
```

Figure 4.8 Horcm1.conf Example

12. Enter the following in the command prompt to verify the connection between CCI and the subsystem.

Example:

```

C:\>cd horcm\etc

C:\horm\etc>echo hdl-7 | .\inraid
Harddisk 1 -> [ST] CL1-A Ser =75000174 LDEV = 0 [HITACHI ] [DF600F-CM ]
Harddisk 2 -> [ST] CL1-A Ser =75000174 LDEV = 2 [HITACHI ] [DF600F ]
                HORC = SMPL HOMRCF[MU#0 = SMPL MU#1 = NONE MU#2 = NONE]
                RAID5[Group 2- 0] SSID = 0x0000
Harddisk 3 -> [ST] CL1-A Ser =75000174 LDEV = 3 [HITACHI ] [DF600F ]
                HORC = SMPL HOMRCF[MU#0 = SMPL MU#1 = NONE MU#2 = NONE]
                RAID5[Group 3- 0] SSID = 0x0000
Harddisk 4 -> [ST] CL1-A Ser =75000174 LDEV = 2 [HITACHI ] [DF600F ]
                HORC = SMPL HOMRCF[MU#0 = NONE MU#1 = SMPL MU#2 = NONE]
                RAID5[Group 2- 1] SSID = 0x0000
Harddisk 5 -> [ST] CL1-A Ser =75000174 LDEV = 4 [HITACHI ] [DF600F ]
                HORC = SMPL HOMRCF[MU#0 = NONE MU#1 = SMPL MU#2 = NONE]
                RAID5[Group 4- 0] SSID = 0x0000
Harddisk 6 -> [ST] CL1-A Ser =75000174 LDEV = 2 [HITACHI ] [DF600F ]
                HORC = SMPL HOMRCF[MU#0 = NONE MU#1 = NONE MU#2 = SMPL]
                RAID5[Group 2- 2] SSID = 0x0000
Harddisk 7 -> [ST] CL1-A Ser =75000174 LDEV = 5 [HITACHI ] [DF600F ]
                HORC = SMPL HOMRCF[MU#0 = NONE MU#1 = NONE MU#2 = SMPL]
                RAID5[Group 5- 0] SSID = 0x0000

C:\horm\etc>
  
```

For details on the configuration definition file, refer to the *Hitachi TagmaStore Adaptable Modular Storage and Workgroup Modular Storage Command Control Interface (CCI) User and Reference Guide (MK-95DF701)*.

4.5 Setting the Environment Variable

To perform Snapshot operations, you must set the environment variable. An example of the system configuration with two instances within the same server (Windows® 2000) follows.

1. Set the environment variable for each instance. Enter the following from the command prompt.

Example:

```
C:\HORCM\etc>set HORCMINST=0
```

2. To use Snapshot, you must set the environment variable shown below.

Example:

```
C:\HORCM\etc>set HORCC_MRCF=1
```

3. Execute the horcmstart script and execute pairdisplay to verify the configuration.

Example:

```
C:\HORCM\etc>horcmstart 0 1
starting HORCM inst 0
HORCM inst 0 starts successfully.
starting HORCM inst 1
HORCM inst 1 starts successfully.

C:\HORCM\etc>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 )75000174 2.SMPL ----,----- ---- -
VG01   oradb1(R)   (CL1-A , 1, 3-0 )75000174 3.SMPL ----,----- ---- -
VG01   oradb2(L)   (CL1-A , 1, 2-1 )75000174 2.SMPL ----,----- ---- -
VG01   oradb2(R)   (CL1-A , 1, 4-0 )75000174 4.SMPL ----,----- ---- -
VG01   oradb3(L)   (CL1-A , 1, 2-2 )75000174 2.SMPL ----,----- ---- -
VG01   oradb3(R)   (CL1-A , 1, 5-0 )75000174 5.SMPL ----,----- ---- -
```

4.6 Configuring Snapshot

4.6.1 Setting the POOL and V-VOL

To prepare for using Snapshot, set the POOL and the V-VOL.

- Setting the POOL:
 - One POOL can be set for each controller. The POOL shares and manages two or more P-VOLs and the differential data of the same number of V-VOLs for each controller.
 - Assign formatted LU(s) to the POOL. Up to 64 LUs can be assigned to each POOL. It is possible to add an LU during an online operation.
 - The capacity of the first LU to be assigned to the POOL must be 2 GB or larger.

Notes:

- Assign an LU larger than 20 GB.
- An LU with an FC drive and an LU with a SATA drive cannot coexist in the POOL.
- Setting V-VOL:
 - Up to 15(14: the micro program earlier than version 0750/A) V-VOLs can be assigned per P-VOL.
 - The V-VOL can be assigned after creating the LU to be the P-VOL and assigning the POOL. When assigning the V-VOL, specify the LU to be the P-VOL and the POOL.
 - The controller that has ownership of the V-VOL (a controller that controls the V-VOL) is the same one that controls the P-VOL and POOL.

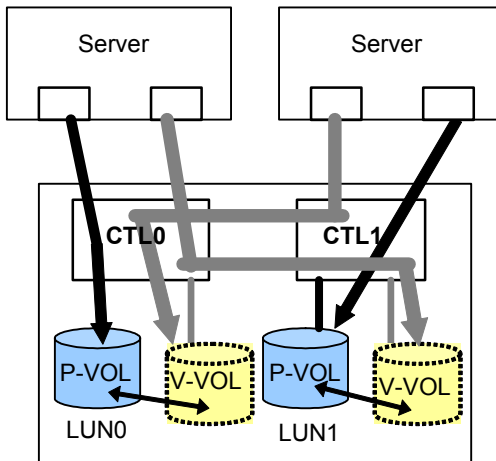
4.6.2 Path Connection

- AMS500

Configuration A is not recommended (the V-VOL is run by the CTL of the other system). For the host side, use two ports per CTL (a single port per CTL causes an FC bottleneck on the host side and disables the process to be performed).

Configuration A:

Accessing the V-VOL from another CTL (cross-connecting).




Cross-use in data share mode
(heavy load on CPU)

Not Recommended

Configuration B:

Accessing the P-VOL and V-VOL from local CTL port.




P-VOL and V-VOL uses the same CTL port.

Recommended

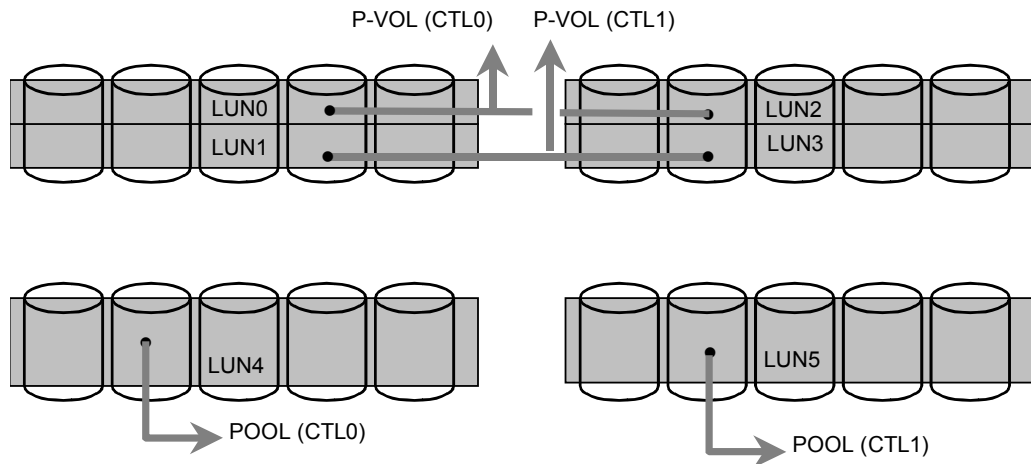
4.6.3 Locating P-VOLs and POOLs

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

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



- **Notes on locating multiple LUs within the same drive column**



If multiple LUs are set within the same drive and each pair state differs, it is difficult to estimate the performance in order to design the system operational settings. For example, when the LU0 and LU2 are both P-VOLs and exist within the same group in the same drive (V-VOLs are located in different drive group), and when the LU0 is in COPY (RS-R) status and the LU2 is in PSUS status.

- **Pair status differences when setting multiple pairs**

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

- **Choosing the drive types of a RAID group in which a P-VOL/POOL will be located**

For optimal performance, a P-VOL should be located in a RAID group which contains FC drives. When a P-VOL or POOL is located in a RAID group made up of SATA drives, the host I/O performance is lessened due to the decreased performance of the SATA drive. You should assign a primary volume to a RAID group consisting of FC drives (refer to section 4.6.3).

- **Setting multiple command devices**

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

4.6.4 P-VOLs and POOLs in a RAID Configuration

- You must use a RAID level with redundancy for both the P-VOL and POOL as RAID 0 is not supported.
- A RAID level and/or a number of drives in RAID (N of $ND+1P$ or $ND+NP$) of a P-VOL are not identical to those of a POOL. However, to improve performance, designate the RAID levels and numbers of drives to be identical respectively.

Table 4.17 P-VOL and POOL RAID Configuration

P-VOL	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 (see Note 1) ($N = 4$)	1	$2+1.25 = 3.25$	1/3.25
RAID 5 (see Note 1) ($N = 4$)	RAID 1+0 ($N = 1$ to 8)	1	$1.25+2 = 3.25$	1/3.25
RAID 5 (see Note 1) (When $N = 4$)	RAID 5 (see Note 1) (When $N = 4$)	1	$1.25+1.25 = 2.5$	1/2.5
RAID 5 (see Note 1) (When $N = 8$)	RAID 5 (see Note 1) (When $N = 8$)	1	$1.125+1.125 = 2.25$	1/2.25
RAID 6 (see Note 2) (When $N = 4$)	RAID 6 (see Note 2) (When $N = 4$)	1	$1.5+1.5 = 3$	1/3
RAID 6 (see Note 2) (When $N = 8$)	RAID 6 (see Note 2) (When $N = 8$)	1	$1.25+1.25 = 2.5$	1/2.5
RAID 6 (see Note 2) (When $N = 4$)	RAID 5 (see Note 1) (When $N = 4$)	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.

4.6.5 Monitoring POOL Capacity

When the ratio of the used POOL capacity to total POOL capacity becomes larger than the threshold value, all the V-VOLs that use the POOL are placed in PSUE because the differential data cannot be saved (70% is the default, but it can be set as low as 50% and as high as 80%).

Even if the contract for hardware maintenance (including a free warranty period) has been made, a user must monitor the data pool so that the capacity will not exceed its capacity. The user is required to manage the data pool capacity.

Table 4.18 Operation Performed when the Threshold Value is Exceeded

Item	Operation
Navigator	A message is displayed in the Information Message field.
	Data pool status is changed to PFUS.
CCI	The status of the pair that uses the data pool concerned is changed from PSUS to PFUS.

- Method for informing a user that the threshold value of the used data pool capacity is exceeded:
 - Because this system has no function which informs the user of the excess threshold value, a host system is required to watch CCI. Prepare a script, etc. for informing the user that the pair status has been changed to PFUS.
 - You can get the pair status as a returned value using the CCI `pairvolchk -ss` command. When the status is PFUS, the returned value is 28. (When the LU is specified, the values for the P-VOL and V-VOL are 28 and 38 respectively.)
 - Monitoring on the used data pool capacity is necessary for each data pool.
 - The used POOL capacity (rate of use) can be accessed through CCI or Navigator. It is recommended not only to monitor the data pool threshold value but also to monitor and manage the hourly transition of the used capacity of the data pool. For details of procedure for referring to the rate of data pool capacity used, refer to the *Hitachi TagmaStore Adaptable Modular Storage and Workgroup Modular Storage Command Control Interface (CCI) User and Reference Guide (MK-95DF701)*.

Example: A script for informing the user of the watching result in the case of Windows® host.

The following is a script to inform the user about the result of the data monitoring process (oradb1) using the data pool of controller 0 and pair (oradb2) of controller 1. The following script is activated every few minutes. You must activate CCI (HORCM) before this script can be activated.

```
echo OFF
set HORCC_MRCF=1
set HORCMINST=0
pairvolchk -g vg01 -d oradb1 -ss
if errorlevel 29 go to pool1
    if errorlevel 28 go to pfusp10
    go to pool1
:pfusp10
<The procedure for informing a user>Note

:pool1
pairvolchk -g vg02 -d oradb2 -ss
if errorlevel 29 go to end
    if errorlevel 28 go to pfusp11
    go to end
:pfusp11
<The procedure for informing a user>Note

