Hitachi Data Systems

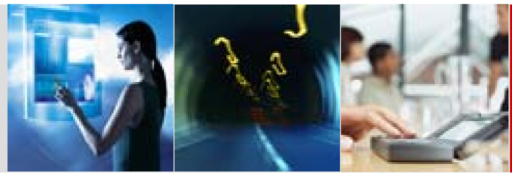

Student Guide for Replication Fundamentals for Hitachi Modular Storage

CSI0157

Courseware Version 4.0

NOTICE:

Notational conventions: 1KB stands for 1,024 bytes, 1MB for 1,024 kilobytes, 1GB for 1,024 megabytes, and 1TB for 1,024 gigabytes, as is consistent with IEC (International Electrotechnical Commission) standards for prefixes for binary and metric multiples.

HDS Academy 0079

Contact Hitachi Data Systems at www.hds.com.

Product Names mentioned in courseware:

Enterprise Storage Systems

- Hitachi Universal Storage Platform® V
- Hitachi Universal Storage Platform® VM
- Hitachi Universal Storage Platform®
- Hitachi Network Storage Controller

Legacy Products:

- Hitachi Lightning 9900™ Series enterprise storage systems
- Hitachi Lightning 9900™ Series enterprise storage systems

Modular Storage Systems

- Hitachi Adaptable Modular Storage system
- Hitachi Adaptable Modular Storage 200
- Hitachi Adaptable Modular Storage 500
- Hitachi Adaptable Modular Storage 1000
- Hitachi Adaptable Modular Storage 200 with iSCSI
- Hitachi Adaptable Modular Storage 200 with NAS Option
- Hitachi Adaptable Modular Storage 500 with iSCSI
- Hitachi Adaptable Modular Storage 500 with NAS Option
- Hitachi Adaptable Modular Storage 1000 with iSCSI
- Hitachi Adaptable Modular Storage 1000 with NAS Option
- Hitachi Adaptable Modular Storage 2000 Family
- Hitachi Adaptable Modular Storage 2100
- Hitachi Adaptable Modular Storage 2300
- Hitachi Adaptable Modular Storage 2500
- Hitachi Workgroup Modular Storage system
- Hitachi Workgroup Modular Storage 100
- Hitachi Workgroup Modular Storage 100 with iSCSI
- Hitachi Workgroup Modular Storage 100 with NAS Option
- Hitachi Simple Modular Storage

Legacy Products:

- Hitachi Thunder 9500™ Series modular storage systems
- Hitachi Thunder 9200V™ entry-level storage

NAS Storage Systems

- Hitachi Essential NAS Platform®
- Hitachi Essential NAS Platform™
- Hitachi High-performance NAS Platform, powered by BlueArc®
- Hitachi NAS Blade for Universal Storage Platform™ and Network Storage Controller

- Hitachi High-performance NAS Platform, powered by BlueArc® 2000 family
    - Hitachi High-Performance NAS Platform 2000
    - Hitachi High-Performance NAS Platform 2100
    - Hitachi High-Performance NAS Platform 2200
- Hitachi High-performance NAS Platform, powered by BlueArc® 3000 family
    - Hitachi High-Performance NAS Platform 3100
    - Hitachi High-Performance NAS Platform 3200

Management Tools

- Hitachi Basic Operating System
- Hitachi Basic Operating System V
- Hitachi Resource Manager™ utility package
    - Module Volume Migration Software
    - LUN Manager/LUN Expansion
    - Network Data Management Protocol (NDMP) agents
    - Logical Unit Size Expansion (LUSE)
    - Cache Partition Manager feature
    - Cache Residency Manager feature
    - Storage Navigator program
    - Storage Navigator Modular program
    - Storage Navigator Modular 2 program
- Hitachi NAS Blade Manager software
- Hitachi NAS Manager Suite of software

Replication Software

Remote Replication:

- Hitachi Universal Replicator software
- Hitachi TrueCopy® Heterogeneous Remote Replication software bundle (for enterprise systems)
- Hitachi TrueCopy® Remote Replication software bundle (for modular systems)
- Hitachi TrueCopy® Synchronous software
- Hitachi TrueCopy® Asynchronous software
- Hitachi TrueCopy® Extended Distance software

Hitachi In-System Replication software bundle:

- Hitachi ShadowImage® Heterogeneous Replication software (for enterprise systems)
- Hitachi ShadowImage® Replication software (for modular systems)
- Hitachi Copy-on-Write Snapshot software

Hitachi Storage Command Software Suite

- Hitachi Chargeback software
- Hitachi Device Manager software

- Hitachi Dynamic Link Manager software
- Hitachi Global Link Availability Manager software
- Hitachi Global Reporter software
- Hitachi Path Provisioning software
- Hitachi Protection Manager software
- Hitachi QoS for File Servers software
- Hitachi QoS for Oracle software
- Hitachi Replication Monitor software
- Hitachi Storage Services Manager software
- Hitachi Tiered Storage Manager software
- Hitachi Tuning Manager software

Other Software

- Hitachi Backup and Recovery software, powered by CommVault®
- Hitachi Backup Services Manager software, powered by APTARE®
- Hitachi Business Continuity Manager software
- Hitachi Command Control Interface (CCI) Software
- Hitachi Dynamic Provisioning software
- Hitachi Storage Resource Management Solutions
- Hitachi Volume Migration software
- Hi-Track® Monitor "call home" service/remote monitoring tool

Other Solutions and Terms

- Hitachi Content Archive Platform
- Hi-Star™ crossbar switch architecture
- Hitachi Universal Star Network™ V

Contents

Introduction

- Student Introductions
  - Name
  - Position
  - Experience
  - Expectations

2

Intended Audience

- Hitachi Data Systems employees
  - Pre-sales
  - Architect
  - Implementation and Support
  - Storage Manager
- Hitachi Channel Partners
  - Pre-sales
  - Architect
  - Implementation and Support
  - Storage Manager
- Hitachi Customers
  - Architect
  - Implementation and Support
  - Storage Manager

3

Prerequisites

- Knowledge of Hitachi modular storage systems or completion of one of the following courses:
  - Hitachi Adaptable Modular Storage 2000 Family Architecture and Operations (TCI1830)
  - Hitachi Adaptable Modular Storage 2000 Family Installation, Configuration, and Support (TCI1835)
  - Hitachi Adaptable Modular Storage 2000 Family Hardware Differences (TCI1780)

4

Course Description

- This five day instructor-led course provides product fundamentals and tasks necessary to size and implement in-system and remote replication for the Adaptable Modular Storage 2000 family of systems.
- Participants are introduced to Hitachi ShadowImage® Replication software, Hitachi Copy-on-Write Snapshot software and Hitachi TrueCopy® Remote Replication software, including Hitachi TrueCopy® Extended replication software, with details on software specifications and theory of operations. Participants will have the opportunity to install, configure, use, and troubleshoot the replication software applications in hands-on labs.

5

Course Objectives

- Upon completion of this course, the learner should be able to:
  - Describe the key features and operations of Hitachi replication software, including Hitachi TrueCopy, TrueCopy Extended Distance, ShadowImage, and Copy-on-Write Snapshot software
  - Compare the functionality of TrueCopy to TrueCopy Extended Distance, and ShadowImage to Copy-on-Write Snapshot replication software
  - Recognize key Hitachi replication software specifications and performance factors
  - Install, configure, use, and troubleshoot Hitachi replication software managing key features, commands, and operations
  - Use Hitachi Command Control Interface (CCI) operations and commands
  - Develop simple CCI scripts to automate software operations
  - Use Hitachi Storage Navigator Modular 2 GUI and CLI operations and commands
  - Explain key considerations for SAN deployment
  - Compare differences between iSCSI and Fibre Channel environment for implementing replication

6

Course Topics

Content	Learning Activities
1. Introduction to Hitachi ShadowImage Replication Software 2. Hitachi ShadowImage Replication Software Performance 3. Hitachi Storage Navigator Modular 2 Basic Operations	1. CCI Install for Sun Solaris and Microsoft Windows 2. ShadowImage Software Operations
4. Hitachi ShadowImage Replication Software with Hitachi Command Control Interface 5. Hitachi Storage Navigator Modular 2 ShadowImage Replication Operations	3. ShadowImage Software and RAID Manager (CCI) for Sun Solaris 4. ShadowImage Software and RAID Manager (CCI) for Microsoft Windows
6. Hitachi Copy-on-Write Snapshot Software Introduction 7. Hitachi Copy-on-Write Snapshot Software with Command Control Interface 8. Hitachi Copy-on-Write Snapshot Software Storage Navigator Modular 2 Operations	5. Copy-on-Write Operations

7

Content	Learning Activities
9. Hitachi TrueCopy Remote Replication Software Overview 10. Hitachi TrueCopy Remote Replication Configurations Using CCI	6. TrueCopy Software Operations 7. TrueCopy with Brocade Switches (Optional) 8. TrueCopy and RAID Manager (CCI) for Sun Solaris
11. Hitachi TrueCopy Remote Replication with Hitachi Storage Navigator Modular 2 12. Transport Infrastructure and WAN Connectivity 13. Hitachi TrueCopy Extended Distance Software	9. TrueCopy and RAID Manager (CCI) for Windows 10. TrueCopy and Swap Takeover Function (Optional) 11. TrueCopy Extended Distance

8

Learning Paths

- Are for customers, partners and employees
  - Available on HDS.com, Partner Xchange and HDSnet
- Enable career advancement
- Are a path to professional Certification
- Are available with the instructor
  - Details or copies

9

HDS.com: http://www.hds.com/services/education/

Partner Xchange: https://extranet.hds.com/http://aim.hds.com/portal/dt/

HDSnet: http://hdsnet.hds.com/hds_academy/

Please contact your local training administrator if you have any questions regarding Learning Paths or visit your applicable website.

1. Introduction to Hitachi ShadowImage® Replication Software

Module Objectives

- Upon completion of this module, the learner will be able to:
  - Describe the key features and operations of Hitachi ShadowImage® Replication software
  - List the rules associated with the operations of ShadowImage software
  - Describe the key competitive advantages of ShadowImage software
  - Identify ShadowImage commands and the operations of each command

2

Overview

- In-system hardware-based copy facility that provides:
  - Full copies of logical units within Hitachi family storage systems
  - Replicates information with a minimum impact to service levels

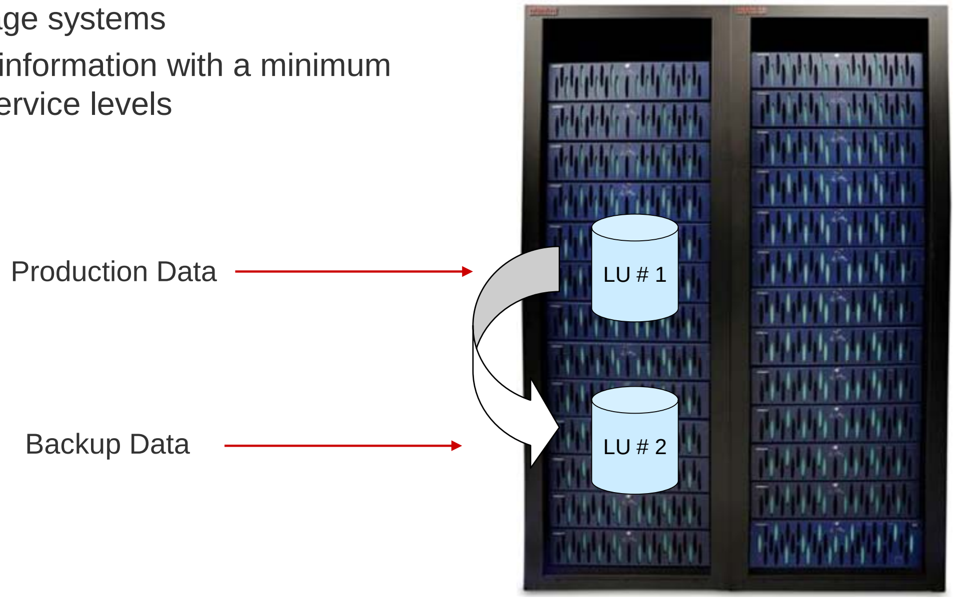

3

ShadowImage software is the in-system copy facility for Adaptable Modular Storage 2000 family. It enables server-free backups, which allows customers to exceed service level agreements (SLAs). It fulfills two primary functions:

- Copy open-systems data
- Backup data to a second LUN

ShadowImage software allows information to be split away and used for system backups, testing and data mining applications while the customer's business continues to run. It uses either graphical or command line interfaces to create a copy and then control data replication and fast resynchronization of logical volumes within the system.

- What does it do?
  - Enables online backup
  - Allows application development testing or tape archival from the ShadowImage mirror copy
- How does it work?
  - Creates copy (S-VOL) of any active application volume or LU (P-VOL)
  - Allows the new copy to be used by another application or system

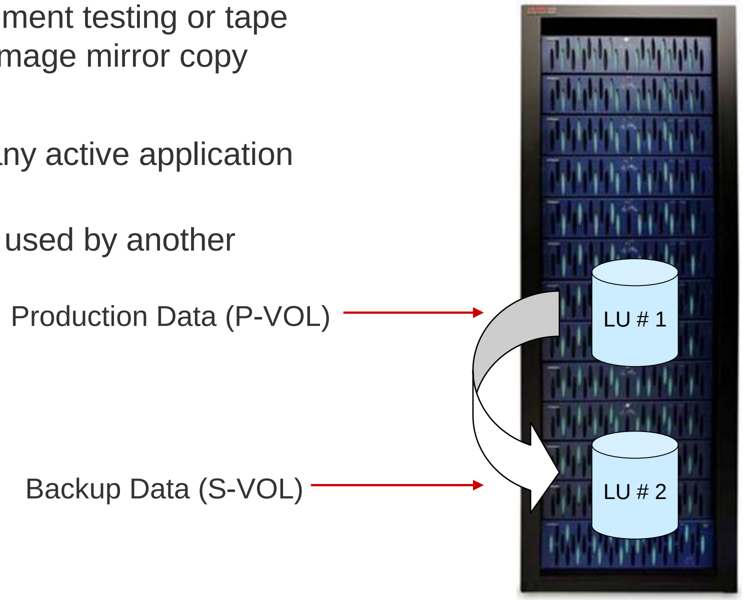

4

ShadowImage software allows you to create a single, local copy of any active application volume while benefiting from full RAID protection. This mirrored copy can be used by another application or system for a variety of purposes, including data mining, full volume batch cycle testing and backups.

It can provide one duplicate, secondary volume (S-VOL) per primary volume (P-VOL) within the same system to maintain redundancy of the primary volume. It allows you to split and combine duplex volumes and provides you the contents of static LUs without stopping the access.

ShadowImage operations are non-disruptive and allow the primary (main) volume of each volume pair to remain online for all hosts for both read and write I/O operations. ShadowImage operations continue unattended to provide asynchronous internal data backup.

Applications

- Backup and recovery
- Data warehousing and data mining applications
- Application development
- Run benchmarks and reports

5

ShadowImage software is replication/backup and restore software that delivers the copy flexibility customers need for meeting today's unpredictable business challenges. With ShadowImage software, customers can:

- Execute logical backups at faster speeds and with less effort than previously possible
- Easily configure backups to execute across a storage area network (SAN)
- Manage backups from a central location
- Increase the speed of applications
- Expedite application testing and development
- Keep a copy of data for backup or testing
- Ensure data availability

Benefits

- Faster time to market
- Quicker decision making
- Surpassed service level agreements (SLAs)
- Reduced management complexity and costs
- Increased efficiency

Benefits of ShadowImage software include:

Faster time to market: Rapidly testing and deploying new applications creates a competitive advantage. ShadowImage software enables production information to be replicated and used to develop, install, test and modify applications swiftly, in real time.

Quicker decision making: ShadowImage software let customers replicate production data, and then to populate their data warehousing/data mining applications without disrupting online business. Always at their fingertips, they have the data needed to make better, more profitable decisions.

Surpassed service level agreements (SLAs): Combined with the all architecture of Hitachi Adaptable Modular Storage 2000 systems, ShadowImage software enables customers to exceed today's SLAs with high availability and performance that improves employee productivity and customer loyalty.

Reduced management complexity and costs: ShadowImage software uses resynchronization to copy only the data that has changed since the last data update. This reduces management complexity and costs.

Increased efficiency: ShadowImage software allows a business to remain online during basic data center activities, eliminating the need for expensive, around-the-clock resources to perform these tasks. Consequently, customers can increase the efficiency of their organizations by refocusing these resources onto other revenue-driven and time-critical projects.

General Replication Terminology

- Command Device (CMD-DEV)
  - Dedicated logical volume that is used only by management software such as CCI Raid Manager, to interface with the arrays
    - Not used by ordinary applications
    - Can be shared between several hosts

- Differential Management Logical Unit (DM-LU)
  - This volume is used to manage differential data in an array. For ShadowImage, the DM-LU is an exclusive volume used for storing data when the storage system is powered down.

7

Specifications

- General Replication Program Product Specifications

Item	Adaptable Modular Storage 2000 Family
Maximum number of command devices	128 with a minimum of 40MB each
Maximum number of DM-LUs	2 with a minimum of 10GB each

8

Item	Adaptable Modular Storage 2000 Family
User interface for replication operations and configuration	Storage Navigator Modular 2 (GUI, CLI) CCI (RAID Manager)
Controller configuration	Dual controllers
Maximum number of pairs	Model 2100: 1023 Models 2300 & 2500: 2047 * Volume Migration and ShadowImage share number of pairs.
Maximum number of copy jobs running concurrently in the array	Models 2100 & 2300: 4/CTL Model 2500: 8/CTL (4/Core) * Volume Migration and ShadowImage share copy jobs

9

* Volume Migration stands for Hitachi Volume Migration software.

Item	Adaptable Modular Storage 2000 Family
P-VOL: S-VOL	1:8
RAID level	RAID-6 , RAID-5, RAID-1, RAID-1+0, RAID-0 * P-VOL cannot be RAID-0.
RAID level of P-VOL to S-VOL	Any to any
RAID Groups that P-VOL and S-VOL resides	It is recommended that P-VOL and S-VOL reside in different RAID Groups.
Drive type for a P-VOL and S-VOL	Any
Quick Resync/Split	Supported
Split by CTG	Supported

10

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

Quick mode enabled Advantages

Access from the host to the S-VOL becomes possible without waiting for data copy completion while creating or updating the copying volume.

Quick mode enabled Disadvantages

Since the P-VOL data is used when accessing the SVOL, the load for the S-VOL affects the I/O performance of the P-VOL. The I/O performance of the S-VOL is affected by the load of the P-VOL and is more deteriorated than the case where Quick Mode is not used.

Quick mode disabled Advantages

Since access from the host to the S-VOL is independent of the P-VOL, the I/O performance is less affected.

Quick mode disabled Disadvantages

Access to the S-VOL cannot start unless the data copy is completed while creating or updating the copying volume. Therefore it is required to access the S-VOL after waiting for data copy completion.

Item	Adaptable Modular Storage 2000 Family
Pair operation and format • Pair operation during format • Format during pair operation	Not supported
Combination with another ShadowImage pair • P-VOL of a ShadowImage pair is S-VOL of another ShadowImage pair	Not supported
Combination with Copy-on-Write Snapshot pair • P-VOL of a ShadowImage pair is P-VOL of a Snapshot pair • S-VOL of a ShadowImage pair is P-VOL of a Snapshot pair • P-VOL of a ShadowImage pair is V-VOL of a Snapshot pair	Not supported
Combination with TrueCopy pair (See TrueCopy section)	Supported (including Swap-resync)
Combination with TrueCopy Extended Distance pair (See TrueCopy Extended section)	Not supported

11

Hitachi TrueCopy® Extended Distance software:

Management Resources

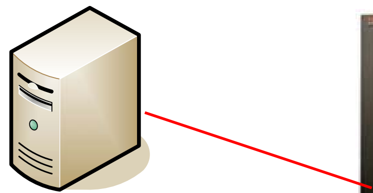

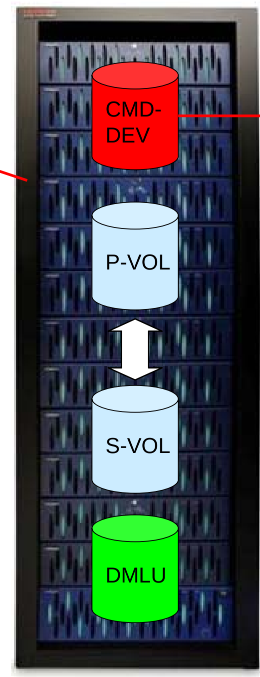

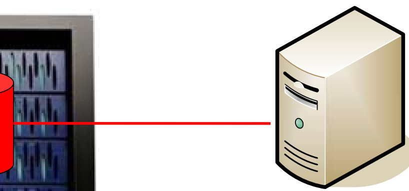

Management Server with SNM2

- Create pairs
- Split pairs
- Resynchronize pairs
- Delete pairs
- No CMD-DEV necessary

Management Server with CCI

- Create pairs
- Split pairs
- Resynchronize pairs
- Delete pairs
- CMD-DEV necessary

12

The ShadowImage system components include:

- ShadowImage volume pairs (P-VOLs and S-VOLs) in the Adaptable Modular Storage
- Storage Navigator Software Modular 2 and/or CCI Raid Manager software

Rules

- Two LUNs that compose a ShadowImage paired volume must have same LUN size in blocks.
- ShadowImage software supports LUSE volumes if they are exact the same size.
- Hitachi Cache Residency Manager feature volumes are not supported.

13

Internal ShadowImage Operation

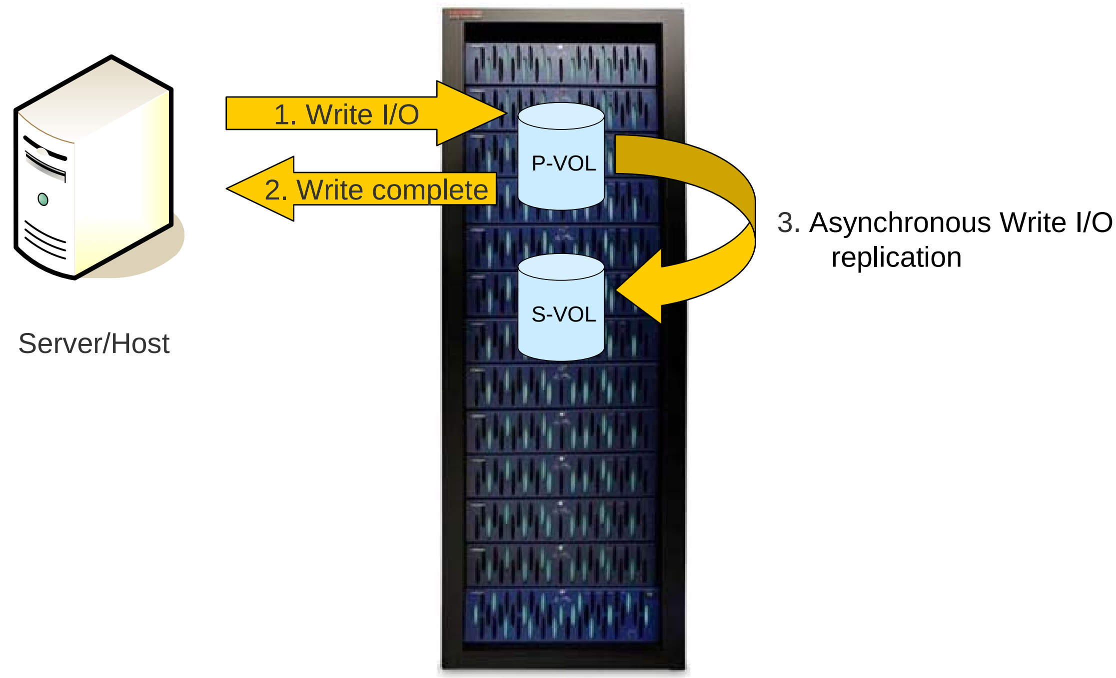

Hitachi Adaptable Modular Storage (AMS) 2000 Family storage array

14

Operations

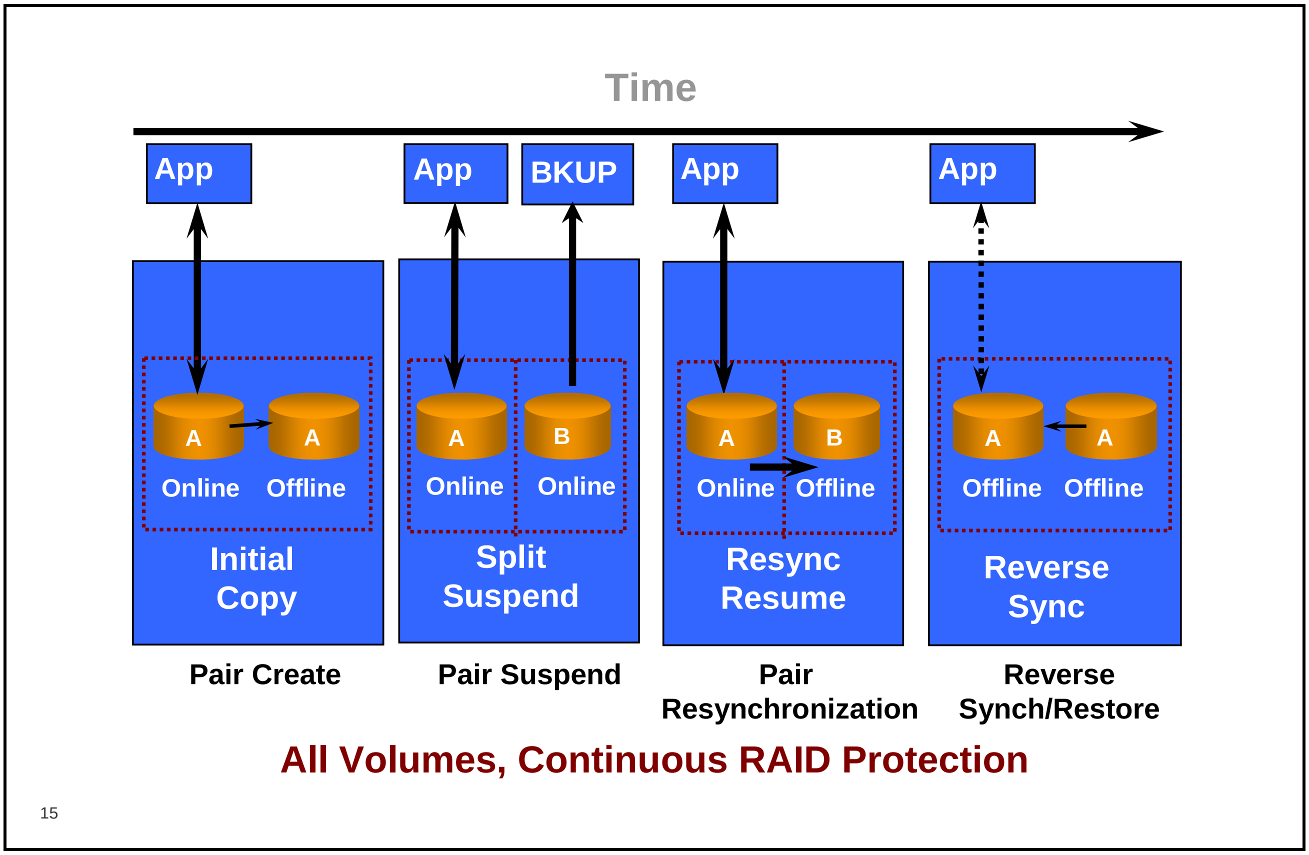

ShadowImage operations include:

- PAIR-CREATE
- PAIR-SPLIT
- PAIR-RESYNCHRONIZE

paircreate Command

- Establishes a new ShadowImage pair

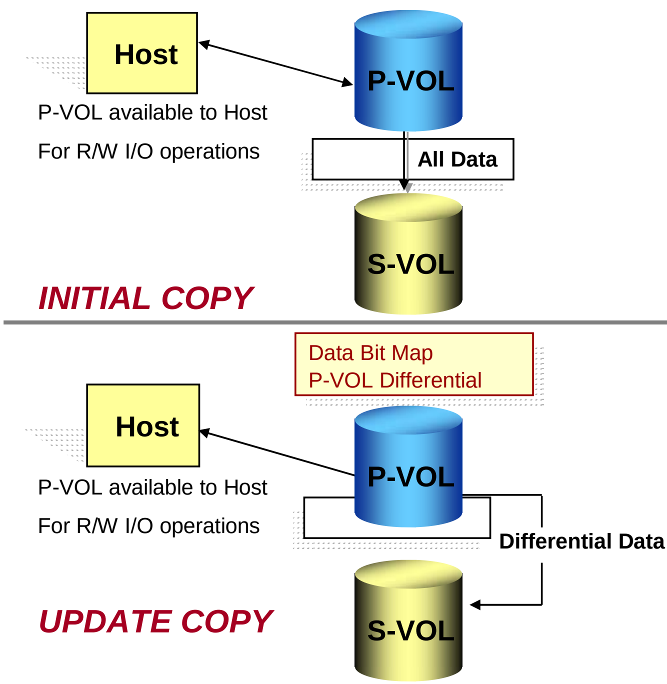

- Updates S-VOL after initial copy
- Write I/O to P-VOL during initial copy — duplicated at S-VOL by update copy after initial copy

16

The ShadowImage paircreate operation establishes the newly specified ShadowImage pair.