:end
```

Note: Describe the following processes in the procedure for informing a user of the pair status, as necessary.

- E-Mail notification process
- Screen display process
- SNMP notification process
- Event log notification process
- A process to add the LU(s) to the data pool

4.6.6 Monitoring Transition to PSUE Status

When there is insufficient data pool capacity or a hardware failure occurs, the Snapshot pair is placed in PSUE status. Even if the hardware maintenance contract (including a free warranty period) has been made, the operation for recovering the status of Snapshot pair must be created by the user. The obligation of service personnel is limited to a recovery from the hardware failure. Therefore, the user must monitor the transition of the Snapshot pair status to PSUE.

Table 4.19 Operation Performed when Pair Status Changed to PSUE

Item	Operation
Navigator	A message is displayed in the Information Message field.
	The pair status is changed to PSUE.
CCI	The status of the pair designated is changed to PSUE.
	An error message is output to the system log file. (For UNIX system and the Windows 2000 system, the syslog file and eventlog file are shown respectively.)

- Informing a user when the pair status is changed to PSUE:
 - Because this system has no function to inform a user of the occurrence of the PSUE status, it is necessary to monitor the system log from a host system. Prepare a script for informing the user that the pair status is changed to PSUE using the `swatch` command, etc. of the monitoring tool.
 - For a UNIX system, or an application program for monitoring the eventlog file in the case of Windows system, the following message is output to the event log during PSUE status.

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

4.7 Recommended Pair Assignment

- Do not assign a frequently updated LU to a pair.

When the pair status is PSUS, the old data is copied to a POOL volume when writing to a primary volume. Because the load on the processor in the controller is increased, the writing performance becomes limited. When the writing load becomes heavier due to: a large number of write operations, the writing of data with a large block size, frequent write I/O instructions, and continuous writing, the effect becomes greater. Therefore, be restrictive in the selection of the LU to which Snapshot is applied. When applying Snapshot to an LU bearing a heavy writing load, it is necessary to consider making loads on the other LUs lighter.
- Use a small number of volumes within the same RAID group.

When volumes are assigned to the same RAID group and used as primary volumes, there may be situations where the host I/O for one of the volumes causes a restriction on host I/O performance of the other volume(s) due to drive contention. Therefore, it is recommended that you assign few (one or two) primary volumes to the same RAID group (refer to section 4.6.3). When creating pairs within the same RAID group, standardize the controllers that control LUs in the same RAID group.
- Make an exclusive RAID group for a POOL volume.

When another volume is assigned to a RAID group to which a POOL volume has been assigned, loads on drives are increased and their performance is restricted because primary volumes correspond to the POOL volume in common. Therefore, use a RAID group, to which a POOL volume is assigned, for the POOL volume only. There can be multiple POOL volumes in a subsystem. Please use different RAID groups for each POOL (refer to section 4.6.3).
- For Snapshot, use FC drives.

When a P-VOL and POOL are located in a RAID group made up of SATA drives, the performance of a host I/O is reduced because of the lower performance of the SATA drive. Therefore, you should assign a primary volume to a RAID group consisting of FC drives (refer to section 4.6.3).
- Assign four or more disks to the data disks.

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

4.7.1 Windows 2000/2003 and Dynamic Disk

4.7.1.1 Environments

You cannot make a P-VOL and a V-VOL into a dynamic disk using the Windows Server 2000 environment, but you can use a P-VOL and a V-VOL as a dynamic disk.

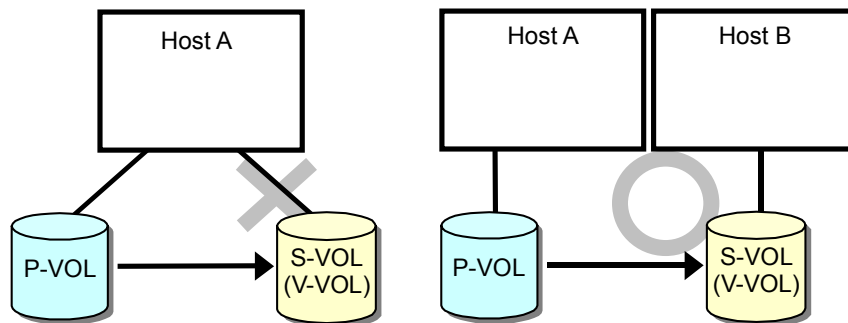
Note: A dynamic disk is a Windows Server 2000/2003 feature.

4.7.1.2 Restrictions


Observe the following when you use a dynamic disk with Windows Server 2003.

Follow the restrictions of the copy functions (ShadowImage, Snapshot, TrueCopy, and TCE) in addition to the following restrictions. For the restrictions of the each copying function, refer to the corresponding manual.

- When a secondary host uses a V-VOL, verify that the host recognizes it after making sure that the pair status is PSUS after the pair is created.
- After the pair is created in the secondary host, PSUS
- One host cannot recognize both a P-VOL and a V-VOL.



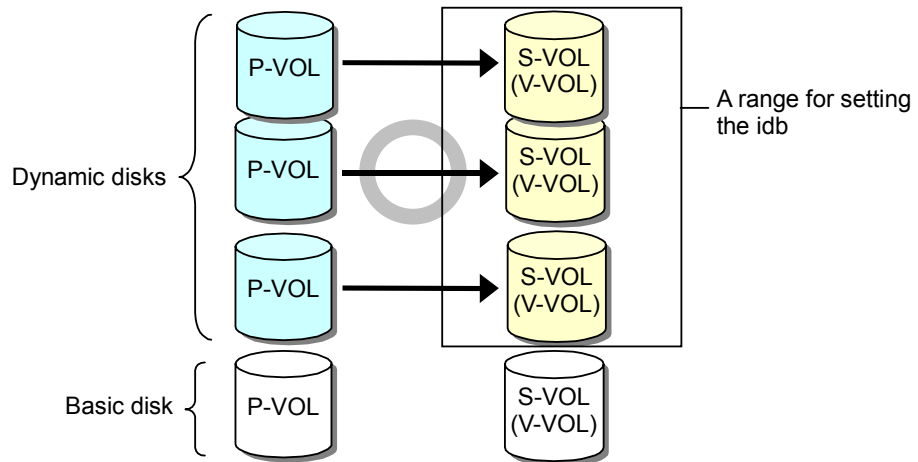
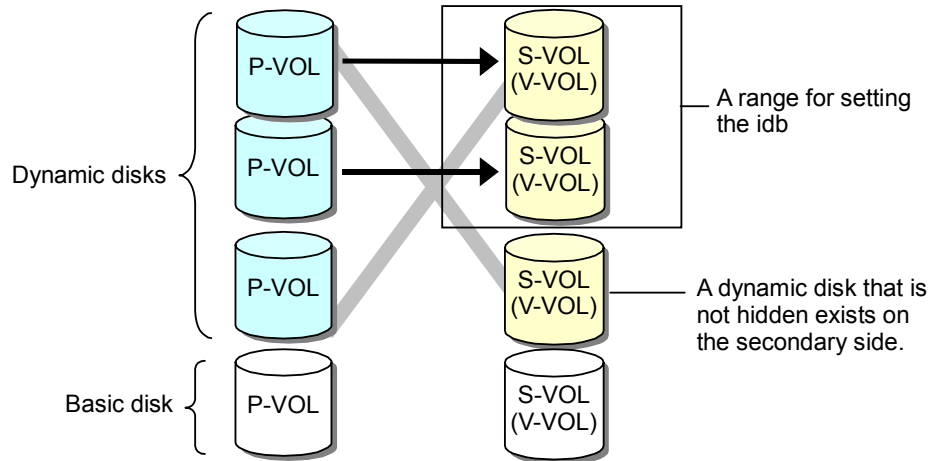
- An LU, in which two or more dynamic disk volumes coexist, cannot be copied.



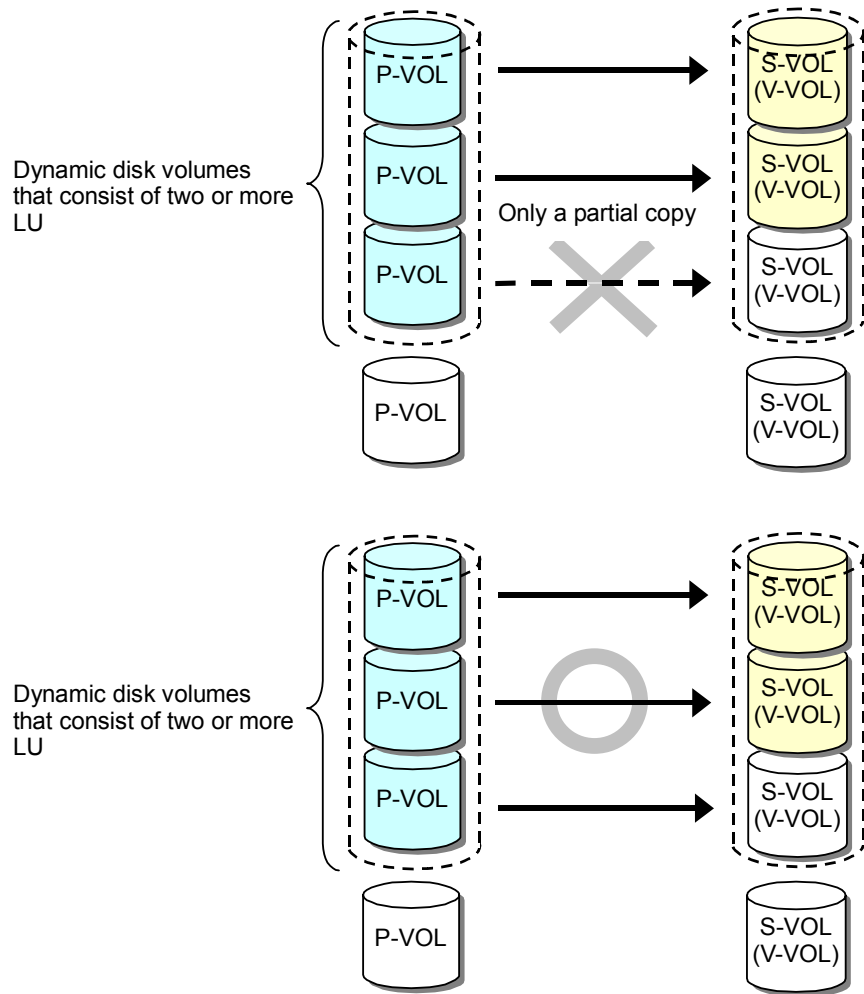
- Do not use a dynamic disk function for volumes other than an S-VOL (V-VOL) on the secondary host side.

When copying, hide all the dynamic disks exist on the primary side using `raidvchkset -vg idb`. 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.)

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.



- When copying dynamic disk volumes that consist of two or more LUs, do this after hiding all LUs composing the dynamic disk at the same time. After the copy is completed, release LUs from being hidden and have them recognized by a host.



- A dynamic disk cannot be used with a cluster (MSCS, VCS, etc.)
- A dynamic disk cannot be used with VxVM and HDLM.

4.7.2 Applying Snapshot to an Existing System

When applying Snapshot to an existing system, follow the description provided in this chapter and assign the appropriate configuration.

4.8 Configuration Restrictions

- LU Mapping and Snapshot Configuration
 - When Mapping mode is active the P-VOL and V-VOL cannot be paired. Use the LUN Manager if you do not want them recognized by a host.
- Cluster Software, Path Switching Software, and Snapshot Configuration
 - Do not make the V-VOL an object of the cluster software and the path switching software.
- MSCS and Snapshot Configuration
 - When setting the V-VOL to be recognized by the host, use the `CCI mount` command instead of using the Disk Administrator.
 - The same host cannot recognize both the P-VOL and V-VOL. Have the host recognize the P-VOL and let another host recognize the V-VOL.
 - Do not place the MSCS Quorum Disk in CCI
 - Shutdown MSCS before executing `sync`
- AIX and Snapshot Configuration
 - The same host cannot recognize both P-VOL and V-VOL. Have the host recognize only P-VOL and let another host recognize V-VOL.
 - Multiple V-VOLs per P-VOL cannot be recognized from the one host. Make V-VOL limit recognition from a host to only one V-VOL per P-VOL.
- VxVM and Snapshot Configuration
 - If you have set the P-VOL and the V-VOL to be recognized by the same host, the VxVM will not operate properly. Set only the P-VOL to be recognized by the host and let another host recognize V-VOL.
- Windows 2000 and Snapshot Configuration
 - Multiple V-VOLs per P-VOL cannot be recognized from the one host. Make V-VOL recognized to limit recognition from a host to only one V-VOL per P-VOL.
- Windows 2003 and Snapshot Configuration
 - Multiple V-VOLs per P-VOL cannot be recognized from the one host. Make V-VOL recognized to limit recognition from a host to only one V-VOL per P-VOL.
 - When mounting a volume, use `Volume{GUID}` as an argument of the mount command of CCI. The `Volume{GUID}` can be used by CCI 01-13-03/00 or later. (For more detail, see the *Hitachi TagmaStore Adaptable Modular Storage and Workgroup Modular Storage Command Control Interface (CCI) User and Reference Guide (MK-95DF701)*).
 - When describing a command device in the configuration definition file, specify it as `Volume{GUID}`. (For more detail, see the *Hitachi TagmaStore Adaptable Modular Storage and Workgroup Modular Storage Command Control Interface (CCI) User and Reference Guide (MK-95DF701)*).

- 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 “re-scan the disks” in Windows to make the recovery. Restart CCI if Windows cannot access the command device even if CCI is able to recognize it.
- Linux and LVM Configuration
 - If you have set the P-VOL and the V-VOL to be recognized by the same host, the LVM will not operate properly. Set only the P-VOL to be recognized by the host and let another host recognize V-VOL.
- Tru64 UNIX and Snapshot Configuration
 - When rebooting the host, it may take some time before host is up if V-VOL is in a state other than PSUS. When rebooting the host, be sure that V-VOL is PSUS, or do not have the host recognize V-VOL that is not PSUS by using LUN Manager.

Chapter 5 Performing SnapShot Operations

This chapter provides examples of Snapshot commands using the Windows 2000 systems. The icons of the pairs, which indicate pair status in Navigator, are also displayed.

- Snapshot Operations (section 5.1)
- Setting the Command Device (section 5.2)
- Setting the Differential Management LU (section 5.3)
- Setting the Pool (section 5.4)
- Setting P-VOL and V-VOL (section 5.5)
- Setting Mapping Information (section 5.6)
- Differential Data (section 5.7)
- Snapshot Pair Status (section 5.8)
- Concurrent Use of LUN Expansion (section 5.9)
- Cascade Configurations (section 5.10)
- Cascade Restrictions (section 5.11)
- Cascade Connection of TCE with SnapShot (section 5.12)
- Creating Snapshot Pairs (section 5.13)
- Updating the V-VOL (section 5.15)
- Restoring a V-VOL to the P-VOL (section 5.16)
- Releasing Snapshot Pairs (section 5.17)
- Confirming Pairs Status (section 5.18)

5.1 Snapshot Operations

This chapter provides examples of Snapshot commands using the Windows 2000 systems. The icons of the pairs, which indicate pair status in Navigator, are also displayed.

To execute Snapshot commands, from the host where CCI is installed, display the command prompt.

Snapshot operations can be performed from the UNIX/PC server host using CCI software. For further information about CCI, refer to the *Hitachi TagmaStore Adaptable Modular Storage and Workgroup Modular Storage Command Control Interface (CCI) User and Reference Guide (MK-95DF701)*.

Table 5.1 Snapshot Functions

No.	Function	Contents
1	Pair Creation	Supported For P-VOL: Read/Write For V-VOL: Read only
2	Updating V-VOL data	Supported For P-VOL: Read/Write For V-VOL: Read/Write
3	Restoration	Supported For P-VOL: Read/Write For V-VOL: Does not accept I/O operations (Read/Write)
4	Pair releasing	Supported

5.2 Setting the Command Device

The Command Device is a user-selected, dedicated logical volume on the TagmaStore subsystem, which functions as the interface to the CCI software. The Snapshot commands are issued by CCI (HORCM) to the TagmaStore Command Device.

A Command Device must be designated in order to issue Snapshot commands.

The Command Device must be defined in the HORCM_CMD section of the configuration definition file for the CCI instance on the attached host (see section 4.4). Two Command Devices can be designated for the TagmaStore subsystem. You can designate Command Devices using Navigator.

Notes:

- LUs set for Command Devices must be recognized by the host. The Command Device LU size must be greater than or equal to 33 MB.
- The following restrictions apply when either a pair of ShadowImage, Snapshot or TrueCopy exists or a TrueCopy path is defined.
 - When two command devices are set, only one command device can be released.
 - When only one command device is set, the command device cannot be released.
- Locating Command Devices
 - When two command devices are set within the one disk subsystem, assign them to the respective RAID groups; if they are assigned to the same RAID group, both command devices become unavailable due to an error such as a drive failure.

To designate Command Device(s):

1. Start Navigator and change the operation mode to **Management Mode** (administrator mode).
2. Connect to the subsystem. The **Array System Viewer** window (see Figure 5.1) opens displaying the connected subsystem.
3. Click **Logical Status** tab.
4. Click **Command Devices**.

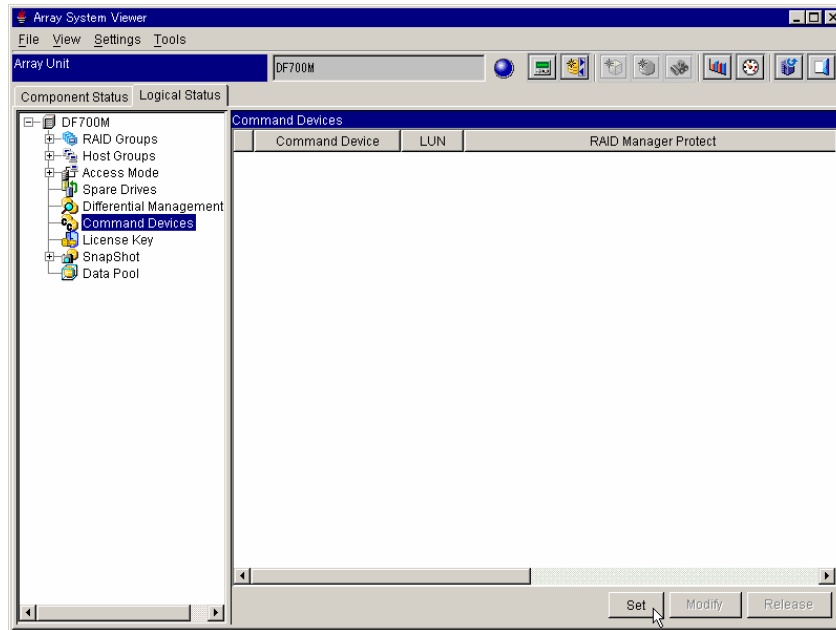


Figure 5.1 Array System Viewer Window (Command Device Page: Before Setting)

5. Click **Set**.

The **Command Devices Settings** dialog box displays.

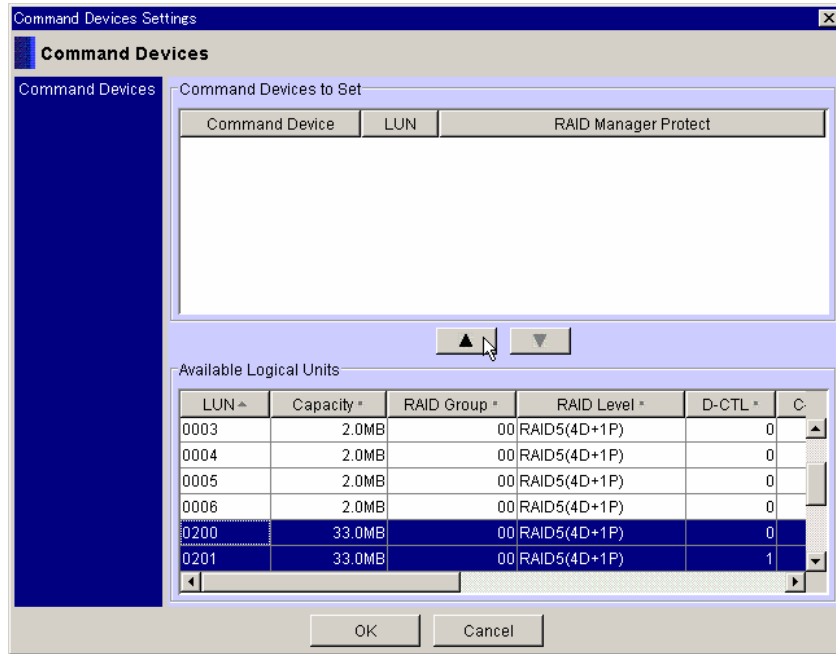



Figure 5.2 Command Devices Settings Dialog Box (Before Setting)

- In **Available Logical Units** list, select the LUN you want to set on the command devices and click .

The selected LUN moves to the **Command Devices to Set** list.

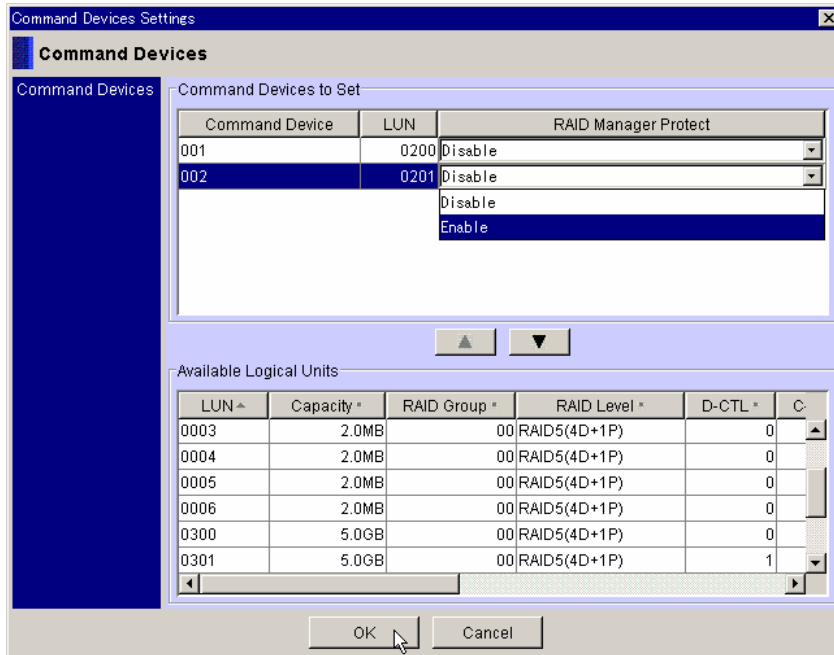



Figure 5.3 Command Devices Settings Dialog (After Setting)

- When you want to use the protection function of CCI, in the **RAID Manager Protect** drop-down list, click **Disable** or **Enable**.

When you want to change the command devices, select the LUN on the **Command Devices to Set** list and click .

The selected LUN moves to the **Available Logical Units** list.

8. Click **OK**.
The setting information displays.

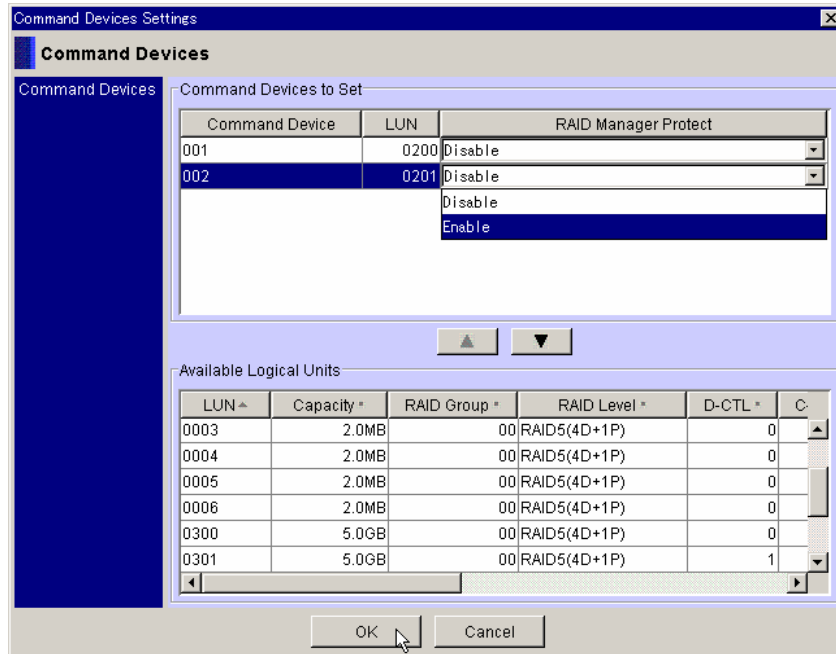


Figure 5.4 Array System Viewer Window (Command Device Page: After Setting)

Note: To use the alternate Command Device function, or to avoid data loss and subsystem downtime, designate two Command Devices. When two command devices are set within one disk subsystem, assign them to the respective RAID groups. If they are assigned to the same RAID group, both command devices become unavailable because of difficulties like drive failure. For details on the alternate Command Device function, refer to the *Hitachi TagmaStore Adaptable Modular Storage and Workgroup Modular Storage Command Control Interface (CCI) User and Reference Guide (MK-95DF701)*.

9. To release an already set command device, select the LUN off or the command device you want to release and click **Release**.
10. A message displays. Click **OK**.

Designating the Command Device is now complete. Leave the **Array System Viewer** window open and go to the next section (section 5.3).

5.3 Setting the Differential Management LU

The Differential Management LU is an exclusive logical unit for storing the differential data when the volume is copied. The Differential Management LU in the disk subsystem is treated the same way as the other logical units. However, a logical unit that is set as the Differential Management LU is not recognized by the host (it's hidden).

You must set a logical unit with at least 5 GB as the Differential Management LU. Up to two Differential Management LUs can be set and the second is used for mirroring. We recommend that you set two Differential Management LUs.

To designate Differential Management LUs:

1. Click **Logical Status** tab on the **Array System Viewer** window.
2. Click **Differential Management LU**.

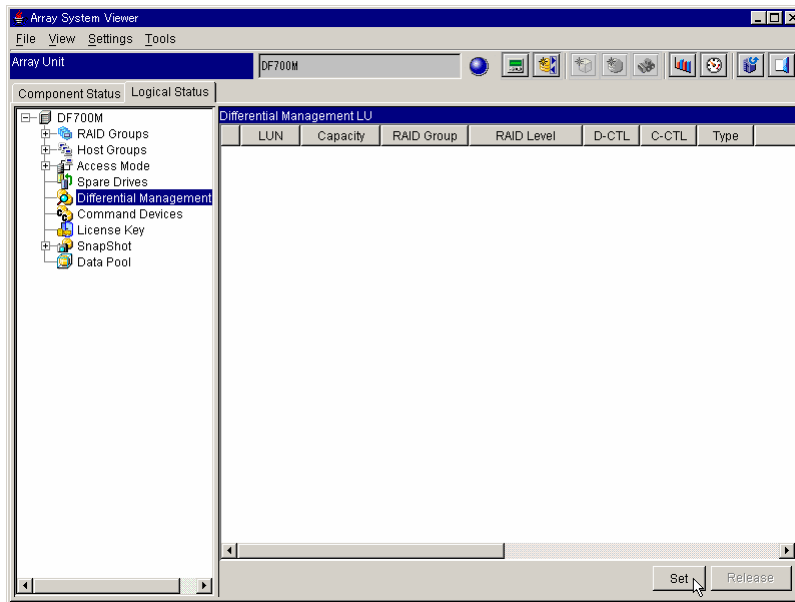


Figure 5.5 Array System Viewer Window (Differential Management LU Page: Before Setting)

3. Click **Set** and the **Select Logical Unit** dialog box is displayed.

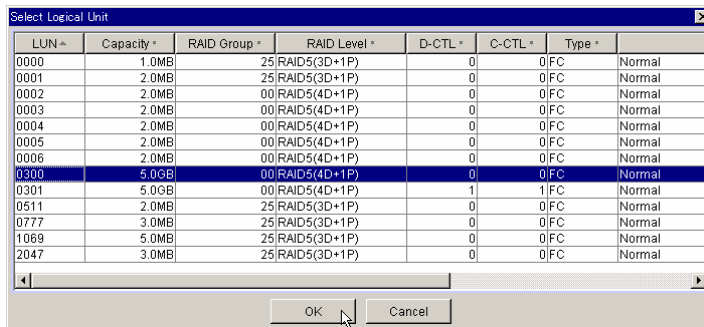


Figure 5.6 Select Logical Unit Dialog Box

4. Select the LUN you want to set differential management LU and click **OK**.
 5. A message appears. Click **OK**.
- The setting information displays.

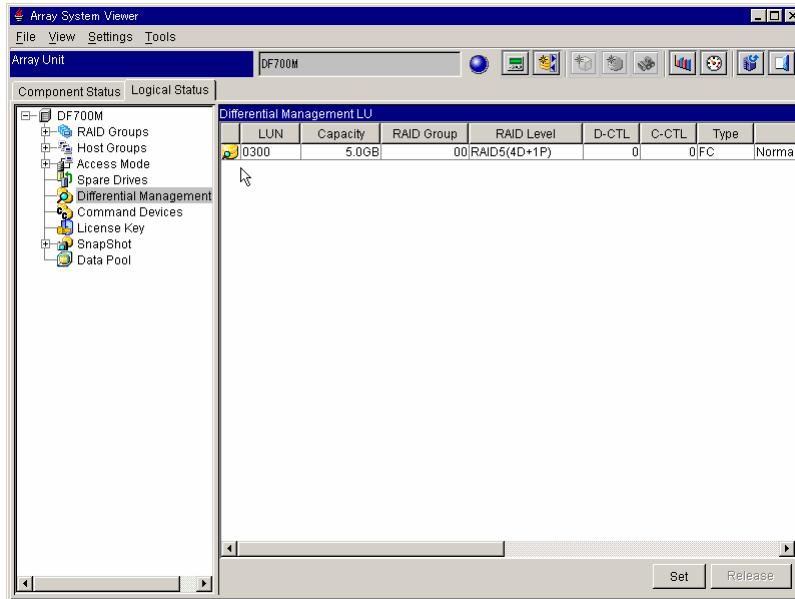


Figure 5.7 Array System Viewer Window (Differential Management LU Page: After Setting)

Designating differential management LU is now complete.

Note: The following restrictions apply when either a pair of ShadowImage, Snapshot or TrueCopy, or TCE exists, the path of True Copy or TCE is defined, or POOL of Snapshot is defined.

- When two differential management LUs are set, only one differential management LU can be released.
- When only one differential management LU is set, the differential management LU cannot be released.

5.4 Setting the POOL

A single data pool must be created for each controller, 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.
- An LU with an FC 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. On the **Array System Viewer** window, click the **Logical Status** tab. The **Logical Status** window displays in front.
2. Click the **Data Pool** icon in the left pane.
3. Select **Data Pool 0** or **Data Pool 1**. Select **Add Logical Unit** from the context menu when an icon for the **Data Pool** is clicked with the right mouse button. (**Data Pool 0** is for Controller 0 and **Data Pool 1** is for Controller 1.)

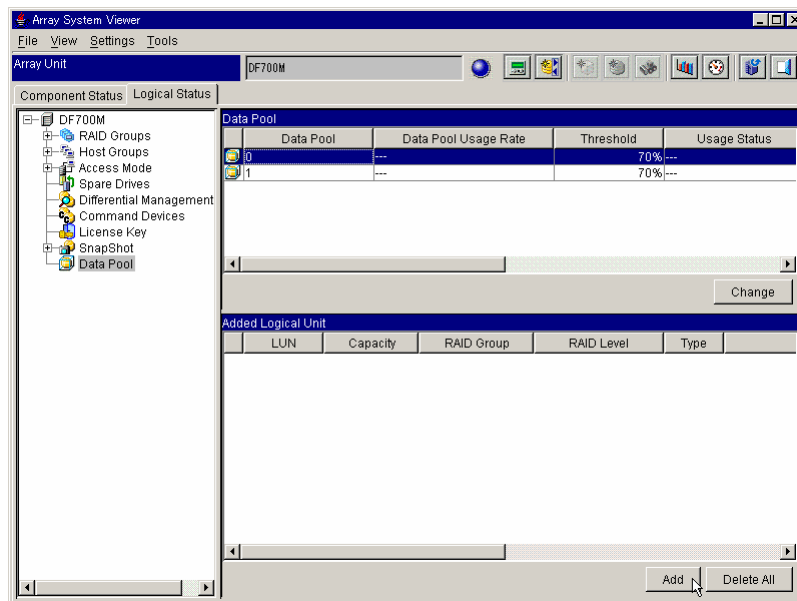


Figure 5.8 Setting the POOL

The Add Logical Unit dialog box displays.

4. Select a logical unit for the POOL and click **OK**.

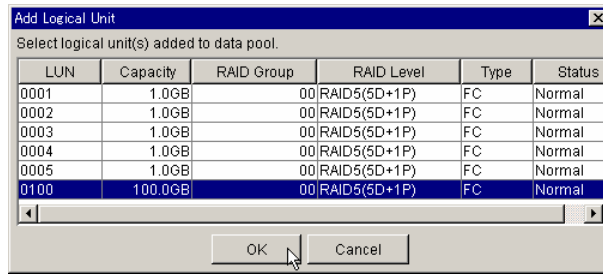


Figure 5.9 Add Logical Unit Dialog

5. Click **OK** in the confirmation dialog box to accept the POOL configuration. The added POOL updates and the following window displays.

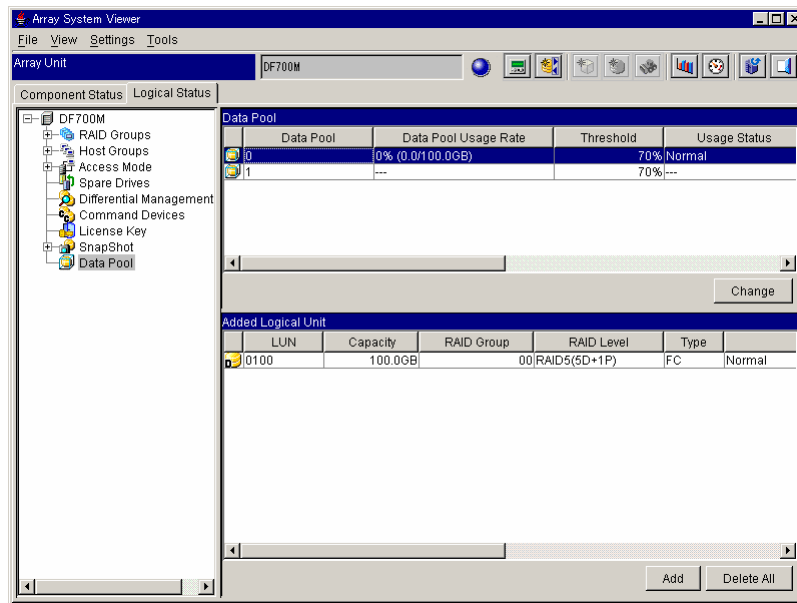


Figure 5.10 Added POOL Updates

6. To add a logical unit to the POOL, repeat Step 3.
7. To delete a logical unit that has been assigned to the Data Pool, select the LUN on the **Added Logical Unit** and click **Delete All**.

Note: When deleting the logical unit set as the POOL, it is necessary to delete all Snapshot images (V-VOLs).

8. A message appears, asking you to confirm that you want to delete the POOL. Click **OK**.

- To change the threshold value of the Data Pool, click **Change** (see Figure 5.10).
The **Property** dialog box displays.

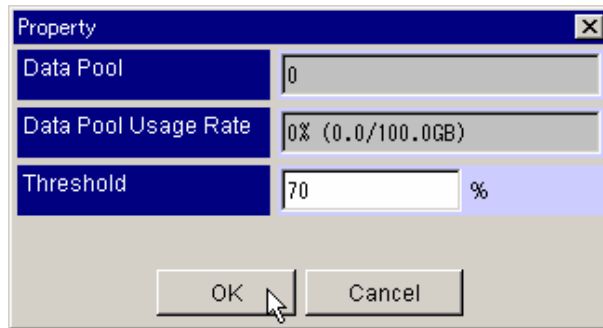


Figure 5.11 Property Dialog Box

Note: The subsystem will warn you when the used capacity of the POOL is Threshold +1% and the pair status will become PFUS.

- The default threshold value is 70. Specify any value from 50 through 80 and click **OK**.

5.5 Setting P-VOL and V-VOL

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

To set the V-VOL, specify the logical unit that will be assigned to a P-VOL. If a specification for the logical unit assigned to a V-VOL is omitted when setting the V-VOL, Navigator assigns the smallest undefined number to the logical unit.

Note: The controller that has ownership of the V-VOL (a controller that controls the V-VOL) is the same controller that controls the P-VOL and POOL.

To set the P-VOL and V-VOL:

1. On the **Array System Viewer** window, click **Logical Status** tab. The **Logical Status** window displays.
2. Click the **Pair** icon in the left pane. Click **Add**.

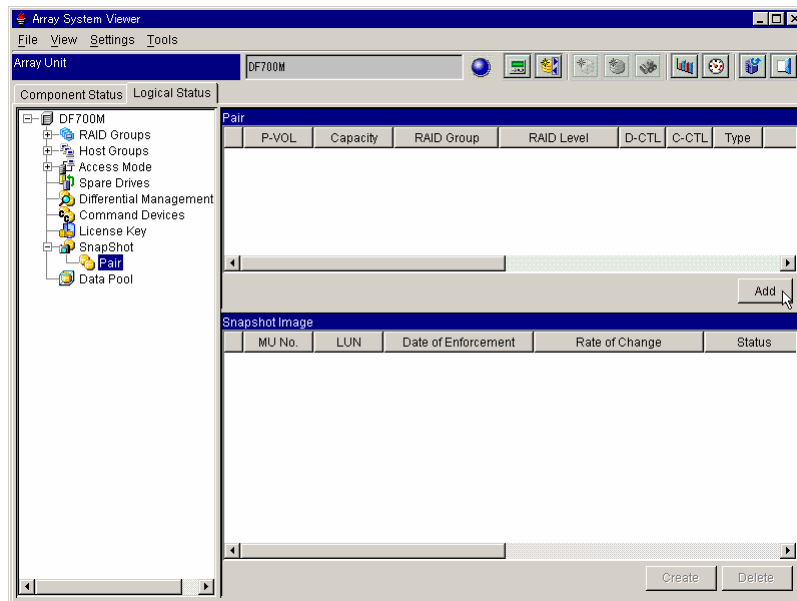


Figure 5.12 Setting the P-VOL and SnapShot Image (V-VOL)

The **Add Pair** dialog box displays.

3. Select a LUN for the P-VOL and click **OK**.

Regarding the number of the logical unit for the V-VOL, Navigator assigns the smallest undefined number. If you want to specify a logical unit, enter the logical unit number.

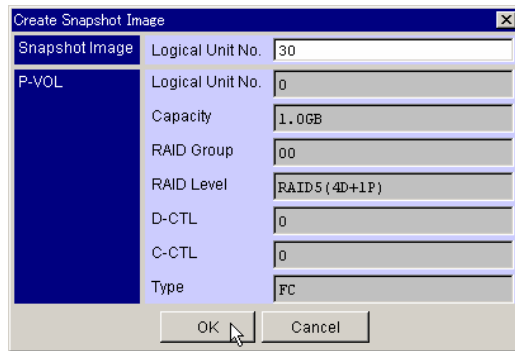
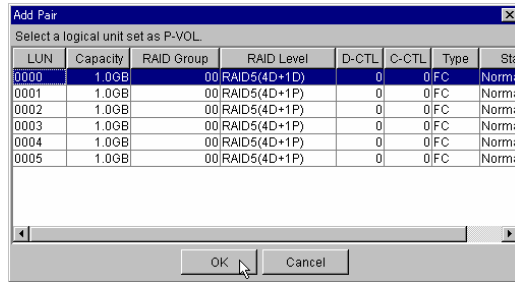


Figure 5.13 Add Pair Dialog Box

- Click **OK** in the confirmation dialog box, asking you to confirm that you want to create the Snapshot Image.

The created Snapshot Image updates and the following window displays.

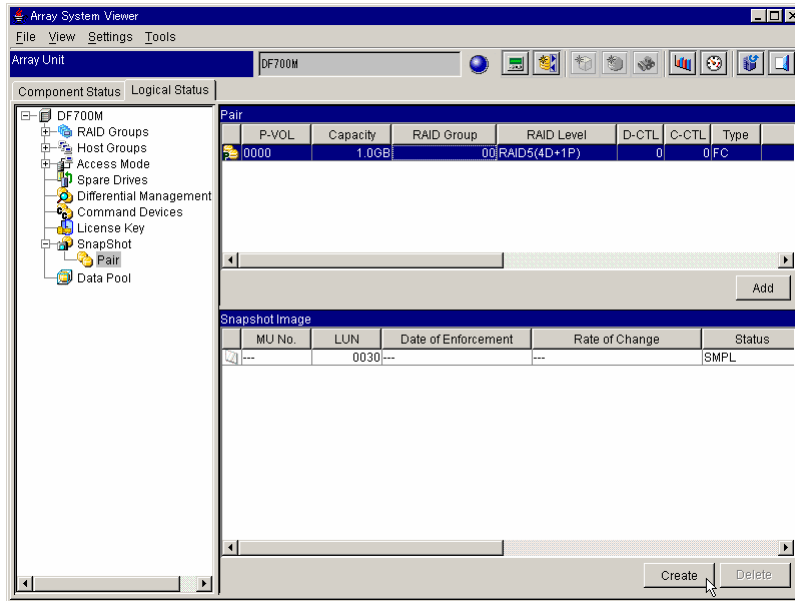


Figure 5.14 Adding the P-VOL and SnapShot Image (V-VOL)

- To add a V-VOL, select the P-VOL icon in the right pane, then click **Create**. The **Create Snapshot Image** dialog box displays.
- Navigators assigns the smallest undefined number; if you approve, click **OK**. When you want to specify a number, enter the logical unit number.

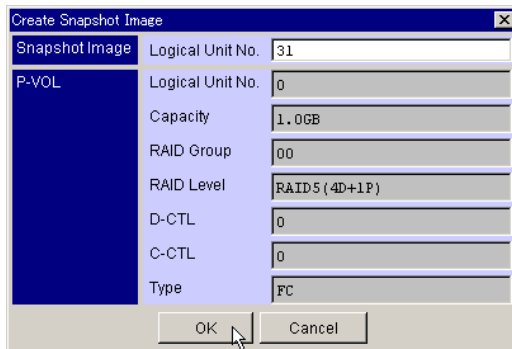


Figure 5.15 Create Snapshot Image Dialog

The added Snapshot image updates and the following window displays.

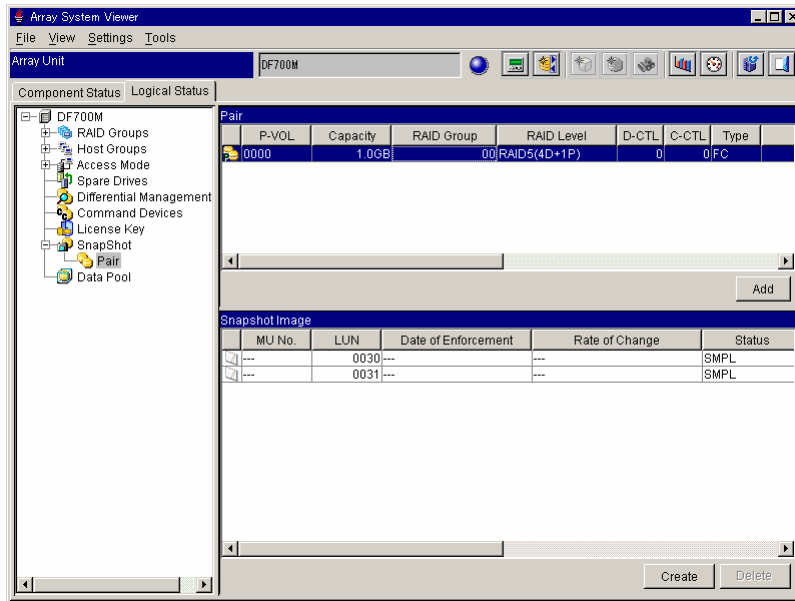


Figure 5.16 Added Snapshot Image Updates

- To delete a logical unit that has been assigned to the Snapshot Image, select the **P-VOL** on the **Pair** and select **LUN** on the Snapshot Image. Click **Delete**.

Note: In order to delete the V-VOL, the pair state must be **SMPL**.

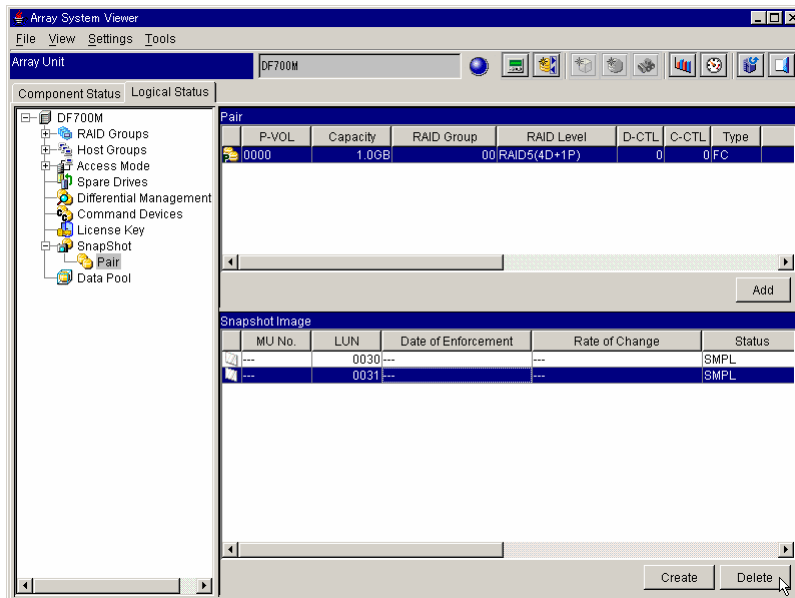


Figure 5.17 Deleting the SnapShot Image (V-VOL)

- Click **OK** in the confirmation dialog box to delete Snapshot Image.

5.6 Setting Mapping Information

This section describes the steps necessary for setting the Mapping Information.

Note: When Mapping mode is active the P-VOL and V-VOL (whose mapping information is not set for the port that was set in the configuration definition file and is not recognizable from a host) cannot be paired. When you do not want to have them recognized by a host, use the LUN Manager.

5.6.1 Selecting Mapping Mode

1. On the **Array System Viewer** window, click the **Logical Status** tab. The **Logical Status** window is displayed in front.
2. Click the **Access Mode** plus signs next to the **Mapping Mode**.

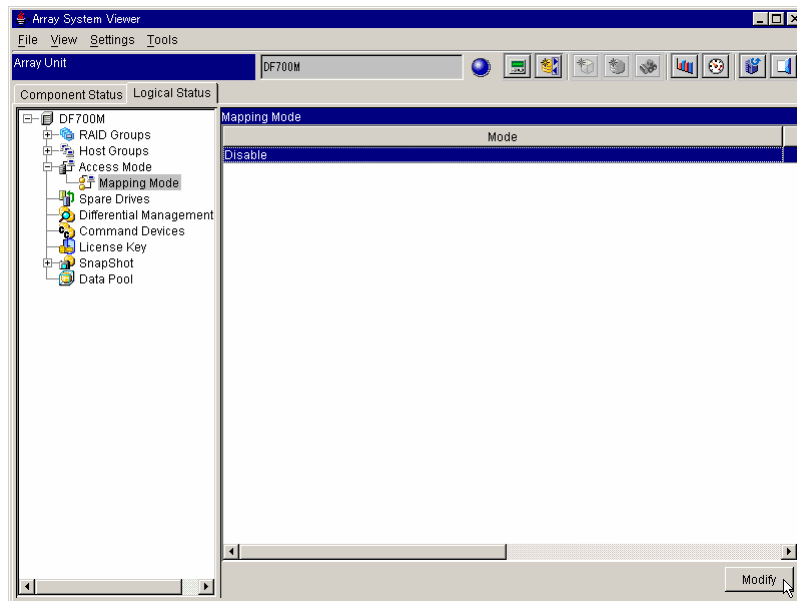


Figure 5.18 Array System Viewer Window (Mapping Mode)

3. Select **Disable** in **Mapping Mode** list and click **Modify**.




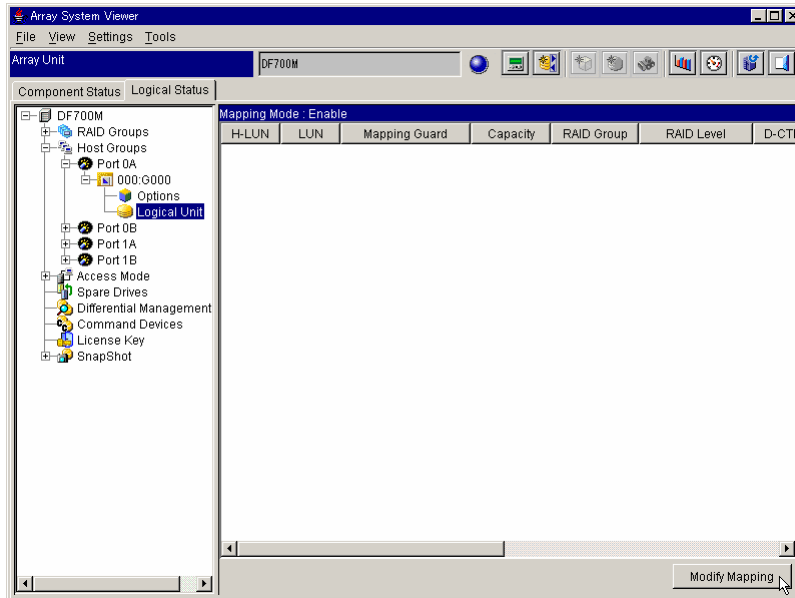
Figure 5.19 Mapping Mode Dialog Box

4. Select **Enable** and click **OK**.
5. Click **OK** in the confirmation dialog box.

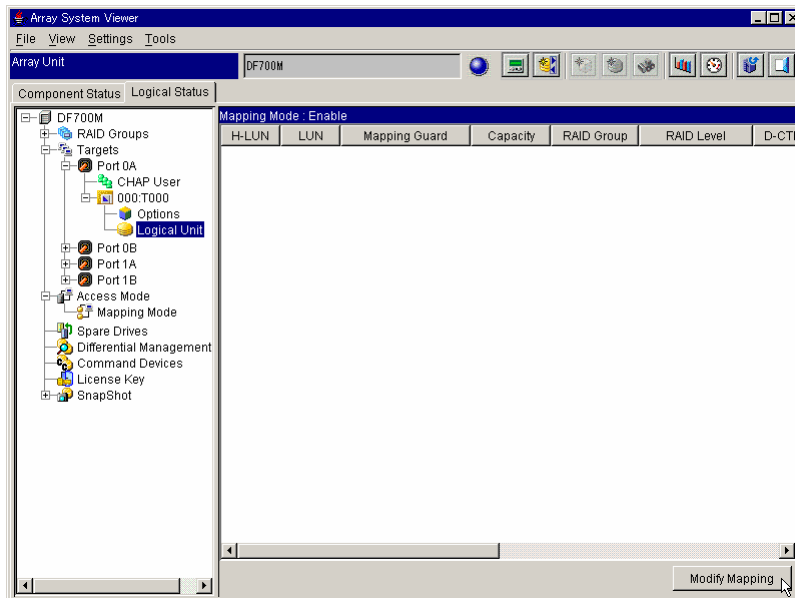
5.6.2 Setting Mapping Information

1. Click **Port**, then select either of the signs next to **000:G000** or **000:T000**.
2. Select **Logical Unit**. Click **Modify Mapping**.

The **Mapping Property** window displays.

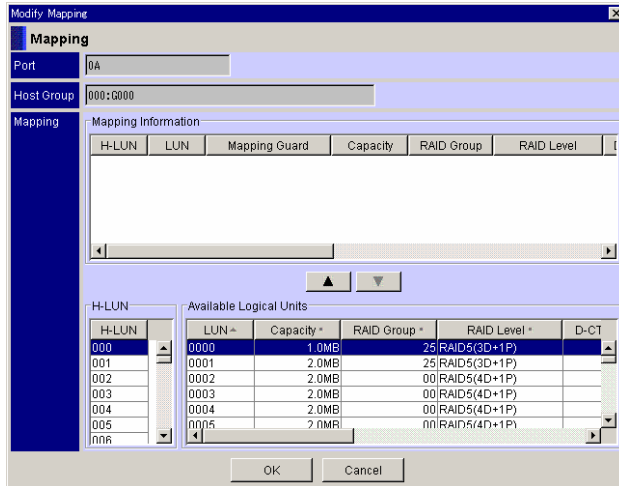


(Host Group)

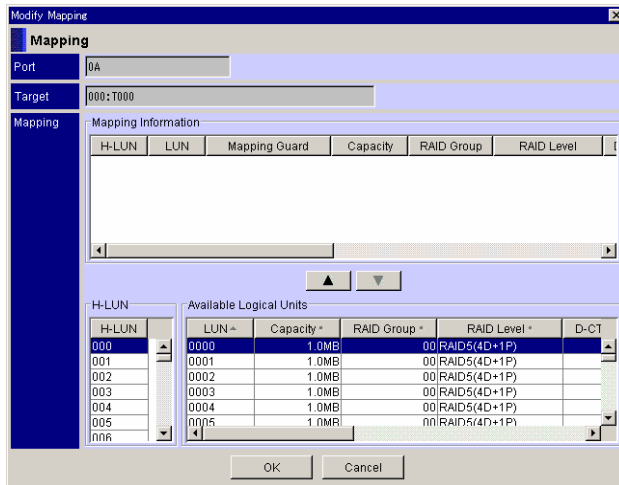


(Target)

Figure 5.20 Array System Viewer Window (Setting Mapping Information)




(Host Group)



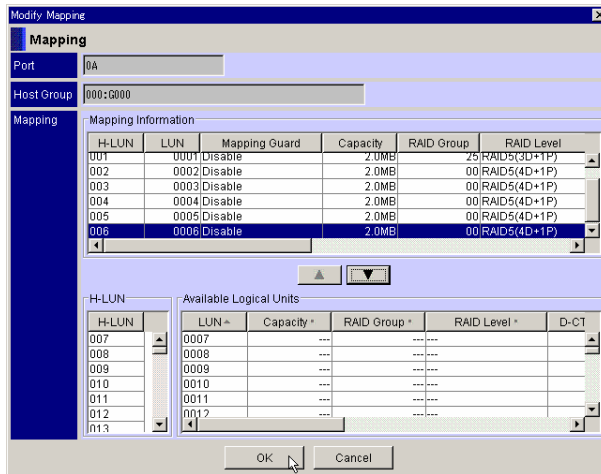
(Target)

Figure 5.21 Mapping Property Window (Before Setting)

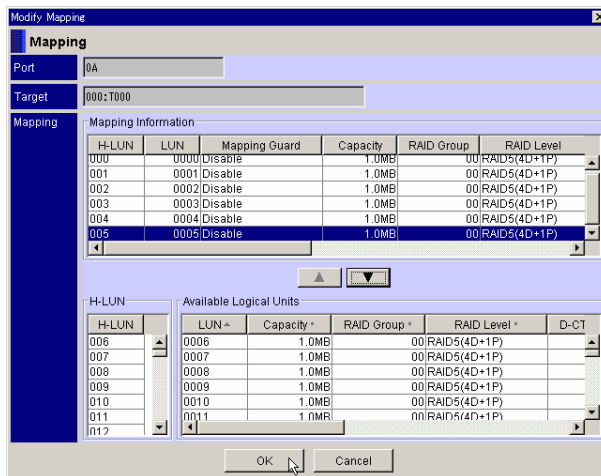
3. Select one H-LUN to add. Select LUN and click . Added contents display in Mapping Information.

Select the following items:

- For LUN, select the subsystem LU number.
- For Host LUN, select an LU number the host can recognize.



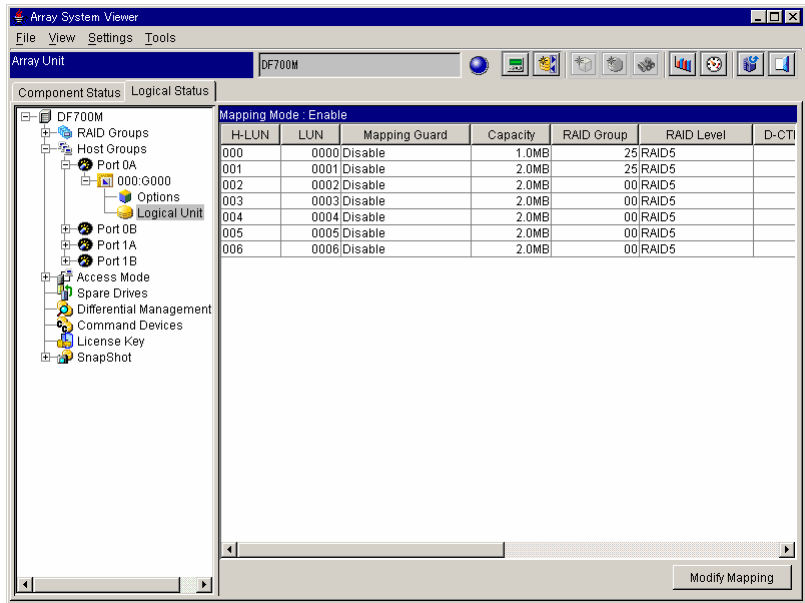
(Host Group)



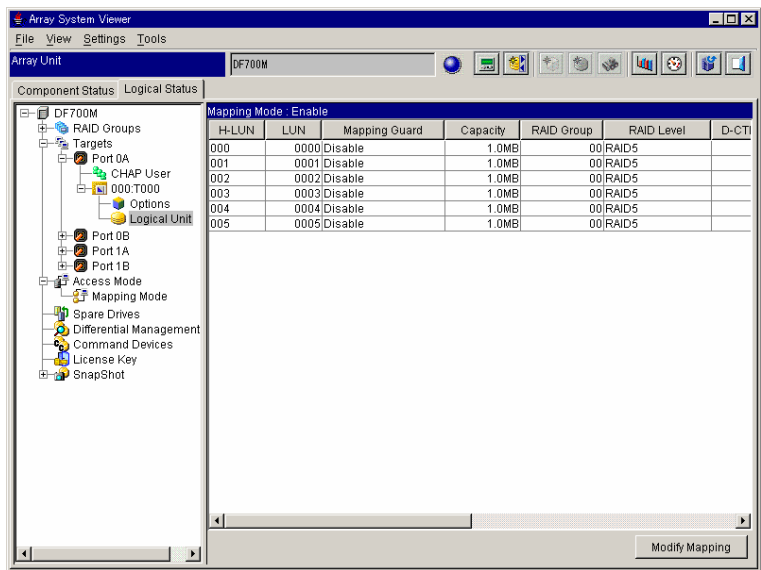
(Target)

Figure 5.22 Mapping Property Window (After Setting)

4. Click OK.
 5. Click OK in the confirmation dialog box.
- The set mapping information updates and the following window displays.



(Host Group)



(Target)

The mapping information settings are complete.

5.6.3 Paircreate Operation

Use the `paircreate` command to create Snapshot pair(s). A Snapshot pair is created the moment `paircreate` is issued and the pair status becomes PAIR. When issuing `paircreate`, a volume assigned to the V-VOL must be in the SMPL (simplex) status.

When `paircreate` is issued with the `-split` option added to it, the pair status becomes PSUS.




Figure 5.23 Creating a Snapshot Pair

Note: For details on the `paircreate` operation, see section 5.1.

5.6.4 Updating V-VOL

When executing the Snapshot instruction (a data update) to the created Snapshot pair, change the pair status from PSUS to PAIR using the `pairresync` command; then change it to PSUS using the `pairsplit` command. The P-VOL data that exists when `pairsplit` is issued is retained in the V-VOL in the same way as when `paircreate -split` is issued.

For details on `pairresync` operation, see section 5.15.

5.6.5 Restoration

A `pairresync -restore` command is used when restoring backup data retained in the V-VOL to the P-VOL. When a restoration instruction is given, the P-VOL data is immediately replaced with the backup data retained in the V-VOL. However, while the backup data in the V-VOL is being physically duplicated in the P-VOL, the pair status is changed to **COPY (RS-R)**. For a V-VOL being restored while the pair status is **COPY (RS-R)**, no Snapshot or Read/Write instruction can be executed. When the ratio P-VOL: V-VOL is 1: n ($n > 1$), Read/Write operation from/to the other V-VOL cannot be performed in the same way as the V-VOL that is being restored. When the restoration is completed, reading/writing from/to the other V-VOL in the **PSUS** status becomes possible again. When the V-VOL in the **PSUS** status becomes readable and writable again, it retains the data it had at the time of the Snapshot instruction, and is not affected by the restoration of the V-VOL data to the P-VOL.

The pair can be split while its status is **COPY (RS-R)**, however, the P-VOL data being restored cannot be used logically. The V-VOLs correlated to the P-VOL and with a status other than **SMPL** are placed in **PSUE** status. Do not split a pair while the pair status is **COPY (RS-R)**, except when an emergency exists.

When the restoration instruction of the V-VOL data is issued to the P-VOL, the pair status is not changed to **PAIR** immediately but is changed to **COPY (RS-R)**. The data of the P-VOL, however, is replaced promptly with the backup data retained in the V-VOL. When the Snapshot instruction is issued to the other V-VOL after the issue of the restoration instruction, the V-VOL retains the data it had at the time of the Snapshot instruction based on the P-VOL data. That data is replaced with the backup data, even before the restoration completes.

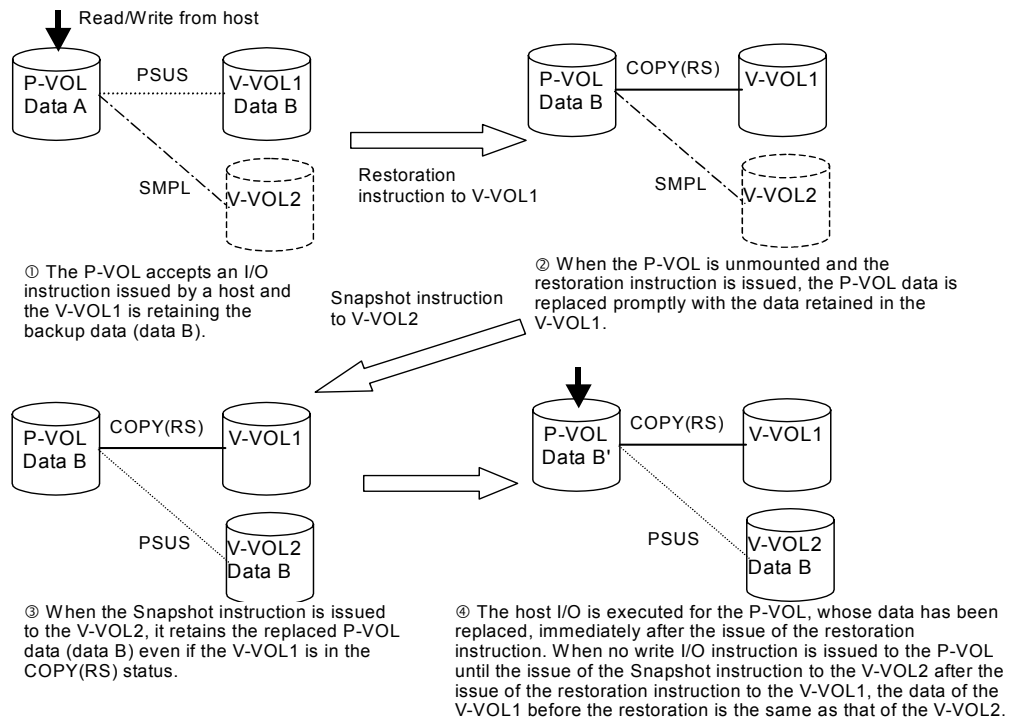


Figure 5.24 Operation Example when Snapshot Instruction is Issued to the Other V-VOL During the Restoration

5.6.6 Pair Failures

When the state of V-VOL data is unstable as a result of a failure (such as the double failure of drives or the POOL capacity has exceeded the limit within the disk subsystem), the pair is placed in PSUE status. When an event such as the above occurs during a restoration process, the P-VOL being restored becomes unable to accept Read/Write instructions and the V-VOL data becomes invalid. To resume the pair that was not retained, it is necessary to split the pair once using `pairsplit -S` and execute the Snapshot instruction `paircreate -split` again. However the V-VOL that is created does not have the previously invalidated data but contains the P-VOL data at the time of the new Snapshot instruction. When the disk subsystem places a pair in the PSUE status, CCI is output to the file of the system log or event log in order to inform a host of the operation.

In the following situation, the disk subsystem places a pair in the PSUE status:

- When a POOL in the disk subsystem cannot be accessed due to a double drive failure as a result of a disk subsystem failure, the disk subsystem changes a pair in **PSUS** status to **PSUE** status..
- When the differential data cannot be saved because no additional free capacity from the Data Pool is available, the disk subsystem places all pairs in **PSUE** status.

5.6.7 Pair Splitting

When the `-S` option is used with `pairsplit`, any pair in the PAIR, PSUS, COPY(RS-R), or PSUE status can be split at any time after it is placed in the SMPL status.

When the command is executed, the V-VOL data is annulled immediately and invalidated. Therefore, if you access the V-VOL after the pair is split, the data retained before the pair is split is not available.

For details on `pairsplit -S`, see section 5.17.

5.7 Differential Data

When a Snapshot operation is performed, the V-VOL is assured through management of differential data (locations of Write data from a host contained in the P-VOL and V-VOL) and reference to the bit map performed in a Read operation by the host (Figure 5.25 provides differential data). The extent of one bit in the bit map is equivalent to 64 KB. Therefore:

- Even an update of a single KB update requires a data transfer as large as 64 KB to copy from the P-VOL to the POOL.
- The copy amount becomes less when the locality of the host access pattern is high.

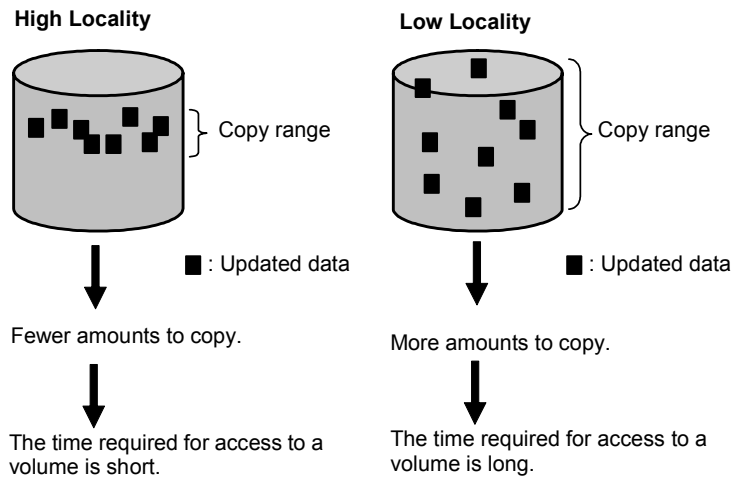


Figure 5.25 Differential Data

5.8 Snapshot Pair Status

Snapshot displays the pair status of all Snapshot volumes (LUs). Figure 5.26 status and the Snapshot operations.

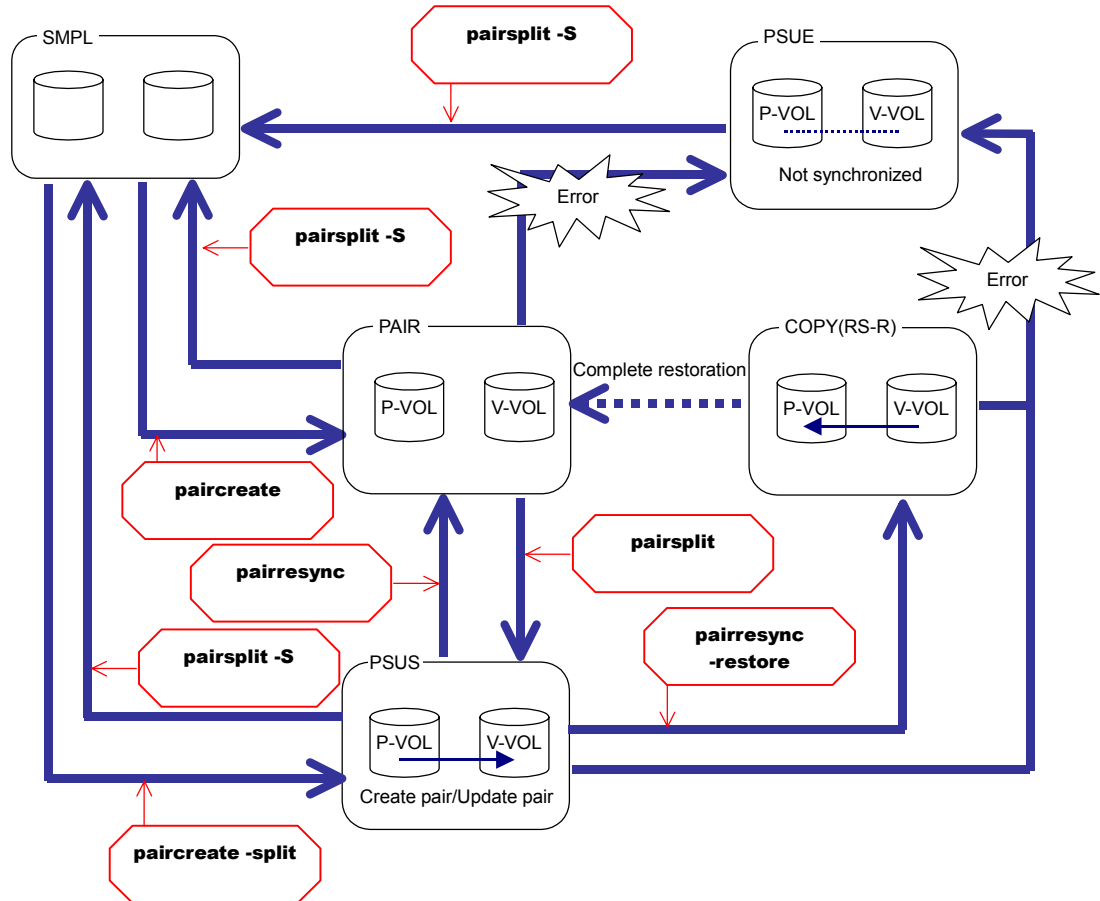


Figure 5.26 Snapshot Pair Status Transitions

Table 5.2 lists and describes the Snapshot pair status conditions.

If a volume is not assigned to a Snapshot pair, its status is SMPL. When you create a Snapshot pair by executing the `paircreate -split` or `pairsplit` operation after `paircreate` has been issued, statuses of the P-VOL and the V-VOL change to PSUS.

It is possible to access the P-VOL or V-VOL in the PSUS state. The pair status changes to PSUE (interruption) when the V-VOL cannot be created or updated, or when the V-VOL data cannot be retained due to a disk subsystem failure. When the `-S` option is specified by executing `pairsplit`, the pair is split and the pair status changes to SMPL.

Table 5.2 Snapshot Pair Status

Pair Status	Description	P-VOL	V-VOL
SMPL	This is a state in which no volume is assigned to a Snapshot pair. The P-VOL in the SMPL status accepts I/O operations of Read/Write. The V-VOL does not accept any Read/Write I/O operations.	Read and write	Does not accept I/O operations (read/write)
PAIR	PAIR is a pseudo status that exists in order to give interchangeability with the ShadowImage command system. The actual status is the same as the PSUS . In this state, access from a host to a P-VOL is lowered.	Read and write	Does not accept I/O operations (read/write)