- The volumes, which will become the P-VOL and S-VOL, must both be in the SMPL (simplex) state before becoming a ShadowImage pair.
- ShadowImage initial copy operation copies all data from the P-VOL to the associated S-VOL.
- P-VOL remains available to all hosts for read and write I/Os throughout the initial copy operation.
- Write operations performed on the P-VOL during the initial copy operations will always be duplicated to the S-VOL after the initial copy is complete.
- Status of the pair is COPY(PD) (PD = pending duplex) while the initial copy operation is in progress. The pair status changes to PAIR when the initial copy is complete.

You can select the pace for the initial copy operation when creating pairs. The following pace options are available:

- Slower
- Medium
- Faster

The slower pace minimizes the impact of ShadowImage operations on system I/O performance, while the faster pace completes the initial copy operation as quickly as possible. The best timing is based on the amount of write activity on the P-VOL and the amount of time elapsed between update copies.

pairsplit Command

- Splits P-VOL and S-VOL pairs in the PAIR state

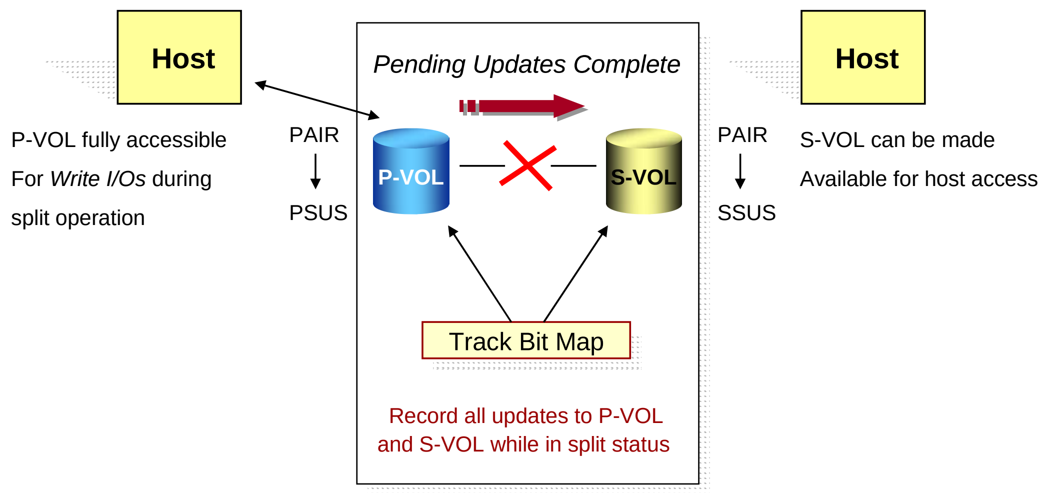

17

The ShadowImage pairsplit operation performs all pending S-VOL updates (those issued prior to the split command and recorded in the P-VOL track map) to make the S-VOL identical to the state of the P-VOL when the suspend command was issued and then provides full read/write access to the split S-VOL.

You can split existing pairs as needed and you can use the paircreate operation to create and split pairs in one-step. This feature provides point-in-time backup of your data, and facilitates real data testing by making the ShadowImage copies (S-VOLs) available for host access.

When the split operation is complete, the pair status changes to PSUS and you have full read/write access to the split S-VOL. While the pair is split, the Thunder 9500 V Series system establishes a track map for the split P-VOL and S-VOL, and records all updates to both volumes. The P-VOL remains fully accessible during the pairsplit operation. Pairsplit operations cannot be performed on suspended (PSUE) pairs.

- pairsplit Illustration

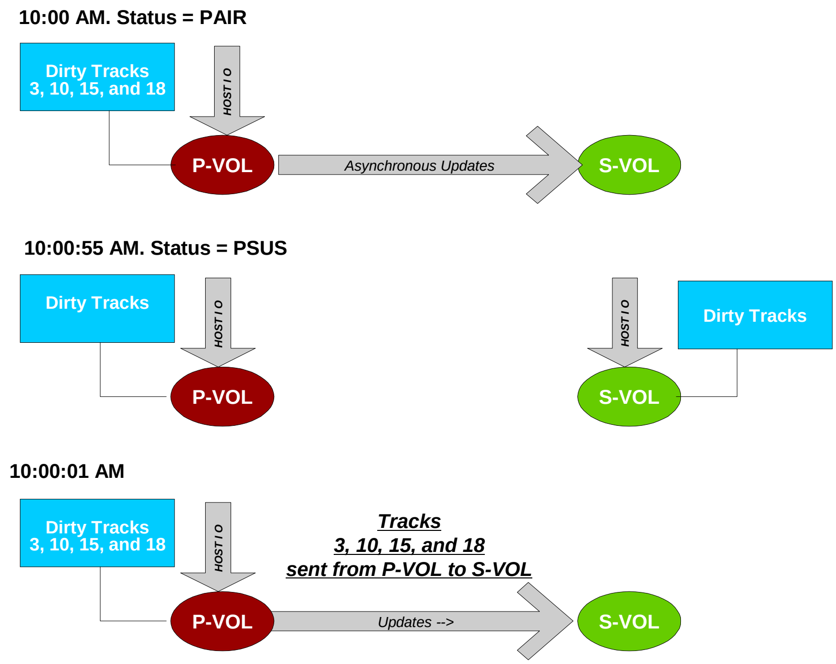

18

1. The P-VOL and S-VOL are in PAIR status as of 10:00 AM. Tracks 3, 10, 15 and 18 are marked as dirty because of host I/O.
2. The status of the P-VOL and S-VOL is changed to PSUS. During this state there are track bitmaps attached to both the P-VOL and the S-VOL. These bitmaps keep track of changes on both the P-VOL and the S-VOL.
3. Tracks 3, 10, 15 and 18 are sent across to the S-VOL from the P-VOL

pairresync Command

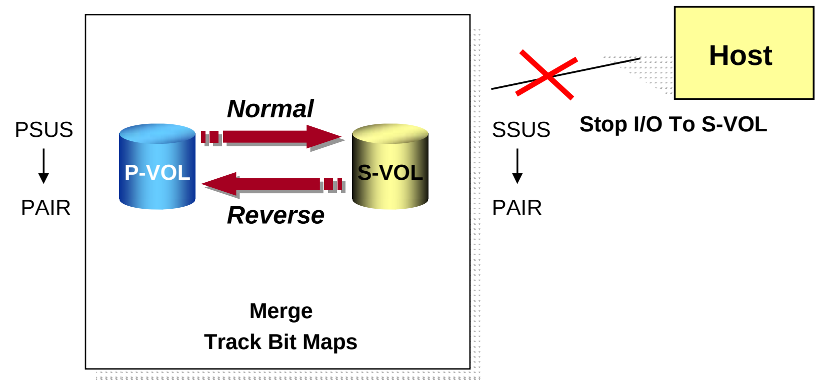

Sync S-VOL with P-VOL	Sync P-VOL with S-VOL
NORMAL	REVERSE

Sync S-VOL with P-VOL

NORMAL

- Copy direction = P-VOL to S-VOL
- S-VOL track map merged into P-VOL track map — copies all “flagged” tracks from P-VOL to S-VOL
- P-VOL remains available to hosts for R/W I/O operations

Sync P-VOL with S-VOL

REVERSE

- Copy direction = S-VOL to P-VOL
- P-VOL track map merged into S-VOL track map — copies “flagged” tracks from S-VOL to P-VOL
- P-VOL is unavailable for R/W I/Os

19

The ShadowImage **pairresync** operation resynchronizes the suspended pairs (PSUS) or the suspended on error pairs (PSUE). When the pairresync operation starts the pair status changes to COPY(RS) or COPY(RS-R). The pair status changes to PAIR when the pairresync operation completes.

ShadowImage software allows you to perform two types of pairresync operations:

Normal. The normal **pairresync** operation resynchronizes the S-VOL with the P-VOL. The copy direction for a normal pairresync operation is P-VOL to S-VOL. The pair status during a normal resync operation is COPY(RS). The S-VOL becomes inaccessible to all hosts for write operations, and the P-VOL is accessible to all hosts for both read and write operations during a normal pairresync. The normal pairresync operation can be executed for pairs with the status PSUS and PSUE.

Reverse. The reverse pairresync (**pairresync –restore**) operation synchronizes the P-VOL with the S-VOL. The copy direction for a reverse pairresync operation is S-VOL to P-VOL. The pair status during a reverse resync operation is COPY(RS-R), and the S-VOL becomes inaccessible to all hosts for write operations during a reverse pairresync operation. The P-VOL is inaccessible for both read and write operations, and the write operations on P-VOL will always be reflected to S-VOL

When a pairresync operation is performed on a suspended pair (status = PSUS), the storage system merges the S-VOL differential track map into the P-VOL differential track map and then copies all flagged data from the P-VOL to the S-VOL. When a reverse pairresync operation is performed on a suspended pair, the storage system merges the P-VOL differential track map into the S-VOL differential track map and then copies all flagged tracks from the S-VOL to the P-VOL. This ensures that the P-VOL and S-VOL are properly resynchronized in the desired direction.

pairresync Command

- Normal resync Illustration

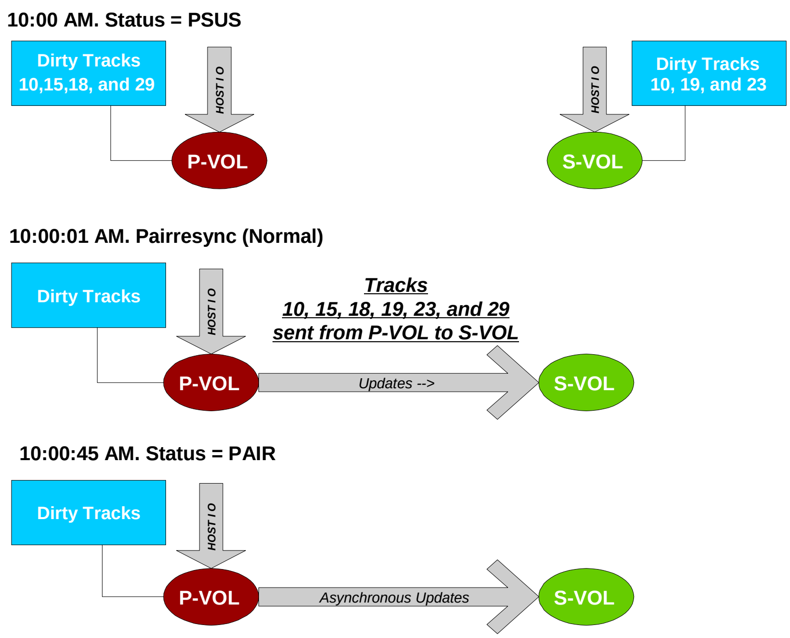

20

1. The status of the P-VOL and the S-VOL is PSUS as of 10:00 AM. Tracks 10, 15, 18 and 19 are marked as dirty on the track bitmap for the P-VOL. Tracks 10, 19 and 23 are marked as dirty on the track bitmap for the S-VOL.
2. At 10:00 AM, a pairresync (normal) command is issued. The track bitmaps for the P-VOL and S-VOL are merged. The resulting track bitmap has tracks 10, 15, 18, 19, 23 and 29 marked as dirty. These tracks are sent from the P-VOL to the S-VOL as part of an update copy operation.
3. Once the update copy operation in step 2 is complete, the P-VOL and S-VOL are declared as a PAIR.

- Reverse resync Illustration

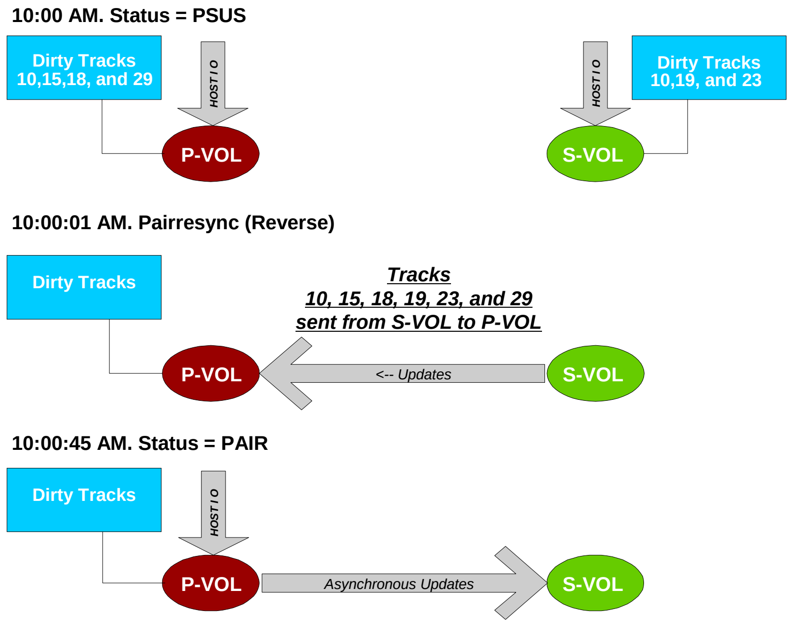

21

- The status of the P-VOL and the S-VOL is PSUS as of 10:00 AM. Tracks 10, 15, 18 and 19 are marked as dirty on the track bitmap for the P-VOL. Tracks 10, 19 and 23 are marked as dirty on the track bitmap for the S-VOL.
- At 10:00 AM, a pairresync (normal) command is issued. The track bitmaps for the P-VOL and S-VOL are merged. The resulting track bitmap has tracks 10, 15, 18, 19, 23 and 29 marked as dirty. These tracks are sent from the S-VOL to the P-VOL as part of an update copy operation.
- Once the update copy operation in step 2 is complete, the P-VOL and S-VOL are declared as a PAIR.

pairsplit -E Command

- Suspends the ShadowImage copy operations to the S-VOL of the pair

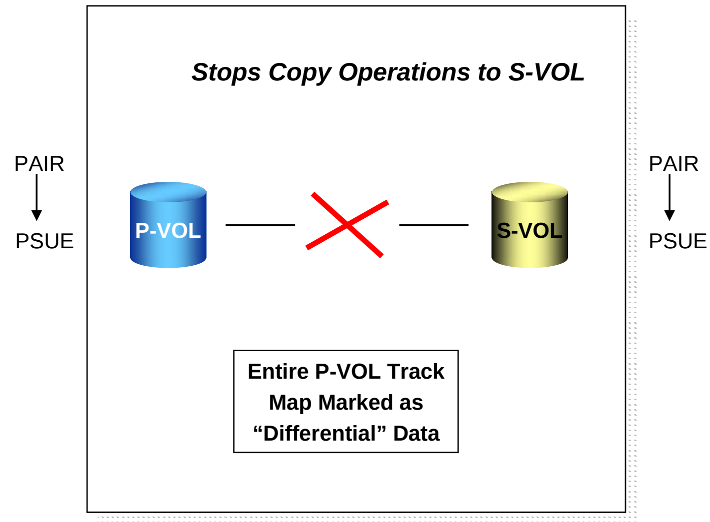

22

When a ShadowImage pair is suspended (status PSUE), the storage system stops performing ShadowImage copy operations to the S-VOL, continues accepting write I/O operations to the P-VOL and marks the entire P-VOL track map as difference data.

The storage system will automatically suspend a ShadowImage pair when it cannot keep the pair mirrored for any reason. When the storage system suspends a pair, a file is output to the system log or event log to notify the host.

The storage system will automatically suspend a pair under the following conditions:

- When the ShadowImage volume pair has been suspended or deleted
- When the storage system detects an error condition related to an initial copy operation.

Resynchronize Suspended Pair

- pairresync on a suspended pair (PSUE)

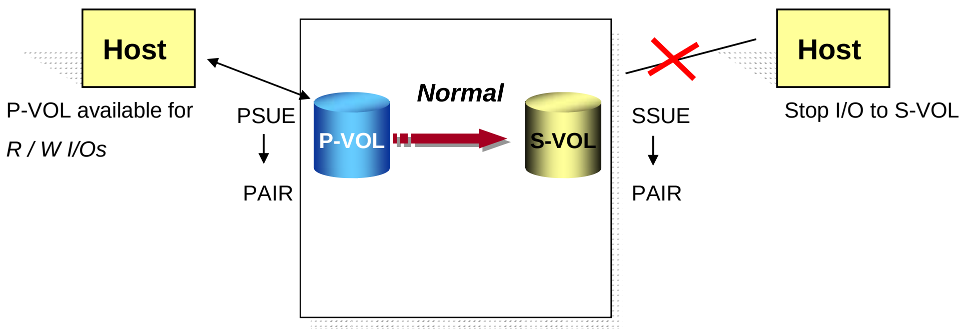

NORMAL
Copy Direction = P-VOL to S-VOL

Syncs S-VOL with P-VOL

- Resynchronizes entire volume

Equivalent to and takes as long as *Initial copy* operation

23

When a **pairresync** operation is performed on a suspended pair, the entire P-VOL is copied to the S-VOL. When a **pairresync –restore** operation is performed, the entire S-VOL is copied to the P-VOL. While the pairresync operation for a split ShadowImage pair can be very fast, the pairresync operation for a suspended on error pair will take as long as the initial copy operation.

pairsplit -S Command

- Stops copy operations and changes status to simplex (SMPL)

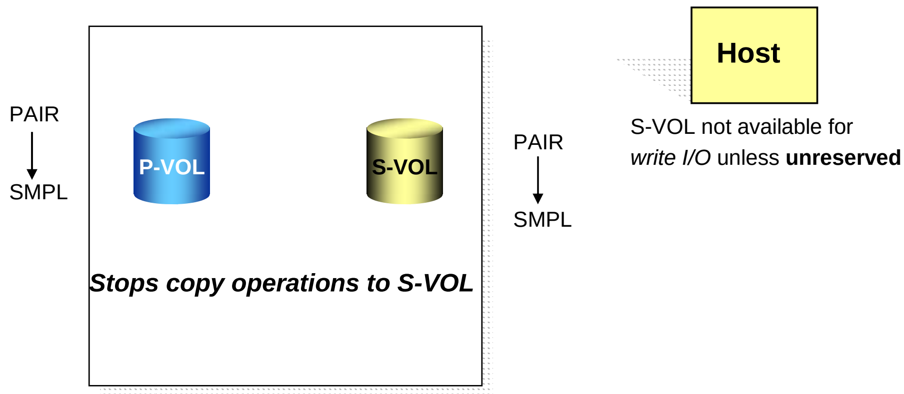

24

The ShadowImage **pairsplit –S** (delete pair) operation stops the ShadowImage copy operations to the S-VOL of the pair and changes the pair status of both volumes to SMPL.

Volume Grouping

- Perform operations for pairs or a group of pairs

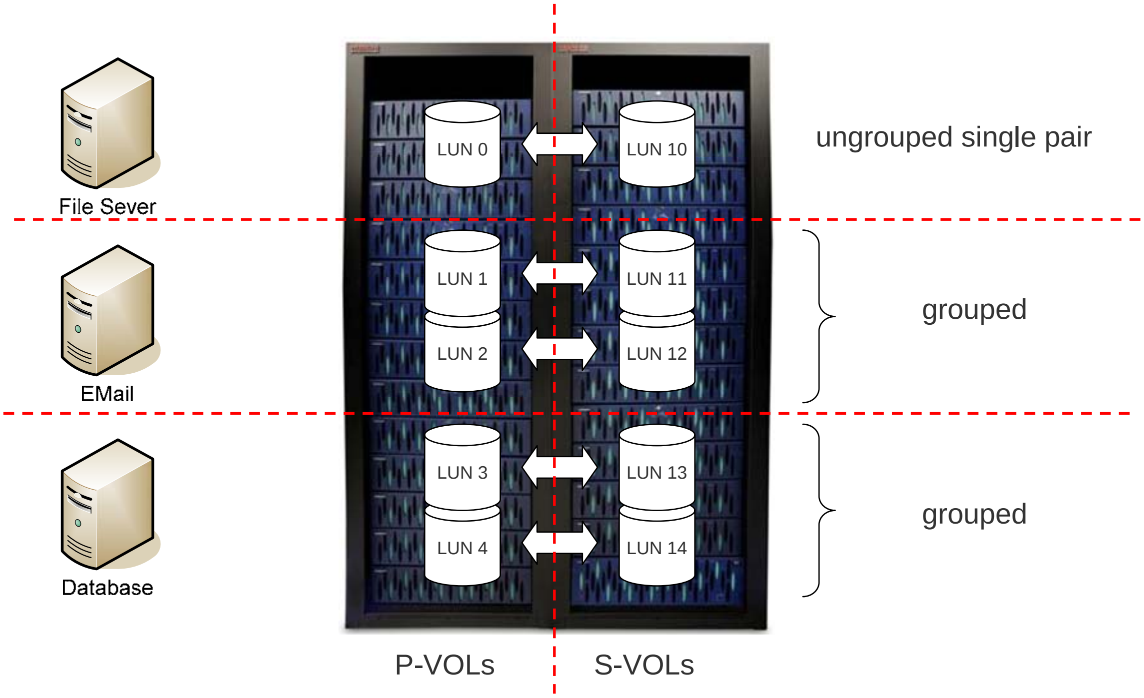

25

You can define or set up ShadowImage pairs in groups which enables you to issue commands or do operations for single pairs or a group of pairs.

Volume Status and Availability Summary

Pair Status	Description	S-VOL Access
SMPL	The volume is not assigned to a ShadowImage pair. The storage system accepts read and write I/Os for all SMPL volumes.	N/A (there is no S-VOL yet.)
COPY(PD)	The initial copy operation is in progress. The storage system continues to accept read and write operations for the P-VOL, but does not accept write operations for the S-VOL.	Read only
PAIR	The initial copy operation is complete, and the storage system starts asynchronously copying the write operation taken to the P-VOL data onto the S-VOL. The P-VOL and S-VOL of a duplex pair (PAIR status) is identical. The storage system rejects all write I/Os for S-VOLs with the status PAIR.	Read only
PSUS	The storage system starts accepting write I/Os for PSUS S-VOLs. The storage system keeps track of all updates to the suspended P-VOL and S-VOL, so that the pair can be resynchronized quickly.	Read and Write. The S-VOL can be mounted.
COPY(RS)	The storage system does not accept write I/Os for COPY(RS) S-VOLs. When a suspended pair is resynchronized in normal mode, the storage system copies only the P-VOL differential data to the S-VOL. When a PSUE (suspended under error) pair is resynchronized, the storage system copies the entire P-VOL to the S-VOL.	Read only.