COPY (RS-R)	COPY(RS-R) is a state in which the backup data retained in the V-VOL is being restored to the P-VOL. In this state, Read/Write I/O operations are accepted for the P-VOL as before (in the PSUS status). The V-VOL will not accept Read/Write I/O operations. Snapshot instruction cannot be executed. Pair status will be returned to PAIR after the restoration is completed. When a failure occurs or a pair is split during restoration, the V-VOLs status is correlated with the P-VOL status. Any status other than SMPL being restored becomes PSUE.	Read and write	Does not accept I/O operations (read/write)
PSUS (SSUS)	PSUS is a state in which the P-VOL data at the time of the Snapshot instruction is retained in the V-VOL. When a change of the P-VOL data occurs, the P-VOL data at the time of the Snapshot instruction is retained as the V-VOL data. The P-VOL and V-VOL in the PSUS status accept Read/Write I/O operations. However, the V-VOL does not accept any read/write instruction while the P-VOL is being restored. In this state, access from a host to a P-VOL is lowered.	Read and write	Read and write (A read/write instruction is not acceptable while the P-VOL is being restored.)
PFUS	PFUS is a state in which the used rate of POOL reaches the threshold of POOL. However, PFUS usually operates as PSUS . Only when -fc option is added in the <code>pairdisplay</code> command and -ss option is added in the <code>pairvolchk</code> command, are you able to recognize PFUS. (<code>pairvolchk</code> is recognized as returned values.)	Read and write	Read and write (A read/write instruction is not acceptable while the P-VOL is being restored.)
PSUE (Error)	PSUE is a state in which the P-VOL data at the time of the Snapshot instruction cannot be retained in the V-VOL due to a failure in the disk subsystem. In this state, I/O operations of Read/Write concerning the P-VOL are accepted as before (in the PSUS status). However, when a failure occurs during restoration, the P-VOL does not accept any Read/Write instruction. The V-VOL data has been invalidated at this point of time. To resume the split pair, it is required to execute the Snapshot instruction (<code>paircreate -split</code>) again after splitting the pair (using the <code>pairsplit -S</code>) once. However, data of the V-VOL created is not the former version that was invalidated, but the P-VOL data at the time of the new Snapshot instruction.	Read and write (The P-VOL does not accept a Read/Write instruction either when the pair status is PSUE due to a failure that occurred during restoration.)	Does not accept I/O operations (read/write)

For details on `pairdisplay` operation, see 4.7.2.

5.9 Concurrent Use of LUN Expansion

In Snapshot, a unified LU can be used as a P-VOL.

Preconditions and restrictions of concurrent use with Snapshot and LUN Expansion are as follows:

The unified LU cannot be assigned to a POOL. A V-VOL cannot be used to compose a unified LU because it is a virtual logical volume.

- Capacities of the P-VOL and POOL
 - Capacities of the P-VOL and POOL are not required to be the same.
 - A unified LU composed of 17 or more LUs cannot be assigned to the P-VOL. Each of all LUs composing the unified LU must have capacity of 1 GB or larger.

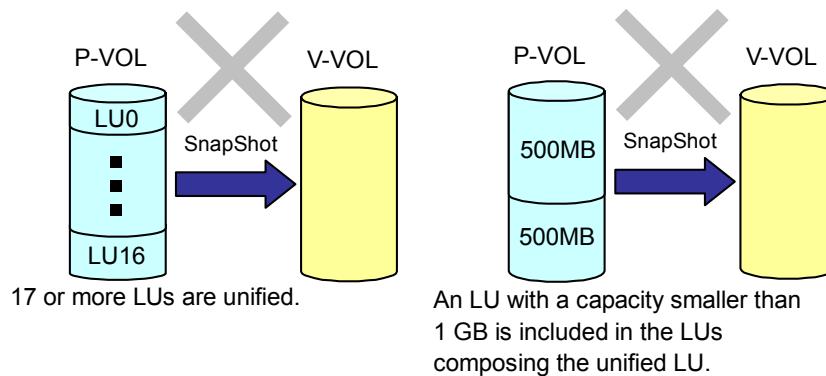


Figure 5.27 Example: Restriction Placed on Composing Unified LU

- RAID Levels of the P-VOL
 - A P-VOL, which is made of a unified LU consists of LUs on various RAID levels or with various numbers of data disks, can be paired. It is not required to make its RAID level and number of data disks the same as those of the POOL.

- Others
 - A P-VOL belonging to a Snapshot pair cannot be used to compose a unified LU. Unify the P-VOL after splitting the pair once and deleting all V-VOLs belonging to the P-VOL.
 - Although an LU assigned to a POOL cannot be unified, the POOL capacity can be increased by setting of up to 64 usual LUs. The LU(s) can be added to the POOL online. However, an LU with an FC drive and an LU with a SATA drive cannot coexist in a POOL.
 - An LU with an FC drive and an LU with a SATA drive cannot be unified. Unify LUs with drives of the same type.




Figure 5.28 Example: Restriction Placed on Uniting LUs with Different Drive Types

5.10 Cascade Configurations

The following cascade configurations are discussed in this section:

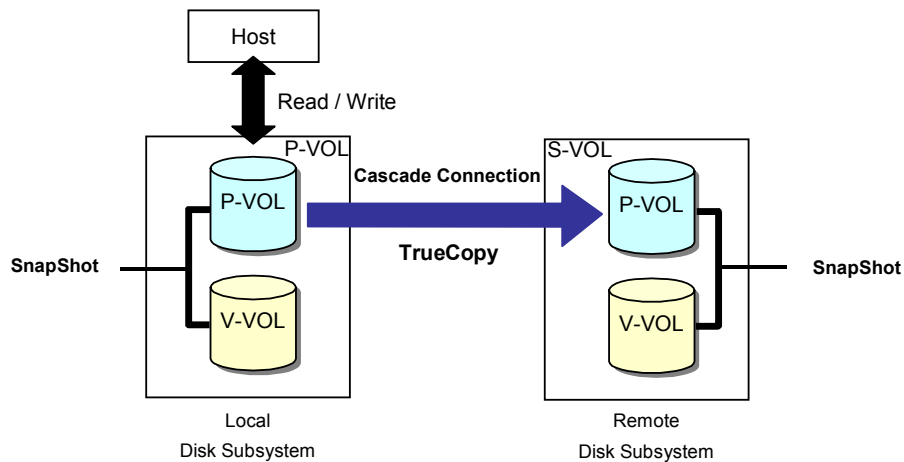
- Cascading TrueCopy with Snapshot
- Cascading with a P-VOL of Snapshot
- Cascading with a V-VOL of Snapshot

5.10.1 Cascading TrueCopy with Snapshot

TrueCopy volumes can be cascaded with those of Snapshot as shown in Figure 5.29. Because cascading Snapshot with TrueCopy lowers performance, use it only when it is necessary.

Note: Snapshot cannot be cascaded with ShadowImage.

- Cascade with a P-VOL of SnapShot



- Cascade with a V-VOL of SnapShot

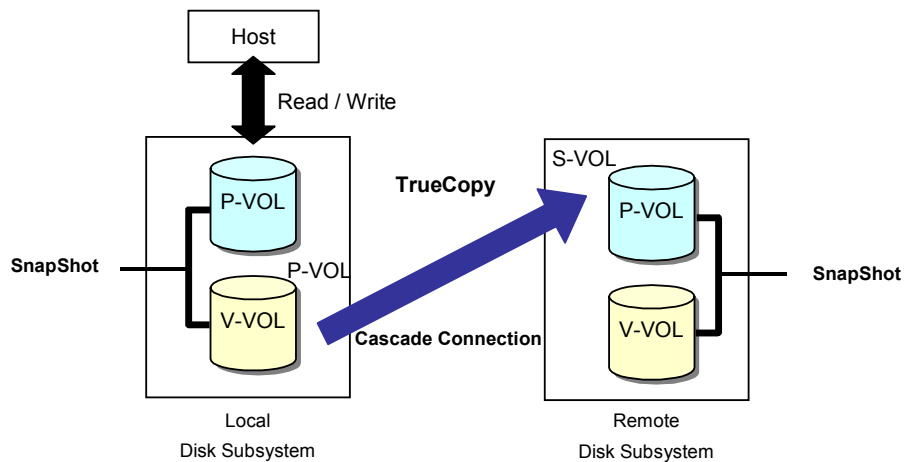


Figure 5.29 Cascade Connection of TrueCopy with Snapshot

5.10.2 Cascading with a Snapshot P-VOL

When both TrueCopy and Snapshot pairs are placed in PAIR status, host performance on the local side is lowered. Therefore, set a PSUS status for TrueCopy and Snapshot pairs when host I/Os are frequently instructed.

Figure 5.30 shows an example of a cascade structure for Snapshot P-VOL.

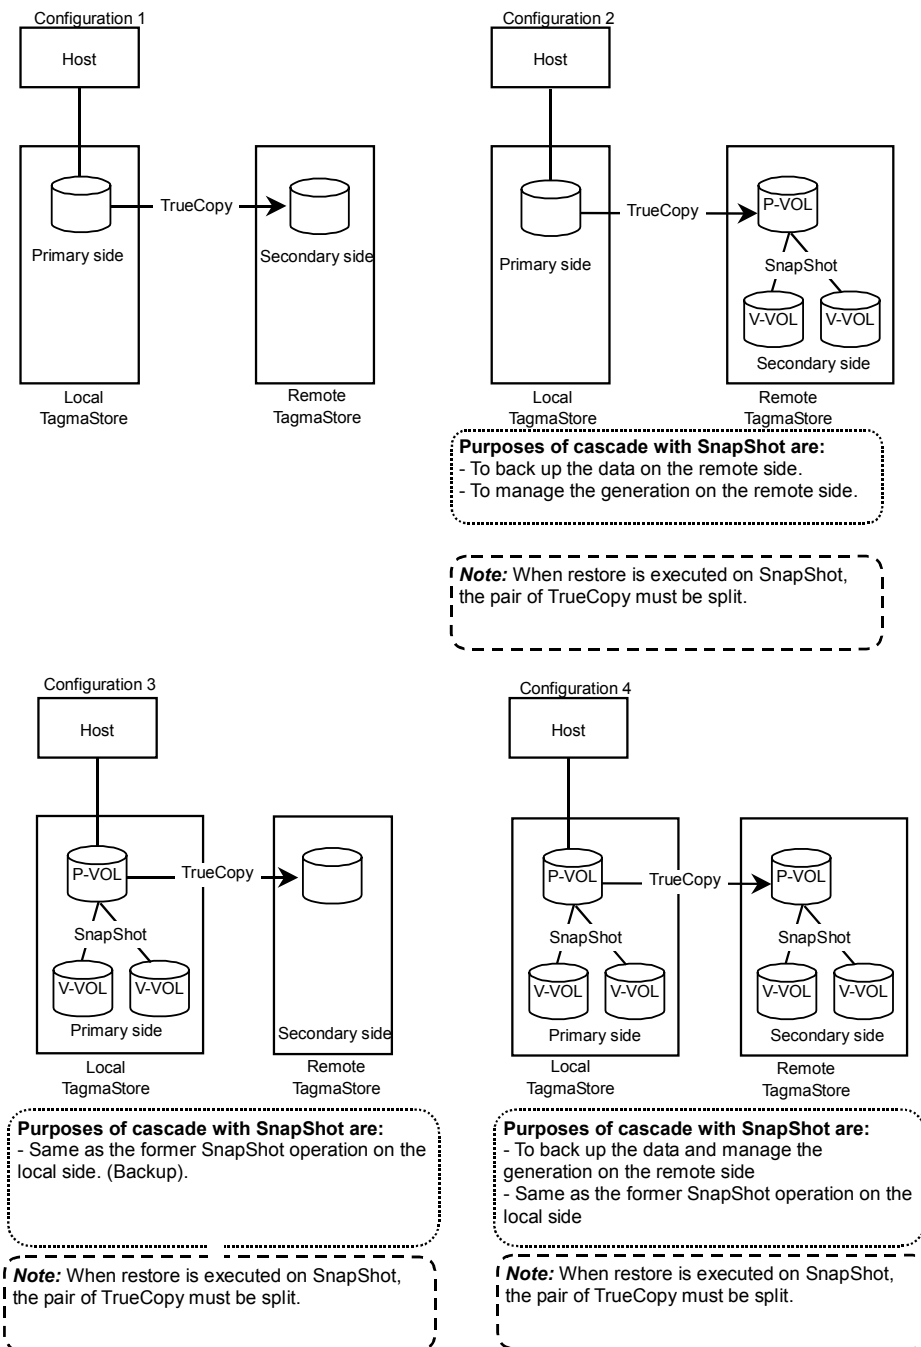


Figure 5.30 Cascade Structure: Snapshot P-VOL

5.10.2.1 Cascading with a V-VOL

A cascade of a Snapshot V-VOL is used when synchronously performing a backup on the remote side. Though the cascade of Snapshot can decrease an S-VOL (V-VOL) capacity differently from a cascade of ShadowImage, the performance on the local side (a P-VOL of Snapshot) is affected by the backup. Snapshot can make two or more V-VOLs for a P-VOL, however, a single V-VOL only can be cascaded with TrueCopy.

Note: A TrueCopy pair must have PSUS status when resynchronizing (giving the Snapshot instruction) a volume of Snapshot on the local side.

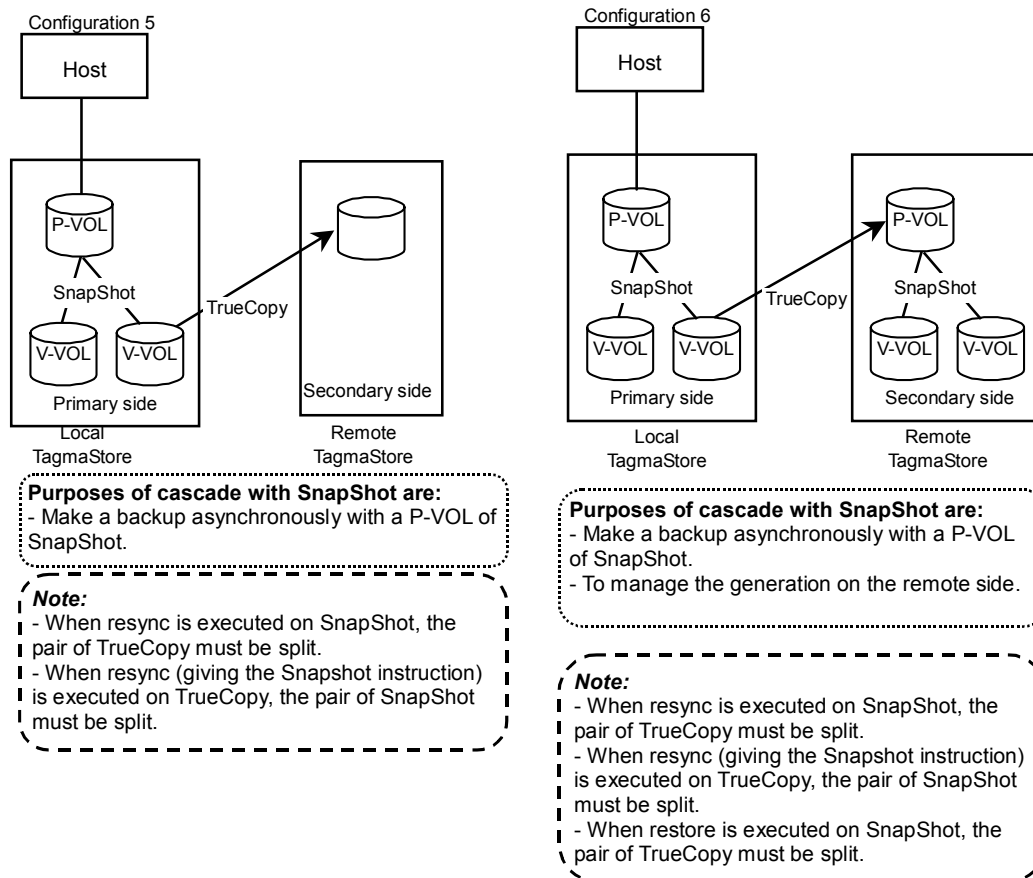


Figure 5.31 Cascade Structure: Snapshot V-VOL

5.11 Cascade Restrictions

The following cascade restrictions are discussed in this section:

- Cascade Restrictions with Snapshot
- Cascade Restrictions on TrueCopy with Snapshot
- Cascade Restrictions with P-VOL
- Cascade Restrictions with V-VOL
- Cascade Restrictions with POOL of Snapshot
- Cascade Restrictions with ShadowImage

5.11.1 Cascade Restrictions with Snapshot

A cascade connection with Snapshot is not available.

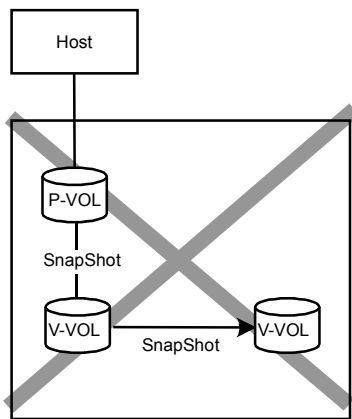


Figure 5.32 Restrictions in Cascade Connection-1

5.11.2 Cascade Restrictions on TrueCopy with Snapshot

The following is an example of a configuration in which restrictions are placed on cascading TrueCopy with Snapshot.

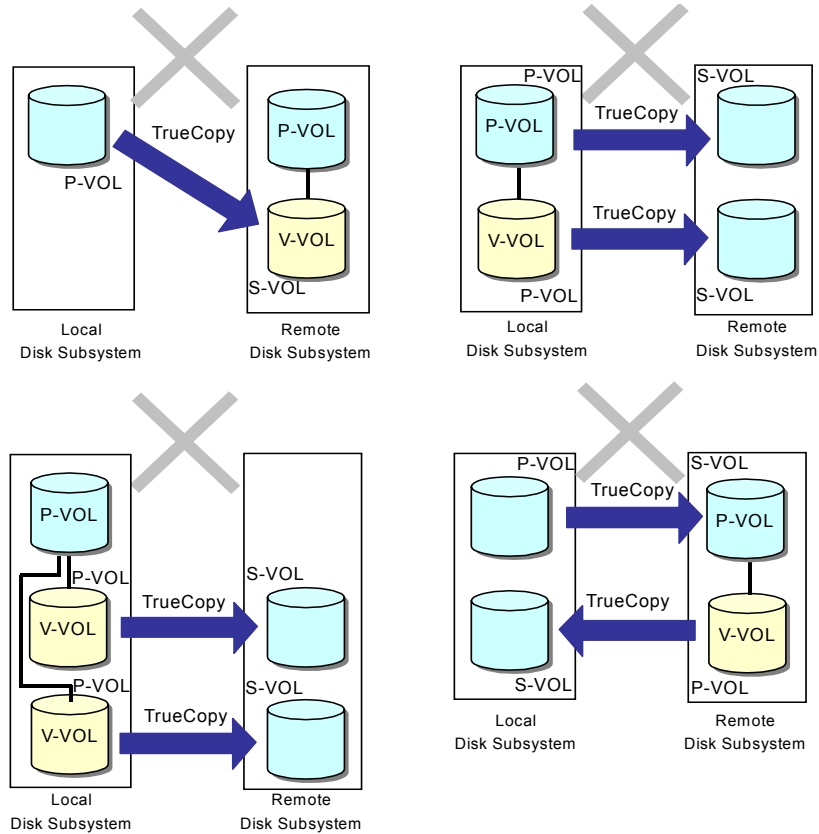


Figure 5.33 Configuration Restrictions on Cascade of TrueCopy with Snapshot

When a P-VOL of Snapshot connects with an LU of TrueCopy and the restore operation is executed on Snapshot, a TrueCopy pair must be split (PSUS).

5.11.3 P-VOL cascade restrictions

The TrueCopy status must be PSUS when `-restore` is executed. If `-restore` is executed and the status of TrueCopy is COPY or PAIR, 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 a P-VOL on Snapshot and a P-VOL on TrueCopy

Table 5.3 displays 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 TrueCopy are the same LU.

Table 5.3 Read/Write Instruction to a P-VOL on the Local Side

		Snapshot P-VOL				
		PAIR	RCPY	PSUS	PSUE	PSUE (Restore)
TrueCopy P-VOL	PAIR	○ R/W	x	○ R/W	○ R/W	x
	COPY	○ R/W	x	○ R/W	○ R/W	x
	PSUS R/W	○ R/W	○ R/W	○ R/W	○ R/W	△ R/W
	PSUS R	○ R	x	○ R	○ R	x
	PSUE R/W	○ R/W	△ R/W	○ R/W	△ R/W	△ R/W
	PSUE R	○ R	x	○ R	△ R	x
	PSUE R/W	○ R/W	x	○ R/W	△ R/W	x

○: Possible, ×: Not Possible

△: When a pair operation causes an error (can occur as a result of a change of the pair status to PSUE)

~~R/W~~: Read/Write is unavailable

Note: PSUE in this table excludes a condition in which LU access is not possible (for example, LU blockage).

- One LU is used for a P-VOL on Snapshot and S-VOL on TrueCopy.

Table 5.4 displays whether 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 TrueCopy share the same LU.

Table 5.4 Read/Write Instruction to P-VOL on the Remote Side

		Snapshot P-VOL				
		PAIR	RCPY	PSUS	PSUE	PSUE (Restore)
TrueCopy P-VOL	PAIR	○ R	x	○ R	○ R	x
	COPY	○ R	x	○ R	○ R	x
	PSUS R/W	○ R/W	○ R/W	○ R/W	○ R/W	△ R/W
	PSUS R	○ R	x	○ R	○ R	x
	PSUE R	○ R	x	○ R	△ R	x

○: Possible, ×: Not Possible

△: When a pair operation causes an error (can occur as a result of a change of the pair status to PSUE)

-R/W: Read/Write is unavailable

Note: PSUE in this table excludes a condition in which LU access is not possible (for example, LU blockage).

- Number of Snapshot V-VOLs

V-VOLs of up to 15(14: the micro program earlier than version 0750/A) generations can be made even when the P-VOL of Snapshot is cascaded with the P-VOL and S-VOL of TrueCopy in the same way as when no cascade connection is made.

5.11.4 V-VOL Cascade Restrictions

- Transitioning the Status of TrueCopy and Snapshot Pairs

Cascading a TrueCopy LU with a Snapshot V-VOL is supported only when the Snapshot V-VOL and a TrueCopy P-VOL are the same LU. Operations including Snapshot and TrueCopy pairs are restricted, depending on the pair status.

When cascading TrueCopy volumes with a Snapshot V-VOL, create a Snapshot pair first. If a TrueCopy pair was created earlier, split the TrueCopy pair once and create a pair using Snapshot.

When changing the status of a Snapshot pair, the status of a TrueCopy pair must be PSUS or PSUE. When changing the status of a TrueCopy pair, the status of a Snapshot pair must be PSUS.

Table 5.5 displays whether a read/write from/to a V-VOL of Snapshot on the local side is possible when a V-VOL of Snapshot and a V-VOL of TrueCopy are the same LU.

Table 5.5 V-VOL Read/Write Instruction on the Local Side (TrueCopy)

		Snapshot P-VOL				
		PAIR	RCPY	PSUS	PSUE	PSUE (Restore)
TrueCopy P-VOL	PAIR	x	x	○ R/W	x	x
	COPY	x	x	○ R/W	x	x
	PSUS R/W	○ R/W	○ R/W	○ R/W	△ R/W	△ R/W
	PSUS R	○ R/W	○ R/W	○ R	△ R/W	△ R/W
	PSUE R/W	○ R/W	○ R/W	○ R/W	△ R/W	△ R/W
	PSUE R	○ R/W	○ R/W	○ R	△ R/W	△ R/W
	PSUE R/W	○ R/W	○ R/W	○ R/W	△ R/W	△ R/W

○: Possible, ×: Not Possible

△: When a pair operation causes an error (can occur as a result of a change of the pair status to PSUE)

~~R/W~~: Read/Write is unavailable

Note: PSUE in this table excludes a condition in which LU access is not possible (for example, LU blockage).

5.11.5 Cascade Restrictions with a Snapshot POOL

Neither TrueCopy nor ShadowImage can be created using a POOL.

5.11.6 Cascade Restrictions with Snapshot and ShadowImage

Snapshot cannot be cascaded with ShadowImage (see Figure 5.34).


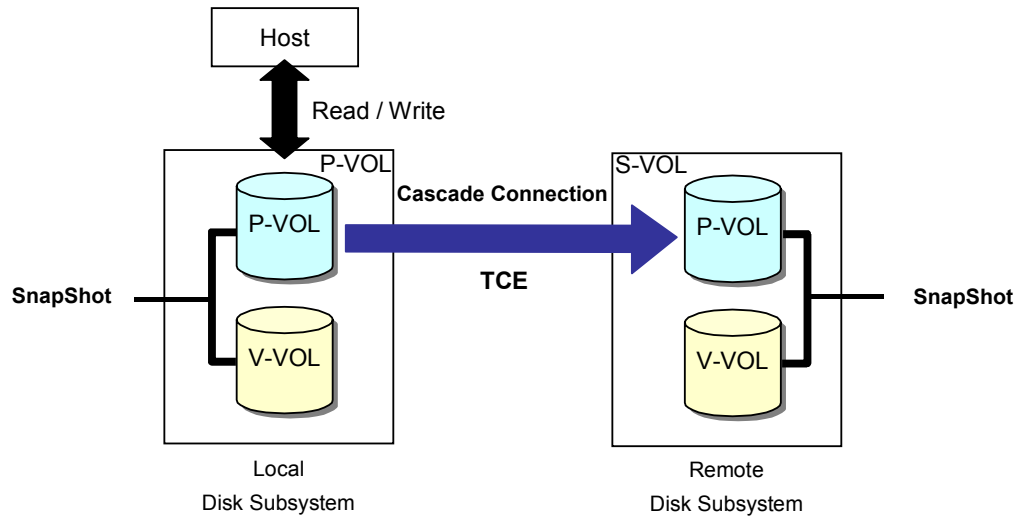


Figure 5.34 Restrictions in Cascade Connection-2

5.12 Cascade Connection of TCE with SnapShot

Volumes of TCE can be cascaded with those of SnapShot P-VOL as shown in Figure 5.35. Cascading of SnapShot with TCE lowers performance, however.

- Cascade with a P-VOL of SnapShot



- Cascade with a V-VOL of SnapShot

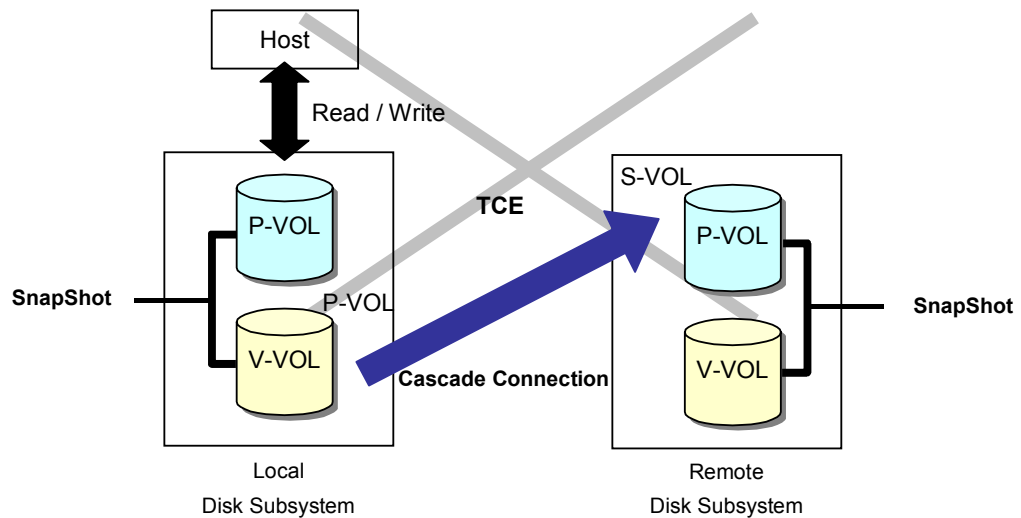


Figure 5.35 Cascade Connection of TCE with SnapShot

The TCE 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 TCE P-VOL can be done only when the status of the TCE pair is PSUS, SMPL, or PFUS.
- When the status of the TCE pair is PAIR, the splitting operation of the SnapShot pair, which is cascaded with the TCE S-VOL and in the PAIR or PSUS status, can be done using the instruction (pairsplit -mscas command) issued to the TCE pair from the host connected to the local side subsystem.
- When restoring the SnapShot pair cascaded with the TCE S-VOL, the status of the TCE pair must be SMPL or PSUS. The restoration can be done in the SSWS status, but it cannot be done when the status is SSWS(R) in which the S-VOL is being restored using the POOL data.
- When the TCE S-VOL is in the SSWS(R) 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 TCE S-VOL.

5.12.1 Cascade Restrictions with P-VOL of TCE

- LU Shared with P-VOL on SnapShot and P-VOL on TCE

Table 5.6 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 TCE are the same LU.

Table 5.6 A Read/Write Instruction to a P-VOL of SnapShot on the Local Side (TCE)

		SnapShot P-VOL				
		PAIR	RCPY	PSUS	PSUE	PSUE (Restore)
TCE P-VOL	PAIR	○ R/W	×	○ R/W	○ R/W	×
	COPY	○ R/W	×	○ R/W	○ R/W	×
	PSUS	○ R/W	○ R/W	○ R/W	○ R/W	Δ R/W
	PFUS	○ R/W	○ R/W	○ R/W	○ R/W	Δ R/W
	PSUE	○ R/W	Δ R/W	○ R/W	Δ R/W	Δ R/W

○: Allowed ×: Not allowed

Δ: A case where a pair operation causes an error (a case that can occur as a result of changing the pair status to PSUE)

R/W: Read/Write by a host is allowed.

R: Read by a host is allowed but write is not allowed.

W: Write by a host is allowed but read is not allowed.

~~R/W~~: Read/Write by a host is not allowed.

Note: PSUE 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 TCE

Table 5.7 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 TCE are the same LU.

Table 5.7 A Read/Write Instruction to a P-VOL of SnapShot on the Remote Side (TCE)

		SnapShot P-VOL				
		PAIR	RCPY	PSUS	PSUE	PSUE (Restore)
TCE S-VOL	PAIR	○ R	×	○ R	○ R	×
	COPY	○ R	×	○ R	○ R	×
	PSUS R/W	○ R/W	○ R/W	○ R/W	○ R/W	△ R/W
	PSUS R	○ R	×	○ R	○ R	×
	PSUS(N) R/W	△ R/W	×	△ R/W	△ R/W	×
	SSWS	○ R/W	○ R/W	○ R/W	○ R/W	△ R/W
	SSWS(R)	△ R/W	×	○ R/W	○ R/W	×
	PFUS	○ R	×	○ R	△ R	×
	PSUE	○ R	×	○ R	△ R	×

○: Allowed ×: Not allowed

△: A case where a pair operation causes an error (a case that can occur as a result of changing the pair status to PSUE)

R/W: Read/Write by a host is allowed.

R: Read by a host is allowed but write is not allowed.

W: Write by a host is allowed but read is not allowed.

~~R/W~~: Read/Write by a host is not allowed.

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

- V-VOLs of SnapShot

V-VOLs of up to 15(14: the micro program earlier than version 0750/A) generations can be made even when the P-VOL of SnapShot is cascaded with the P-VOL and S-VOL of TCE in the same way as in the case where no cascade connection is made.

5.13 Notes and Restrictions

- Restrictions on the SnapShot at the time of the NAS OS boot

When SnapShot is used in the NAS system, do not execute reverse resynchronization during the operations such as installation, boot, shutdown, node stop, and cluster stop of the NAS OS. Conversely, do not perform the operations such as installation, boot, shutdown, node stop, and cluster stop of the NAS OS during the execution of reverse resynchronization of SnapShot.

If the reverse resynchronization of a SnapShot pair and the operations of the NAS OS such as installation, boot, shutdown, node stop, and cluster stop are performed at the same time, a failure may occur because the NAS OS is overloaded.

5.14 Creating Snapshot Pairs

To create Snapshot pairs:

1. If the group name in the configuration definition file is VG01, follow these steps:
Execute `pairdisplay` to verify that the status of the Snapshot volumes is `SMPL` (see section 4.5).
2. Execute `paircreate`. Then, execute `pairevtwait` to verify that the status of each volume is `PSUS`.