26

Pair Status	Description	S-VOL Access
COPY(RS-R)	The storage system does not accept write I/Os for COPY(RS-R) S-VOLs. When a split pair is resynchronized in reverse mode, the storage system copies only the S-VOL differential data to the P-VOL.	Read only.
PSUE	The storage system continues accepting read and write I/Os for a PSUE (suspended under error) P-VOL (however, if the status transits from COPT(RS-R), all access to P-VOL is disabled). The storage system marks the entire P-VOL track map as difference data, so that the entire P-VOL is copied to the S-VOL when the PSUE pair is resumed. Use the pairresync command to resume a PSUE pair.	Read only. (If the status transits from COPY(RS-R), all access to P-VOL is disabled.)
(S-VOL Switch)	This is a state in which a double failure of drives occurred in a P-VOL and P-VOL was switched to an S-VOL internally. This state is displayed as PSUE with CCI.	Read/Write is not available

27

Pair Status Transitions

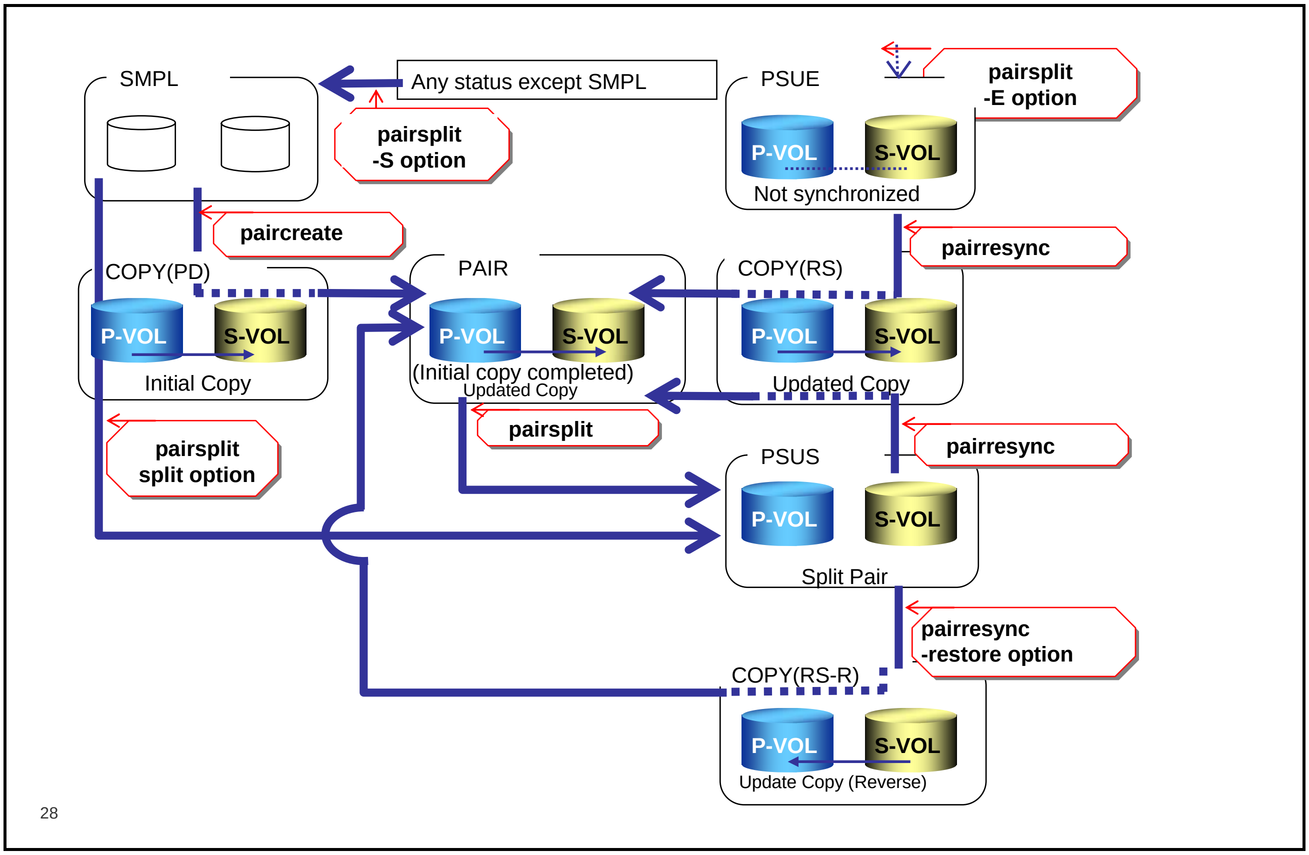

This illustration shows the ShadowImage pair status transitions and the relationship between pair status and ShadowImage operations.

Starting in the upper left of the illustration, if a volume is not assigned to a ShadowImage pair, its status is SMPL.

When you create a ShadowImage pair, the status of the P-VOL and S-VOL changes to COPY(PD). When the initial copy operation is complete, the pair status becomes PAIR. If the Adaptable Modular Storage and Workgroup Modular Storage cannot maintain PAIR status for any reason, or if you suspend on error the pair (**pairsplit –E**), the pair status changes to PSUE. When you suspend a pair (**pairsplit**), the pair status changes to COPY(SP). When the **pairsplit** operation is complete, the pair status changes to PSUS to enable you to access the suspended S-VOL.

When you start a **pairresync** operation, the pair status changes to COPY(RS). When you specify reverse mode for a **pairresync** operation (**pairresync –restore**), the pair status changes to COPY(RS-R) (data is copied in the reverse direction from the S-VOL to the P-VOL). When the **pairresync** operation is complete, the pair status changes to PAIR. When you split/release a pair (**pairsplit -S**), the pair status changes to SMPL.

Module Review

1. Describe the difference between initial copy and update operations for ShadowImage software.
2. What rules must you adhere to when performing ShadowImage operations?
3. Describe the various ShadowImage commands.

29

2. Hitachi ShadowImage® Replication Software Performance

Module Objectives

- Upon completion of this module, the learner should be able to:
  - Describe key ShadowImage performance factors for Adaptable Modular Storage 2000 family

2

Performance Factors

- Location recommendations for P-VOLs and S-VOLs

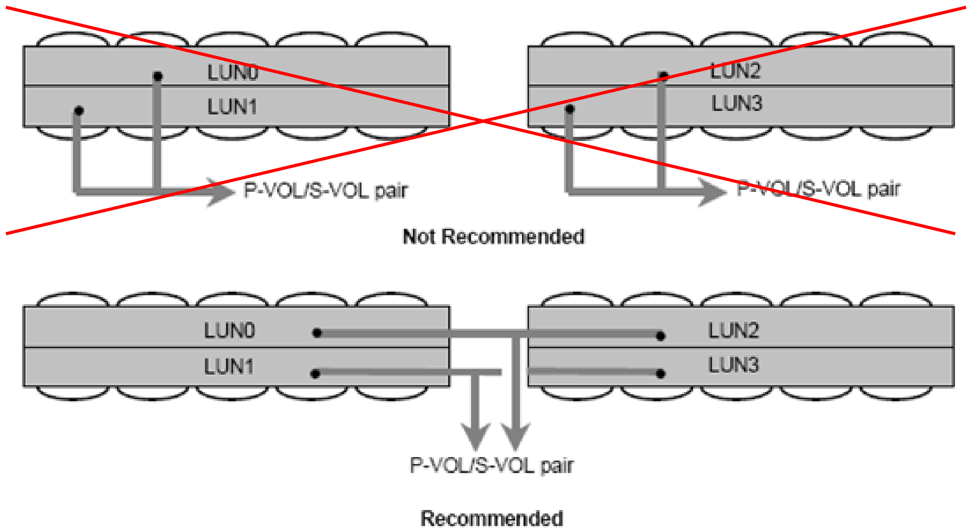

3

Do **not** locate P-VOLs and S-VOLs within the same ECC group of the same RAID Group. If P-VOLs and S-VOLs are located within the same ECC group of the same RAID Group the following issues might arise:

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

- Locating multiple LUs within same drive column

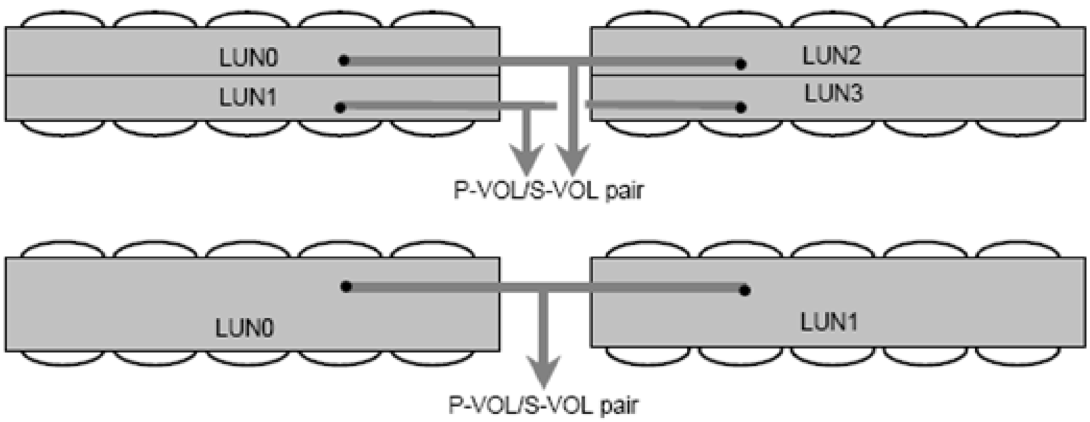

4

It is difficult to estimate performance and to design the system operational settings if multiple LUs with differing states per pair are set within the same parity group. A different state for pairs in parity groups is not recommended. Keep related LUNs in a parity group to keep states the same.

- Pair status differences when setting multiple pairs
  - Keep the status of pairs the same when setting multiple ShadowImage pairs
- P-VOLs and S-VOLs in a RAID configuration
  - Because serial attached SCSI (SAS) drives performance exceeds SATAII drive performance:
- Setting multiple command devices
  - When two command devices are set within the one disk system, assign them to the respective RAID Groups

5

- Pair status differences when setting multiple pairs
  - Even with a single LU per parity group, keep the status of pairs the same (such as SMPL, coupling, and split) when setting multiple ShadowImage pairs
    - If each ShadowImage pair status differs, it is difficult to estimate the performance when designing the system operational settings
- P-VOLs and S-VOLs in a RAID configuration
  - Because SAS drives performance exceeds SATAII drive performance:
    - Locate a P-VOL in a RAID group consisting of SAS drives
    - Conduct a thorough investigation beforehand when locating an S-VOL in a RAID group consisting of the SATAII drives
- Setting multiple command devices
  - When two command devices are set within the one disk subsystem, assign them to the respective RAID Groups
    - Both command devices become unavailable due to difficulties such as a drive failure if they are assigned to the same RAID Group

- P-VOLS and S-VOLs in a RAID configuration

No.	P-VOL	S-VOL	Amount of User Data	Total Amount of ShadowImage	Share of User Data
1	RAID 1+0 (*N* = 1 to 8)	RAID 1+0 (*N* = 1 to 8)	1	4	1/4
2	RAID 1+0 (*N* = 1 to 8)	RAID 5 (*N* = 1 to 8)	1	2 + 1.25 = 3.25	1/3.25
3	RAID 1+0 (*N* = 2 to 8)	RAID 0 (*N* = 2 to 8)	1	2+1=3	1/3
4	RAID 5	RAID 5 (When *N* = 4)	1	1.25+1.25 = 2.5	1/2.5
5	RAID 5	RAID 0 (When *N* = 4)	1	1.25+1 = 2.25	1/2.25

6

The recommended RAID configuration is RAID-5 (4D+1)/RAID-5 (4D+1).

- 4D+1 has outstanding performance in RAID.
- Large number of 4D+1 has been adopted based on past sales records.
- Cost can be reduced due to the large share of user data.

Usability capacity/usable production is the amount of user data.

Total Amount of ShadowImage is the raw capacity (terabytes).

RatioShare of User Data is the ratio of the usable capacity to the raw capacity used up when ShadowImage is used.

- Relationship between copy pace parameters and resync process

No.	Parameters	Resync Pace	Resync Process	Host's I/O Performance
1	1 to 5	Slow *(do not select)*	Single copy unit (stripe size) is processed per copy every 10 ms. The process is performed in longer intervals if the I/O performance from the host is heavy. The time of resync completion cannot be guaranteed because the pace differs depending on the I/O performance from the host.	Read: Executing approximately 90 of host limitation is possible. Write: Executing approximately 90 of host limitation is possible.
2	6 to 10	Normal	The resync process is performed continuously, but the time of resync completion cannot be guaranteed. The pace differs depending on the I/O performance from the host.	Same as above.
3	11 to 15	Prior	The resync process is performed continuously and takes prior. Therefore, the I/O performance from the host is restricted. The time of resync completion is guaranteed.	Random Read: Random Write: Sequential Read: Sequential Write:

7

Set the copy pace for initial copying and restoring based on the following criteria:

1. It is recommended to prioritize the resync operation. Enter a value 15 for <size> to set the copy pace.
2. Select **Normal** only when placing importance on the performance to the P-VOL during a resync operation. This can be specified when there is no update processes during the resync operation. The copy time for the resync operation is not guaranteed. Copy time for resync operation delays depends on the I/O performance.
3. **DO NOT SELECT** "Slow".

Pair Assignment

- Assign two different RAID Groups to each of the primary and secondary volumes.
- Assign a small number of volumes within the same RAID Group.
- For a primary volume, use the SAS drives.
- Assign four or more disks as the data disks to the Raid Group.

8

- Assign two different RAID groups to each of the primary and secondary volumes.
  - When a secondary volume is assigned to a RAID group in which the primary volume has been assigned, the reliability of data is lowered because a failure that occurs in a single drive affects both of the primary and secondary LUs. The performance becomes limited because the load of writing applied on a drive is doubled. Therefore, it is recommended to assign primary and secondary volumes to respective RAID groups
- Assign a small number of volumes within the same RAID group.
  - When volumes are assigned to the same RAID group and used as pair volumes, there may be a case where a pair creation or resynchronization for one of the volumes causes a restriction to be placed on performance of a host I/O, pair creation, resynchronization, etc. for the other volumes because of contention between drives. It is recommended that you assign a small number (one or two) of volumes to be paired to the same RAID group. When creating two or more pairs within the same RAID group, standardize the controllers that control LUs in the same RAID group and pay attention to make the pair creation or resynchronization timely.
- For a primary volume, use the FC drives.

- When a P-VOL is located in a RAID group consisting of the SATA drives, performance of a host I/O, pair creation, and pair resynchronization, etc. is lowered because of the lower performance of the SATA drive. Therefore, it is recommended to assign a primary volume to a RAID group consists of the FC drives.

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

Pair Creation and Resynchronization

- Pair creation
  - Create a pair when the I/O utilization is low
  - Try to avoid the creation of multiple pairs simultaneously
- Resynchronization
  - Resynchronize a pair at a time when the I/O utilization is low
  - Give priority to resynchronization

- Pair creation
  - Create a pair when the I/O load is low.
    Pair creation initiates data copying (initial copying) from a primary volume to a secondary volume and therefore affects the performance of the host I/O. It affects the same parity group to which the paired LUs belong. Therefore, it is best to create a pair when the I/O load is low. Perform an operation such as a batch processing at a different time from when a pair is created.
- Resynchronization
  - Resynchronize a pair at a time when the I/O load is low.
    Resynchronization initiates data copying between the primary and secondary volumes and therefore affects the performance of the host I/O. It affects the same parity group to which the paired LUs belong. Therefore, it is best to resynchronize a pair when the I/O load is low. Perform an operation such as a batch processing at a different time from when a pair is resynchronized.
- Give priority to resynchronization.
  - It is possible to specify a priority to the resynchronization process and the host process. Give priority to the resynchronization process because the time required for it to occur can be estimated. When priority is given to the host process, the influence of resynchronization on the host I/O performance can be reduced; it is difficult to estimate when the resynchronization process is completed because the progress of the resynchronization depends on the load of the host I/O.

ShadowImage and TrueCopy Configurations

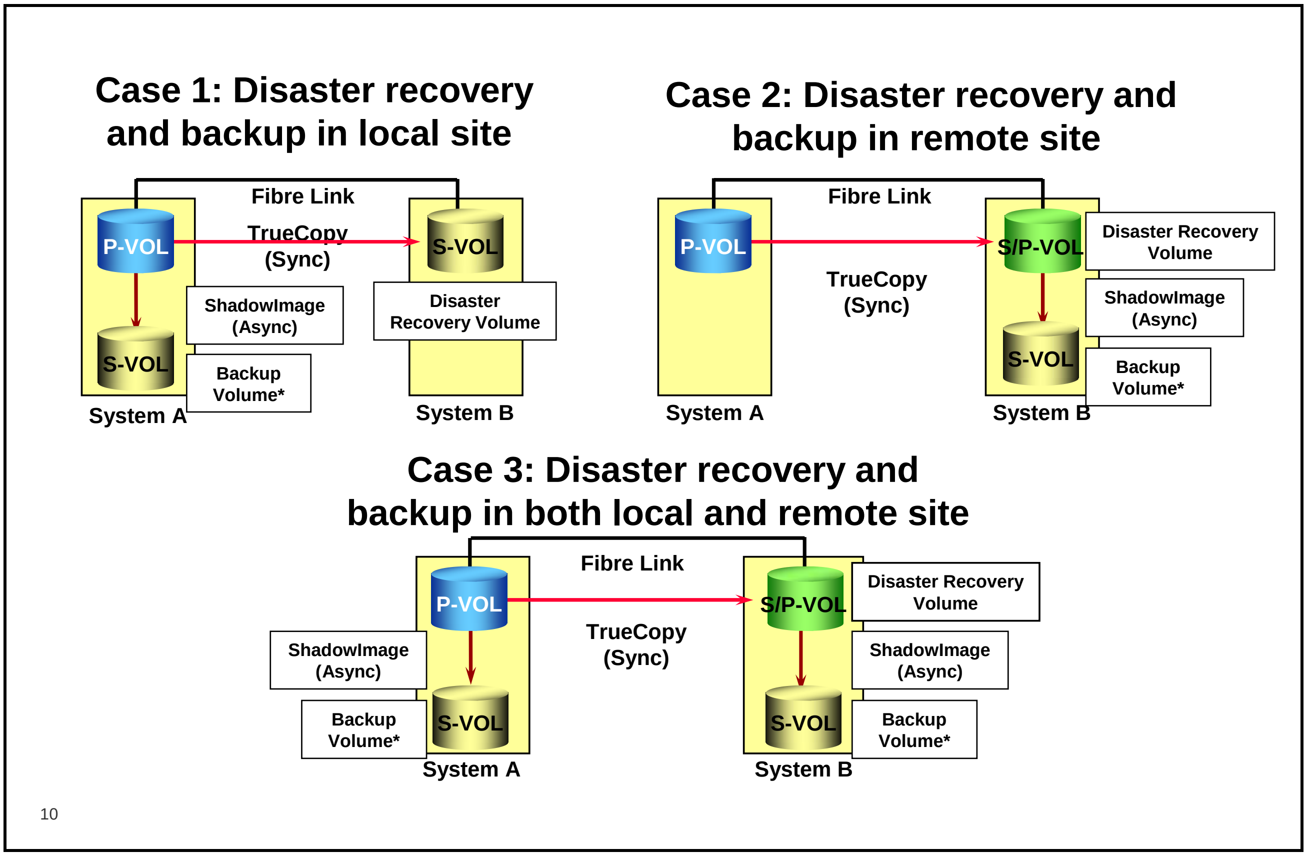

* The backup volume is typically mounted on a backup server for data backup to tape.

TrueCopy software and ShadowImage software can be combined in a variety of ways (cascading) which allow you to create unique solutions based upon customer's requirements.

Cascade Connection of TrueCopy with ShadowImage Software

- Cascade with a P-VOL of ShadowImage software

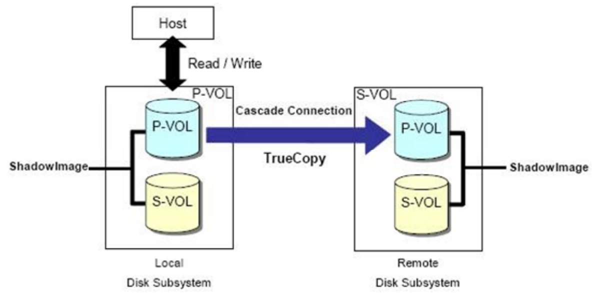

- Cascade with an S-VOL of ShadowImage software

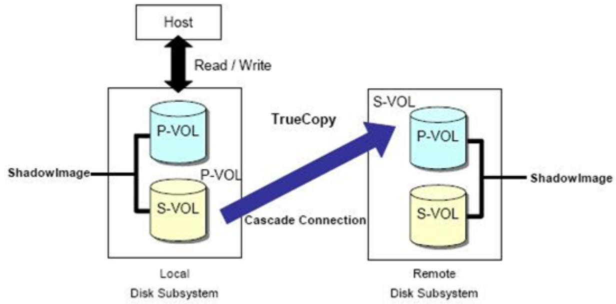

11

- Volumes of TrueCopy can be cascaded with those of ShadowImage software as shown.
- TrueCopy does not always need ShadowImage.
- Snapshot/TrueCopy Extended Distance cannot be cascaded with ShadowImage software.

Simultaneous Cascade Connection of TrueCopy with ShadowImage Software

- Cascade with a P-VOL and an S-VOL of ShadowImage software (P-VOL: S-VOL=1: 1)

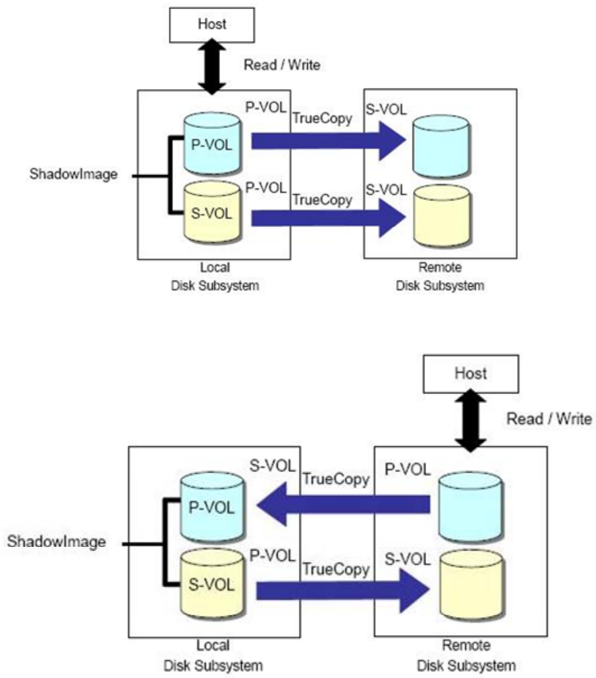

12

Module Review

1. What parameter needs to be setup to give ShadowImage operations priority over front-end I/O activity?
2. What does the pairsplit time depend upon for ShadowImage operations?
3. What does the initial paircreate time depend upon for ShadowImage operations?
4. What does the pair resynchronization time depend upon for ShadowImage operations?

14

3. Hitachi Storage Navigator Modular 2 Basic Operations

Module Objectives

- Upon completion of this module, the learner will be able to:
  - Use the Hitachi Storage Navigator Modular 2 GUI and CLI to prepare an Adaptable Modular Storage 2000 family storage system for replication operations

2

Overview

- Storage Navigator Modular 2 GUI and CLI enables you to:
  - Set up and manage Command Devices
  - Set up and manage Differential Management Logical Units (DM-LUs)

3

The following slides display screenshots of the ShadowImage software installation.

Setting up Command Devices in SNM2

1. Select **Command Devices** and click **Add Command Device.**

4

The command device is a user-selected, dedicated logical volume on the Adaptable Modular Storage 2000 family system that functions as the interface to the CCI software.

The Command Device must be mapped to the server which runs the CCI Raid Manager software in order to operate and manage ShadowImage software.

It is recommended to set up at least two Command Devices in two different Raid groups because of redundancy.

2. Select the **LUN or LUNs** and click **Add.**

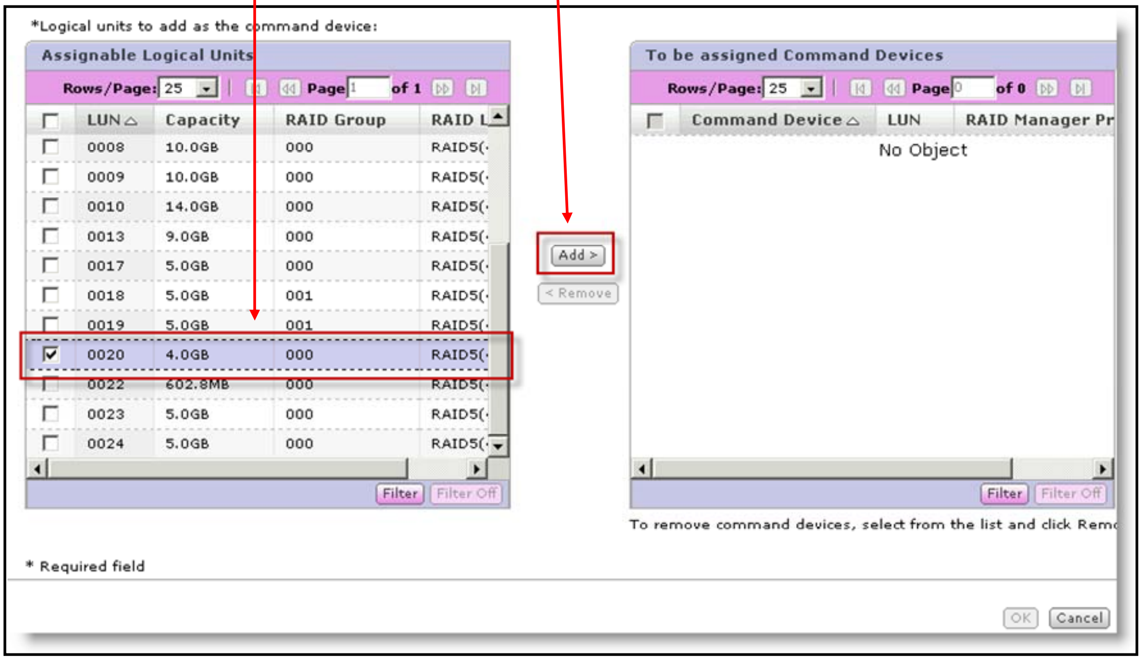

5

3. Click **OK** to finish the operation.

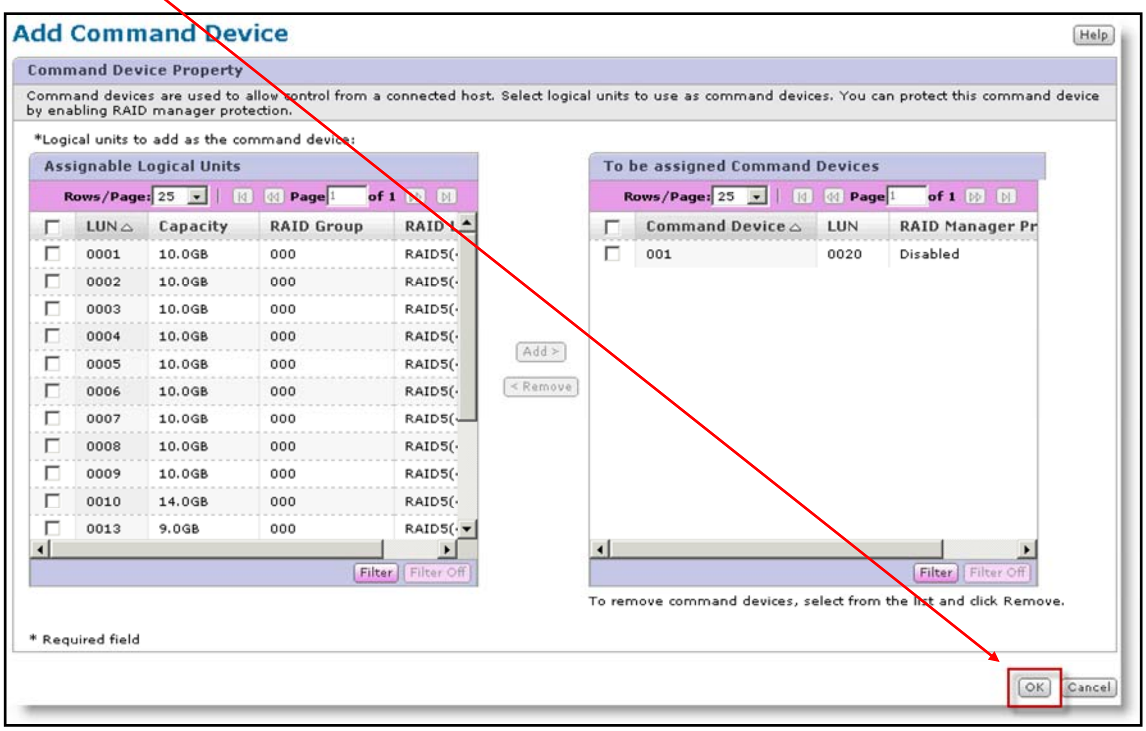

6

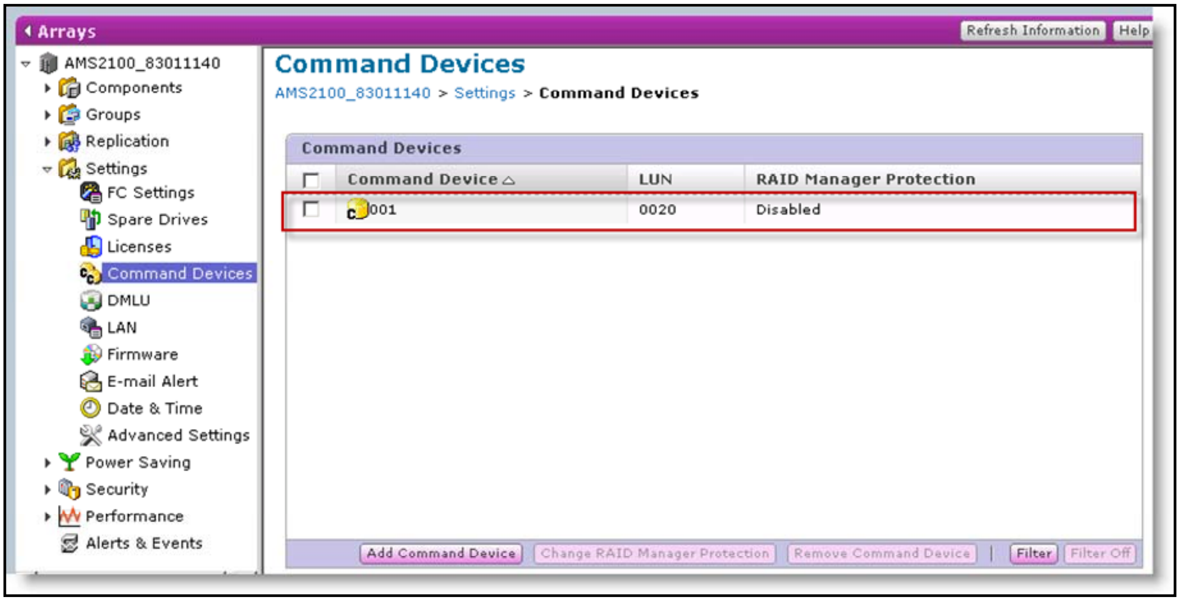

To set up the Command Device with the SNM2 CLI, you have to run the command as follows:
aucmddev -unit ‘unit_name’ -set -dev 1 20

7

You will find more detailed information on the **aucmddev** CLI command in the *Storage Navigator Modular 2 Command Line Interface (CLI) User‘s Guide* on page 4-140

Setting up DM-LUs in SNM2

1. Select **DMLU** and click **Add DMLU.**

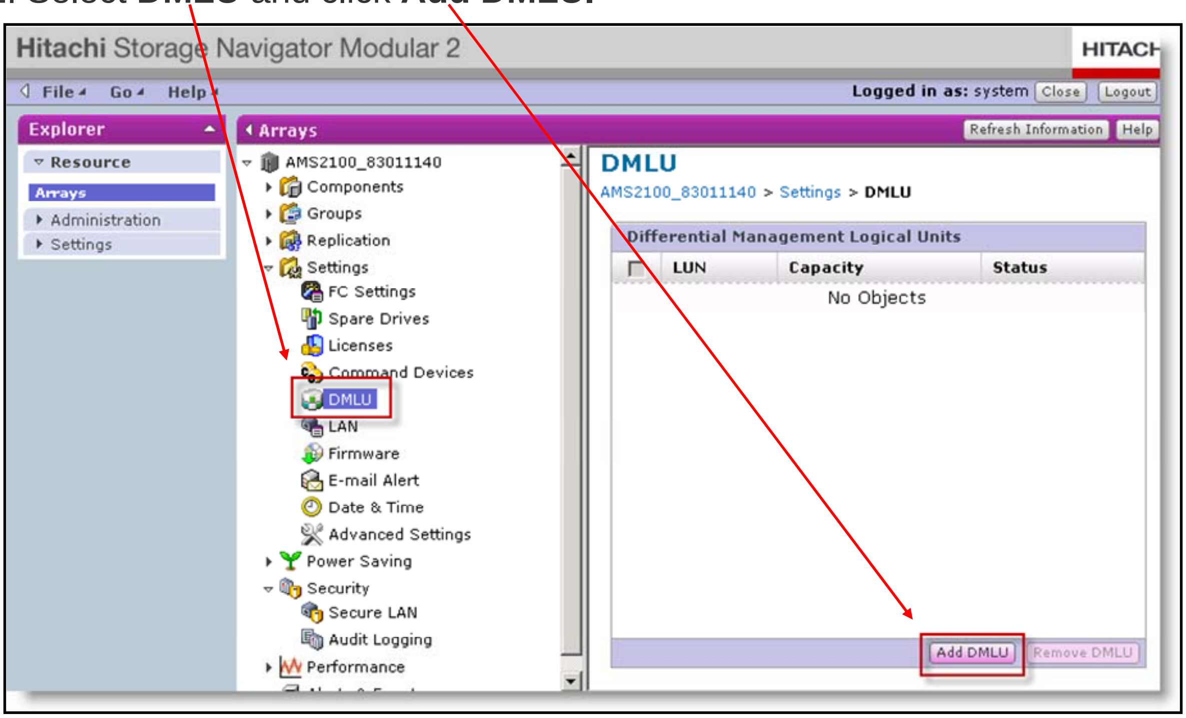

8

The DM-LU is an exclusive volume used for storing ShadowImage, Copy-on-Write Snapshot, TrueCopy and TrueCopy Extended information when the array is powered down.

You must set up at least one DM-LU before you can use any replication product.

It is recommended to set up two DM-LUs in two different Raid groups because of redundancy.

2. Select the **LUN or LUNs** and click **OK.**

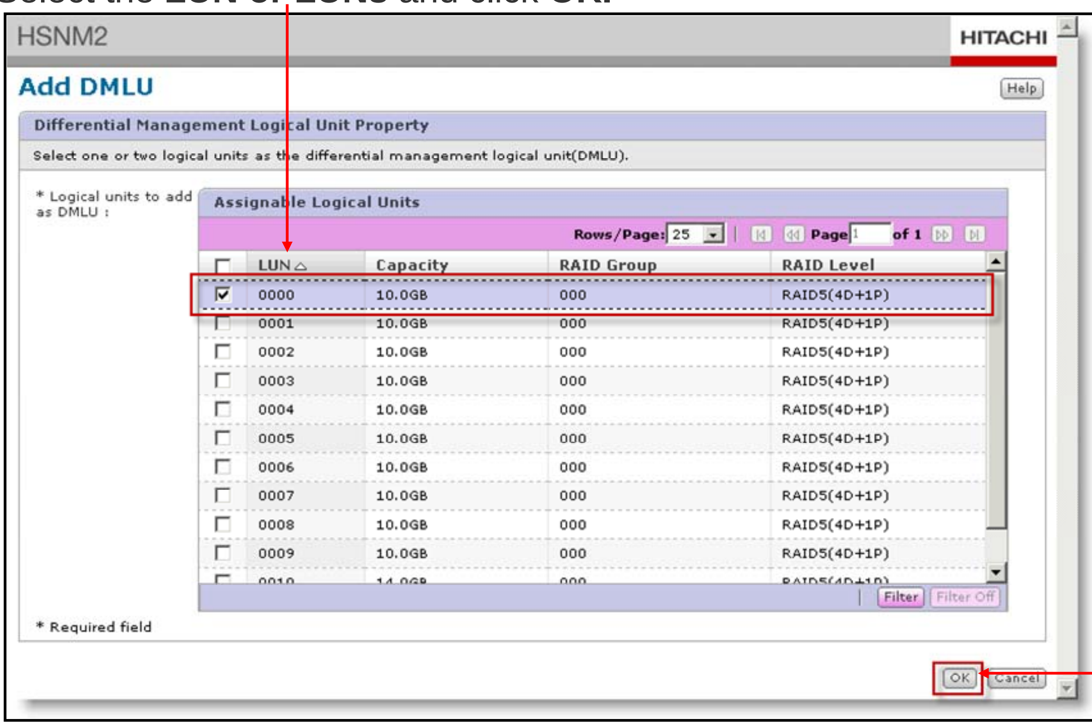

9

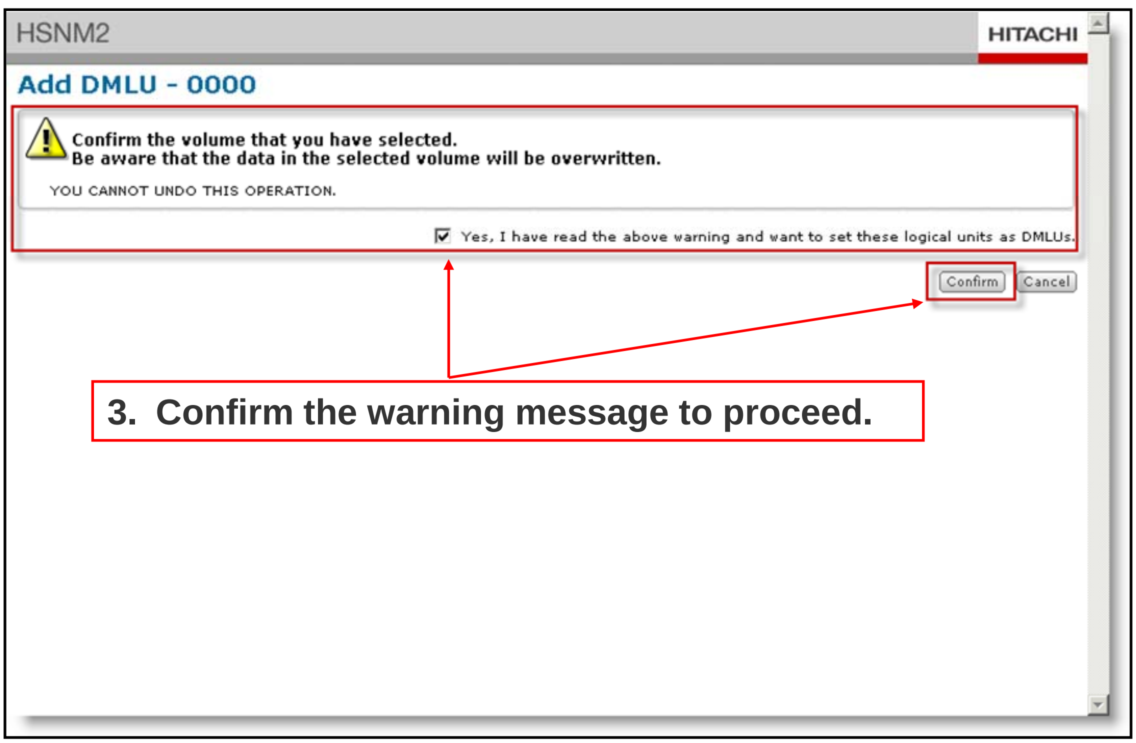

3. Confirm the warning message to proceed.

10

 HDS Confidential: For distribution only to authorized parties. Hitachi Data Systems

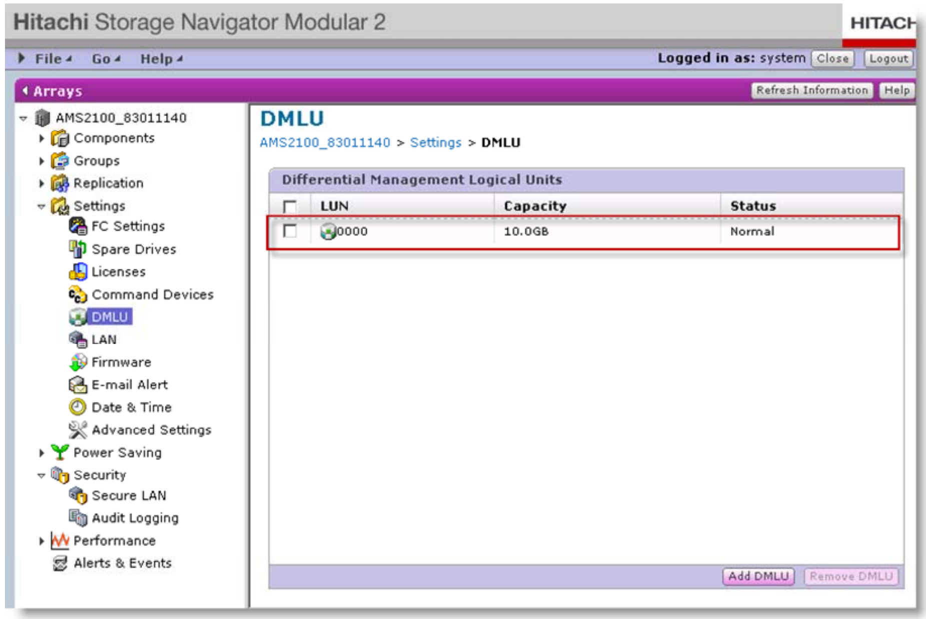

To set up the DM-LU with the SNM2 CLI, you have to run the command as follows:

audmlu -unit ‘unit_name’ -set -lu 0

11

You will find more detailed information on the **audmlu** CLI command in the *Storage Navigator Modular 2 Command Line Interface (CLI) User's Guide* on page 4-147

4. Hitachi ShadowImage Replication Software with Hitachi Command Control Interface

Module Objectives

- Upon completion of this module, the learner should be able to:
  - Identify the capabilities of Hitachi Command Control Interface (CCI) software relating to Hitachi ShadowImage® Replication software
  - Describe the steps to configure ShadowImage software with CCI
  - List key considerations when creating or revising the ShadowImage configuration file
  - Describe key troubleshooting scenarios and corrective actions

2

Overview of CCI

- Command Control Interface (CCI) Software:
  - Enables Hitachi replication software commands from hosts to the storage system
  - Interfaces with the system software and high-availability software on the system hosts
  - Provides failover and operation commands, which support mutual hot standby in conjunction with industry-standard failover products
  - Supports a scripting function for defining multiple replication operations in a script file
- CCI is sometimes called RAID Manager.

3

- CCI software must be installed on a host.
- Users can run CCI scripts to control and automate ShadowImage operations.
- Using CCI scripting, you can quickly set up and execute a large number of commands while integrating host-based high-availability control over remote copy operations.

Hitachi replication software includes:

- ShadowImage
- TrueCopy Remote Replication software
- Copy-on-Write Snapshot software

Components

- CCI software on the host
- HORCM Instance
  - “Service or Daemon”
  - Used to communicate with storage system and remote server
- HORCM Configuration File
  - Defines communication paths
  - Defines LUNs (volumes) to be controlled
- HORCM Commands
  - Monitor and control remote copy operations
- HORCM Command Device
  - Used to accept commands to and from HORCM
- DM-LU Modular Storage System
  - Used to store differential information

4

CCI software is loaded on O/S server. Supported platforms are Sun Solaris, HP-UX, AIX, Linux, Tru64 UNIX, IRIX, Microsoft Windows NT/2000/2003/2008.

HORCM “instance” is the CCI software that provides communication between the storage system and “remote” server. Users run two instances of HORCM – one is the ‘sending’ instance, which manages the P-VOLs, while the other HORCM instance is the ‘receiving’ instance, which manages the S-VOL.

HORCM Configuration Files – each HORCM instance requires its own configuration file

HORCM CMD DEVICE

The size of the CMD device should be 35MB to 50MB and should be accessible to the host running CCI. The command device should not have any user data on it.

A command device is a host-attached LUN that facilitates communication between the CCI and the storage system. The storage system port on which the command device is mapped should have the “Standard Inquiry Data Expand” option selected in host connection mode 2 settings.

Basic Connectivity Setup

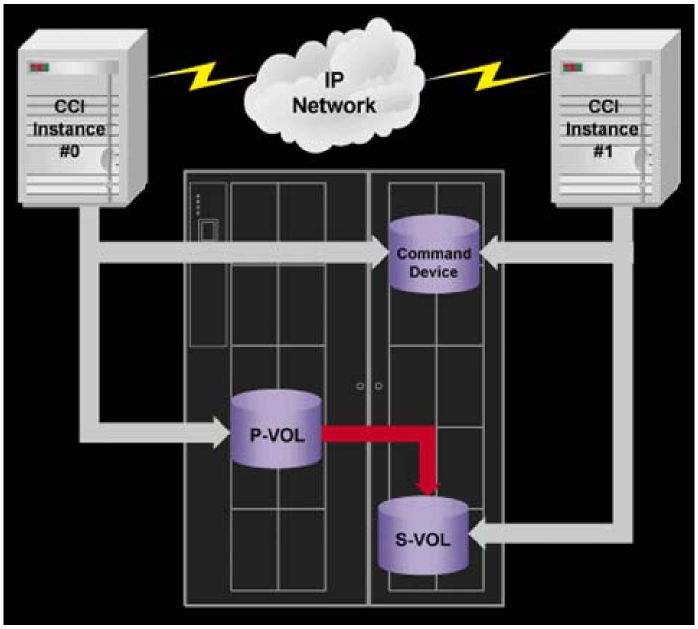

- CCI environment establishes a "conversation"
- Instances negotiate out-of-band with TCP/IP
- In-band communication using the command device

5

Basic connectivity setup includes:

1. A CCI installation on the hosts that are assigned the P-VOL and the S-VOL
2. IP connections between hosts running CCI instances
3. FC connection between the host running CCI, the production/backup host and the storage system where the P/S-VOL and command device are defined

CCI can also run on a separate dedicated host (not the production or the backup host) with Fibre Channel connectivity to the storage system where the command device is configured.

Description

- How CCI Works

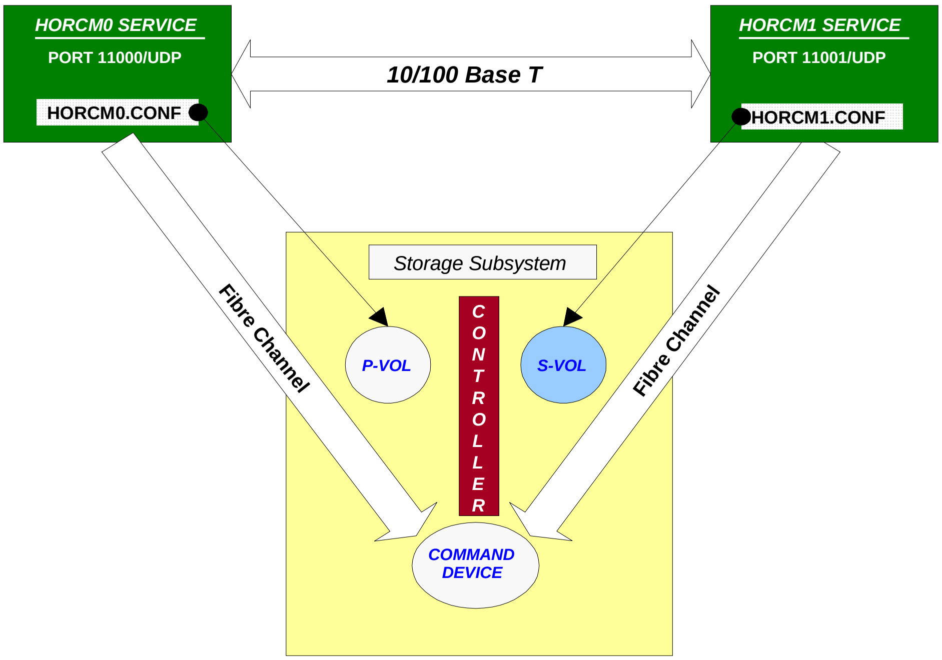

6

How CCI Works:

1. A command is issued either from the console (root prompt) on a Sun Solaris Server or through a script, such as, **paircreate**.
2. HORCM0 service traps the command if the variable HORCMINST=0 is set. This service expands the commands sent by the script/operator from the console and converts them to SCSI-3 commands, which are placed on the command device. This service also monitors the P-VOLs listed in HORCM0.CONF.
3. The HORCM0 service then communicates with the HORCM1 service over the LAN and informs the service to monitor the S-VOLs.
4. The controller on the storage system constantly polls the command device. Once available commands are detected, it picks up and executes the commands. For example, if the command was a **paircreate** command then the DKA pair copies data from the P-VOL to the S-VOL.

- One Server — Two HORCM Instances

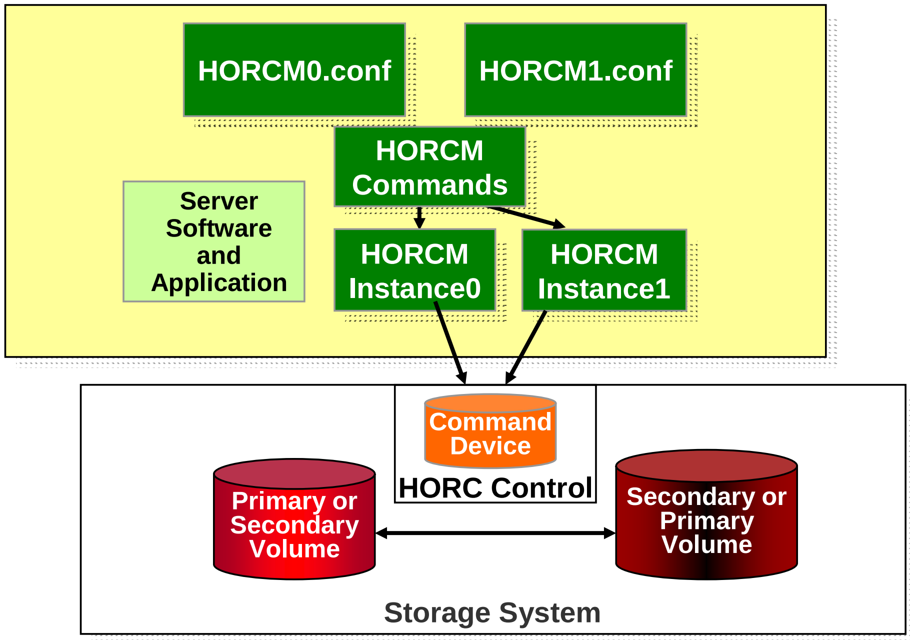

7

- Two Servers — Two HORCM Instances

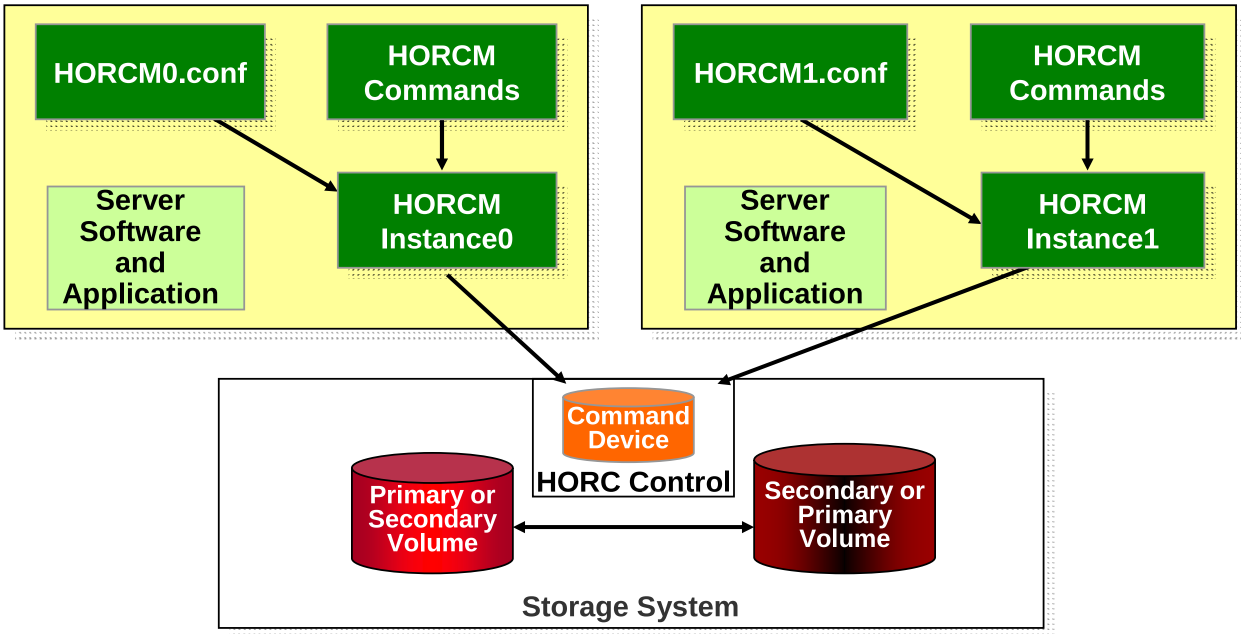

8

Above are the components of the Command Control Interface.

The relationship between components includes:

- There are always at least two HORCM instances, a sending instance and a receiving instance. Instance0 is the sending instance and Instance1 is the receiving instance.
- Each instance relies on a configuration file in order to communicate with other instance and the subsystem.
- The configuration file defines the volumes that will be "paired up".

If there are two instances, there must be two corresponding configuration files.

When a command is issued (usually via a script), the instance sends the command to the CMD device and then the storage system will execute the command.

Configuration Files

- CCI Configuration files include:
  - Services files
    - /etc/services (UNIX)
    - /WINNT/system32/drivers/etc/services (Microsoft Windows NT and Server 2000)
    - /Windows/system32/drivers/etc/services (Microsoft Windows Server 2003)
  - horcm0.conf and horcm1.conf
    - /etc Directory (UNIX)
    - WINNT Directory (Microsoft Windows NT and Server 2000)
    - Windows Directory (Microsoft Windows Server 2003)

9

- Service files include:
  - /etc/services (UNIX)
  - /WINNT/system32/drivers/etc/services (Microsoft Windows NT and Server 2000)
  - /Windows/system32/drivers/etc/services (Microsoft Windows Server 2003)
- horcm0 11000/udp #HORCM instance 0 port
- horcm1 11001/udp #HORCM instance 1 port

Any number between 5000 and 65535. Do not duplicate previously assigned ports.

udp - User Datagram Protocol.
A connection-less method of data packet delivery. The application is responsible for data transmission monitoring, error detection, and retransmission request.