```
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.
```

Figure 5.36 Paircreate -split Command Example

3. Execute `pairdisplay` to verify the pair status and the configuration.

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 )75000174 2.P-VOL PSUS,----- ---- -
VG01   oradb1 (R)   (CL1-A , 1, 3-0 )75000174 3.S-VOL SSUS,----- ---- -
VG01   oradb2 (L)   (CL1-A , 1, 2-1 )75000174 2.SMPL ----,----- ---- -
VG01   oradb2 (R)   (CL1-A , 1, 4-0 )75000174 4.SMPL ----,----- ---- -
VG01   oradb3 (L)   (CL1-A , 1, 2-2 )75000174 2.SMPL ----,----- ---- -
VG01   oradb3 (R)   (CL1-A , 1, 5-0 )75000174 5.SMPL ----,----- ---- -
```

To assure that the data of two or more Snapshot images included in a group are of the same time, the CTG is used. The method of creating a pair using the CTG is explained below.

1. If the group name in the configuration definition file is VG01, follow these steps:
Execute `pairdisplay` to verify that the status of the Snapshot volumes is `SMPL` (see section 4.5).
2. Execute `paircreate -m grp`. Then, execute `pairevtwait` to verify that the status of each volume is `PAIR`.

```
C:\HORCM\etc>paircreate -g VG01 -vl -m grp
C:\HORCM\etc>pairevtwait -g VG01 -s pair -t 300 10
pairevtwait : Wait status done.
```

Figure 5.37 Paircreate -m grp Command Example

3. Next, execute `pairsplit`. Then, execute `pairevtwait` to verify that the status of each volume is `PSUS`.

```
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.

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 )75000174 2.P-VOL PSUS,----- ---- -
VG01   oradb1 (R)   (CL1-A , 1, 3-0 )75000174 3.S-VOL SSUS,----- ---- -
VG01   oradb2 (L)   (CL1-A , 1, 4-0 )75000174 4.P-VOL PSUS,----- ---- -
VG01   oradb2 (R)   (CL1-A , 1, 5-0 )75000174 5.S-VOL SSUS,----- ---- -
VG01   oradb3 (L)   (CL1-A , 1, 6-0 )75000174 6.P-VOL PSUS,----- ---- -
VG01   oradb3 (R)   (CL1-A , 1, 7-0 )75000174 7.S-VOL SSUS,----- ---- -
```

Note: When using the CTG, it is required to specify the `-m grp` option. However, the `-split` option and the `-m grp` option cannot be used at the same time. When collecting the Snapshot image using the CTG, split a pair after changing a pair status to RAIR with the `paircreate` command.

The Snapshot pair is created.

5.15 Updating the V-VOL

To update the V-VOL:

Note: Differential data stored in the POOL is deleted when the pair status is changed to PSUS using `pairsplit`. It is not deleted when the pair status is changed to PSUS by `pairresync`. The deletion of differential data is not completed immediately after the pair status is changed to PSUS; it is completed after a few moments. The time required for the deletion process is proportional to the P-VOL capacity, it takes about five minutes to delete a 100 GB P-VOL.

1. For example, if the group name in the configuration definition file is VG01:

Change the PSUS status of the Snapshot pair to PAIR status using `pairresync`. Then, change the status to PSUS using `pairsplit`.

```
C:\HORCM\etc>pairresync -g VG01 -d oradb1
C:\HORCM\etc>pairwait -g VG01 -s pair -t 300 10
pairwait : Wait status done.
C:\HORCM\etc>pairsplit -g VG01 -d oradb1
```

Figure 5.38 Pairresync Command Example

2. Execute `pairdisplay` to update the pair status and the configuration.

```
C:\HORCM\etc>pairdisplay -g VG01
Group  PairVol (L/R) (Port#,TID, LU-M) ,Seq#,LDEV#.P/S,Status, Seq#,P-LDEV# M
VG01   oradb1 (L)   (CL1-A , 1, 2-0 )75000174 2.P-VOL PSUS,----- ----
VG01   oradb1 (R)   (CL1-A , 1, 3-0 )75000174 3.S-VOL SSUS,----- ---- -
VG01   oradb2 (L)   (CL1-A , 1, 2-1 )75000174 2.SMPL ----,----- ---- -
VG01   oradb2 (R)   (CL1-A , 1, 4-0 )75000174 4.SMPL ----,----- ---- -
VG01   oradb3 (L)   (CL1-A , 1, 2-2 )75000174 2.SMPL ----,----- ---- -
VG01   oradb3 (R)   (CL1-A , 1, 5-0 )75000174 5.SMPL ----,----- ---- -
```

The V-VOL was updated.


5.16 Restoring a V-VOL to the P-VOL

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

Note: When no differential exists between a P-VOL and V-VOL to be restored, restoration is not completed immediately. It takes time to examine the differential between the P-VOL and V-VOL.

The method of search for the differentials is “Search All”. Because the search is terminated when all the differentials are found, the time required for completion of the search varies depending on the positions of the differentials.

A pattern that requires a short search time A pattern that requires a long search time



The differential management area on the cache memory

→ : Search operation

■: Bit that indicates differential data

The standard time required for the examination of the differential is shown below. When no I/O is issued by a host, capacity per LU is 100 GB, and the whole differential management area is searched (when the rate of coincidence is approximately 100%).

Table 5.8 Time Required to Examine Differential

Equipment Type	Copy Speed	Searched Time (minutes)
AMS500	11 to 15 (fast copy)	Approximately 4
	6 to 10 (normal copy)	Approximately 4
	1 to 5 (slow copy)	Approximately 25
AMS200	11 to 15 (fast copy)	Approximately 6
	6 to 10 (normal copy)	Approximately 6
	1 to 5 (slow copy)	Approximately 35
WMS 100	11 to 15 (fast copy)	Approximately 6
	6 to 10 (normal copy)	Approximately 6
	1 to 5 (slow copy)	Approximately 35

Notes:

- When executing `pairresync -restore`, assign the “-c value”; the mode is set as the default value (slow copy) if it is not assigned. Although the copy speed is set as a specified value such as `-c 15` when `paircreate` is executed, the mode is also set as “slow copy” if the “-c value” is not given when `pairresync -restore` is executed. It is recommended that you select “fast” (11 to 15) for the `-c` option. It is recommended to select “fast” (11 to 15) when copying pace is preceded. When “fast” (11 to 15) is selected, the host I/O performance may be lowered. To prevent lowering performance for the host I/O, select “normal” or “slow” (1 to 10).
- The restoration command can be issued up to 128 P-VOLs at the same time. However, the number of P-VOLs to which the physical copying (background copying) from a V-VOL can be done at the same time is up to four per controller. The first to 4th background copies are completed in the order the command was issued. The 5th and the following background copies are completed in the ascending order of LU numbers, after the preceding restoration is completed.

1. For example, if the group name in the configuration definition file is VG01:

Execute `pairresync` to restore the V-VOL to the P-VOL.

```
C:\HORCM\etc>pairresync -restore -g VG01 -d oradb1 -c 15
```

Figure 5.39 Pairresync -restore command example

2. Execute `pairdisplay` to restore the pair status and the configuration.

```
C:\HORCM\etc>pairdisplay -g VG01
Group  PairVol (L/R) (Port#,TID, LU-M) ,Seq#,LDEV#.P/S,Status, Seq#,P-LDEV# M
VG01   oradb1 (L)   (CL1-A , 1, 2-0 )75000174 2.P-VOL COPY,----- ---- -
VG01   oradb1 (R)   (CL1-A , 1, 3-0 )75000174 3.S-VOL COPY,----- ---- -
VG01   oradb2 (L)   (CL1-A , 1, 2-1 )75000174 2.SMPL -----,----- ---- -
VG01   oradb2 (R)   (CL1-A , 1, 4-0 )75000174 4.SMPL -----,----- ---- -
VG01   oradb3 (L)   (CL1-A , 1, 2-2 )75000174 2.SMPL -----,----- ---- -
VG01   oradb3 (R)   (CL1-A , 1, 5-0 )75000174 5.SMPL -----,----- ---- -
```

3. Execute the `pairsplit` command. Replace pair status PAIR with PSUS.

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

Figure 5.40 Pairsplit Command Example

V-VOL to P-VOL is restored.

5.17 Releasing Snapshot Pairs

To release the Snapshot pair and change the status to SMPL:

1. For example, if the group name in the configuration definition file is VG01:

Execute the `pairdisplay` command to verify that the status of the Snapshot pair is PSUS or PSUE.

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 )75000174 2.P-VOL PSUS,----- ---- -
VG01   oradb1 (R)   (CL1-A , 1, 3-0 )75000174 3.S-VOL SSUS,----- ---- -
VG01   oradb2 (L)   (CL1-A , 1, 2-1 )75000174 2.SMPL ----,----- ---- -
VG01   oradb2 (R)   (CL1-A , 1, 4-0 )75000174 4.SMPL ----,----- ---- -
VG01   oradb3 (L)   (CL1-A , 1, 2-2 )75000174 2.SMPL ----,----- ---- -
VG01   oradb3 (R)   (CL1-A , 1, 5-0 )75000174 5.SMPL ----,----- ---- -
```

2. Execute the `pairsplit -S` command to release the Snapshot pair.

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

Figure 5.41 Pairsplit -S Command Example

3. Execute the `pairdisplay` command to verify that the pair status changed to SMPL.

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 )75000174 2.SMPL ----,----- ---- -
VG01   oradb1 (R)   (CL1-A , 1, 3-0 )75000174 3.SMPL ----,----- ---- -
VG01   oradb2 (L)   (CL1-A , 1, 2-1 )75000174 2.SMPL ----,----- ---- -
VG01   oradb2 (R)   (CL1-A , 1, 4-0 )75000174 4.SMPL ----,----- ---- -
VG01   oradb3 (L)   (CL1-A , 1, 2-2 )75000174 2.SMPL ----,----- ---- -
VG01   oradb3 (R)   (CL1-A , 1, 5-0 )75000174 5.SMPL ----,----- ---- -
```

The Snapshot pair is released.

5.18 Confirming Pairs Status

The status of the logical units is displayed in icon format in the **Logical Status** pane of the **Array System Viewer** window. P designates a primary volume and V designates a virtual volume.

Table 5.9 Icons for Pair Status

Volume Status				Paired Status	Meaning
Formatted Yellow	Unformatted Gray	Regressed Pink	Blockaded Red		
				PAIR	Indicates paired volume.
				COPY	Indicates a paired status, but resync operation is not complete.
				PSUS	In paired status, but updates to S-VOL data are suspended due to user-requested pairsplit.
				PSUE	In paired status, but updates to S-VOL data are suspended forcibly or due to an error condition.

Chapter 6 Troubleshooting

Two types of pair failures occur when Snapshot is used. Type one is caused by a hardware failure and type two is caused by a data pool capacity shortage. Recovery is also required when the accessed data pool capacity exceeds its threshold value (a transition of pair status to PFUS). The procedures for recovery from these failures differ; they are described in this chapter.

- Hardware Failure (section 6.1)
- Data Pool Capacity Shortage (section 6.2)
- Backup Operation for Quick Recovery (section 6.3)
- Online Backup Operation Using an Inexpensive Configuration (section 6.4)
- Restoring Backup Data (section 6.5)
- Previously Designed Procedure for Recovering Data during Malfunction (section 6.6)

6.1 Hardware Failure

When a hardware error occurs while operating Snapshot, CCI user operation is necessary in addition to system maintenance performed by Hitachi Data Systems personnel.

For example, when formatting is necessary because of an LU error and the LU is used for Snapshot, the pair must be released by the user before formatting the LU. Contact Hitachi Data Systems personnel because maintenance requires that the user issue CCI commands. Note that only Hitachi Data Systems personnel can remove errors caused by hardware. Operations such as recovering a Snapshot pair status (e.g. resynchronizing) must be done by the user.

If you have a problem with operating Snapshot, refer to the troubleshooting section in *Hitachi TagmaStore Adaptable Modular Storage and Workgroup Modular Storage Command Control Interface (CCI) User and Reference Guide (MK-95DF701)*. Snapshot error messages and procedures for recovering PSUE errors are described in detail.

6.2 Data Pool Capacity Shortage

There are two situations in which data pool capacity shortages occur:

- When the threshold value of the usable data pool capacity is exceeded:

When the accessed data pool capacity exceeds its threshold value, pair status is changed to PFUS. The subsystem operates the same with the pair in PFUS status as with the pair in PSUS status, however, it is necessary to ensure that a sufficient data pool capacity is available early because data pool capacity is very likely to run short. The operation to ensure a sufficient data pool capacity must be done by a user. The method required for this process involves splitting the V-VOL from the pair or increasing the data pool capacity.

To split the pair, refer to section 5.17. Split the pair after saving the V-VOL data to a tape drive; the data of the V-VOL to be split from the pair will be invalidated.

To increase the data pool capacity, add a logical unit to the data pool; refer to section 4.3. Note that an addition cannot be made when 64 LUs have already been assigned to the data pool.

- When data pool capacity is insufficient:

When differential data cannot be saved to the data pool caused by a data pool capacity shortage, all the V-VOLs, which use the data pool and whose statuses are other than SMPL, are placed in the PSUE status. This occurs because the V-VOL data cannot be retained. If this occurs, split all pairs whose statuses have been changed to PSUE. Then check the propriety of the system configuration (data pool capacity, number of V-VOLs, etc.) because the system may be inappropriately configured. When this is completed, perform the Snapshot operation for correcting the status. A series of operations regarding a shortage of data pool capacity must be performed by a user.

If you have a problem with Snapshot operation, refer to the troubleshooting section in the *Hitachi TagmaStore® Adaptable Modular Storage and Workgroup Modular Storage Command Control Interface (CCI) User and Reference Guide (MK-95DF701)*. Snapshot error messages and procedures for recovering the PSUE errors are described in detail.

6.3 Backup Operation for Quick Recovery

Note the following regarding backup operations for quick recovery:

- Because the V-VOL that can be used to restore can be retained 24/365, recovery can occur without doing a restore from a tape device. (However, because of the possibility of becoming PSUE by overflow of POOL capacity or hardware failure, tape backup is mandatory.)
- Make two V-VOLs per P-VOL and issue Snapshot instruction every 0:00 and 12:00.
- Backup onto tape at night (when few I/O instructions are issued by a host). (As a general guide, a time of the day when host I/O is below 100 IOPS.)
- Backup onto tapes via two ports simultaneously.
- When backing up onto tapes, the total capacity of the V-VOL concerned must be 1.5 TB or smaller.

Example: When total V-VOL capacity is 1 TB, time required for backing up (at a speed of 100 MB/sec) is 3 hours minutes.

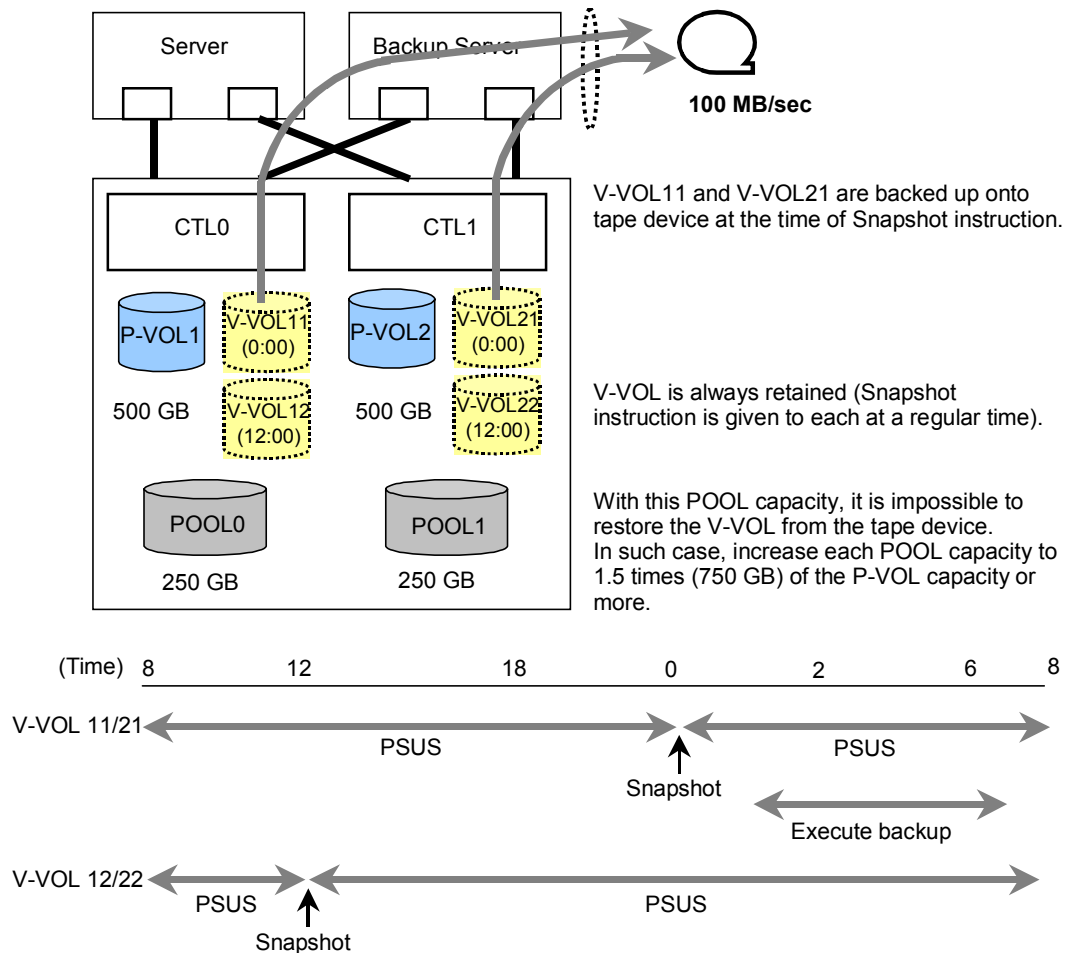


Figure 6.1 Ordinary Quick Recovery Operation

6.4 Online Backup Operation Using an Inexpensive Configuration

Note the following while doing an online backup operation:

- A period during which the V-VOL must be retained is limited to the period required for backup to a tape device.
- The configuration cost is minimal because the data pool capacity can be managed by storing only data updated during a backup to a tape device.
- Back up onto a tape device at night (when few I/O instructions are issued by a host). (In general, a time of the day when the host I/O is below 100 IOPS.)
- Backup onto tapes via two ports simultaneously.
- When backing up onto tapes, the total capacity of the V-VOL concerned must be 1.5 TB or smaller.

Example: When the total V-VOL capacity is 1 TB, time required for back up (at a speed of 100 MB/sec) is 3 hours.

To restore from a tape device, only “Direct restoration of the P-VOL” is possible.

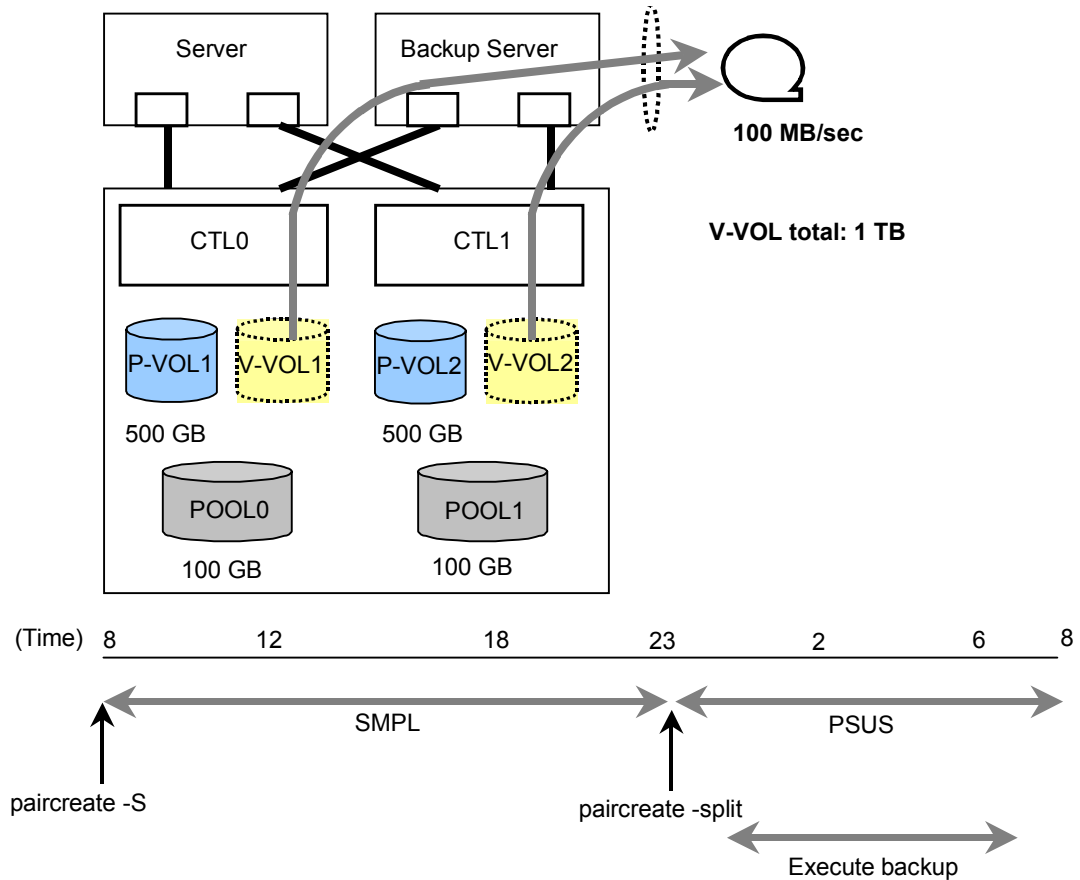


Figure 6.2 Ordinary Operation

6.5 Restoring Backup Data

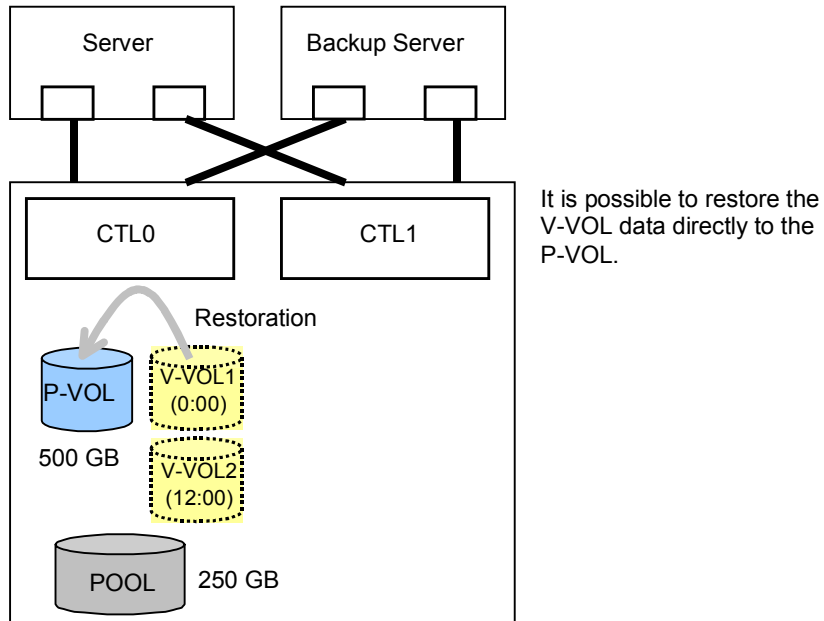
There are two ways to restore backup data. One is restoration using backup data within the same disk array and the other is restoration using backup data stored in a tape device.

- Restoration using backup data within the same disk array
 - This is the restoration method used for backup operation for quick recovery.
- Restoration using backup data stored in a tape device (via V-VOL)
 - This is a restoration method used when the V-VOL has a PSUE status in backup operation for quick recovery.
 - Free POOL capacity should be larger than the P-VOL capacity (recommendation is 1.5 times of the P-VOL capacity or more).
- Restoration using backup data stored in a tape device (direct P-VOL)
 - This is a restoration method used for the operation when the V-VOL is a PSUE status in backup operation for a quick recovery.
 - This is a restoration method used for an operation whose free POOL capacity is smaller than the P-VOL capacity, like online backup operation using an inexpensive configuration.