10

- horcm0.conf
  - A text file located in the:
    - /etc Directory (UNIX)
    - /WINNT Directory (Microsoft Windows NT and Server 2000)
    - /WINDOWS Directory (Microsoft Windows Server 2003)
  - Provides a definition of hosts, groups, volumes, and command device to the CCI instance
  - Contains four areas that define the CCI environment
    - HORCM_MON
    - HORCM_CMD
    - HORCM_DEV
    - HORCM_INST

11

horcm.conf

- HORCM_MON – Defines monitoring and communication parameters. The host IP address or DNS name, service name, poll rate, and timeout value are all set in this section
- HORCM_CMD – Defines the command devices on the storage system
  - More than one command device can be used to provide failover functionality should the original command device become inoperable
  - More than one path can be used in order to multipath the command device

12

- HORCM_DEV – Defines the device group (names a group of paired logical volumes), the device names (names the paired volume in the device group), the port ID, target ID and LUN are set in the DEV section of the configuration file
  - Pairs are group and device-name dependent.
- HORCM_INST – Defines the network settings that correspond to the defined group names in the DEV section on the configuration file
  - The group name, IP address of the other host, and the function name are set in the INST section of the configuration file.

13

horcm0.conf Managing P-VOL

- **HORCM_MON**

```
#ip_address  service   poll(ms)  timeout(ms)
SVR1         horcm0    6000      3000
```

- **HORCM_CMD**

```
#dev_name
/dev/rdsk/c2t1d1s2                                 # Solaris
\\.\Physicaldrive2                                 # Windows NT, 2000, and 2003
\\.\Volume{f66c6208-6da0-11da-912a-505054503030}   # Windows 2000 and 2003
\\.\CMD-77010047-2-CL1-C                           # Windows 2000 and 2003
```

- **HORCM_DEV**

```
#dev_group  dev_name  port #  target ID  LUN#  MU#
oradb1      disk1     CL1-A   3          1     0
```

- **HORCM_LDEV**

```
#dev_group  dev_name  Serial#   LUN#  MU#
SIgroup1    p0        77010047  0     0
```

- **HORCM_INST**

```
#dev_group  ip_address  service
oradb1      SVR1        horcm1
```

14

HORCM_MON describes:

- **IP address** (host name) or the number of the server running instance 0.
- **service** (local service) is the /etc/services file port name line entry for instance 0. The port can be explained as a "socket" number to communicate to instance 1, which is also located in the /etc/services file and vice-versa.
- **poll** interval in milliseconds (1000 milliseconds = 1 second). This indicates the number of times HORCM daemon will "look at" the command device for status about the pairs. When this number is higher, HORCM daemon overhead on the running server is reduced. 1000ms is the default value.
- **timeout** value in milliseconds (1000 milliseconds = 1 second). This indicates the time for which the HORCM daemon will wait for status from instance 1 before timing out. In ShadowImage mode, this will apply to communication between the two instances running on one server when applicable.

HORCM_CMD describes the path to the raw device serving as the command device.

HORCM_DEV describes the source LUNs

- **dev group** name associates all LUNs to be controlled as a group for manipulation from one command.
- **dev name** must be unique for all devices within a group.

- Source to target pairs are group and device name associated and not line number associated. **PORT#, TID and LUN#** are self explanatory
- **MU#** is mirror unit number when creating one source to multiple targets.
  - 0 = first copy and is implied
  - 1 = second copy of same source
  - 2 = third pair of same source.
    Each MU number must be specified as a unique group name and device name. The MU field is valid in the horcm0.conf file only and not horcm1.conf. MU is always 0 for the modular arrays since one source can only have one target.

HORCM_INST

- **dev group** field requires one entry per group specified in the HORCM_DEV definitions.
- **ip address** describes the ip address or the name of instance 1.
- **service** is the /etc/services file port name line entry for instance 1.

horcm1.conf Managing S-VOL

- **HORCM_MON**

```
#ip_address    service    poll(ms)        timeout(ms)
SVR1           horcm1     6000            3000
```

- **HORCM_CMD**

```
#dev_name
/dev/rdsk/c2t1d1s2                                  # Solaris
\\.\Physicaldrive2                                  # Windows NT, 2000 and 2003
\\.\Volume{f66c6208-6da0-11da-912a-505054503030}    # Windows 2000 and 2003
```

- **HORCM_DEV**

#dev_group	*dev_name*	*port #*	*target ID*	*LUN#*	*MU#*
oradb1	disk1	CL1-A	0	2	0

- **HORCM_INST**

```
#dev_group     ip_address   service
oradb1         SVR1         horcm0
```

15

Setting Up ShadowImage Replication Software

1. Create the CMD device and differential management LUN
2. Map P-VOL and S-VOLs as LUNs on the storage system
3. Create HORCM configuration files
4. Edit the **`services`** file
5. Find the command device
6. Update HORCM with the command device info
7. Set the environment variables
8. Find absolute LUN numbers and shutdown HORCM
9. Update HORCM with the LUN info
10. Start instances
11. Issue basic commands

16

Configuration

Step 1 – Create command device

- Create a 36MB volume.
- Use Storage Navigator Modular 2.
  - Map the volume to the correct port.
  - Configure the volume as the command device.

17

- Example
  - The command device accepts TrueCopy and ShadowImage read and write commands executed by the system.
  - The volume designated as the command device is used only by the system and is a raw device.
  - The command device can be any LUN that is accessible by the host.
  - A LUN as small as 36MB can be used as a command device.
  - A LUSE volume cannot be used as a command device.

```
HORCM_CMD
#dev_name          dev_name          dev_name
/dev/rdsk/c0t0d11s2                                   # Solaris
\\.\Physicaldrive2                                    # Windows NT, 2000 and 2003
\\.\Volume{f66c6208-6da0-11da-912a-505054503030}      #Windows 2000 and 2003
\\.\CMD-77010047-2-CL1-C                              #Windows 2000 and 2003
```

18

The command device accepts TrueCopy and ShadowImage read and write commands executed by the system and returns read requests to the UNIX/PC host. The volume designated as the command device is a raw device used only by the system. Make sure that the volume to be selected as the command device does not contain any user data; any user data on that volume will become inaccessible to the UNIX/PC server host. A command device can be any OPEN-x device (e.g., OPEN-3, OPEN-8) that is accessible by the host.

A virtual LVI/LUN volume as small as 36MB (for example, OPEN-3-CVS) can be used as a command device. A LUSE volume cannot be used as a command device. For Sun Solaris operations, the command device must be labeled. To enable dual pathing of the command device under Sun Solaris systems, make sure to include all paths to the command device on a single line in the HORCM_CMD section of the configuration file. Putting the path information on separate lines may cause parsing issues and failover may not occur unless the HORCM startup script is restarted on the Sun Solaris system.

Step 2 – Create S-VOLs

- Open Storage Navigator Modular 2
- Record the Port, LUN, Target ID, and Host Number of both the P-VOLs and S-VOLs

19

LUNs cannot be defined through CCI. S-VOLs need to be defined through LUN Manager prior to using CCI. Ensure that S-VOL LUNs are mapped to the appropriate ports.

For example, a backup scenario:

- The P-VOL is mapped to the port of the production server and S-VOL is mapped to the port of the backup server.
- Record the PORT/LUN information.
- This information will be added to the configuration files.

- Proposed Configuration

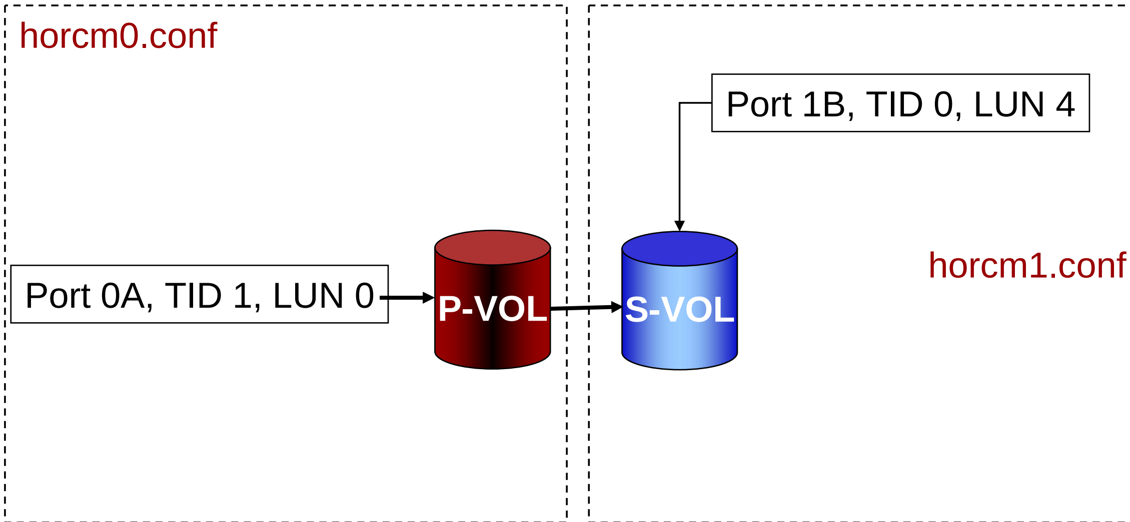

20

Step 3 – Create HORCM files

- Four sections
  - HORCM_MON
  - HORCM_CMD
  - HORCM_DEV
  - HORCM_INST
- Two HORCM files are required
- Non-numbered instance configured by `horcm.conf`
- Numbered instances configured by `horcm`*`n`*`.conf` (for example, **`horcm0.conf`**)
- Location of configuration files:
  - UNIX: /etc/
  - Microsoft Windows NT and 2000: C:\WINNT\
  - Microsoft Windows 2003: C:\Windows

21

- Create HORCM Files — HORCM_MON Section

```
HORCM_MON
#IP          service     poll        timeout
hosta        horcm0      1000        3000

HORCM_CMD
\\.\Volume{f66c6208-6da0-11da-912a-505054503030}

HORCM_DEV
#dev_group         dev_name    port     TID   LUN   MU#
group1             dev01       CL2-A    5     0     0

HORCM_INST
#dev_group         IP          service
group1             hosta       horcm1
```

22

HORCM_MON section provides network identity and timing data.

The monitor parameter (**HORCM_MON**) defines the following values:

- **ip_address:** The IP address of the local host.
- **service:** The port name assigned to the CCI service (registered in the **/etc/services** file). The service parameter defines the CCI instance that runs on the local host. If a port number is specified instead of a port name, the port number will be used.
- **poll:** The interval for monitoring paired volumes. To reduce the HORCM daemon load, make this interval longer. If set to -1, the paired volumes are not monitored. The value of -1 is specified when two or more CCI instances run on a single machine.
- **timeout:** The time-out period of communication with the remote server.

- Create HORCM Files — HORCM_CMD Section

```
HORCM_MON
#IP         service     poll        timeout
hosta       horcm0      1000        3000

HORCM_CMD
\\.\Volume{f66c6208-6da0-11da-912a-505054503030}

HORCM_DEV
#dev_group        dev_name    port     TID   LUN   MU#
group1            dev01       CL2-A    5     0     0

HORCM_INST
#dev_group        IP          service
group1            hosta       horcm1
```

23

HORCM_CMD section identifies the command device LUN.

- The command device is a user-selected, dedicated logical volume on the system, which functions as the interface to the CCI software on the UNIX/PC host.
- The command device is dedicated to CCI communications and cannot be used by any other applications.
- The command device accepts ShadowImage read and write commands that are executed by the system.
- The volume designated as the command device is used only by the system and is blocked from the user.
- The command device uses 16MB while the remaining volume space is reserved for CCI and its utilities. A virtual LVI/LUN volume as small as 36MB (e.g., OPEN-3-CVS) can be used as a command device.

- Create HORCM Files — HORCM_DEV Section

```
HORCM_MON
#IP         service     poll        timeout
hosta       horcm0      1000        3000

HORCM_CMD
\\.\Volume{f66c6208-6da0-11da-912a-505054503030}

HORCM_DEV
#dev_group        dev_name    port     TID   LUN   MU#
group1            dev01       CL2-A    5     0     0

or:

HORCM_LDEV
#dev_group        dev_name    Ser #        LUN   MU#
group1            dev01       7701004    5     0


HORCM_INST
#dev_group        IP          service
group1            hosta       horcm1
```

24

HORCM_DEV section identifies groups of LUNs that are maintained by this instance. These can be either P-VOLs or S-VOLs.

The device parameter (HORCM_DEV) defines the device addresses for the paired logical volume names.

- **group**: Names a group of paired logical volumes; commands are executed for all corresponding volumes according to this group name.
- **dev_name**: Names the paired logical volume within a group (i.e., name of the special file or unique logical volume). The name of paired logical volume must be different from the “dev name” on another group.
- **port #**: Defines the Adaptable Modular Storage/Workgroup Modular Storage port number of the volume that corresponds with the dev_name volume.
- **target ID**: Defines the SCSI/fibre target ID number of the physical volume on the specified port.
- **LU #**: Defines the SCSI/fibre logical unit number (LU#) of the physical volume on the specified target ID and port.
- **MU #**: Defines the mirror unit number (0-2) of ShadowImage volumes. If this number is omitted, it is assumed to be zero (0).

- Create HORCM Files — HORCM_INST Section

```
HORCM_MON
#IP          service     poll        timeout
hosta        horcm0      1000        3000

HORCM_CMD
\\.\Volume{f66c6208-6da0-11da-912a-505054503030}

HORCM_DEV
#dev_group         dev_name    port     TID   LUN   MU#
group1             dev01       CL2-A    5     0     0

HORCM_INST
#dev_group         IP          service
group1             hosta       horcm1
```

25

HORCM_INST section identifies the remote HORCM instance that manages the alternate half of each group's mirror set.

Instance parameter (**HORCM_INST**) defines the network address (IP address) of the remote server. The following values are defined in the HORCM_INST parameter:

- **dev_group**: The server name described in dev_group of HORCM_DEV.
- **ip_address**: The network address of the specified remote server.
- **service**: The port name assigned to the HORCM communication path.

- Case 1 – Create HORCM files on system on same host

```
HORCM_MON
#IP          service             poll  timeout
hosta        horcm0              1000  3000

HORCM_CMD
#LEAVE COMMENTED FOR NOW
HORCM_DEV
#-----------------------------------------------------
#dev_group   dev_name   port    TID LUN   MU#
#-----------------------------------------------------
#LEAVE COMMENTED FOR NOW
#-----------------------------------------------------
HORCM_INST
#-----------------------------------------------------
#dev_group        IP          service
group1            hosta       horcm1
```

26

Add the HORCM.conf files to the local host connected to the system.

Two HORCM files must be created, one for P-VOLs and one for S-VOLs.

Each file name should start with “HORCM” and end with [0-9].conf.

Leave most information blank since it will be filled in later.

Note: The IP column in the HORCM_INST and HORCM_MON section are the same if both HORCM files exist on the same host.

- Case 2 – Create HORCM files on different hosts

horcm0.conf

```
HORCM_MON
#IP             service      poll      timeout
hosta           horcm0          1000    3000

HORCM_CMD
#LEAVE COMMENTED FOR NOW
HORCM_DEV
#---------------------------------------------
#dev_group   dev_name   port    TID LUN   MU#
#---------------------------------------------
#LEAVE COMMENTED FOR NOW
#---------------------------------------------
HORCM_INST
#---------------------------------------------
#dev_group      IP              service
group1          hosta           horcm1
```

horcm1.conf

```
HORCM_MON
#IP             service         poll    timeout
hosta           horcm1          1000    3000

HORCM_CMD
#LEAVE COMMENTED FOR NOW
HORCM_DEV
#---------------------------------------------
#dev_group   dev_name   port    TID LUN   MU#
#---------------------------------------------
#LEAVE COMMENTED FOR NOW
#---------------------------------------------
HORCM_INST
#---------------------------------------------
#dev_group      IP              service
group1          hosta           horcm0
```

27

Note: If both HORCM files are on different hosts, then the IP column in the HORCM_INST and HORCM_MON section are the interchanged. The service column in the HORCM_INST group are is interchanged.

Best Practice: Always configure P-VOLs in horcm0.conf and S-VOLs in horcm1.conf regardless of whether or not both instances are on the same server.

- Key Configuration File Considerations
  - The HORCM instance must be stopped and restarted after changes have been made to the configuration file.
  - Be aware of extra spaces and tabs in the file.
    - Separate columns with one tab only.
  - Group and device names:
    - Case sensitive
    - Must be unique within the file
    - Must match across both files
      - For example: Device name “DeviceOne” in HORCM0.conf must be listed as “DeviceOne” in HORCM1.conf for the corresponding device.
  - Save the HORCM configuration files in the correct OS directory.

28

If you use the device name “DeviceOne” in HORCM0.conf, then it must be “DeviceOne” in HORCM1.conf. Variance in the name such as “deviceone” or “Deviceone” or “ DEVone” are not accepted. Remember that this rule applies when referring to that device at the command prompt or in a script; users tend to forget this when issuing commands via the command prompt.

Step 4 – Find command device

- To find command device
  - Microsoft Windows: raidscan –x findcmddev hdisk 0,20
    - Resulting output >
      \\.\PhysicalDrive5
      \\.\Volume{b9b31c79-240a-11d5-a37f-00c00d003b1e}
  - Unix: Type command “format”
    - Resulting Output >
      c2t0d10 <HITACHI – DF600F – 000 – CM -0112 cyl >
- Although raidscan is a CCI-specific command, the instances do not need to be started on Microsoft Windows NT/2000 host
- Record this information for insertion into HORCM file at a later time
- Enter these commands on both local and remote host

Sun Solaris:

ls /dev/rdsk/* | ./inqraid

HP-UX:

ioscan –fun | grep rdsk | ./inqraid

Microsoft Windows:

raidscan –x findcmddev hdisk0,20

If you are on a Windows system, you do not need the HORCM instances running in order to execute a raidscan command. On Unix systems, the HORCM instances need to be running in order to execute a raidscan command.

Look for a Hitachi disk with a “-CM” in the device identifier.

Note:

The "\\.\Volume{GUID}" must be made by setting a partition used by the disk management without file system format, and you will be able to know "\\.\Volume{GUID}" for the command device using "inqraid $Volume -CLI -fv" or "raidscan -x findcmddev 0,X".

The "\\.\Volume{GUID}" will be kept as the same command device even though the physical drive numbers are changed on every reboot in a SAN environment, and HORCM will convert to "\\.\PhysicalDriveX" by finding the PhysicalDrive number through "\\.\Volume{GUID}" if "\\.\Volume{GUID}" is specified to HORCM_CMD.

Step 5 – Edit the services file

```
# Copyright (c) 1993-1995 Microsoft Corp.
#
# This file contains port numbers for well-known services as defined by
# RFC 1060 (Assigned Numbers).
#
# Format:
#
# <service name>  <port number>/<protocol>  [aliases...]   [#<comment>]
#

echo                7/tcp
echo                7/udp
discard             9/tcp    sink null
discard             9/udp    sink null
systat             11/tcp
systat             11/tcp    users
daytime            13/tcp
daytime            13/udp
netstat            15/tcp
qotd               17/tcp    quote
qotd               17/udp    quote
chargen            19/tcp    ttytst source
chargen            19/udp    ttytst source
ftp-data           20/tcp
ftp                21/tcp
telnet             23/tcp
smtp               25/tcp    mail
time               37/tcp    timserver
time               37/udp    timserver
rlp                39/udp    resource      # resource location
name               42/tcp    nameserver
```

30

- Example: Service file with HORCM instances added

```
nfs              2049/udp                    # sun nfs
knetd            2053/tcp                    # Kerberos de-multiplexor
eklogin          2105/tcp                    # Kerberos encrypted rlogin
rmt              5555/tcp    rmtd
mtb              5556/tcp    mtbd            # mtb backup
man              9535/tcp                    # remote man server
w                9536/tcp
mantst           9537/tcp                    # remote man server, testing
bnews           10000/tcp
rscs0           10000/udp
queue           10001/tcp
rscs1           10001/udp
poker           10002/tcp
rscs2           10002/udp
gateway         10003/tcp
rscs3           10003/udp
remp            10004/tcp
rscs4           10004/udp
rscs5           10005/udp
rscs6           10006/udp
rscs7           10007/udp
rscs8           10008/udp
rscs9           10009/udp
rscsa           10010/udp
rscsb           10011/udp
qmaster         10012/tcp
qmaster         10012/udp

horcm0          11000/udp                 # horcm instance 0
horcm1          11001/udp                 # horcm instance 1
```

31

Step 6 – Update HORCM files

– Enter the command device you have located into the horcm files.

– Be sure to add the command device in BOTH locations

```
HORCM_MON
#IP         service     poll  timeout
hosta       horcm0            1000  3000

HORCM_CMD
\\.Volume{b9b31c79-240a-11d5-a37f-00c00d003b1e}

HORCM_DEV
#dev_group   dev_name   port    TID LUN   MU#
#COMMENT THIS OUT

HORCM_INST
#dev_group       IP          service
group1           hosta       horcm1
```

32

At this time the device information should be commented out, but enter the command device information on the local and remote hosts. If the information that was collected from LUN Management is added, the instances will not correctly recognize the volumes due to an incorrect/missing LU#. The LU# entered must be the ABSOLUTE LUN number.

Use system raw device name to enter the command device.

Sample formats:

- Sun Solaris: /dev/rdsk/c1t1d2s2
- HP-UX: /dev/rdsk/c1t1d2
- Microsoft Windows: \\.\PhysicalDrive4

At this time, if multiple command devices are required, they should be listed on a single line.

Step 7 – Locate absolute LUN number on both the local and remote systems and perform an update on the HORCM files

Important: TID and LUN numbers are based on the system's view, not the host's view

– Set environment variables
  - Set HORCMINST=<select the configuration that contains P-VOL and enter the decimal digit at the end of horcm[0-9] here>

– Start instance
  - Execute **horcmstart.s**h or **horcmstart** <select the configuration that contains P-VOL and enter the decimal digit at the end of horcm[0-9] here>

– Execute command

raidscan –p <*port that you connected desired volume to*> **–fx**

33

Execute the following command to determine the TID associated with each LUN once the "shell" HORCM instances are running:

raidscan –p CL1-A –CLI

- Locate Absolute LUN Number — Storage versus host TID and LUN

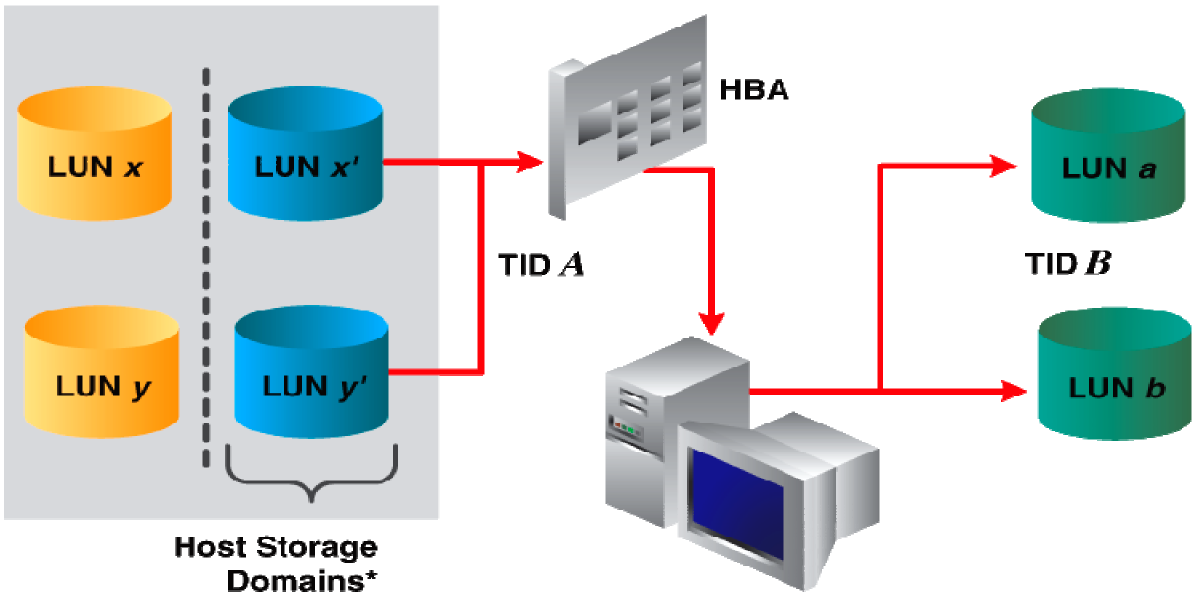

34

- Locate Absolute LUN Number
  - Information resulting from portscan
  - Record the absolute LUN numbers (required)
  - Stop the horcm instances using the following commands. Ensure that the absolute LUN#s are recorded prior to stopping the instances.
    - Horcmshutdown
      or
      horcmshutdown.sh [horcm instance number]

RAIDSCAN –p CL1-A –CLI –f

```
PORT# /ALPA/C TID# LU#  Seq#  Num LDEV#  P/S  Status Vol.Type
CL1-A   ef  5  1   0 0   995   1     0  SMPL     -   DF600F
CL1-A   ef  5  1   1 0   995   1     1  SMPL     -   DF600F-CM
CL1-A   ef  5  1  10 0   995   1    10  SMPL     -   DF600F
CL1-A   ef  5  1  11 0   995   1    11  SMPL     -   DF600F
CL1-A   ef  5  1  12 0   995   1    12  SMPL     -   DF600F
CL1-A   ef  5  1  13 0   995   1    13  SMPL     -   DF600F
CL1-A   ef  5  1  14 0   995   1    14  SMPL     -   DF600F
```

35

- Locate Absolute LUN Number

```
HORCM_MON
#IP_ADDR                SERVICE          POLL        TIMEOUT
10.0.0.10               horcm0           1000        3000

HORCM_CMD
#dev_name
\\.\PhysicalDrive5
\\.\Volume{fab8dec0-0618-11d8-b67a-000103dc105e}

HORCM_DEV
#GROUP             DEVNAME     PORT#       TID   H-LU   MU
VG                 SI1         CL1-A-001   1     0      0

HORCM_INST
#DEV_GROUP  IP_ADDR     SERVICE
VG          10.0.0.10   horcm1
```

It's possible to use the mapped addresses (H-LUN and not the I-LUN – 001 is the Host Group 001)
You need the latest Microcode and CCI

36

- Start Naming the P-VOLs
  - Each DEV_GROUP is composed of one or more DEV_NAMEs
  - Users choose their DEV_GROUP and DEV_NAME field names
    - When creating field names it is helpful to make them be descriptive
  - DEV_GROUP and DEV_NAME can use up to 31 characters for each name

37

- Match Names to the S-VOLs
  - Device names match for both the P-VOLs and S-VOLs
  - Device groups match for both the P-VOLs and S-VOLs
  - See the sample HORCM conf handout for a complete view

38

- Step 8 – Update HORCM files

```
HORCM_MON
#IP          service           poll  timeout
hosta        horcm0            1000  3000

HORCM_CMD
\\.\PhysicalDrive5
\\.\Volume{b9b31c79-240a-11d5-a37f-00c00d003b1e}

HORCM_DEV
#dev_group  dev_name  port         TID   LUN   MU#
group1      dev01     CL1-A        1     10    0
group1      dev02     CL1-A        1     11    0

HORCM_INST
#dev_group        IP          service
group1            hostb       horcm1
```

39

Step 9 – Set environmental variables

- HORCMINST = <last number with horm_.conf>
- HORCC_MRCF = 1
  (or use –IMX or -SIX in command, where X = <last number with horm_.conf>)
-

Step 10 – Start instances

- Microsoft Windows
  - Horcmstart <last number used in horcm_conf>
- UNIX
  - horcmstart.sh <last number used in horcm_conf>

Step 11 – Issue commands

- pairdisplay
- paircreate
- pairsplit
- pairresync

On ShadowImage software, an environment must also be set; it is required to set the environment variable HORCC_MRCF to 1. The HORCMINST environment variable should be designated to the number you have chosen when you named your HORCM configuration file. This will direct CCI in the right direction when it needs to find a configuration file to read.