6.5.1 Backup Operation for Quick Recovery

- This restoration method uses backup data within the same disk array
- When a software failure (caused by a wrong operation by a user or an application program bug) occurs, perform restoration, selecting backup data you want to return from the V-VOL being retained.

Note: It is necessary to unmount the P-VOL once before restoring the P-VOL from the V-VOL.



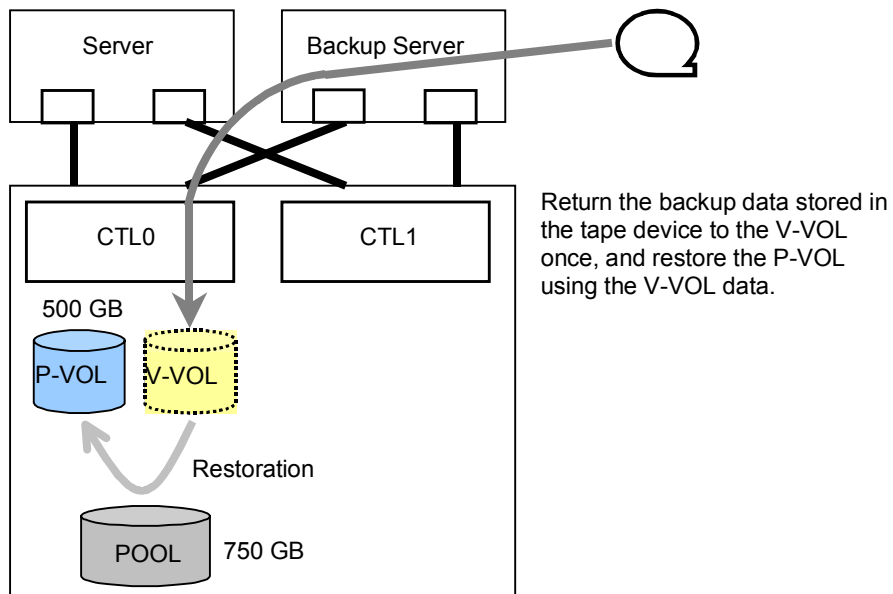
6.5.2 Restoring Backup Data from a Tape Device

Return backup data stored in the tape device to the V-VOL from which the backup data was executed and restore P-VOL using V-VOL data.

- Restoration via V-VOL
 - Return backup data stored in tape device to V-VOL once and restore P-VOL using V-VOL data.

Notes:


- When returning the backup data to the V-VOL, free POOL capacity larger than the P-VOL capacity is required. More than 1.5 times the P-VOL capacity is recommended.
- It is necessary to unmount the P-VOL once, before restoring the P-VOL from the V-VOL.



- Direct restoration of P-VOL
 - Use it when free POOL capacity is insufficient or V-VOL has entered PSUE status.
 - Restore data from a backup server to the P-VOL via a LAN or directly from tape device to the P-VOL.

Notes:

- When copying the backup data from tape device onto P-VOL, lift all pairs on that P-VOL (= SMPL). If backup data on tape device is restored to paired P-VOL (= PSUS or PAIR), then data copy will be initiated from P-VOL to POOL in order to retain data on V-VOL. This causes a drop in restore performance and open capacity of POOL has to be larger than P-VOL capacity.
- It is necessary to suspend host access during restoration of P-VOL.



6.6 Previously Designed Procedure for Recovering Data during Malfunction

For Snapshot, V-VOL data cannot be retained (PSUE status) because of a hardware failure or a data pool capacity shortage. Even if the user has a hardware maintenance contract (including free warranty period) with Hitachi, recovering Snapshot statuses (such as regeneration of the Snapshot pair) is beyond maintenance personnel's task. Maintenance personnel's tasks include removing the malfunction cause and recovering the hardware only. The procedure for recovering Snapshot statuses must be specified by the user during system design.

Chapter 7 Specifications and Function Comparison between Snapshot and ShadowImage

7.1 Differences between Snapshot and ShadowImage

7.1.1 External Specifications

This section describes basic specifications for Snapshot and ShadowImage.

Table 7.1 External Specifications

No.	Item	Snapshot	ShadowImage
1	Applicable model	AMS1000/AMS500/AMS200/WMS100 (For dual configuration only.)	AMS500/AMS200/WMS100 (For dual configuration only.)
2	Cache memory	WMS100: 1 GB/CTL AMS200: 1, 2 GB/CTL AMS500: 2, 3, 4 GB/CTL is supported (see Note 1) AMS1000: 2, 4, 6, 8 GB/CTL is supported (see Note 1)	WMS100: 512 MB, 1 GB/CTL AMS200: 1, 2 GB/CTL AMS500: 1, 2, 3, 4 GB/CTL is supported AMS1000: 2, 4, 6, 8 GB/CTL is supported
3	Unit of pair management	Logical unit	
4	Number of pairs	AMS200/WMS100: 254 (maximum) AMS500: 1,022 (maximum) AMS1000: 2,046 (maximum)	AMS200/WMS100: 255 (maximum) AMS500: 1,023 (maximum) AMS1000: 2,047 (maximum)
5	Number of command devices	Up to two per disk subsystem can be set (see Note 2).	
6	Pair structure (number of S-VOLs per P-VOL)	1:15(14: the micro program earlier than version 0750/A)	1:3(1: the micro program earlier than version 0750/A)
7	P-VOL and POOL (S-VOL) arrangement	Within the same CTL (see Note 3)	Within the same CTL
8	RAID level	RAID 1 (1D+1D) RAID 5 (2D+1P to 15D+1P) RAID 6 (2D+2P to 28D+2P) RAID 1+0 (2D+2D to 8D+8D)	RAID 0 (2D to 16D) RAID 1 (1D+1D) RAID 5 (2D+1P to 15D+1P) RAID 6 (2D+2P to 28D+2P) RAID 1+0 (2D+2D to 8D+8D)
9	RAID configuration for P-VOL and POOL (S-VOL)	P-VOL and S-VOL can be paired between different RAIDs. Note that <i>N</i> for ND of P-VOL and POOL can be different. (See Note 4)	P-VOL and S-VOL can be paired between different RAIDs. Note that <i>N</i> for ND of P-VOL and S-VOL can be different.
10	Consistency Group (CTG) number	Max 128/subsystem AMS200/WMS100: 254 (maximum) AMS500: 1,022 (maximum) AMS1000: 2,046 (maximum)	Max 128/subsystem AMS200/WMS100: 255 (maximum) AMS500: 1,023 (maximum) AMS1000: 2,047 (maximum)

No.	Item	Snapshot	ShadowImage
11	Mixing Snapshot and non-Snapshot	Available (mixing within CTL and between CTLs) (See Note 5)	Available (mixing within CTL and between CTLs)
12	Types of drive for a P-VOL/POOL	Both LUs with FC drives and SATA drives can be assigned to any P-VOL and a POOL. However, an LU with FC drives assigned to a P-VOL and a POOL is recommended. When creating a pair with the LUs configured in the SATA drive, the use conditions of the SATA drive may be different. (See Table 7.2)	
13	Maximum supported capacity value of P-VOL and S-VOL (TB)	Maximum Snapshot single capacity $-(2 \times \text{Total S-VOL capacity of ShadowImage} \div 17) - (\text{Total P-VOL and S-VOL capacity of TrueCopy} \div 17) - (\text{Total P-VOL and S-VOL capacity of TCE} \times 3)$ (See Note 6) Those with restriction (See Note 7)	Maximum supported capacity value of S-VOL (Maximum ShadowImage single capacity $-(\text{Total P-VOL and S-VOL capacity of TrueCopy}) - 17 \times (\text{Total P-VOL capacity of Snapshot}) - (\text{Total P-VOL and S-VOL capacity of TCE} \times 51)) \div 2$ (See Note 6)
14	Host interface	Fibre or iSCSI	Fibre, iSCSI, or NAS
15	Management size of difference data	64 KB	1 MB
16	Number of POOLs that can be set	One per controller. (POOL number "0" is for Controller 0 and POOL number "1" is for Controller 1.)	—
17	LU number for POOL that can be set	Up to 64 per disk subsystem can be set.	—
18	Expansion of POOL capacity	Possible (The capacity is expanded through an addition of LU(s) to the POOL. The extension of a POOL can be made while a pair is formed.) However, an LU with an FC drive and an LU with a SATA drive cannot coexist in a POOL.	—
19	Reduction of POOL capacity	Possible only when all the V-VOLs have been deleted.	—
20	RAID configuration for the same POOL	POOL can be paired between different RAIDs. Note that <i>N</i> for ND of POOL can be different (see Note 4).	—
21	Deleting during Coupling (RAID group, P-VOL, V-VOL) Formatting during coupling	Not available	
22	Deleting RAID group/P-VOL, of the LU for the POOL Formatting of the LU for the POOL	Not available.	
23	Potential effect caused by installation of optional feature function	Reboot is required.	Reboot is not required.
24	Potential effect caused by a detachment of one of the controllers	Detaching one of the controllers has no effect on the V-VOL data.	Detaching one of the controllers has no effect on the S-VOL data.

No.	Item	Snapshot	ShadowImage
25	Potential effect caused by a P-VOL failure	V-VOL data exists in the P-VOL; therefore, if the P-VOL fails, the V-VOL also fails.	The entire P-VOL is copied to the S-VOL; therefore, a P-VOL failure has no effect.
26	Concurrent use of LUN Expansion	Available.	Available.
27	Concurrent use of Data Retention Utility	When the S-VOL Disable is set for an LU, a pair formation using the LU as an S-VOL (POOL) is suppressed. A setting of the S-VOL Disable of a volume that has already become an S-VOL (V-VOL or POOL) is not suppressed only when the pair status is PSUS. When the S-VOL Disable is set for a P-VOL, restoration of SnapShot, reverse resynchronization of ShadowImage is suppressed.	
28	Concurrent use of Cache Residency Manager	The LU specified for Cache Residency (LU cache residence) cannot be set to P-VOL, V-VOL, or POOL.	The LU specified for Cache Residency (LU cache residence) cannot be set to S-VOL or P-VOL.
29	Concurrent use of SNMP Agent	Available.	
30	Concurrent use of LUN Manager	Available.	
31	Concurrent use of Password Protection	Available.	
32	Concurrent use of TrueCopy	TrueCopy can be cascaded with Snapshot. For details, see section 5.10.1.	Available.
33	Concurrent use of ShadowImage	Snapshot and ShadowImage can be used together at the same time, but cascade between Snapshot and ShadowImage is not supported. The number of CTG at the time of using SnapShot and ShadowImage together is limited to the maximum of 128 combining that of SnapShot and ShadowImage.	—
34	Concurrent use of Volume Migration	Available. A P-VOL, an S-VOL, and reserved LU of Volume Migration cannot be specified as the Snapshot image (V-VOL) and creating SnapShot.	Available. A P-VOL, an S-VOL, and reserved LU of Volume Migration cannot be specified as a P-VOL or an S-VOL of ShadowImage.
35	Concurrent use of Cache Partition Manager	Available.	
36	License	ShadowImage and SnapShot become usable through entries of their respective key codes.	
37	Location for storing a bit map	1 GB area at the top of the Differential Management LU.	
38	Number of Differential Management LUs	One or two (One of them is for mirroring) The Differential Management LU size must be greater than or equal to 5 GB. We recommend that you set two Differential Management LUs according to following conditions. <ul style="list-style-type: none"> ▪ To be created in different RAID group ▪ To be allocated in different controllers 	
39	Point in time Split	All pairs can be defined as one group.	
40	Events that make a differential to be	Transition of pair status from PSUS to PAIR; releasing a pair.	

No.	Item	Snapshot	ShadowImage
	released		

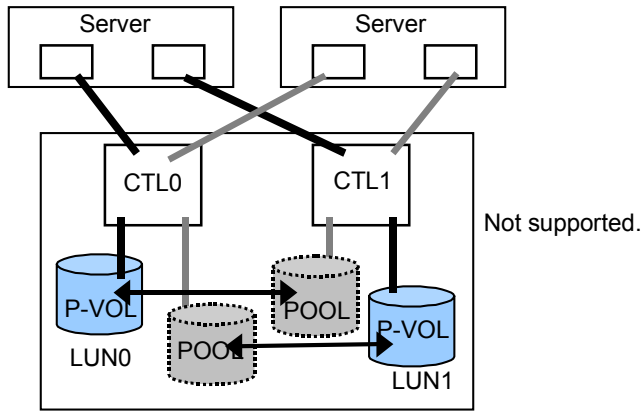
Table 7.2 Characteristics of FC Drive and SATA Drive

Items	SATA Drive		FC Drive
	ATE250R/ATE400R	Others	
Drives presumed operation environment	The Read/Write access takes 330 hours/month or less on the average.	The Read/Write access takes 720 hours/month.	The Read/Write access takes 720 hours/month.

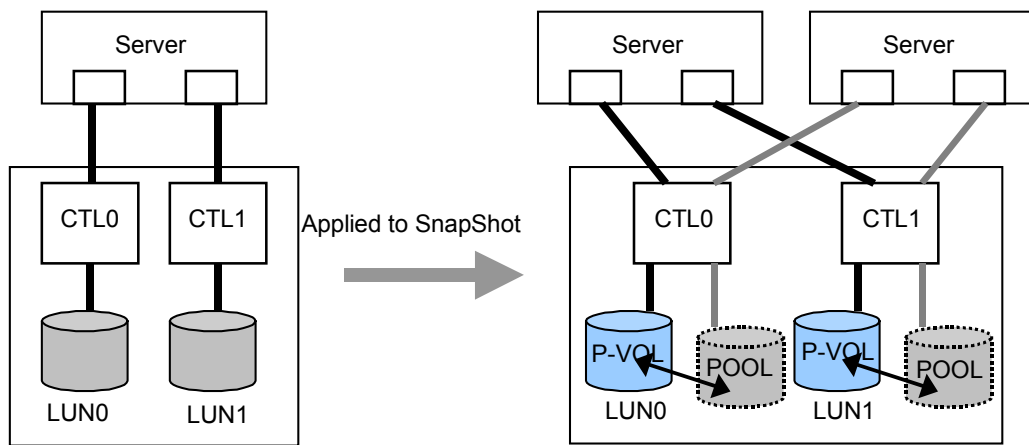
Note 1: Part of cache memory is used to retain Snapshot differential data according to capacity of installed cache memory. Note this when estimating cache.

Note 2: For details on command devices, see the *Hitachi TagmaStore Adaptable Modular Storage and Workgroup Modular Storage Command Control Interface (CCI) User and Reference Guide (MK-95DF701)*. Command devices can be shared between multiple servers. However, this does not mean the data is shared. Command devices must not be used by ordinary applications. Set the command device with caution. For details on setting the command device, see the *Hitachi TagmaStore Adaptable Modular Storage and Workgroup Modular Storage Command Control Interface (CCI) User and Reference Guide (MK-95DF701)*.

Note 3: The configuration shown below (the P-VOL and S-VOL arranged to cross two different CTLs) is not supported.



Concept:



Configuration for ordinary system operation (without SnapShot):
 Usually, LUs are assigned to each dual controller (LU ownership).

SnapShot Configuration:
 The P-VOLs and POOL are defined to each LU within and under the same controller.

Figure 7.1 Concept

Note 4: A RAID level and/or “N” of the term ND (*N* = Number of data drives in RAID) of a P-VOL cannot be identical to those of an LU in a POOL respectively.

An example of a valid combination is a configuration consisting of RAID 5 (7D+1P) for P-VOL and a configuration of RAID 5 (4D+1P) and RAID 1+0 (4D+4D) LUs in a POOL.

Note 5: Mixing LUs (P-VOL, POOL, and V-VOL) of Snapshot and LUs of non-Snapshot are available within the AMS200, AMS500, WMS100, and AMS1000 subsystem. However, note that there may be some effects to the performance.

Note 6: The maximum capacity supported by Snapshot function can be calculated from the following formula. The single maximum capacity of Snapshot is shown in Table 7.3. The capacity of all the P-VOLs must be less than or equal to the maximum supported, and it must meet the requirements shown in Table 7.4 to Table 7.7.

$$\begin{aligned} &\text{Snapshot: Maximum supported capacity value of P-VOL (TB)} \\ &= \text{Maximum Snapshot single capacity} \\ &- (2 \times \text{Total S-VOL capacity of ShadowImage} \div 17) \\ &- (\text{Total P-VOL and S-VOL capacity of TrueCopy} \div 17) \\ &- (\text{Total P-VOL and S-VOL capacity of TCE} \times 3) \end{aligned}$$

Note 7: Because TrueCopy and TCE cannot be used together, the capacity supported by TrueCopy or TCE in the above formula is zero.

Table 7.3 Single Maximum Supported Capacity of SnapShot Function (TB)

Equipment Type	Mounted Memory Capacity	Snapshot
WMS100	512 MB/CTL	—
	1 GB/CTL	11
AMS200	1 GB/CTL	7
	2 GB/CTL	22
AMS500	1 GB/CTL	—
	2 GB/CTL	27
	3 GB/CTL	33
	4 GB/CTL	48
AMS1000	2 GB/CTL	47
	4 GB/CTL	69
	6 GB/CTL	84
	8 GB/CTL	100

Note8: The Snapshot function restricts the capacity of the P-VOL/POOL. The supported maximum capacity value of the P-VOL/POOL varies, depending on the capacity ratio of P-VOL to POOL and cache memory. Because the POOL used by SnapShot is the same as that used by TCE, sufficient POOL capacity must be designed into the system when the SnapShot function is used together with TCE. Table 7.4 to Table 7.7 show the maximum supported capacities of the P-VOL, the POOL for each cache memory capacity, and the formula for calculating the previous capacities. The capacity of all the P-VOLs must be less than or equal to the maximum supported capacity that is required in **Note 6**, and it must meet the requirements shown in Table 7.4 to Table 7.7.

Table 7.4 Formula for Calculating Maximum Supported Capacity Value for P-VOL/POOL (AMS200)

Capacity of Cache Memory Installed	Differential Information used Capacity
1 GB/CTL	Total P-VOL capacity ÷ 4 + Total POOL capacity < 1.0 TB

Table 7.5 Formula for Calculating Maximum Supported Capacity Value for P-VOL/POOL (AMS200)

Capacity of Cache Memory Installed	Differential Information used Capacity
1 GB/CTL	Total P-VOL capacity ÷ 4 + Total POOL capacity < 1.0 TB
2 GB/CTL	Total P-VOL capacity ÷ 4 + Total POOL capacity < 3.7 TB

Table 7.6 Formula for Calculating the Maximum Supported Capacity Value for P-VOL/POOL (AMS500)

Capacity of Cache Memory Installed	Differential Information used Capacity (shared by SnapShot and TCE)
2 GB/CTL	Total P-VOL of SnapShot and P-VOL (S-VOL) of TCE capacity ÷ 4 + Total POOL capacity < 1.8 TB
3 GB/CTL	Total P-VOL of SnapShot and P-VOL (S-VOL) of TCE capacity ÷ 4 + Total POOL capacity < 3.7 TB
4 GB/CTL	Total P-VOL of SnapShot and P-VOL (S-VOL) of TCE capacity ÷ 4 + Total POOL capacity < 7.0 TB

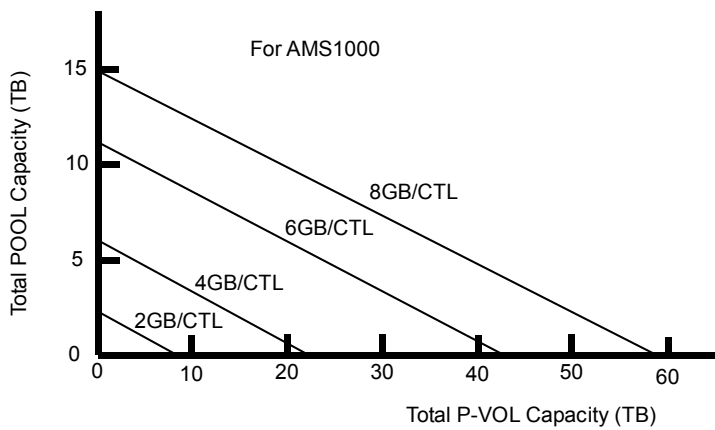
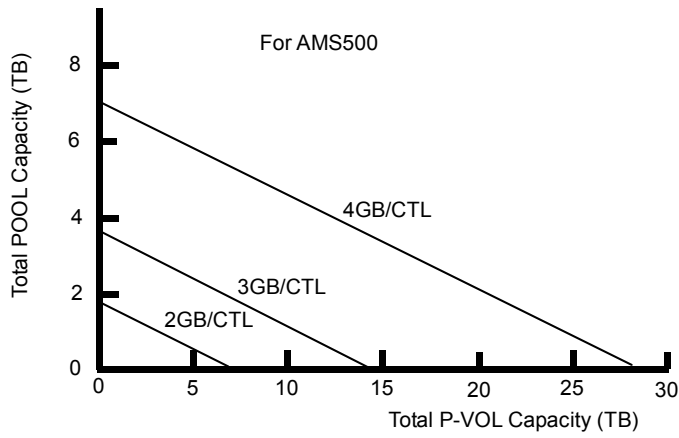

Table 7.7 Formula for Calculating the Maximum Supported Capacity Value for P-VOL/POOL (AMS1000)

Capacity of Cache Memory Installed	Differential Information Used Capacity (Shared by SnapShot and TCE)
2 GB/CTL	Total P-VOL of SnapShot and P-VOL (S-VOL) of TCE capacity ÷ 4 + Total POOL capacity < 1.8 TB
4 GB/CTL	Total P-VOL of SnapShot and P-VOL (S-VOL) of TCE capacity ÷ 4 + Total POOL capacity < 5.5 TB
6 GB/CTL	Total P-VOL of SnapShot and P-VOL (S-VOL) of TCE capacity ÷ 4 + Total POOL capacity < 11.0 TB
8 GB/CTL	Total P-VOL of SnapShot and P-VOL (S-VOL) of TCE capacity ÷ 4 + Total POOL capacity < 15.0 TB

Note 9: The capacity of each P-VOL is managed in units of 15.75 GB. When the P-VOL capacity is 17 GB, the P-VOL is regarded as using a capacity of 31.5 GB. When there are two P-VOLs; each has a capacity of 17 GB. They use a total capacity of 63 GB (= 31.5 GB×2), though the actual capacity is 34 GB (= 17 GB×2).

Note 10: The capacity of each LU, which has been registered to POOL, is managed in units of 4 GB. When the LU capacity is 5 GB, the LU is regarded as using a capacity of 8 GB. When there are two LUs has been registered to POOL; each has a capacity of 5 GB. They use a total capacity of 16 GB (= 8 GB×2), though the actual capacity is 10 GB (= 5 GB×2).

Note 11: For P-VOL/POOL supported volume of each cache memory, refer to the following graphs.



7.1.2 Redundancy

Snapshot and ShadowImage are identical functions from the viewpoint of producing a duplicate within a disk subsystem. However, the duplicated volume (S-VOL) of ShadowImage is a copy of the entire P-VOL data to a single LU; the duplicated volume (V-VOL) of Snapshot consists of the P-VOL data and differential data saved in the POOL. Therefore, when a hardware failure, such as a double failure of drives occurs in the P-VOL, a similar failure also occurs in the V-VOL and the pair status is changed to PSUE (see section 5.7).

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

When hardware failures occur in the POOL and S-VOL during a restoration for both Snapshot and ShadowImage, the P-VOL being restored accepts no Read/Write instruction. The difference between Snapshot and ShadowImage in redundancy is shown below.

- P-VOL failures

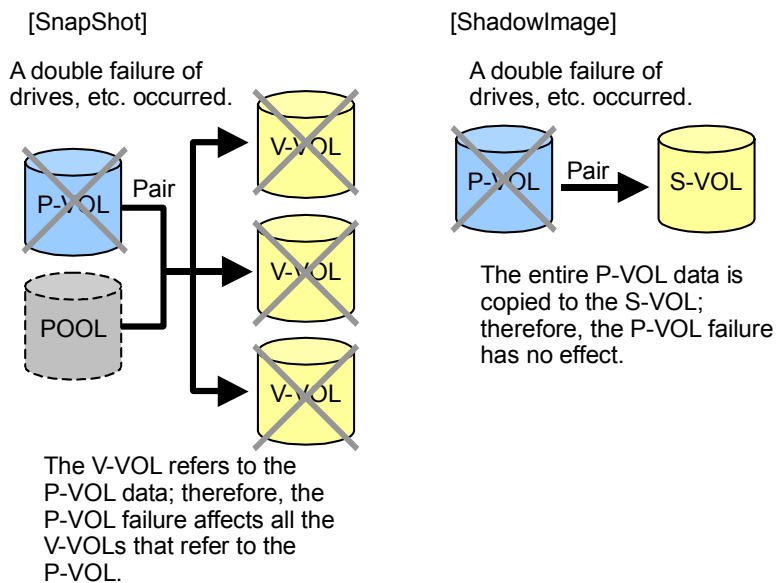


Figure 7.2 P-VOL Failures

- POOL (S-VOL) failures

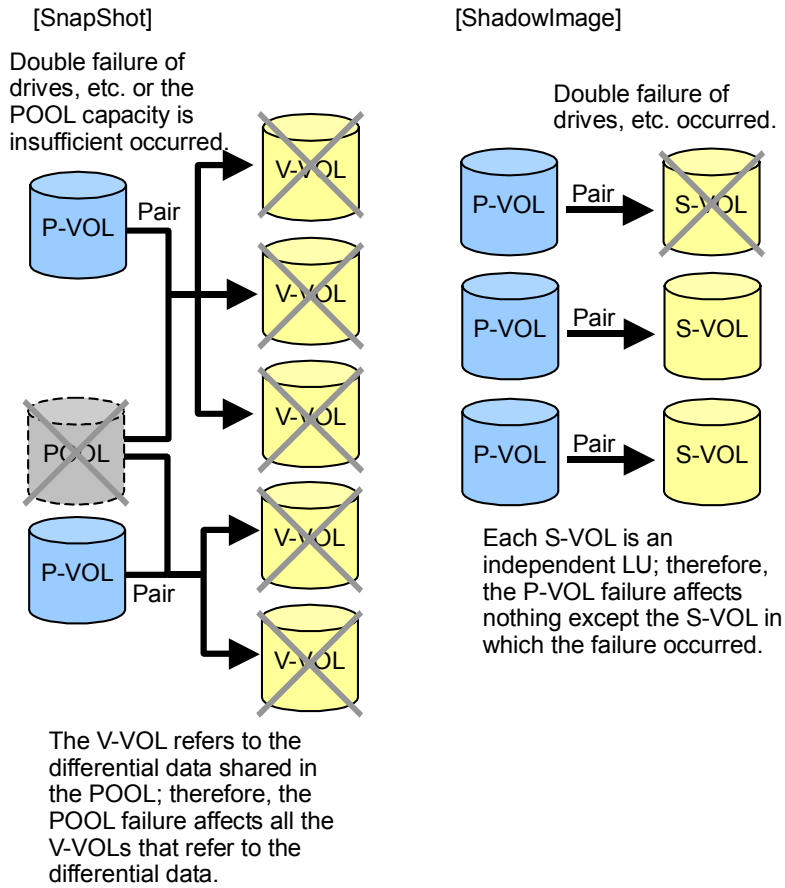


Figure 7.3 POOL (S-VOL) Failures

- POOL (S-VOL) failures during a restoration

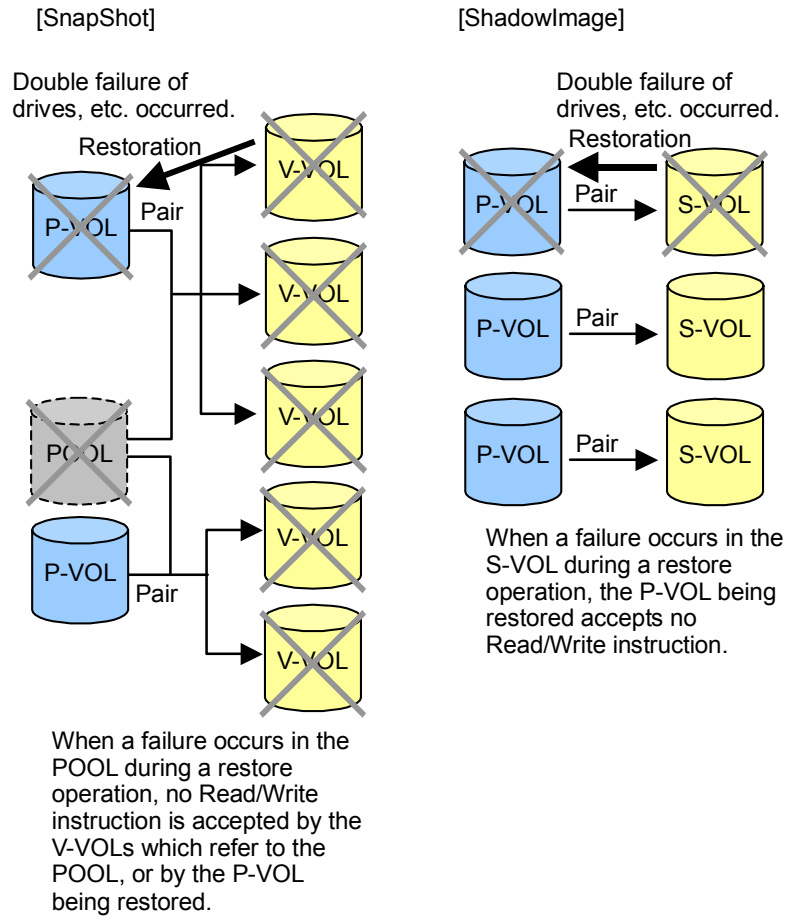


Figure 7.4 POOL (S-VOL) Failures during Restore Operation

7.1.3 Snapshot and ShadowImage Functions

Snapshot and ShadowImage both create a duplication within a disk subsystem; however, their uses differ because their external specifications and data assurance measures vary from each other. Advantages, disadvantages, and the functions of Snapshot and ShadowImage are shown below.