Horcminst=#

You must specify the instance number. By default, this environment variable is blank. This number depends on the ending number labeled on your HORCM file. For example, if your HORCM file is labeled as horcm0.conf, the horcminst variable must be equal to 0. If a blank environment variable is chosen, the horcm file can be named “horcm.conf”, but this configuration will not work if both horcm files are on the same host since two horcm files are required for any ShadowImage configuration.

Both the command execution environment and the command activation environment require an instance number.

Horcc_mrcf=1

The CCI command, TrueCopy, or ShadowImage is selected by the environment variable HORCC_MRCF. The default command execution environment for CCI is TrueCopy. You must set this environmental variable when performing ShadowImage.

Note: The Horcm_Log environmental variable allows you to set the directory where the log files will be located. By default, the log files are stored in Horcm/log/curlog.

CCI Commands

- CCI Commands Overview
  - pairdisplay Displays pair information
  - paircreate Creates ShadowImage pairs
  - pairsplit Splits a ShadowImage pair
  - pairresync Resynchronizes a ShadowImage pair
  - pairevtwait Waits for a event transition
  - pairmon Monitors and reports on pair events
  - raidscan Displays port configuration
  - raidqry Displays CCI host configuration
  - inqraid Displays relation between special files and PDEVs

41

- pairdisplay example:

```
# pairdisplay -g oradb
Group Pair Vol(L/R) (Port#,TID,LU-M), Seq#, LDEV#…. P/S, Status, Seq#, P-LDEV# M
oradb oradb1(L)    (CL1-A,  1, 0  ) 30053   18... P-VOL  PAIR   30053  19    -
oradb oradb1(R)    (CL1-D,  1, 0  ) 30053   19... S-VOL  PAIR   ----   18    -
```

42

In Pair Vol (L/R) - L stands for local machine, R stands for remote machine.

In ShadowImage software:

L is the P-VOL

R is the S-VOL

If there is a “W” shown in the M column, the S-VOL is write enabled.

- pairdisplay within a shell script:

```
pairevtwait -g $datagroup -nowait | grep " PSUS" > /dev/null
 if [ $? -ne 0 ]
 then
   pairdisplay -g $datagroup -fcx
   exit 1
 else
   echo "$W Pair '$datagroup' already Split"
 fi
```

43

pairdisplay -g oradb1 –fcx

It is used to verify the pairing of volumes, as well as to verify completion of pair synchronization

-x option is used to display the LDEVs in hex.

-fc option is used to display the pair copy completion percentage

- paircreate examples:

paircreate -g oradb –vl

paircreate -g oradb -d dev1 -vl

paircreate -g oradb -vl -split

44

If horcminst=0

Then if the command **paircreate –g oradb –vl** is used data goes from P-VOL to S-VOL

However if the command **paircreate –g oradb –vr** is used then data goes from S-VOL to P-VOL

If horcminst=1

Then if the command **paircreate –g oradb –vl** is used then data will go from S-VOL to P-VOL

However if the command **paircreate –g oradb –vr** is used then data will go from P-VOL to S-VOL

paircreate –g oradbi -vl

- This command copies data from source volumes to target volumes
  - The **-vl** option copies data from the SOURCE volume to the target volume
  - Using the **-vr** instead of -vl option copies data from the target volumes to the source volumes
  - Pair Status Transition SIMPLEX >>> COPY >>>PAIR

pairsplit -g oradb1

This command stops updating the secondary volume and pair status is maintained

- The pair is put into Suspend Status (PSUS)
- Changed-Data Track Table is maintained in the modular storage system shared memory
- Pair Status Transition PAIR >>> PSUS

pairsplit -g oradb1 –S

- This command immediately stops updating the secondary volume and sets the shadow pairs to simplex status via the **-S** option
- No 'Change Data' bit map is maintained
- **pairresync** is not possible so you must **paircreate**

- pairsplit examples:

```
# pairsplit -g oradb

# pairsplit -g oradb -d dev1

# pairsplit -g oradb -rw -FMRCF

# pairsplit -g oradb -S

# pairsplit -g oradb -E
```

45

pairsplit –g oradb: Splits all the devices in the group oradb

pairsplit –g oradb –d dev1: Splits only the device dev1 in the group oradb

pairsplit –g oradb –rw –FMRCF: Splits a ShadowImage pair that has been cascaded from a TrueCopy S-VOL/P-VOL

pairsplit –g oradb –S: Deletes the pair and switches it to SIMPLEX

pairsplit –g oradb –E: Aborts current action

- pairsplit within a shell script:

```
for datagroup in $VGNAMES
do
  pairsplit -g $datagroup
  if [ $? -ne 0 ]
  then
    echo "Error Pairsplit command failed for '$datagroup' \n"
    exit 1
  else
    echo "Split command '$datagroup' successful"
  fi
  pairevtwait -g $datagroup -s psus -t 32000
  if [ $? -ne 0 ]
  then
    echo "Error Pairsplit not complete for '$datagroup' \n"
    exit 1
  else
    echo "Split '$datagroup' successful"
  fi
done
```

46

- pairresync examples:

```
# pairresync -g oradb

# pairresync -g oradb -d dev1

# pairresync -g oradb -restore
```

47

pairresync -g oradb1: It re-establishes the pair for the entire group oradb1 and copies 'Change Data' only from the primary volume to the secondary volume. Changed-data track table maintains changes made to both source and target volumes. Re-synchronization is always from source to target volume

Pairresync –g oradb –d dev1: It re-establishes the pair only for the device dev1 within the group oradb.

Pairresync –g oradb –restore: It re-establishes the pair for the entire group oradb and copies 'Change Data' only from the secondary volume to the primary volume. Changed-data track table maintains changes made to both source and target volumes. Re-synchronization is always from target to source volume

- pairresync within a shell script:

```
for datagroup in $VGNAMES
do
pairresync -g $datagroup
  if [ $? -ne 0 ]
  then
    echo "Error Pairresync command failed for '$datagroup' \n"
    exit 1
  else
    echo "Start of resynchronization of '$datagroup' successful"
  fi
  echo "Wait for end of resynchronization"
  $pairevtwait -g $datagroup -s pair -t 32000
  if [ $? -ne 0 ]
  then
    echo "Error Resync not complete \n"
    exit 1
  fi
pairdisplay -g $datagroup -fcx
done
```

48

- pairevtwait within a shell script:

```
for datagroup in $VGNAMES
do
  $pairevtwait -g $datagroup -nowait | grep " PSUS" > /dev/null
  if [ $? -ne 0 ]
  then
    echo "$W Pair '$datagroup' not in Split status \n"
    exit 1
  fi
done
```

49

- pairmon example:

```
# pairmon -allsnd -nowait
Group Pair vol Port targ# lun# LDEV# Oldstat code > Newstat code
oradb oradb1   CL1-A 1    5    145    SMPL   0x00 > COPY    0x01
oradb oradb2   CL1-A 1    6    146    PAIR   0x02 > PSUS    0x04
```

50

Pairmon shows the last status change that happened.

- raidscan example:

```
# raidscan -p CL1-A
PORT# /ALPA/C,TID#, LU#..Num(LDEV#....)...P/S, Status, LDEV#,P-Seq#,P-LDEV#
CL1-A  /ef/ 0, 0,   7-0 .1(7).............SMPL  ----  -----,  -----  -----
CL1-A  /ef/ 0, 0,   8-0 .1(8).............SMPL  ----  -----,  -----  -----

bash-2.05# ls /dev/rdsk/* | raidscan -find
DEVICE_FILE        UID  S/F PORT   TARG  LUN    SERIAL  LDEV  PRODUCT_ID
/dev/rdsk/c3t0d0s2    0   F  CL1-B     0    0  75040072     0  DF600F
/dev/rdsk/c3t0d1s2    0   F  CL1-B     0    1  75040072     1  DF600F
/dev/rdsk/c3t0d2s2    0   F  CL1-B     0    2  75040072     2  DF600F
/dev/rdsk/c3t0d3s2    0   F  CL1-B     0    3  75040072     3  DF600F
/dev/rdsk/c3t0d4s2    0   F  CL1-B     0    4  75040072     4  DF600F
/dev/rdsk/c3t0d5s2    0   F  CL1-B     0   20  75040072    20  DF600F
/dev/rdsk/c3t0d6s2    0   F  CL1-B     0   21  75040072    21  DF600F
/dev/rdsk/c3t0d70s2   0   F  CL1-B     0   30  75040072    30  DF600F-CM
```

51

- raidqry example:

```
bash-2.05# raidqry -r SUN01
No Group    Hostname          HORCM_ver  Uid  Serial#    Micro_ver Cache(MB)
 1 SUN01    sunlabq         01-17-03/05    0 75040072  07-26-00/00      2048

bash-2.05# raidqry -l
No Group    Hostname          HORCM_ver  Uid  Serial#    Micro_ver Cache(MB)
 1  ---     sunlabq         01-17-03/05    0 75040072  07-26-00/00      2048
```

52

- inqraid example:

```
bash-2.05# ls /dev/rdsk/* | /HORCM/usr/bin/inqraid
/dev/rdsk/c3t0d0s2 -> [ST] CL1-B Ser =75040072 LDEV =   0 [HITACHI ] [DF600F         ] HORC =
      SMPL  HOMRCF[MU#0 = P-VOL MU#1 = SMPL MU#2 = SMPL] RAID5[Group  0- 0] SSID =
      0x0000
```

53

This example shows the relationship between special files.

- Pair event wait
  - Waits until a specific pair status is achieved before returning control
  - Useful for scripts that need to wait until a specific pair status is achieved

```
pairevtwait -g <group> -s <status> -t <timeout>
```

The **pairresync** command terminates before resynchronization of the secondary (or primary) volume is complete. Use pairevtwait to verify that the resync operation completed successfully (status changes from COPY to PAIR).

- Pair event wait (without waiting)
  - Query the pair status and return immediately
  - Returns a numeric value that indicates the current pair state
  - Useful when scripting actions based on the current pair status

pairevtwait -g <group> -s <status> -nowait

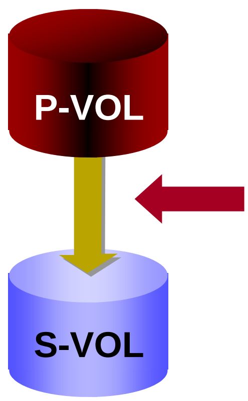

pairevtwait –nowait

Returns the state of the "data flow" between the pairs

55

WIN32 Subcommands

- Provides UNIX-like functionality for Microsoft Windows hosts
  - Can be executed by appending “–x <subcommand>” to any CCI command
  - Best Practice: Use a non-destructive command (such as **pairdisplay**) when using any subcommand

56

- WIN32 sync Subcommand
  - Flushes the system cache to disk
  - Sync does not propagate to volume mount points under a specified drive letter

```
pairdisplay –x sync all
pairdisplay –x sync D:
pairdisplay –x sync D:\mountpoint
pairdisplay –x sync hdisk3 hdisk4
pairdisplay –x sync \Vol2
```

57

- WIN32 mount Subcommand:
  - Mounts the specified volume to a drive letter or volume mount point
  - Partition number may not be required
  - No arguments: list mounted volumes

```
pairdisplay –x mount
pairdisplay –x mount D: hdisk3 p1
pairdisplay –x mount E: \Vol2
pairdisplay –x mount F:\mountpoint \Vol3
```

58

- WIN32 umount Subcommand
  - Unmounts the volume by deleting the drive letter or volume mount point mapping.
  - This subcommand will flush the system cache to disk prior to unmounting

```
pairdisplay –x umount D:
pairdisplay –x umount \Vol2
pairdisplay –x umount F:\mountpoint
```

59

Troubleshooting

- If the instance will not start, then check to see if:
  - The CMD device path is correct in the configuration file.
  - The CMD device is labeled, if not label it.
  - The configuration file is in the correct OS directory.
- If **pairdisplay** fails, then check to see if:
  - The environment variables are set:
    - HORCMINST = 0 or 1
    - HORCC_MRCF = 1

60

- If **paircreate** fails, then check to see if:
  - The target IDs are correct in the configuration file.
    - Use raidscan
  - LUNs are correct in the configuration file.

61

Module Review

1. What are the components of ShadowImage software on a modular storage system?
2. ShadowImage operations are (A) asynchronous or (B) synchronous?
3. Describe the difference between splitting and suspending pairs.
4. What is the purpose of the HORCM configuration files?

63

5. Hitachi Storage Navigator Modular 2 ShadowImage Replication Operations

Module Objectives

- Upon completion of this module, the learner will be able to:
  - Use the Hitachi Storage Navigator Modular 2 GUI and CLI to perform ShadowImage Replication operations on the Adaptable Modular Storage (AMS) 2000 family

2

Overview

- Storage Navigator Modular 2 GUI and CLI enables you to:
  - Create pairs
  - List and view pairs
  - Split pairs
  - Resynchronize pairs
  - Restore pairs
  - Delete pairs

- **No** Command Device, **no** CCI Raid Manager, **no** HORCM script and **no** HORCM instance is necessary when using SNM2 to perform ShadowImage operations on an AMS2000 Family storage system.

3

The following slides display screenshots of the ShadowImage software installation.

SNM2 CLI – CCI/Raid Manager Commands

Action	CLI – Command	CCI - Command
Display pair information	aureplicationlocal -refer	pairdisplay
Create pairs	aureplicationlocal -create	paircreate
Split pairs	aureplicationlocal -split	pairsplit
Resynchronize pairs	aureplicationlocal -resync	pairresync
Restore pairs	aureplicationlocal -restore	pairresync -restore
Delete pairs	aureplicationlocal -simplex	pairsplit -S
Event wait	aureplicationmon -evwait	pairevtwait

4

License

- A license key for ShadowImage is required and must be enabled.

5

The license unlocks the option using the key code for installing ShadowImage software.

SNM2 – ShadowImage Pair Operations

- ShadowImage Pair management with Storage Navigator Modular 2
  - Create pairs
  - View the pair status
  - Split pairs
  - Resynchronize pairs
  - Restore pairs
  - Delete pairs

6

Create ShadowImage Pairs with SNM2

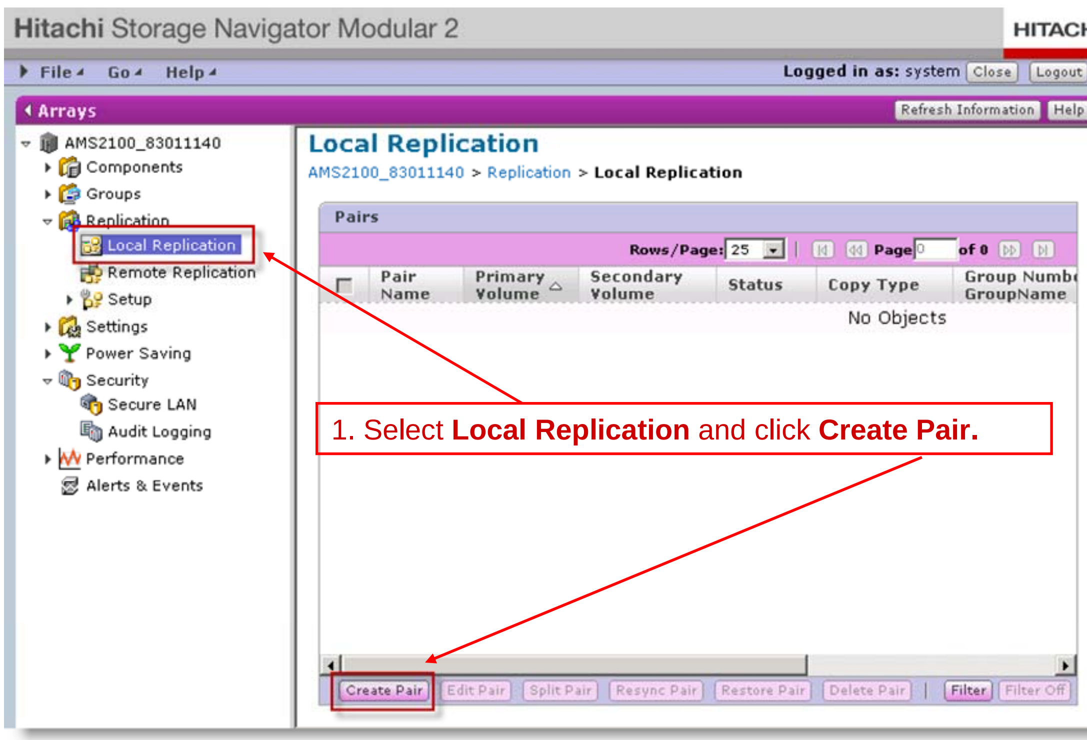

7

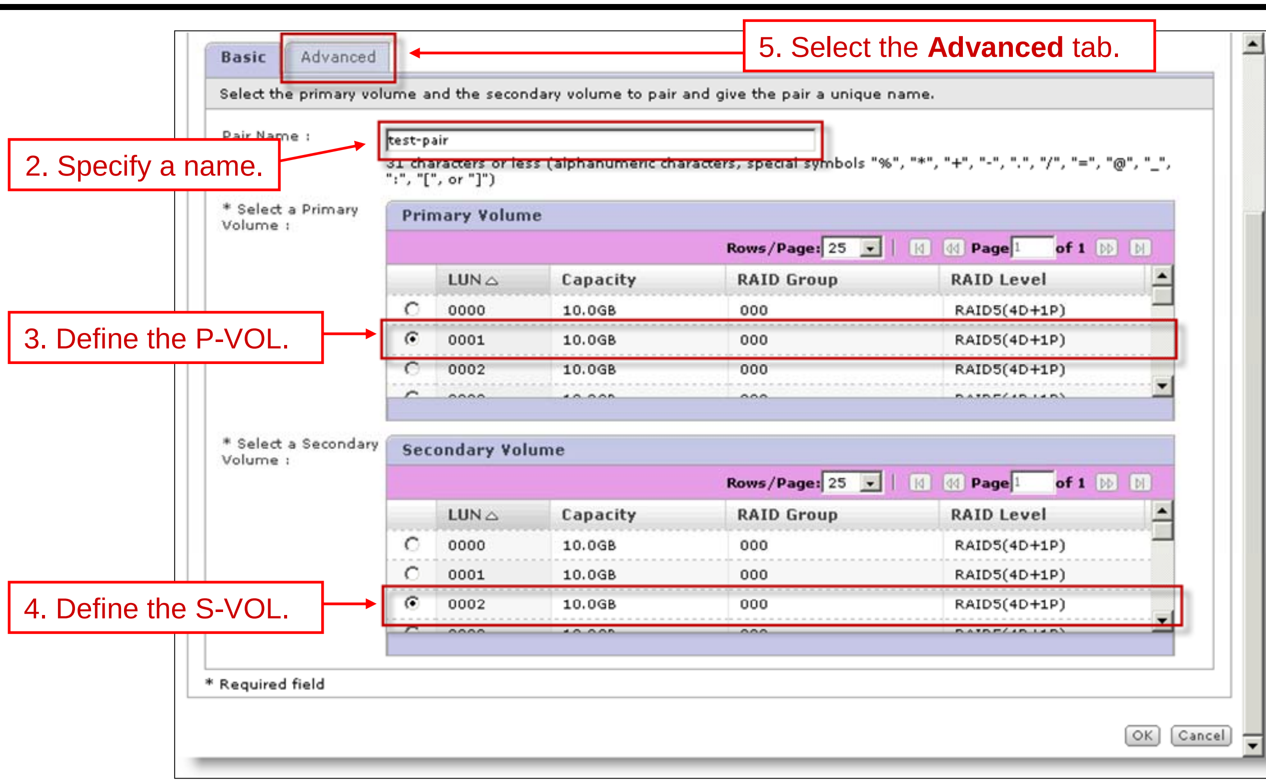

8

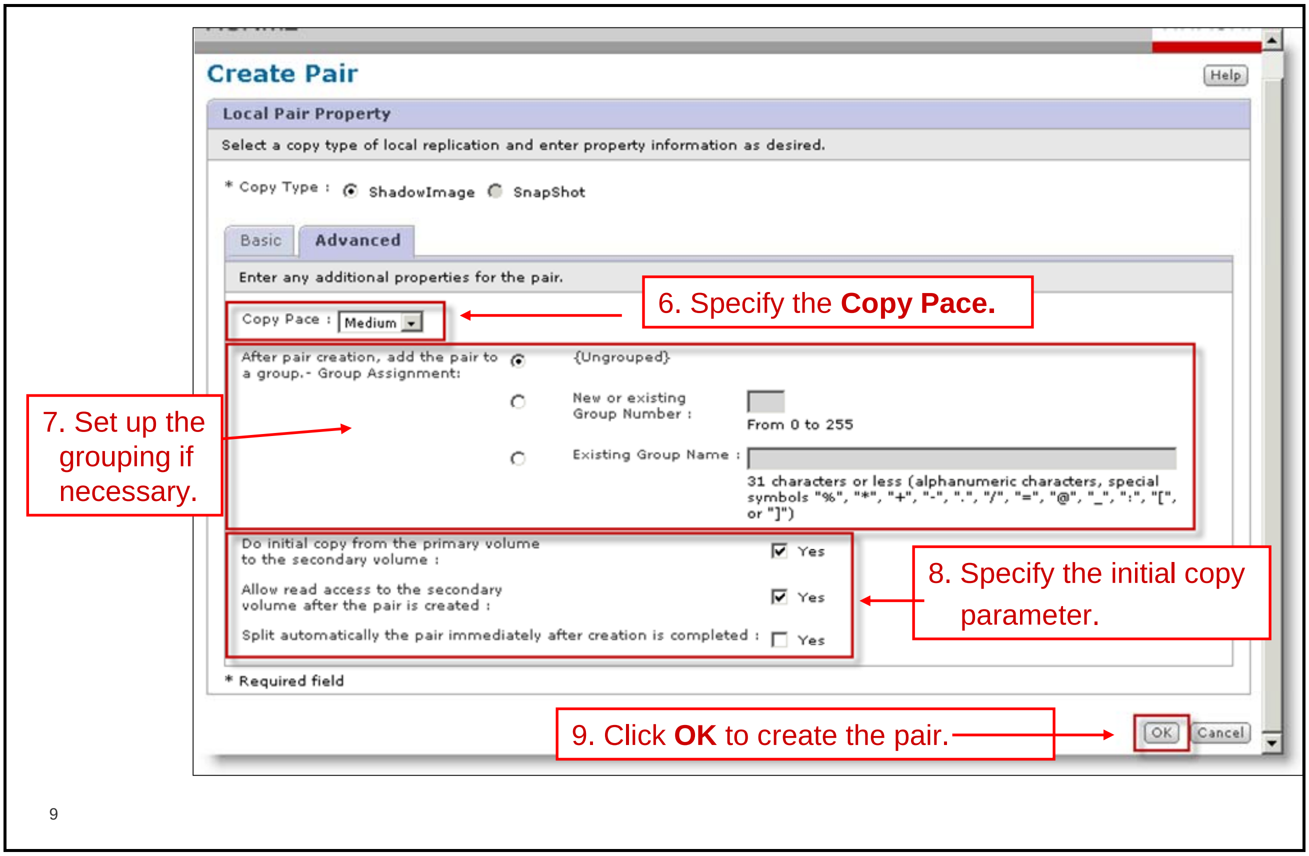
Create Pair
Help
Local Pair Property
Select a copy type of local replication and enter property information as desired.
* Copy Type : ShadowImage SnapShot
Basic
Advanced
Enter any additional properties for the pair.
Copy Pace : Medium
6. Specify the **Copy Pace.**
After pair creation, add the pair to a group.- Group Assignment:
{Ungrouped}
New or existing Group Number :
From 0 to 255
Existing Group Name :
31 characters or less (alphanumeric characters, special symbols "%", "*", "+", "-", ".", "/", "=", "@", "_", ":", "[", or "]")
7. Set up the grouping if necessary.
Do initial copy from the primary volume to the secondary volume : Yes
Allow read access to the secondary volume after the pair is created : Yes
Split automatically the pair immediately after creation is completed : Yes
8. Specify the initial copy parameter.
* Required field
9. Click **OK** to create the pair.
OK
Cancel
9