Table 7.8 Snapshot and ShadowImage Functions

Contents	Snapshot	ShadowImage
Advantages	<p>Amount of physical data to be used for the V-VOL is small because only the differential data is managed.</p> <p>Up to 15(14: the micro program earlier than version 0750/A) V-VOLs per P-VOL can be created.</p> <p>The POOL can be used by the two or more P-VOLs and the same number of V-VOLs by sharing between them; unitary management of its capacity can be done.</p> <p>A pair creation/resynchronization is completed in a moment.</p>	<p>When a hardware failure occurs in the P-VOL, it has no effect on the S-VOL.</p> <p>When a failure occurs in the S-VOL, it has no effect on the other S-VOLs.</p> <p>Access performance is only slightly lowered in comparison with ordinary cases because the P-VOL and S-VOL are independent LUs.</p>
Disadvantages	<p>Due to a hardware failure in the P-VOL, all the V-VOLs correlated to the P-VOL, in which the failure has occurred, are placed in the PSUE status.</p> <p>Due to a hardware failure or a shortage of the POOL capacity in the POOL, all the V-VOLs that use the POOL, in which the failure has occurred, are placed in the PSUE status.</p> <p>Data of the V-VOL, which has been placed in the PSUE status, cannot be restored.</p> <p>When the P-VOL is accessed, the access performance is lowered in comparison with ordinary cases because data is copied to the POOL.</p> <p>When the V-VOL is accessed, the performance regarding access to the POOL is lowered because the V-VOL data is shared among the P-VOL and POOL.</p>	<p>Only one S-VOL can be produced per P-VOL.</p> <p>The S-VOL must have the same capacity as the P-VOL.</p> <p>A pair creation/resynchronization requires time because it accompanies data copying from the P-VOL to S-VOL.</p>
Uses	<p>Backup for quick recovery</p> <p>To make a restoration quickly when software failure occurs, manage multiple backups (for example, by making backups every several hours and managing them according to their generations). It is important to backup onto a tape device due to low redundancy.</p> <p>Online backup</p> <p>Store backup data in a tape device when I/O operations are few (at night, for instance) with a disk capacity as small as possible.</p>	<p>Backup for quick recovery</p> <p>Not recommended</p> <p>Online backup</p> <p>Recommended</p> <p>When many I/O operations are required at night or an amount of data to be backed up is too large to be disposed during the night.</p>

Appendix A Operations Using CLI

This appendix describes the following operation procedure for Snapshot using the Navigator CLI. The following sections are included:

- Installing Snapshot (section A.1)
- Uninstalling Snapshot (section A.2)
- Enabling or Disabling Snapshot (section A.3)
- Setting Command Device (section A.4)
- Setting Differential Management LU (section A.5)
- Setting Pool (section A.6)
- Setting P-VOL and V-VOL (section A.7)
- Setting Mapping Information (section A.8)
- Confirming Status of Pair (section A.9)

For details on Navigator, refer to the *Hitachi TagmaStore Adaptable Modular Storage and Workgroup Modular Storage Storage Navigator Modular Command Line Interface (CLI)* (MK-95DF712).

Note: When you install, uninstall, enable, or disable Snapshot while the disk array subsystem is used on the remote side of TrueCopy or TCE, the following occurs when you restart the subsystem:

- Both paths of TrueCopy or TCE are blocked. When a path is blocked, a TRAP (notification to the SNMP Agent Support Function) occurs. Should this occur, notify all affected departments immediately. A blocked path of TrueCopy or TCE recovers automatically when the subsystem is restarted.
 - When the TrueCopy or TCE pair status is PAIR or COPY, it is changed to PSUE.
-
- **When you restart the disk array subsystem, the subsystem installs, uninstalls, enables, or disables Snapshot after changing the TrueCopy or TCE pair status to PSUS.**

A.1 Installing Snapshot

Since Snapshot is an extra-cost option, Snapshot cannot usually be selected (locked) before it can be used. To make Snapshot available, you must install Snapshot and set its function selectable (unlock).

Note: Before installing/uninstalling Snapshot, verify that the array unit to be operated is functioning normally. If a failure such as a controller blockage has occurred, installation/uninstallation cannot be performed.

To install Snapshot, a key code or key file provided with the optional feature is required. The following describes the installation procedure.

To install Snapshot:

1. From the command prompt, register the subsystem in which Snapshot is to be installed, then connect to the subsystem.
2. Execute `auopt` to install Snapshot. An example is shown below.

Example:

Storage Navigator Modular version is 5.00 or higher and Cache Partition Manager is enabled

```
% auopt -unit subsystem-name -lock off -keycode manual-attached-keycode
Password: manager-password
Are you sure you want to install the option? (y/n[n]): y
When Cache Partition Manager is enabled, if the option using data pool will be
enabled the default cache partition information will be restored.
Do you want to continue processing? (y/n [n]): y
The option is installed successfully.
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 TrueCopy, restarting the remote subsystem will cause both TrueCopy paths to
fail.
TrueCopy pair status will be changed to "PSUE" when pair status is "PAIR" or "COPY".
Please change TrueCopy pair status to "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.
%
```

Example:

Storage Navigator Modular version is 3.00 to 4.05 and Cache Partition Manager is enabled

```
% auopt -unit subsystem-name -lock off -keycode manual-attached-keycode
Password: manager-password
Are you sure you want to unlock the option? (y/n[n]): y
When Cache Partition Manager is enabled, if the option using data pool will be enabled
the default cache partition information will be restored.
Do you want to continue processing? (y/n [n]): y
The option is unlocked.
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.
Also, if you are logging in, the login status will be cancelled when restarting begins.
When using TrueCopy, restarting the remote subsystem will cause both TrueCopy paths to
fail.
TrueCopy pair status will be changed to "PSUE" when pair status is "PAIR" or "COPY".
Please change TrueCopy pair status to "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.
%
```

Example:

Before Storage Navigator Modular version is 3.00 and Cache Partition Manager is enabled

```
% auopt -unit subsystem-name -lock off -keycode manual-attached-keycode
Password: manager-password
Are you sure you want to unlock the option? (y/n[n]): y
The option is unlocked.
Restart the subsystem to apply the setting.
The subsystem stops accepting access from the host while restarting.
Also, if you are logging in, the login status will be canceled when restarting begins.
Do you want to restart the subsystem now? (y/n [n]): y
When using the TrueCopy, restart of the remote subsystem will cause both
TrueCopy paths failure. And when TrueCopy pair status is "PAIR" or "COPY",
TrueCopy pair status will be changed to "PSUE".
Do you want to continue processing? (y/n [n]): y
Now restarting the subsystem. Start Time hh:mm:ss Time Required 4 - 15min.
The subsystem restarted successfully.
%
```

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

3. Execute `auopt` to confirm whether Snapshot has been installed. An example is shown below.

Example:

```
% auopt -unit subsystem-name -refer
Password: manager-password
Option Name      Type      Term      Status
SNAPSHOT        Permanent ---      Enable
%
```

Snapshot is installed and the status is “Enable”. Snapshot installation is complete.

A.2 Uninstalling Snapshot

To uninstall Snapshot, a key code provided with the optional feature is required. Once uninstalled, Snapshot cannot be used (locked) until it is again unlocked using the key code or key file.

Note: The following conditions must be satisfied in order to uninstall Snapshot.

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

The following describes the uninstallation procedure.

To uninstall Snapshot:

1. From the command prompt, register the subsystem in which the Snapshot is to be uninstalled, then connect to the subsystem.
2. Execute the `auopt` command to uninstall Snapshot. An example is shown below.

Example:

Storage Navigator Modular Version 5.00 or higher

```
% auopt -unit subsystem-name -lock on -keycode manual-attached-keycode
Password: manager-password
Are you sure you want to de-install the option? (y/n[n]): y
The option is de-installed successfully.
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 TrueCopy, restarting the remote subsystem will cause both TrueCopy paths to
fail.
TrueCopy pair status will be changed to "PSUE" when pair status is "PAIR" or "COPY".
Please change TrueCopy pair status to "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.
%
```

Example:

Storage Navigator Modular Version 4.03 or higher

```
% auopt -unit subsystem-name -lock on -keycode manual-attached-keycode
Password: manager-password
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.
Also, if you are logging in, the login status will be cancelled when restarting begins.
When using TrueCopy, restarting the remote subsystem will cause both TrueCopy paths to
fail.
TrueCopy pair status will be changed to "PSUE" when pair status is "PAIR" or "COPY".
Please change TrueCopy pair status to "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.
%
```

Example:

Before Storage Navigator Modular Version 4.03

```
% auopt -unit subsystem-name -lock on -keycode manual-attached-keycode
Password: manager-password
Are you sure you want to lock the option? (y/n[n]): y
The option is locked.
Restart the subsystem to apply the setting.
The subsystem stops accepting access from the host while restarting.
Also, if you are logging in, the login status will be canceled when restarting begins.
Do you want to restart the subsystem now? (y/n [n]): y
When using the TrueCopy, restart of the remote subsystem will cause both
TrueCopy paths failure. And when TrueCopy pair status is "PAIR" or "COPY",
TrueCopy pair status will be changed to "PSUE".
Do you want to continue processing? (y/n [n]): y
Now restarting the subsystem. Start Time hh:mm:ss Time Required 4 - 15min.
The subsystem restarted successfully.
%
```

Note: The response time varies, depending on the condition of the array unit. If it does not respond after approximately 15 minutes, check the status and condition of the array unit.

3. Execute the `auopt` command to confirm whether Snapshot has been uninstalled. The example is shown below.

Example:

```
% auopt -unit subsystem-name -refer
Password: manager-password
DMEC002015: No information displayed.
%
```

Snapshot Uninstall is complete.

A.3 Enabling or Disabling Snapshot

Once installed, Snapshot can be enabled or disabled.

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 the Data Pools must be deleted.
- All the Snapshot Images (V-VOL) must be deleted.

The following describes the enabling/disabling procedure.

1. From the command prompt, register the subsystem (array unit) in which the status of the feature is to be changed, then connect to the subsystem.
2. Execute `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.

Example:

Storage Navigator Modular 4.03 and higher

```
% auopt -unit subsystem-name -option SNAPSHOT -st disable
Password: manager-password
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 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 TrueCopy, restarting the remote subsystem will cause both TrueCopy paths to
fail.
TrueCopy pair status will be changed to "PSUE" when pair status is "PAIR" or "COPY".
Please change TrueCopy pair status to "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.
%
```

Example:
(Before Storage Navigator Modular 4.03)

```
% auopt -unit subsystem-name -option SNAPSHOT -st disable
Password: manager-password
Are you sure you want to disable the option? (y/n[n]): y
The option has been set successfully.
Restart the subsystem to apply the setting.
The subsystem stops accepting access from the host while restarting.
Also, if you are logging in, the login status will be canceled when restarting begins.
Do you want to restart the subsystem now? (y/n [n]): y
When using the TrueCopy, restart of the remote subsystem will cause both
TrueCopy paths failure. And when TrueCopy pair status is "PAIR" or "COPY",
TrueCopy pair status will be changed to "PSUE".
Do you want to continue processing? (y/n [n]): y
Now restarting the subsystem. Start Time hh:mm:ss Time Required 4 - 15min.
The subsystem restarted successfully.
%
```

Note: The response time varies, depending on the condition of the array unit. If it does not respond after approximately 15 minutes, check the status and condition of the array unit.

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

Example:

```
% auopt -unit subsystem-name -refer
Password: manager-password
Option Name      Type      Term      Status
SNAPSHOT         Permanent ---      Disable
%
```

Snapshot Enable/Disable is complete.

A.4 Setting Command Device

The Command Device is a user-selected, dedicated logical volume on the TagmaStore subsystem that functions as the interface to the CCI software. The Snapshot commands are issued by the CCI (HORCM) to the TagmaStore Command Device.

In order to accept read and write commands that are executed by the TagmaStore subsystem and return read requests to the UNIX/PC host, the Command Device must be designated. The Command Device must be defined in the HORCM_CMD section of the configuration definition file for the CCI instance on the attached host. Two Command Devices can be designated for the TagmaStore subsystem. You can designate Command Devices using the Navigator.

Note 1: LUs set for Command Devices must be recognized by the host. The Command Device LU size must be greater than or equal to 33 MB.

Note 2: The following restrictions apply when either a pair of ShadowImage, Snapshot or TrueCopy exists or the path of True Copy is defined.

- When two command devices are set, only one command device can be released.
- When only one command device is set, the command device cannot be released.

To designate Command Device(s):

1. From the command prompt, register the subsystem to which you want to create the Command Device. Connect to the subsystem.
2. Execute the `aucmddev` command to create a Command Device.
The following is an example of specifying LU 200 for Command Device 1.

First, display the LUs to be assigned to the Command Device; then, create a Command Device.

To use the protection function of CCI, enter “enable” following the `-dev` option.

Example:

```
% aucmddev -unit subsystem-name -availablelist
Password: manager-password
Available Logical Units
  LUN Capacity  RAID Group RAID Level  D-CTL C-CTL Type Status
  200 35.0 Mbyte      0 5( 4D+1P)    0    0 FC  Normal
  201 35.0 Mbyte      0 5( 4D+1P)    1    1 FC  Normal
%
% aucmddev -unit subsystem-name -set -dev 1 200
Password: manager-password
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 created. The following shows the example.

Note: To set the alternate Command Device function or to avoid data loss and subsystem downtime, designate two Command Devices. For details on alternate Command Device function, refer to the *Hitachi TagmaStore Adaptable Modular Storage and Workgroup Modular Storage Command Control Interface (CCI) User and Reference Guide (MK-95DF701)*.

Example:

```
% aucmddev -unit subsystem-name -refer
Password: manager-password
Command Device LUN RAID Manager Protect
1             200 Disable
%
```

4. To release an already set Command Device, specify as follows.

Following is an example of releasing Command Device 1.

Example:

```
% aucmddev -unit subsystem-name -rm -dev 1
Password: manager-password
Are you sure you want to release the command devices? (y/n [n]): y
The command devices have been released successfully.
%
```

5. To change an already set Command Device, release the already set Command Device first, then change the LU number. Following is an example of specifying LU 201 for Command Device 1.

Example:

```
% aucmddev -unit subsystem-name -set -dev 1 201
Password: manager-password
Are you sure you want to set the command devices? (y/n [n]): y
The command devices have been set successfully.
%
```

A.5 Setting Differential Management LU

When the Differential Management LU is not set, create it. The Differential Management LU is an exclusive logical unit for storing differential data at a time when the volume is copied. The Differential Management LU in the disk subsystem is treated the same way as other logical units. However, a logical unit that is set as the Differential Management LU is not recognized by a host (it is hidden).

It is necessary to set a logical unit of 5 GBs; it must be at the least as large as the Differential Management LU. Up to two Differential Management LUs can be set and the second is used for mirroring. We recommend that you set two Differential Management LUs.

To designate Differential Management LU(s):

1. From the command prompt, register the subsystem to which you want to differential management LU. Connect to the subsystem.
2. Execute the `audmlu` command to create a differential management LU. The following is an example of specifying LU 0 for differential management LU.

First, display the LUs to be assigned to the differential management LU; then, create a differential management LU.

Example:

```
% audmlu -unit subsystem-name -availablelist
Password: manager-password
Available Logical Units
  LUN Capacity   RAID Group  RAID Level  D-CTL  C-CTL Type Status
    0  5.0 GByte         0    5 ( 4D+1P)    0      0 FC  Normal
%
% audmlu -unit subsystem-name -set -lu 0
Password: manager-password
Are you sure you want to set the DM-LU? (y/n [n]): y
The DM-LU has been set successfully.
%
```

3. To release an already set differential management LU, specify it as follows.

Example:

```
% audmlu -unit subsystem-name -rm -lu 0
Password: manager-password
Are you sure you want to release the DM-LU? (y/n [n]): y
The DM-LU has been released successfully.
%
```

Note: The following restrictions apply when either a pair of ShadowImage, Snapshot or TrueCopy exists, the path of True Copy is defined, or POOL of Snapshot is defined.

- When two differential management LUs are set, only one differential management LU can be released.
- When only one differential management LU is set, the differential management LU cannot be released.

A.6 Setting the POOL

A single data pool must be created for each controller, 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.
- An LU with an FC 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 subsystem to which you want to create the Data Pool, then connect to the subsystem.
2. Execute the `aupool` command create a Data Pool.

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

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

Example:

```
% aupool -unit subsystem-name -availablelist -poolno 0
Data Pool      : 0
Available Logical Units
  LUN Capacity   RAID Group RAID Level  Type Status
  100  30.0 Gbyte      0  5( 4D+1P)  FC  Normal
  200  35.0 Gbyte      0  5( 4D+1P)  FC  Normal
%
% aupool -unit subsystem-name -add -poolno 0 -lu 100
Password: manager-password
Are you sure you want to add the logical units to the data pool 0?
(y/n[n]): y
The logical unit has been successfully added.
%
```

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

Example:

```
% apool -unit subsystem-name -refer -poolno 0
Data Pool      : 0
Data Pool Usage Rate: 6% (2.0/30.0 Gbyte)
Threshold      : 70%
Status        : Normal
  LUN Capacity RAID Group RAID Level Type Status
  100 30.0 Gbyte      0      5( 4D+1P) FC Normal

Addable LU
  LUN Capacity RAID Group RAID Level Type Status
%
```

- To delete an existing Data Pool, refer to the following example.

This is an example of deleting Data Pool 0.

Note: When deleting the logical unit set as the Data Pool, it is necessary to delete all Snapshot images (V-VOLs).

Example:

```
% apool -unit subsystem-name -rm -poolno 0
Password: manager-password
Are you sure you want to delete all logical units from the data pool 0?
(y/n/[n]): y
The logical units has been successfully deleted.
%
```

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

An example of changing controller 0.

Example:

```
% apool -unit subsystem-name -cng -poolno 0 -thres 70
Password: manager-password
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.
%
```

A.7 Setting P-VOL and V-VOL

Following is the procedure for setting P-VOL and V-VOL.

To set V-VOL, it is necessary to specify the logical unit that will be assigned to P-VOL. If a specification for the logical unit assigned to a V-VOL is omitted when setting the V-VOL, Navigator assigns the smallest undefined number to the logical unit.

Note: The controller that has ownership of the V-VOL (controller that controls the V-VOL) is the same controller that controls the P-VOL and POOL.

To set P-VOL and V-VOL:

1. From the command prompt, register the subsystem to which you want to set the P-VOL and V-VOL, then connect to the subsystem.
2. Execute the `auvvol` command create a V-VOL.

Following is an example of specifying P-VOL 1.

Example:

```
% auvvol -unit subsystem-name -add -pvol 1
Password: manager-password
Are you sure you want to create the Snapshot image 30 to P-VOL 1?
(y/n[n]): y
The Snapshot image has been successfully created.
%
```

3. To delete an existing Snapshot Image, refer to the following example.

This is an example of deleting Snapshot Image 6.

Note: To delete the V-VOL, the pair state of that V-VOL needs to be SMPL.

Example:

```
% auvvol -unit subsystem-name -rm -lu 1
Password: manager-password
Are you sure you want to delete the Snapshot image 30 to P-VOL 1?
(y/n[n]): y
The Snapshot image has been successfully deleted.
%
```

A.8 Setting Mapping Information

Following is the procedure to specify Mapping Information. The Mapping Information is specified using the Navigator.

Note: When the Mapping mode is valid, the P-VOL and V-VOL cannot be paired. The mapping information is not set for the port that was set in the configuration definition file (not recognizable from a host). When you do not want the P-VOL and V-VOL recognized by a host, use the LUN Manager.

Note: For iSCSI model, use an `autargetmap` command instead of an `auhgmap` command.

1. From the command prompt, register the subsystem to which you want to set the Mapping Information, then connect to the subsystem.
2. Execute the `auhgmap` command to set the Mapping Information. The following is an example of setting the LU 0 in the subsystem to be recognized as 6 by the host. The port is connected via host group 0 of port 0A on controller 0.

Example:

```
% auhgmap -unit subsystem-name -MappingMode on
Password: manager-password
Are you sure you want to set the mapping mode? (y/n [n]): y
when setting starts, the subsystem stops accepting access to the controller from the
host.
Before setting, stop access to the controller from the host.
Do you want to continue processing? (y/n [n]): y
The mapping mode has been set successfully.
%

% auhgmap -unit subsystem-name -add 0 A 0 6 0
Password: manager-password
Are you sure you want to add the mapping information? (y/n [n]): y
The mapping information has been set successfully.
%
```

3. Execute `auhgmap` to verify that the target ID (LU Mapping) has been set.

Example:

```
% auhgmap -unit subsystem-name -refer
Password: manager-password
Mapping mode = ON
Port  Group  H-LUN  LUN
  0A      0      6      0
%
```

A.9 Confirming Status of Pairs

You can display the Snapshot paired logical units using CLI. The following steps describe the procedure.

1. From the command prompt, register the subsystem to which you want to display the status of paired logical volumes. Connect to the subsystem.
2. Execute `aupair` to display the pair status.

Example:

```
% aupair -unit subsystem-name -refer
P-VOL : 1
Capacity      : 5.0 Gbyte
RAID Group    : 0
RAID Level    : 5(4D+1P)
D-CTL        : 0
C-CTL        : 1
Data Pool     : 0
Data Pool Usage Rate : 6% (2.0/30.0 Gbyte)
Snapshot Image
  MU No.  LUN  Date of Enforcement  Rate of Change      Status
    1     5   2003/05/20 12:34:56  20% (1.0/5.0 Gbyte) PSUS
    2     3   2003/05/21 01:23:45  40% (2.0/5.0 Gbyte) PSUS
    --    6   ---                ---                SMPL
%
```

Pair status displays.

Appendix B Installing SnapShot when Cache Partition Manager is Being Used

SnapShot uses a part of the cache area to manage the internal resources. When this happens, the cache capacity that Cache Partition Manager can use therefore decreases.

Ensure that the cache partition information is initialized as shown in Figure B.1 and Figure B.2 when SnapShot is installed and Cache Partition Manager is already in use.

- All the logical units are moved to the master partitions on the side of the default owner controller.
- All sub-partitions are deleted and the size of the each master partition is reduced to half of the user data area after installing SnapShot.

An example of Cache Partition Manager being used is shown in Figure B.1 and an example of SnapShot being installed while Cache Partition Manager is used is shown in Figure B.2.

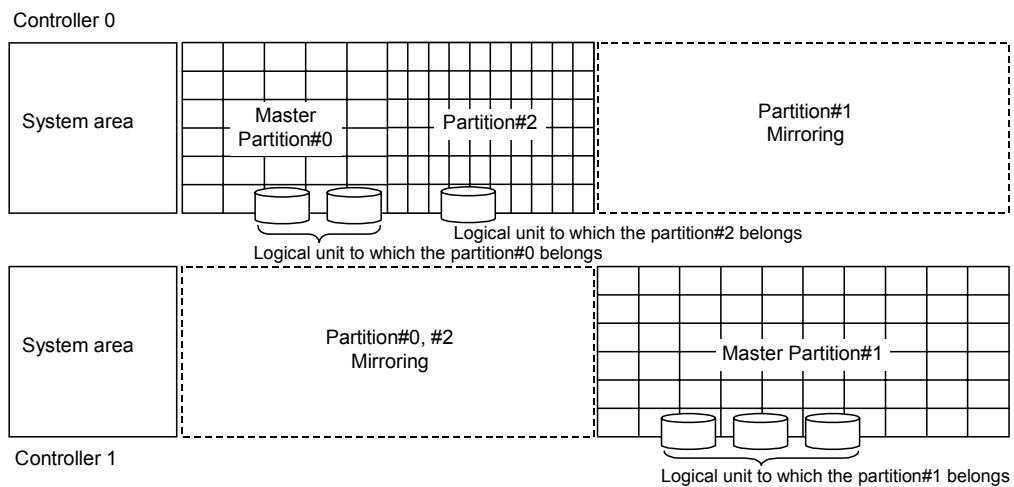


Figure B.1 When Cache Partition Manager is Used

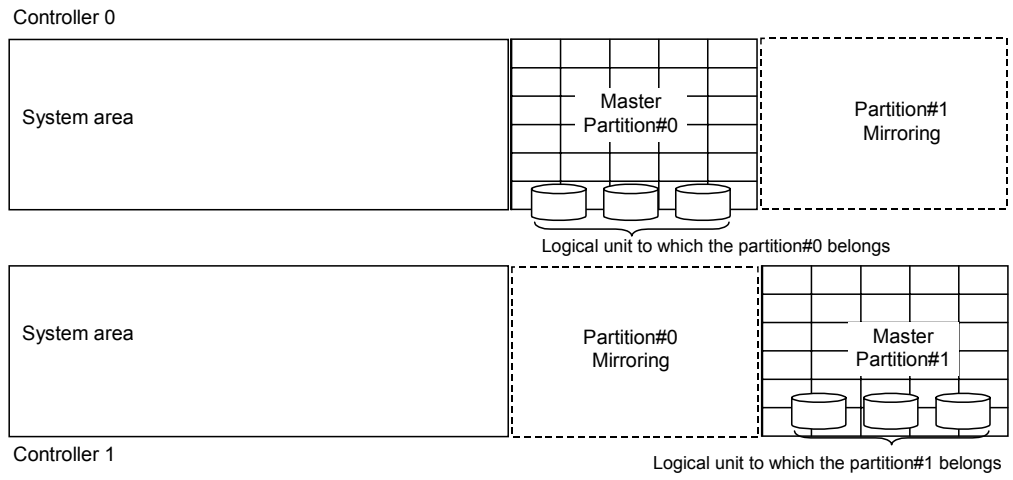


Figure B.2 Where SnapShot is Installed while Cache Partition Manager is Used

Acronyms and Abbreviations

AMS	Adaptable Modular Storage
API	application programming interface
ATA	Advanced Technology Attachment standard
CCI	command control interface
CFW	cache fast write
CIFS	common internet file system
CKD	count-key data
CLI	command line interface
CTG	consistency group
CTL	controller
CU	controller unit
DKC	disk controller unit
DM-LU	differential management logical unit
DWDM	dense wavelength division multiplexer
FC	fibre channel
HA	high availability
HDLM	Hitachi Dynamic Link Manager
HORCM	Hitachi Open Remote Copy Manager
H-LUN	host logical unit
H-RAIN	heterogeneous redundant array of independent nodes
IDE	integrated drive electronics; see also ATA.
IOPS	input output operations per second
LDEV	logical device
LUN	logical unit number
LUSE	LU size expansion
LVI	logical volume image
LVM	logical volume manager
MCU	main control unit
MDB	master directory block
MIB	message information block
MR	magneto-resistive
MU	mirror unit
MVS	multiple virtual storage
NAS	network attached storage
NFS	network file system
NDMP	Network Data Management Protocol
NSC	network storage controller
OSI	open systems interconnection
PFUS	pool full status

POSIX	portable operating system interface
PPRC	peer-to-peer remote copy
PSUE	pair suspended-error status
PSUS	pair suspended-split
PV	physical volume
P-VOL	primary volume
RAID	redundant array of independent disks
RPO	recovery point objective
RTO	recovery time objective
SAN	storage-area network
SATA	serial ATA
SCSI	small computer system interface
SM	shared memory module
SMB	server message block
SMTF	simple mail transfer protocol
SNIA	Storage Networking Industry Association
SNMP	simple network management protocol
SONET	synchronous optical network
SSL	secure socket layer
SSWS	suspend for swapping S-VOL
S-VOL	secondary volume
TID	target identifier
TPOF	tolerable points of failure
USP	Universal Storage Platform
VCS	Veritas Cluster Server™
VIB	volume information block
V-VOL	virtual volume (Snapshot Image)
VxVM	Veritas Volume Manager
WDM	wavelength division multiplexing

Glossary

A

Asynchronous

The term asynchronous is used to describe data communications between computers and devices which occurs intermittently rather than in a steady stream. Communication within a computer, however, is usually synchronous and is governed by the microprocessor clock.

Attribute

As used in this document, an attribute is one or more qualities possessed by an object.

Attribute

An attribute is one or more qualities possessed by an object.

B

background copy

A physical copy of all tracks from the source volume to the target volume

Bind

To bind is to assign a value to a symbolic placeholder. For example, when a program is bound, or linked, the binder replaces the symbolic addresses in the code with real machine addresses.

C

Cache

Cache is a temporary, high-speed storage mechanism. It can be either 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.

Cache fast write (CFW)

CFW is an attribute of record caching in which the cache fast write access function (either Simplex Write or Duplex Write) enables the specified record ID to be placed in the volatile control unit cache when a file-type macro is issued and the cache is available. If the cache is not available, the record is written directly to the DASD surface. A single write is issued to the prime module only or a duplexed write is issued to both the prime and the duplicate modules.

Cache Partition Manager

Cache Partition Manager (CPM) is an optional feature of the Hitachi AMS/WMS disk array subsystem. It enables the user data area of cache to be more finely tuned to the application and enables cache to be divided into sections called partitions to which logical units may be assigned.

Capacity

Capacity is the amount of information (in bytes) that can be stored on a disk drive. The capacity of a hard disk drive is usually expressed in megabytes. Capacity 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.

cascade configuration

A cascade configuration is a connection configuration of volume pairs in which a P-VOL or an S-VOL from a pair belonging to one copy function is used as a P-VOL or S-VOL of the other copy function. TCE supports cascading with ShadowImage and SnapShot.

Channel

A channel is the path data communication follows between two nodes of a network. It is the link between the central processor and the peripherals. A channel can be the physical cabling that connects the nodes on a network, an electronic signal traveling over a pathway, or a sub-channel in a carrier frequency.

channel adapter (CHA)

Provides the channel interface control functions and intercache data transfer functions. It is used to convert the data format between CKD and FBA. The CHA contains an internal processor and 128 bytes of edit buffer memory.

channel extender

A channel extender is a device used to increase the communication distances between channel-connected mainframe computers or between a computer and peripheral devices such as workstations, printers, and storage devices. Optical fiber channel connections are part of the system.

Configuration

Configuration for hardware involves setting various switches and jumpers. For software it means defining the values of parameters. For hardware and software respectively, configuration is the arrangement of the components that make up the system or the set up and set values of the software.

CIFS

Common Internet File System is a protocol used to expose the contents of an archive. CIFS allows clients to access files on a remote Windows computer as if they were part of the local file system.

Cluster

A cluster is 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

Cluster capacity is the total amount of disk space in a cluster, excluding the space required for system overhead and operating system. Cluster capacity is the amount of space available for all archive data, including original file data, metadata, and redundant data.

command devices

Command devices are dedicated logical volumes that are used only by management software such as CCI, to interface with the storage subsystems. Command devices are not used by ordinary applications. Command devices can be shared between several hosts. Up to two command devices can be configured per TCE subsystem.

concurrency of S-VOL

A state resulting when an S-VOL is synchronized, by simultaneously updating S-VOL with P-VOL data and data cached in the primary host memory. There may be discrepancies in S-VOL data, 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 S-VOL, cached data is written onto the P-VOL before subsequent remote copy operations

concurrent copy

Concurrent copy is a combined hardware, license, and software systems management solution that creates data dumps or copies while other applications are updating that data, allowing end-user processing to continue. Concurrent copy allows you to update the data in the files being copied, but the copy or dump of the data it secures does not contain any of the intervening updates.

configuration

Configuration for hardware involves setting various switches and jumpers. For software it means defining the values of parameters. For hardware and software respectively, configuration is the arrangement of the components that make up the system or the set up and set values of the software

configuration definition file

The configuration definition file describes the system configuration for making CCI operational in a TCE 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 TCE pairs on the primary or remote subsystem.

consistency group

A consistency group is a group of two or more logical units in a file system or 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 subsystem. 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.

console (administrative)

The cluster-specific web application that allows monitoring and managing of a HCA cluster and its individual nodes.

control unit (CU)

The control unit is a CPU component that implements microprocessor instructions.

count-key data (CKD)

Format for encoding data on hard drives, typically used in the mainframe environment. It is a physical disc format (Count, Key, Data) introduced by IBM with Series/360 2311 disks in 1964. Count-key-data (CKD) disks format each track as a new file is written on that track (all files have at least one track). CKD disks like SCSI and IDE have sector (count) ID fields and data fields. CKD disks can also have a third kind of field between ID and data called a key field.

cycle time

Cycle time is a user specified time interval used to execute recurring data updates for remote copying. Cycle time cycle updates are set for each subsystem and are calculated based on the number of CTGs.

cycle update

Cycle update processing involves periodically transferring differential data updates from the P-VOL to the S-VOL. TCE remote replication processes are implemented as recurring cycle update operations executed in specific time periods (cycles).

cycle update

Cycle update processing involves periodically transferring differential data updates from the P-VOL to the S-VOL. TCE remote replication processes are implemented as recurring cycle update operations executed in specific time periods (cycles).

D**dark fiber**

A dark fiber is an optical fiber cable that has been physically laid but is still unactivated.

Dark Fiber

The dark fiber is the optical fiber that does not work but it is laid. The optical fibers are generally laid by several dozen fibers to several hundred. Only fibers to be needed are activated and others are left as the dark fiber.

data pools

A data pool is a group of one or more disk volumes, designated to temporarily store untransferred differential data (in the local subsystem) or snapshots of backup data (in the remote subsystem). 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 data volume is a volume that stores database information, whereas other files, such as index files and data dictionaries, store administrative information, known as metadata.

dataset

A dataset is a named collection of data in an IBM mainframe operating system. A dataset in an IBM mainframe is the equivalent of a file in other operating systems, such as Mac OS, Windows and UNIX, for PCs.

device emulation

See logical volume image (LVI).

differential-data

The data to be updated from the suspended status of the pair volume to the primary volume.

differential-data

Differential-data is a set of snapshot data that consists of only the data that changed since a previous snapshot was taken. A backup of differential-data is called an incremental backup.

Differential data control

To control the differential data constantly.

Differential data copy

To copy the updated data to the secondary volume. The data is updated from the differential data control status (the pair volume is under the suspended status) to the primary volume.

Differential Management Logical Unit (DM-LUs)

Differential management logical units are logical units used to manage differential data in a storage subsystem. In a TCE system, there may be up to two DM-LUs configured per subsystem.

direct access storage device (DASD) fast write (DFW)

DFW is an attribute of record caching (while DASD Fast Write Access is a function of record caching) in which a specified record ID is placed in the cache and nonvolatile storage when a file-type macro is issued. If the cache is not available or the nonvolatile storage is not available, the record is written directly to the DASD surface.

disaster recovery

Disaster recovery implies procedures executed to recover critical application data and processing after a disaster recovery processes include failover and failback procedures.

disaster recovery

To recover the data and to keep an operation going when a critical failure.

Disc array device

Disc array subsystem. The disc array subsystem is referred to as “Disc array device” or “Disc subsystem” in this manual.

disk array device

In this manual, a disk array subsystem is sometimes referred to as a disk array device.

disk controller unit (DKC)

A disk controller unit consisting of CHA, CHF, DKA, Cache and other components except DKU.

DNS manager

The Domain Name System (manager) provides host-name resolution services to clients. It also balances requests across all nodes to ensure maximum cluster throughput and availability.

dual copy

Dual copy is the process of simultaneously updating of a P-VOL and S-VOL using a single write operation.

E**emulation**

Emulation is the ability of a program or device to imitate another program or device. Emulation causes a software package to accept that a device it is talking to is really another device. At the system level, emulation is a package of hardware, firmware, and software able to recreate a target machine environment on a new system.

Entire copy

To copy all data in the primary volume to the secondary volume for conforming the data of both volumes completely.

Extender

A converter that changes a signal to other signal.

When the data is transmitted to a distance, an extender is used for changing a signal for Fibre Channel to a signal for a dark fiber or an Ethernet (IP).

Extender

An extender is a converter used to change signals when data is transmitted over long distances. For instance, changing a fiber channel signal to a signal for dark fiber or an Ethernet (IP).

extent

An extent is a contiguous area of storage in a computer file system that is reserved for writing or storing a file.

F

fabric

The hardware that connects workstations and servers to storage devices in a SAN. The SAN fabric enables any-server-to-any-storage device connectivity through the use of Fibre Channel switching technology.

Failover

To take over a process or a data with an alternative host when a failure occurs in a host. By using the High Availability software, an automatically switching the hosts is possible.

Failover

A failover operation involves takeover of critical application processing by an alternate host in the event of a failure at the primary site. Automatic switching of hosts is possible using High Availability software.

fallback

Fallback refers to the process of restarting business operations at a local site using the P-VOL, after the storage subsystems have been recovered.

Fibre channel

Input/output channel using Fibre cable.

fibre channel (FC)

Input/output channel using optical fiber cables. It is the physical media that forms the lowest layer of fibre channel transport.

firmware

Microcode complementing the hardware used to implement the architecture of a system.

fixed block architecture

A model of disks in which storage space is organized as linear, dense address spaces of blocks of a fixed size. Abbreviated FBA. Fixed block architecture is the disk model on which SCSI is predicated. cf. count-key-data.

fixed-content data

Fixed-content data is an exact digital reproduction of a data file as it existed before the file was archived. Fixed-content data cannot be modified or deleted before its retention period expires.

G

gateway

A gateway is a protocol that provides users and applications access to data in a cluster.

granularity of differential data

The 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 P-VOL. The amount of differential data that can be managed per write command is limited by the difference between the number of incoming host write operations (inflow) and outgoing data transfers (outflow).

H

HCA cluster

An HCA cluster is an implementation of Hitachi Content Archiver. An HCA cluster is both a repository that stores terabytes of data and a gateway that enables access to that data.

High Availability (HA) software

High Availability software is used for automatically switching to a stand-by host, in the event of primary host or disk failure. High availability software has to be installed on the primary and secondary hosts.

High Availability (HA) Software

Software that automatically switches a host to a stand-by host (fail-over) during a failure of a host and disks. High Availability software is installed in more than one host respectively.

H-RAIN

Heterogeneous Redundant Array of Independent Nodes. A collection of networked servers that can differ by vendor or model.

HTTP

Hypertext Transfer Protocol is one of the protocols used to expose the contents of an archive. Using HTTP, archived files and directories can be viewed on a web page.

HTTPS

HTTPS is HTTP with SSL security.

ICKDSF

A DSF command used to perform media maintenance.

I

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 MCU port of Fibre Remote Copy function.

iSCSI

iSCSI (Internet-Small Computer Systems Interface) is used as an IP-based standard for carrying SCSI commands over IP networks which link data storage devices and allows the transfer of data.

Java applet

A Java applet is a program written in the Java™ programming language that can be included in an HTML page. When you use a Java technology-enabled browser to view a page that contains an applet, the applet's code is transferred to your system and executed by the browser's Java Virtual Machine (JVM).

J

journal volume

A journal volume is a volume in which the hard disk maintains data integrity in the event of a system crash. The journal volume maintains a log or journal of the activities that have taken place in the volume. This journal allows any lost data to be recreated, because updates to the metadata in directories and bit maps have been written to a serial log.

L

logical

Logical is used to describe a user's view of the way data or systems are organized. The opposite of logical is physical, which refers to the real organization of a system. A logical description of a file is that it is a quantity of data collected together in one place. The file appears this way to users. Physically, the elements of the file could live in segments across a disk.

logical device (LDEV)

A logical device is a group of hardware items that the operating system treats as a single unit.

Logical Unit (LU)

See User Logical Unit (LU)

logical unit number (LUN)

LUN is a three-bit code identifier for a logical unit. LUN0-7 can be assigned.

logical volume

An area on a disk consisting of device files that are logically integrated using a volume manager.

logical volume image (LVI)

LVI (also called device emulation) is a feature used to create virtual LUs that are up to 36 times larger than the standard OPEN-x LUs.

Logical Volume

An area on a disk consisting of device files that are logically integrated using a volume manager.

LU

Logical unit in a device that is connected with Fibre.

LUN

Logical Unit Number. The number of the unit that is connected with Fibre.

M**mainframe**

A mainframe is a large and expensive computer capable of simultaneously supporting thousands of users. In this document the term mainframe is used for IBM computers (zSeries® and S/390®-based systems). This term also marks a distinction between Unix or Windows server computers and the larger more, powerful mainframe. A mainframe is also commonly referred to as the *host*, even though any computer host having a unique IP address can equally be referred to as a *host*.

metadata

Metadata is information about an archived object.

microcode

Microcode is the lowest-level instructions directly controlling a microprocessor. Microcode is generally hardwired and cannot be modified.

Microsoft cluster server

Microsoft Cluster Server is a clustering technology built that supports clustering of two NT servers to provide a single fault-tolerant server.

middleware

Middleware is software that connects two otherwise separate applications. For example, a middleware product can be used to link a database system to a Web server. Using forms, users request data from the database; then, based on the user's requests and profile, the Web server returns dynamic Web pages to the user.

mount

To mount a device or a system means to make a storage device available to a host or platform.

mount point

Mount points are 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.

multiple allegiance support

The multiple allegiance support means that the storage unit can accept concurrent I/O requests for a volume from multiple channel paths. Therefore the storage unit can process requests from separate FICON hosts in parallel, improving throughput and performance.

multiple virtual storage

MVS (including MVS/370, MVS/ESA, MVS/XA) is an operating system that runs on IBM or compatible mainframe computers. The host component works on MVS/ESA (Enterprise Systems Architecture) and MVS/XA (Extended Architecture).

N**NDMP**

Network Data Management Protocol is a protocol used to backup and restore archived objects.

NFS

Network File System is a protocol used to expose the contents of an archive. NFS allows clients to access files on a remote computer as if they were part of a local file system.

node

In networks, a node is a processing location. A node can be a computer or some other device, such as a printer. Every node has a unique network address. In HCA clusters, nodes are Linux-based servers running HCA software and networked to form an HCA cluster.

open device

Collectively refers to the host computer, peripheral control units, and intelligent peripherals that are connected to fibre channel.

P**pair**

See ShadowImage pair.

P-VOL

Primary volume, the volume on the primary side of the pair volume. The primary volume controls the pair status, and the status reflects on the secondary volume.

pair splitting

Pair splitting refers to the termination of a volume pair relationship to temporarily stop update copy processing for the specified volume pair. Pairs may also be split before system reduction tasks.

pair status

A pair status is an 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.

Pair status

The status of the logical volume that is paired.

Paired volume

Primary and secondary volume that are paired in a disk array device.

paired volumes

Paired volumes are primary and secondary volumes comprising a volume pair.

Panel

In this document, a *panel* is equivalent to a *window*.

Parity

The quality of being either odd or even. The fact that all numbers have a parity is commonly used in data communications to ensure the validity of data. This is called parity checking. So parity provides 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 groups

RAID groups can contain single or multiple parity groups where the parity group acts as a partition of that container.

path blockade watch

The path blockade watch setting specifies the time for monitoring blockade in the Fibre Channel paths on the MCU side. The path blockade watch value must be from 0 to 45 seconds. This setting is available for Fibre Channel interface only.

pattern file

A pattern is a table that contains the access attributes of all logical volumes. Pattern files enable administrators to change the access attributes of all logical volumes quickly and easily.

peer-to-peer remote copy (PPRC)

The Peer-to-Peer Remote Copy (PPRC) function is 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).

permission

Using SMB/CIFS, HTTP, WebDAV, and SMTD gateways, permissions are granted from the owner to members of a group or other users to allow access to data files or directories in an archive. Read, write, and execute permissions can be granted for files, directories, or symbolic links.

point-in-time logical copy

A logical copy or snapshot of a volume at a point in time. This allows a backup or mirroring application to run concurrently with the system.

point-to-point

A configuration that allows two ports to be connected serially.

policy

A policy is a process that performs a specific function to aid maintenance of the overall health of a cluster. For example, authentication, cluster balance, garbage collection, protection, retention, and scavenging.

pool volume

A pool volume is used to store backup versions of files, archive copies of files, and files migrated from other storage.

POSIX

Portable Operating System Interface for UNIX is a set of standards that define an application programming interface (API) for software designed to run under heterogeneous operating systems.

primary or local site

A site where the production applications run.

primary volume (P-VOL)

A primary volume is the storage volume in a volume pair, 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).

Q

quota values

Quota values are set for users with write access to volumes providing data storage limits for that user/volume. Quota values are applied to a snapshot and set for the target file system when the snapshot is taken.

R

RCU target port

A port-type used for RCU port of Fibre Remote Copy function. This port allows LOGIN of host computers and MCUs.

recovery point objective

Recovery point objective is the maximum desired time period prior to a disaster, during which changes to data may be lost as a result of recovery. This measure determines up to what point in time data should be recovered in the event of a disaster. Data changes preceding the disaster by at least this time period are preserved by recovery.

recovery time objective

Recovery time objective is the maximum desired time period required to bring one or more applications and associated data back to a correct operational state. It defines the time frame within which specific business operations or data must be restored to avoid any business disruption.

Remote or target site

A site that has the mirrored data of the production site.

remote or target site

A site that has the mirrored data of the primary site.

remote path

A remote path is a route connecting identical ports on the local subsystem and the remote subsystem. Two remote paths must be setup for each subsystem (one path for each of the two controllers built in the subsystem).

remote volume stem

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

repeater

A repeater is a network device used to regenerate or replicate a signal. A repeater relays messages between sub-networks that use different protocols or cable types. A repeater cannot do the intelligent routing performed by bridges and routers.

Resynchronization

Resynchronization is a copy operation performed to make data in the secondary volume consistent with data in the primary volume. This operation involves copying only untransferred differential data to the target secondary volume.

Resynchronization (TrueCopy)

To copy the differential data to the secondary volume for conforming the data of both the primary and secondary volumes. When the resynchronization is completed, the status changes to the pair status.

S**SATA**

Serial ATA is a serial link, a single cable with a minimum of four wires creates a point-to-point connection between devices. ATA is a computer bus technology primarily designed for transfer of data to and from a hard disk. SATA is the successor to the legacy Advanced Technology Attachment standard (ATA, also known as IDE or Integrated Drive Electronics).

secondary volume (S-VOL)

A secondary volume (S-VOL) is a replica of the primary data volume (P-VOL), maintained on the standby subsystem. Recurring differential data updates are performed to keep the data in the S-VOL consistent with data in the P-VOL.

sequential data striping

Sequential data striping refers to writing to multiple disk drives in a pre-planned sequence. Because the processor writes faster than the disk can accept, it has the left-over capacity to locate the next segment of the logically sequential data and prepare to write to it.

Service

A service is the set of functions that one of the seven (7) Open Systems Interconnection (OSI) model layers delivers to the layer above it. For example, the TCP layer provides a reliable byte-stream service to the application layer above it.

server set identifier

The SSID is an alphanumeric name that is 1-32 bytes. The purpose of an SSID is to help hardware clients find and connect to an access point (AP) on the correct network.

session

A session is a series of communications or exchanges of data between two end points that occurs during the span of a single connection. The session begins when the connection is established at both ends and terminates when the connection is ended. For some applications each session is related to a particular port. In this document a session the exchange of data between groups of primary and secondary volumes

ShadowImage

ShadowImage is a software program that replicates user data on TagmaStore AMS/WMS disks, bypassing the host system.

ShadowImage

ShadowImage is a software program that replicates user data on TagmaStore® USP disks, bypassing the host system.

Shadow image Pair

A disk is a Logical Volume Image (LVI). S-VOLs and T-VOLs are ShadowImage volumes for Source and Target. Data is physically copied from the S-VOL to the T-VOL.

shared memory module (SM)

Stores the shared information about the system and the cache control information (director names). This type of information is used for the exclusive control of the system. Like CACHE, shared memory is controlled as two areas of memory and fully non-volatile (sustained for approximately 7 days).

sidefile

A side file is an area of a controller's storage that occupies about 1MB of the controller storage. The sidefile is used to hold data that is changed or updated while being backed up or copied. The sidefile holds it for later integration into the copied data.

SMB

Server Message Block is a protocol used to expose the contents of an archive. SMB allows clients to access files on a remote computer as if they were part of a local file system.

SMTP

Simple Mail Transfer Protocol is a protocol used to receive and store email data directly from email servers.

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.

SNMP

Simple Network Management Protocol is a protocol used to facilitate monitoring and management of clusters through an external interface. SNMP sends notifications to IP addresses whenever certain types of events occur.

source copy

Source copy is the place from which data is taken. The place from which the data is moved is called the source. The source can also indicate the node on a network from which data is sent to its destination.

SSL

Secure Sockets Layer is a key-based Internet protocol for transmitting documents through an encrypted link.

SSL certificate

An SSL certificate is a file containing the cryptographic keys and signatures used with an SSL protocol to verify the authenticity of web sites to protect data sent to or from that site.

Status transition

To change the pair status of the pair volume

Storage Navigator

The TagmaStore Storage Navigator consists of a group of Java™ applet programs that enable users to manage the TagmaStore subsystem. Storage Navigator Java™ applet programs run on a web browser to provide a user-friendly interface for TagmaStore web client functions.

Suspended status

The status when the update operation is suspended with keeping the pair status. Under this status, the differential data control for the updated data is performed in the primary volume.

S-VOL

Secondary volume, the volume on the secondary side of the pair volume. The primary volume controls the pair status, and the status reflects on the secondary volume.

S-VOL determination

S-VOL determination is an internal process on the remote subsystem, which replicates the S-VOL independent of the update operations. This process occurs at the end of each update cycle. This process allows a pre-determined copy of S-VOL data consistent with P-VOL data (as of the previous cycle) to be maintained on the remote site at all times.

sysplex

Sysplex denotes *system complex*. This is a processor complex formed by connecting a number of processors together into a single unit through channel-to-channel adapters or ESCON/FICON fiber optic links. The processors are synchronized using a Sysplex Timer and are managed as a single system image (SSI1). The Sysplex Timer is an invaluable component when systems on multiple CPCs share access to the same data.

system reduction

System reduction refers to maintenance tasks performed to improve system performance. These tasks may include pair deletion, deletion of command devices and data pools.

takeover processing

Takeover processing involves transferring of critical application processing to the S-VOL on the remote standby subsystem. The remote S-VOL is immediately enabled to process subsequent host I/O operations.

target copy

Target copy is a file, device or any type of location to which data is moved or copied.

target port

A port-type which is different from "Initiator Port" and "RCU Target Port". This port is a normal target port which is used without configuration of Fibre Remote Copy. This "Target port" allows LOGIN of host computers. It does not allow LOGIN of MCUs

target site

See Remote or target site.

tier architecture

Tier 1 is fully supported computing expected to be production quality. Tier 2 platforms are not supported by the security officer and release engineering teams. Tier 2 systems are targeted at Tier 1 support, but are still under development. Tier 3 platforms are architectures for which hardware is not or will not be available or which are considered legacy systems unlikely to see broad future use. Tier 4 systems are not supported in any way.

Tier 1: Static content, Tier 2: Application logic, Tier 3: Database

TPOF

Tolerable points of failure consists of the number of concurrent failures beyond with a cluster is no longer viable. A failure is defined as either a node or a disk not functioning properly or responsively.

track

A track is a ring on a disk where data can be written. For hard disks, tracks aggregate into platters and a single track location that cuts through all platters is termed a cylinder. Each track can be subdivided into a number of sectors. The operating system and disk drive find stored information using its track and sector numbers.

truck size

Truck size represents a fixed sector size for each volume type.

Target site

See Remote or target site.

Trap

A program interrupt usually caused by some exceptional situation in a user program. In most cases, the OS performs some action and then returns control to the program.

TrueCopy

TrueCopy is a software program that replicates user data between two TagmaStore USP disks, bypassing the host system.

TrueCopy

The TrueCopy™ feature enables you to create and maintain duplicate copies of all user data stored on a Hitachi TagmaStore™ subsystem for data duplication, backup, and disaster recovery purposes.

TrueCopy

TrueCopy is a software program that replicates data between two TagmaStore disks, independent of the host system. TrueCopy versions are available for TagmaStore AMS/WMS and USP/NSC subsystems. TrueCopy for z/OS is a mainframe version.

U

User Logical Unit (LU)

A user logical unit is a term used to describe any device file located on an external disk subsystem connected to the TagmaStore USP or NSC by a fibre channel.

user logical unit (LU)

A user logical unit is a term used to describe any device file located on an external disk subsystem connected to the TagmaStore subsystem by a fibre channel.

V

Volume

A volume is the basic unit of storage that includes recovery logs and storage pools. A volume can be a logical volume management (LVM) logical volume, a standard file system file, a tape cartridge, or an optical cartridge. The various types of defined volumes include: external, internal, copy source, copy destination, reserve, data, journal, virtual, pool, system, LUSE, copy pair, and USP.

Volume copy

To copy all data of P-VOL into S-VOL.

virtual LVI/ LUN (VLL)

Virtual LVI/LUN is an option that enables the configuration of custom-size logical device images and logical units, which are smaller than standard size devices.

volume

A volume is the basic unit of storage that includes recovery logs and storage pools. A volume can be a logical volume management (LVM) logical volume, a standard file system file, a tape cartridge, or an optical cartridge. The various types of defined volumes include: external, internal, copy source, copy destination, reserve, data, journal, virtual, pool, system, LUSE, copy pair, and USP.

Volume copy

To copy all data of P-VOL into S-VOL.

volume pair

A volume pair is formed by pairing two logical data volumes. It typically consists of one primary volume (P-VOL) on the local storage subsystem and one secondary volume (S-VOL) on the remote storage subsystems. TCE remote copy operations are performed using logical volume pairs.

volume signature

A volume signature is an integer that is an element in the master directory block (MDB) (volume information block (VIB)). For example, for HFS volumes, this field (drSigWord) contains the number \$4244.

WDM

Wavelength Division Multiplexing. Generally, WDM is to multiplex the optical signal of several channels and DWDM (Dense WDM) is to multiplex the optical signal of several dozen channels. Generally, WDM is to multiplex the optical signal of several channels and DWDM (Dense WDM) is to multiplex the optical signal of several dozen channels.

WDM/DWDM

WDM denotes Wave Division Multiplexing technology used to multiplex optical signals from multiple channels. DWDM denotes Dense Wave Division Multiplexing and is used to multiplex optical signals of several dozen channels.

World Wide Name (WWN)

A unique identifier for an open systems host. It consists of a 64-bit physical address (the IEEE 48-bit format with a 12-bit extension and a 4-bit prefix). The WWN is essential for defining the SAnTinel™ parameters because it determines whether the open systems host is to be allowed or denied access to a specified LU or a group of LUs.

WORM

Write once, read many is a data storage technique in which files are protected from being modified, overwritten, or deleted.

write order guarantee

The write order guarantee feature ensures that data is updated in S-VOL in the same order that the host was updated data in the P-VOL, particularly when there are multiple write operations in one update cycle. This feature is critical to maintain data consistency in the remote S-VOL and is implemented by inserting sequence numbers in each update record. Update records are then sorted in the cache within the remote system, to assure write sequencing.