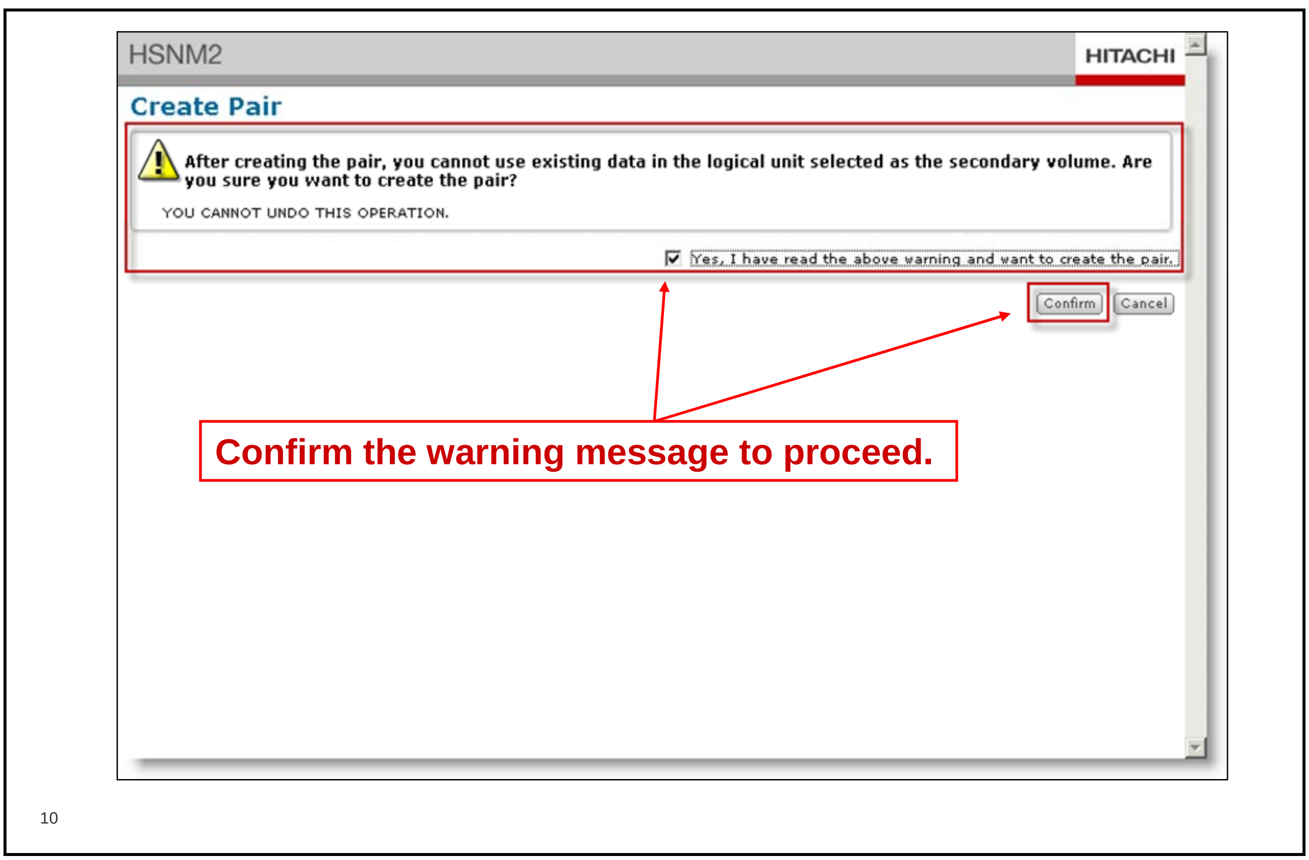
HSNM2
HITACHI
Create Pair
After creating the pair, you cannot use existing data in the logical unit selected as the secondary volume. Are you sure you want to create the pair?
YOU CANNOT UNDO THIS OPERATION.
Yes, I have read the above warning and want to create the pair.
Confirm
Cancel
Confirm the warning message to proceed.
10

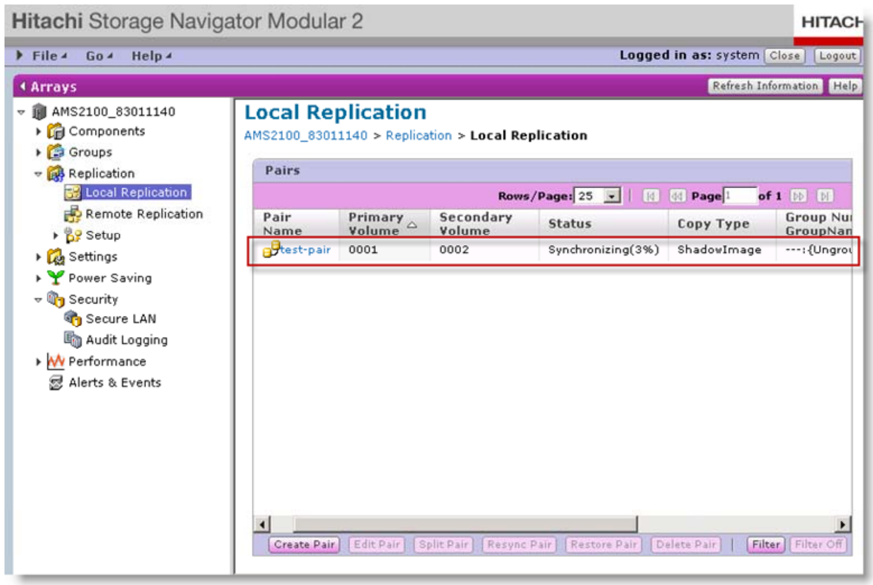

To create the pair with the SNM2 CLI, you can use the CLI command as follows:

aureplicationlocal -unit ‘unit_name’ -create -si -pvol 1 -svol 2 -pairname test_pair -pace normal

11

You will find more information on creating ShadowImage pairs using the SNM2 CLI in the *Hitachi Navigator 2 Command Line Interface (CLI) Reference Guide for Replication* on page 1-9

View the ShadowImage Pair Status with SNM2

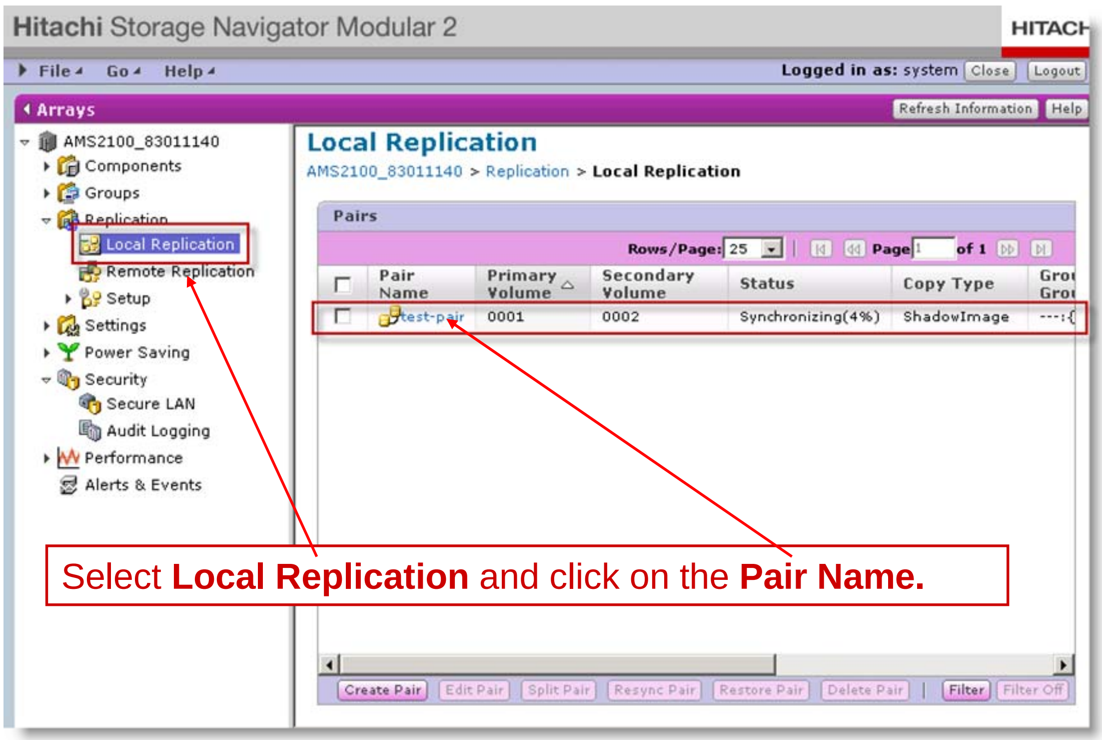

12

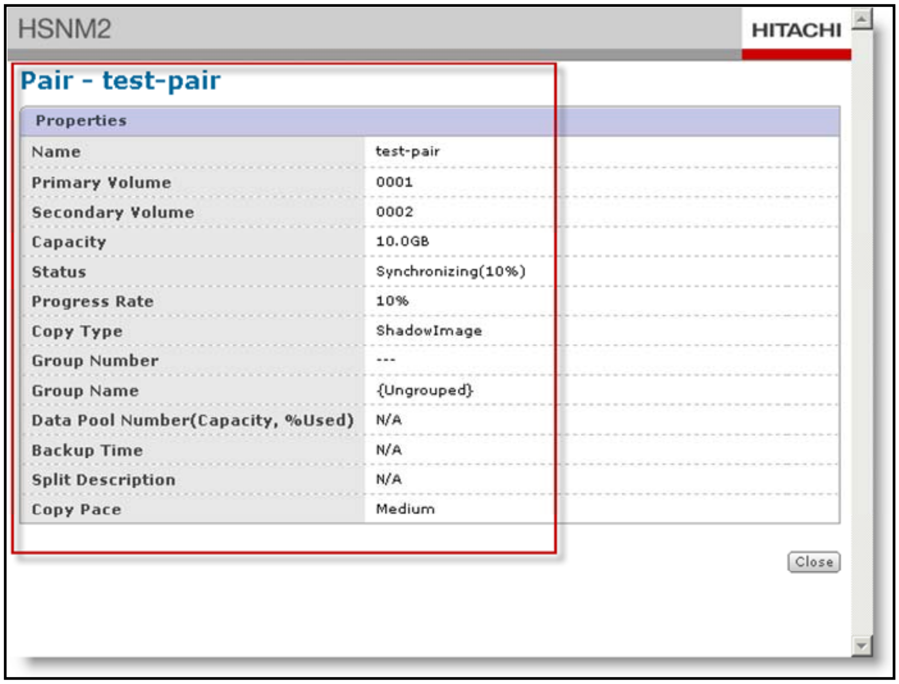

To view the pair with the SNM2 CLI, you can use the CLI command as follows:

aureplicationlocal -unit ‘unit_name’ -refer -detail -pvol 1 -svol 2

13

You will find more information on viewing ShadowImage pairs using the SNM2 CLI in the *Hitachi Navigator 2 Command Line Interface (CLI) Reference Guide for Replication* on page 1-2

Split ShadowImage Pairs with SNM2

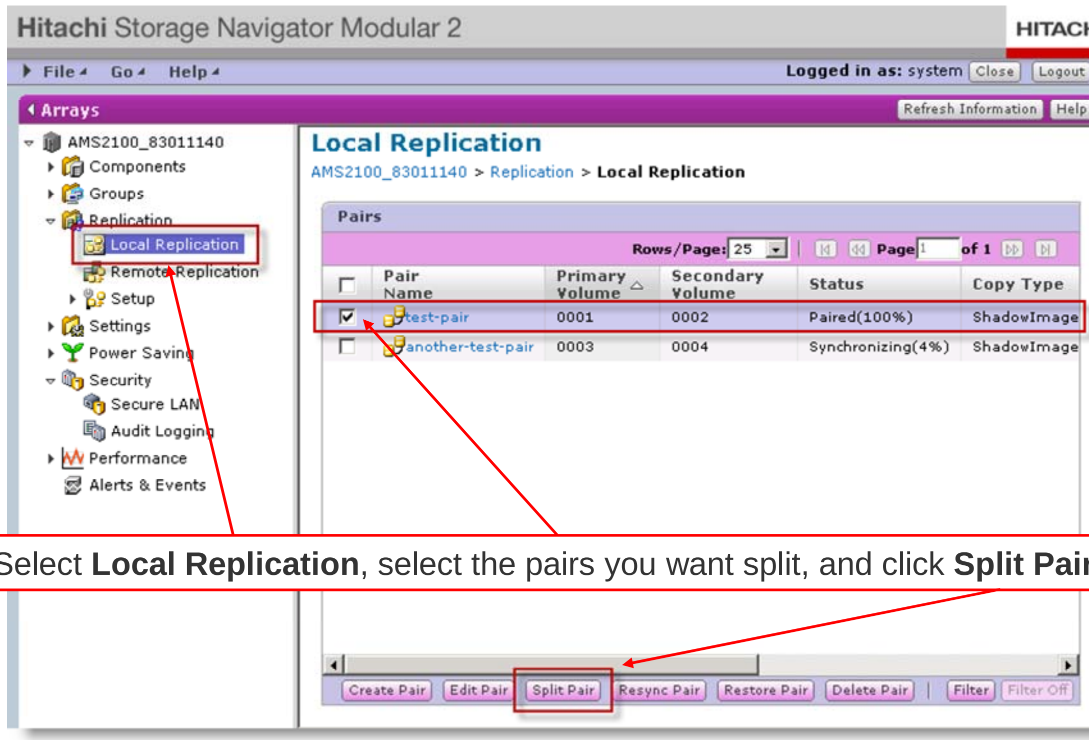

1. Select **Local Replication**, select the pairs you want split, and click **Split Pair.**

14

2. Set the options and click **OK** to proceed.

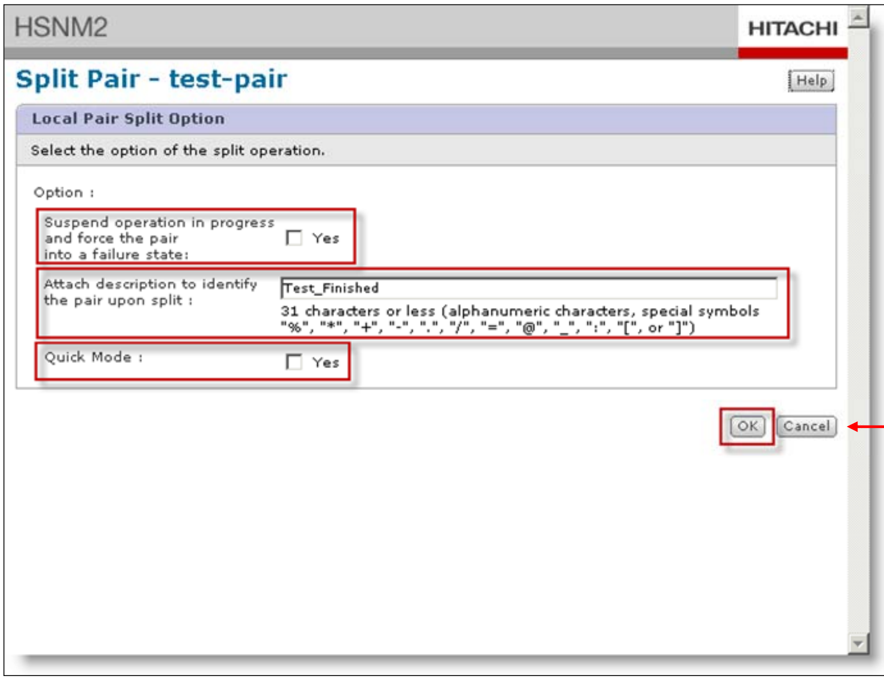

15

- The detailed pair information after the split operation

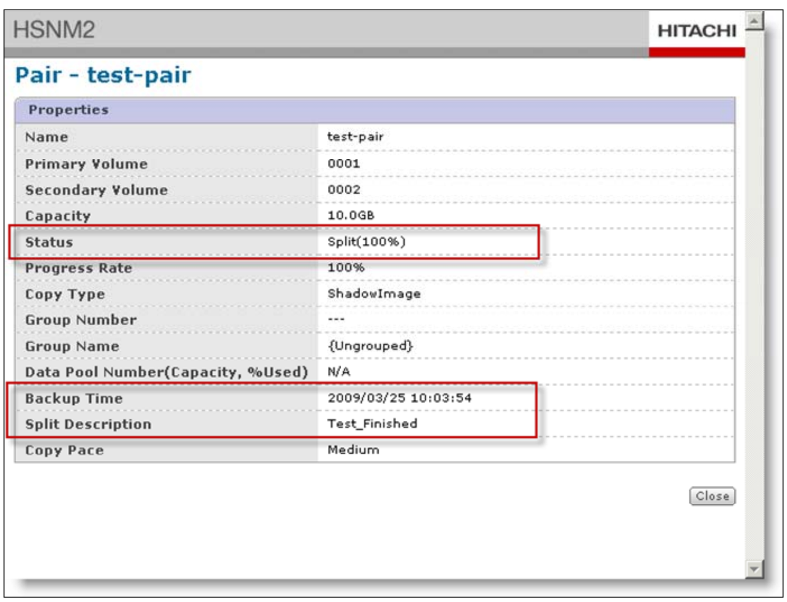

To split the pair with the SNM2 CLI, you can use the CLI command as follows:
aureplicationlocal -unit 'unit_name' -split -si -pvol 1 -svol 2

You will find more information on splitting ShadowImage pairs using the SNM2 CLI in the *Hitachi Navigator 2 Command Line Interface (CLI) Reference Guide for Replication* on page 1-12

Resynchronize ShadowImage Pairs with SNM2

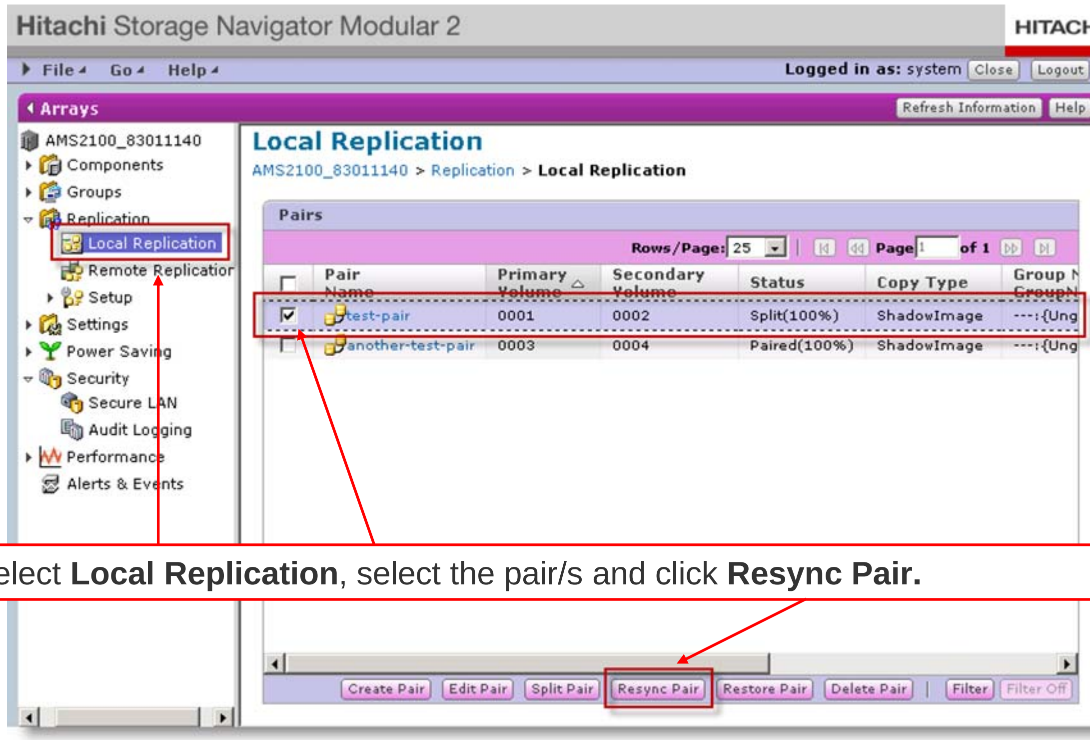

1. Select **Local Replication**, select the pair/s and click **Resync Pair.**

17

2. Specify the **Quick Mode** settings and click **OK** to proceed.

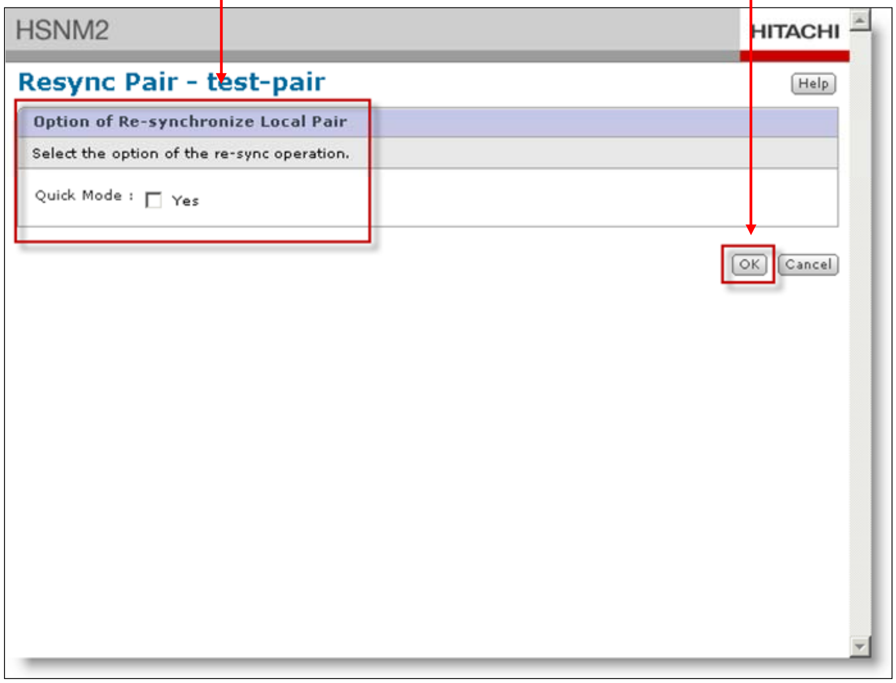

18

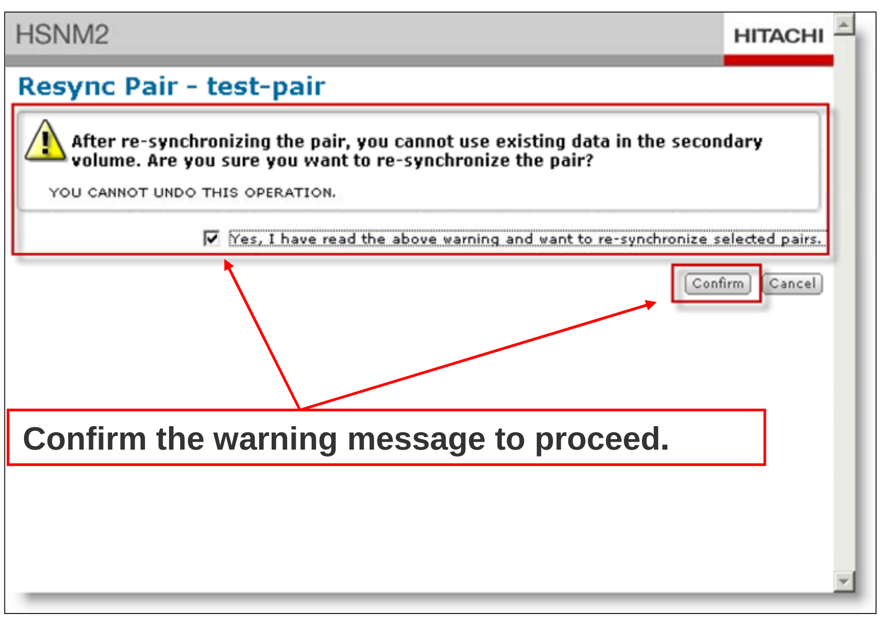

19

- The detailed pair information after the resync operation

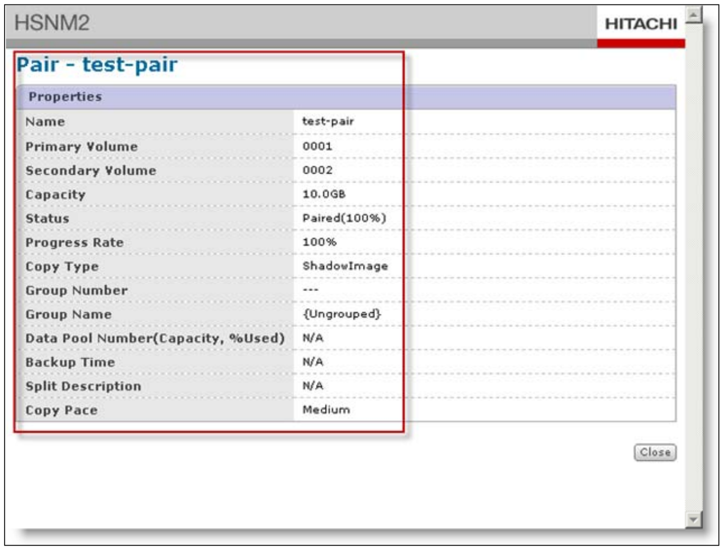

To resynchronize the pair with the SNM2 CLI, you can use the CLI command as follows:

aureplicationlocal -unit ‘unit_name’ -resync -si -pvol 1 -svol 2

20

You will find more information on re-synchronizing ShadowImage pairs using the SNM2 CLI in the *Hitachi Navigator 2 Command Line Interface (CLI) Reference Guide for Replication* on page 1-15

Restore ShadowImage Pairs with SNM2

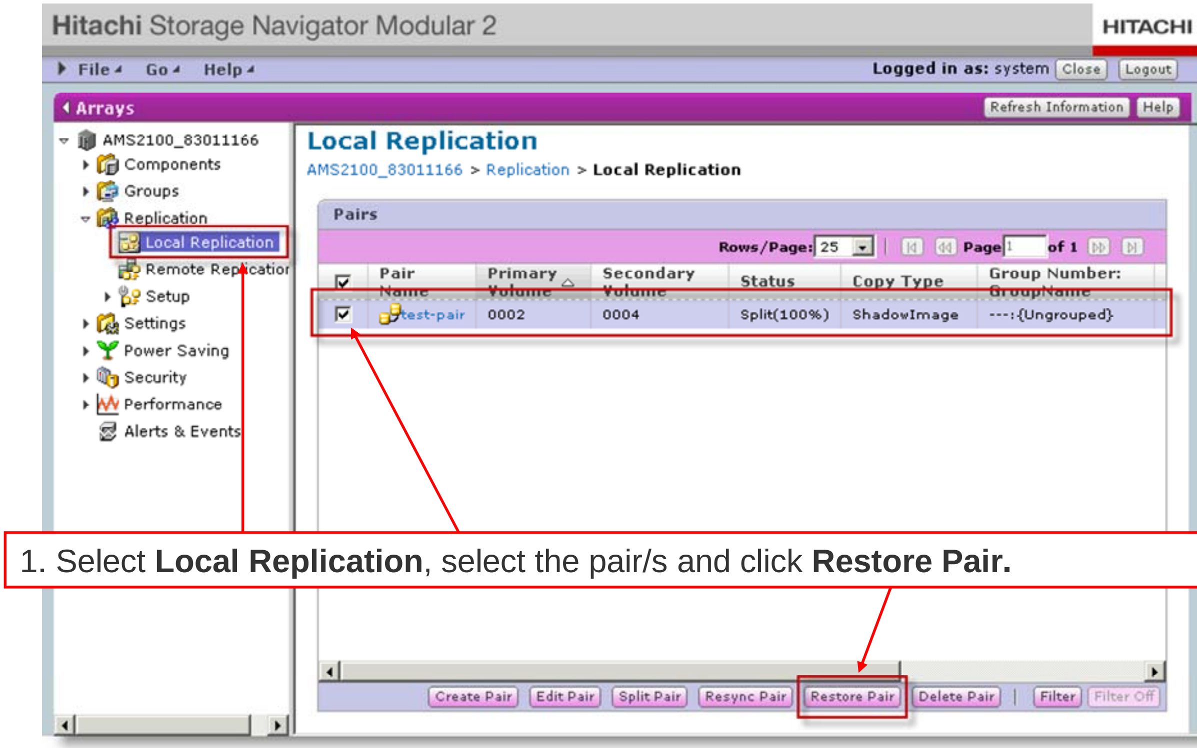

21

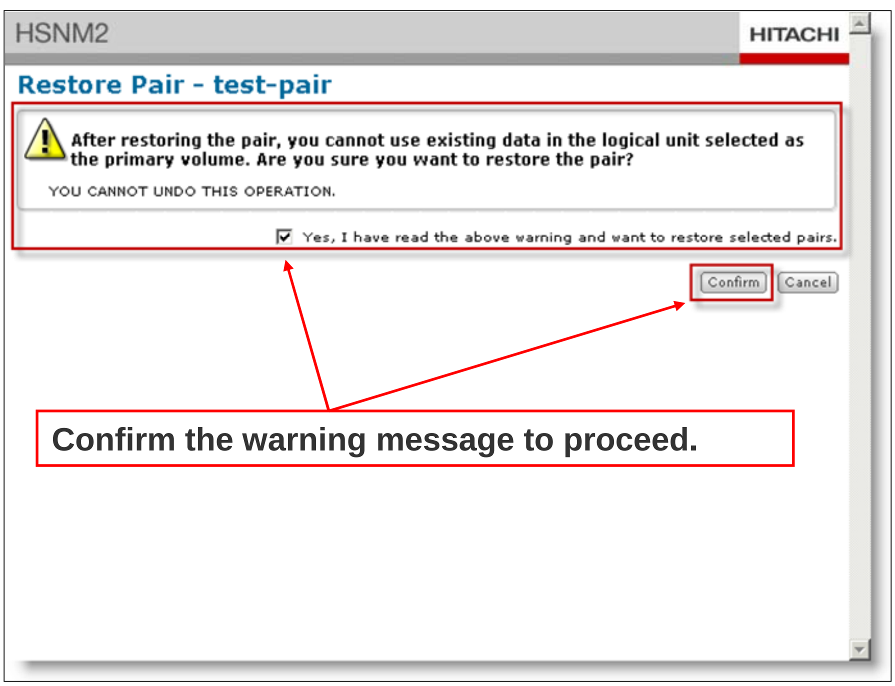

22

- The detailed pair information after the restore operation

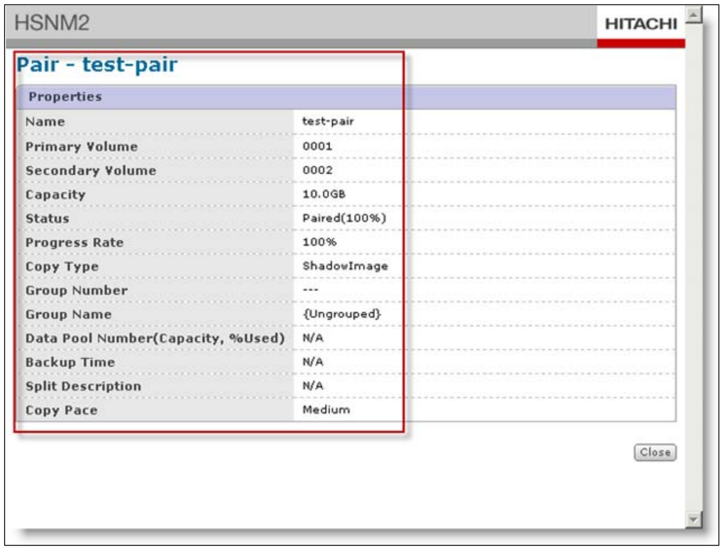

To restore the pair with the SNM2 CLI, you can use the CLI command as follows:
aureplicationlocal -unit ‘unit_name’ -restore -si -pvol 1 -svol 2

23

You will find more information on restoring ShadowImage pairs using the SNM2 CLI in the *Hitachi Navigator 2 Command Line Interface (CLI) Reference Guide for Replication* on page 1-17

Delete ShadowImage Pairs with SNM2

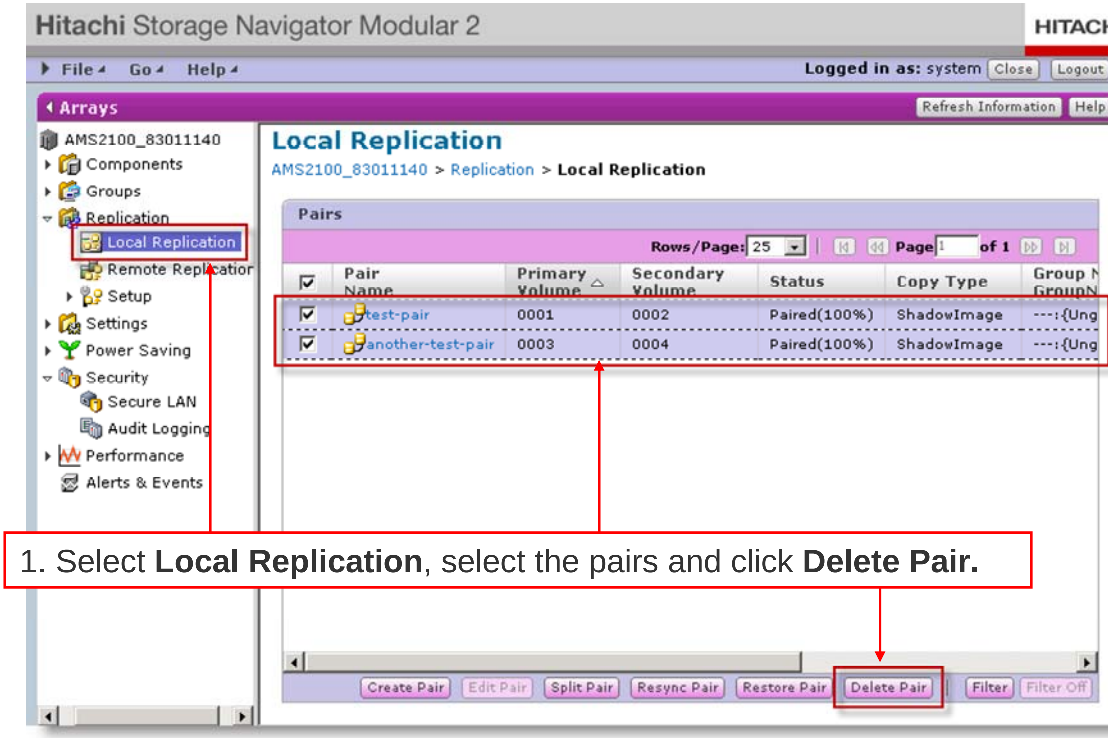

24

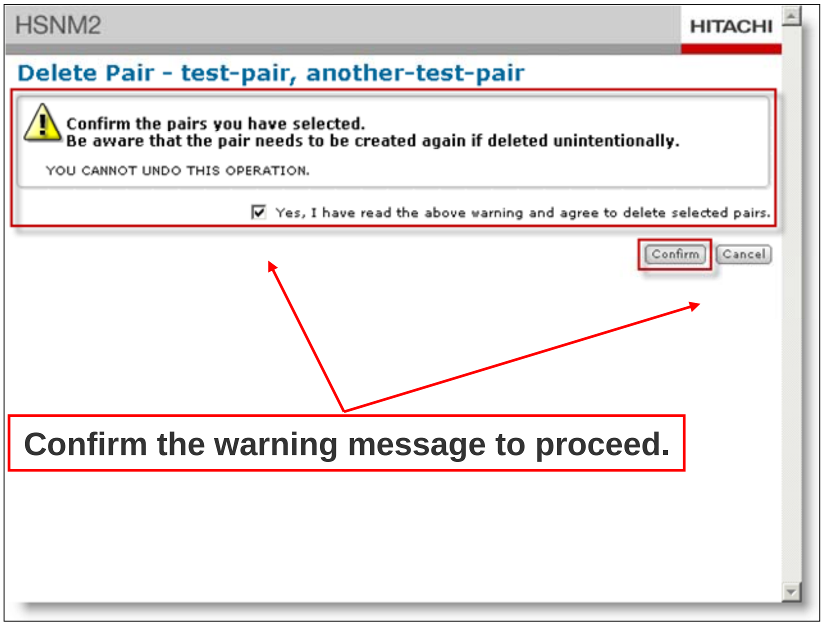

To delete a pair with the SNM2 CLI, you can use the CLI command as follows:

aureplicationlocal -unit 'unit_name' -simplex -si -pvol 1 -svol 2

25

You will find more information on deleting ShadowImage pairs using the SNM2 CLI in the *Hitachi Navigator 2 Command Line Interface (CLI) Reference Guide for Replication* on page 1-19

6. Hitachi Copy-on-Write Snapshot Software Introduction

Module Objectives

- Upon completion of this module, the learner should be able to:
  - Describe the purpose of Hitachi Copy-on-Write Snapshot software
  - List key Copy-on-Write Snapshot specifications
  - Compare the functionality of Copy-on-Write Snapshot software to Hitachi ShadowImage Replication software
  - Describe typical Copy-on-Write Snapshot operations

2

Overview

- Copy-on-Write Snapshot software is a storage-based hardware solution for point-in-time backups, or snapshots.
- A volume by Copy-on-Write Snapshot consists of physical data stored in the primary volume and differential data stored in the Data Pool.
- The Copy-on-Write Snapshot primary volumes (P-VOLs) contain the original data and the snapshot images (V-VOLs) contain the snapshot data.
- As each P-VOL is paired with its V-VOL independently, each volume can be maintained as an independent copy set that can be created (paircreate), given the Copy-on-Write Snapshot instruction (pairsplit) and released (pairsplit –S) separately.

Copy-on-Write Snapshot software internally retains a logical duplication of the primary volume data at the time of command instruction. This software is used for restoration; it stores data at the time of Snapshot instruction when a logical error occurs in the primary volume.

The duplicated volume of the Copy-on-Write Snapshot function consists of physical data stored in the primary volume and differential data stored in the data pool. This differs from the ShadowImage function, where all data is retained in the secondary volume. Although the capacity of the used data pool is smaller than that of the primary volume, a duplicated volume can be created logically when the Snapshot instruction is given. The data pool can share two or more primary volumes and the differential data of two or more duplicated volumes.

Definitions:

Data Pool: A volume that contains differential data only.

Snapshot Image: A logical duplicated volume (V-VOL) of the primary volume. It is an internal volume meant for restoration.

Why use Copy-on-Write Snapshot?

- Because other midrange disk array product vendors have already supported Snapshot.
- Proposing ShadowImage software for backup solutions raises the entire system price, causing tough competition.
- Due to large increases in storage capacity, there is more need for Snapshot function that is comprised of differential data only and is therefore small in capacity.

Terminology

- Command Device (CMD-DEV)
  - Dedicated logical volume that is used only by management software such as Hitachi CCI, to interface with the storage systems. Command Devices are not used by ordinary applications. Command Devices can be shared between several hosts.
- Differential Management Logical Unit (DM-LU)
  - Volume used to manage differential data in a storage system. For replication software, the DM-LU is an exclusive volume used for storing data when the array system is powered down.
- Primary Volume (P-VOL)
  - The storage volume in a volume pair. It is used as the source of a copy operation.
- Virtual Volume (V-VOL)
  - In Copy-on-Write, a secondary volume in which a view of the primary volume (P-VOL) is maintained as it existed at the time of the last snapshot. The V-VOL contains no data but is composed of pointers to data in the P-VOL and the data pool. The V-VOL appears as a full volume copy to any secondary host.
- Data Pool (Pool)
  - One or more disk volumes designated to store the differential data of the V-VOLs.

4

Overview

- Copy-on-Write Snapshot software is used to restore data during the Snapshot instruction, if a logical error occurs in the primary volume.
- Up to 32 V-VOLs (snapshot images) can be created per primary volume and manage data in two or more generations within the disk system.
- Instructions for pair creation or pair splitting can be issued using the Hitachi Storage Navigator Modular 2 GUI, CLI, or CCI.

5

Comparison with ShadowImage Functions

ShadowImage

All data is saved from Primary Volume (P-VOL) to Secondary Volume (S-VOL).

Copy-on-Write Snapshot

Only differential data is saved from Primary Volume (P-VOL) to Data Pool area (Pool).

Pool is shared by multiple Snapshot Images (V-VOL).

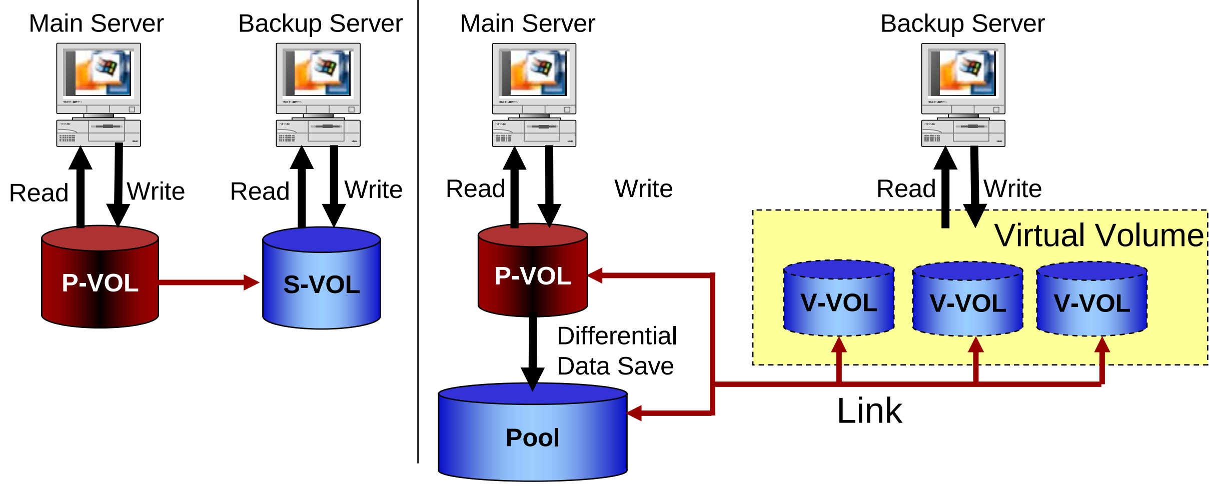

6

ShadowImage Replication Software

- ShadowImage software for Hitachi modular storage is a storage-based hardware solution for duplicating logical volumes that reduces backup time and provides point-in-time backup.
- Primary volumes (P-VOLs) contain original data; the secondary volumes (S-VOLs) contain duplicate data. Since each P-VOL is paired with its S-VOL independently, each volume can be maintained as an independent copy set that can be split (pairsplit), resynchronized (pairresync) and released (pairsplit –S) separately.

Copy-on-Write Snapshot software

- Copy-on-Write Snapshot for Hitachi modular storage is a storage-based hardware solution for duplicating logical volumes that reduces backup time and provides point-in-time backup. The Copy-on-Write Snapshot primary volumes (P-VOLs) contain the original data; Snapshot Images (V-VOLs) contain the Snapshot data. Since each P-VOL is paired with its V-VOL independently, each volume can be maintained as an independent copy set that can be created (paircreate), given the Snapshot instruction (pairsplit) and released (pairsplit –S) separately.

- Each Copy-on-Write Snapshot pair consists of one primary volume (P-VOL) and up to 32 Snapshot Images (V-VOLs), which are located in the same storage system. The Copy-on-Write Snapshot P-VOLs are the primary volumes, which contain the original data. The Copy-on-Write Snapshot V-VOLs are duplicated volumes that contain the data that exists at the time of Snapshot instruction.
- The Copy-on-Write Snapshot Image (V-VOL) contains physical data from the primary volume and differential data stored in the Data Pool (Pool). (This differs from the ShadowImage software where all the data is retained in the secondary volume only; it is actually a pseudo volume with no capacity.)

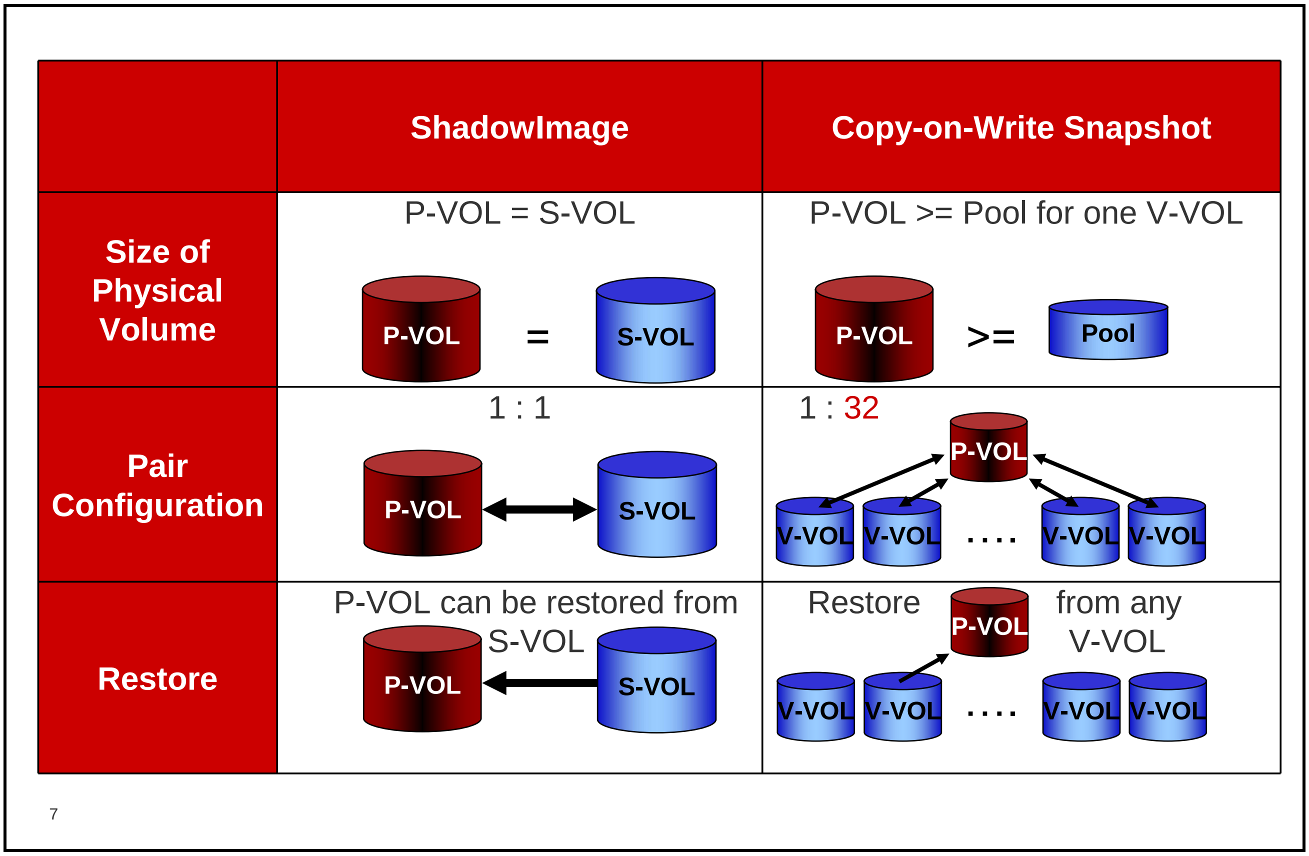

Size of Physical Volume:

The P-VOL and the S-VOL are the same size in ShadowImage software. In contrast, less disk space is required for building a V-VOL image since only part of the V-VOL is on the Pool and the rest is still on the primary volume in Copy-on-Write Snapshot.

Pair Configuration:

Only one S-VOL can be created for every P-VOL in ShadowImage software. In Copy-on-Write Snapshot software, there can be up to 32 V-VOLs per primary volume.

Restore:

A primary volume can only be restored from the corresponding secondary volume in ShadowImage software. With Copy-on-Write Snapshot software, the primary volume can be restored from any Snapshot Image (V-VOL).

- Copy-on-Write Snapshot volume size
  - The capacity needed for one Snapshot Image (V-VOL) the Data Pool area (Pool) is smaller than the P-VOL.

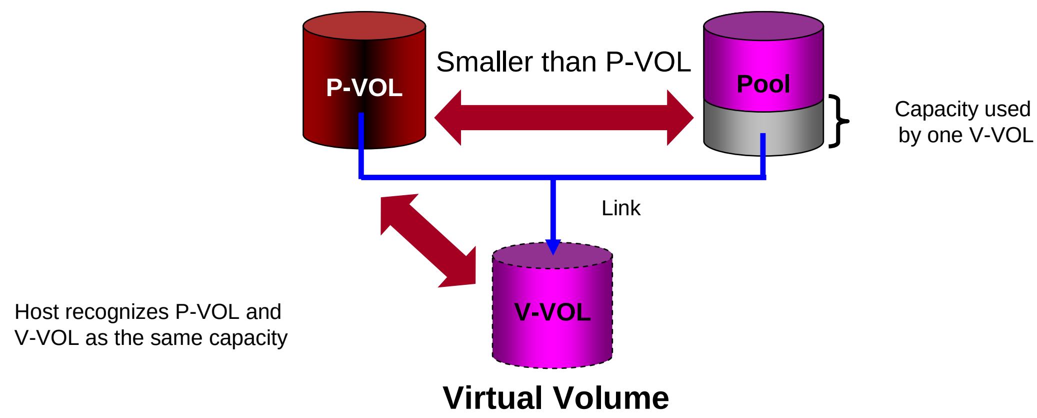

8

In Copy-on-Write Snapshot, the **V-VOL** does not physically exist but is a representation of a **set of pointers** that point to the locations where the data is physically located, partly in the Pool and partly in the P-VOL. Since only part of the data belonging to the V-VOL is located in Pool (and the other part is still on P-VOL), Copy-on-Write Snapshot software does not require twice the disk space to establish a pair, as in ShadowImage software.

However, a host will recognize the P-VOL and the V-VOL as a pair of volumes with identical capacity.

Operation Scenarios

- Before modifying the Data Block on P-VOL
  - V-VOL will get the data from P-VOL

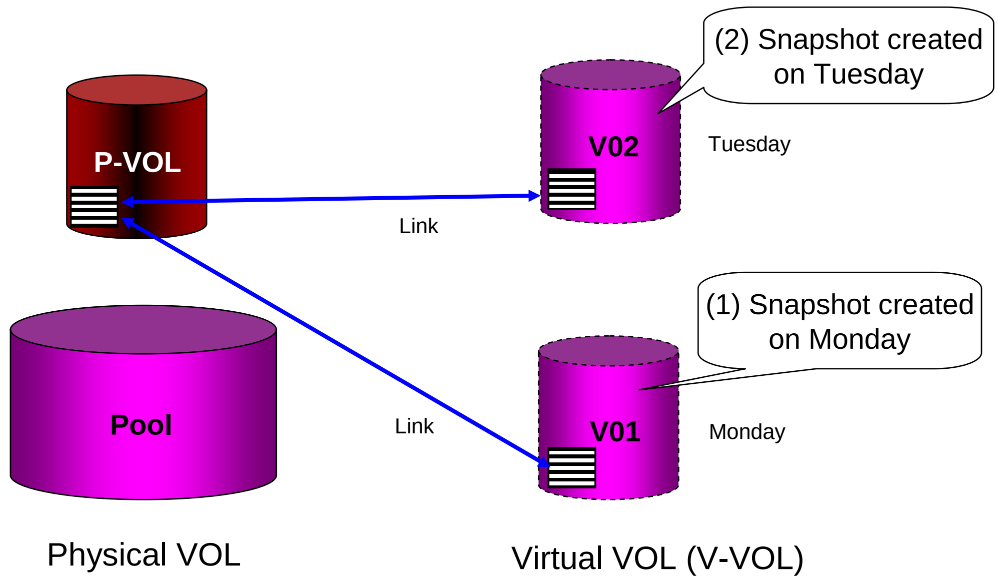

9

This picture shows a situation where two Snapshots have been taken. The highlighted data block in the Snapshots is also available on the primary volume. Requests for this block through the V-VOL would be physically taken from the P-VOL.

This situation will last as long as the corresponding block on the P-VOL is not altered.

- Before writing to the Data Block on P-VOL
  - When there is a write after having created a Snapshot, *old* data is saved on the Data Pool area (Pool) first.
  - This saved data in the Pool is now used by the V-VOLs.
  - Since Pool data can be shared by multiple V-VOLs, only one copy of the data is required.

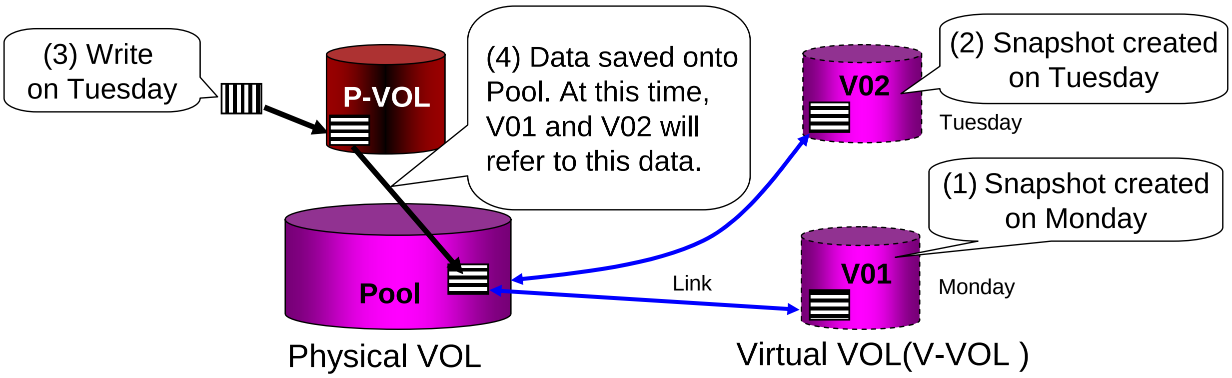

In order to link Pool and Snapshot Images, **Addresses and Generations** are managed by cache.

10

The data block on the P-VOL needs to be written to, but before the write is executed the block is copied to the Pool area. The set of pointers that represent the V-VOL are updated and all requests through a V-VOL for the original block are physically taken from the Pool.

From the host's perspective, the V-VOL (Snapshot Image) has not changed.

- On Wednesday a new snapshot image is created.
  - The new Snapshot Image will refer to the data on P-VOL.

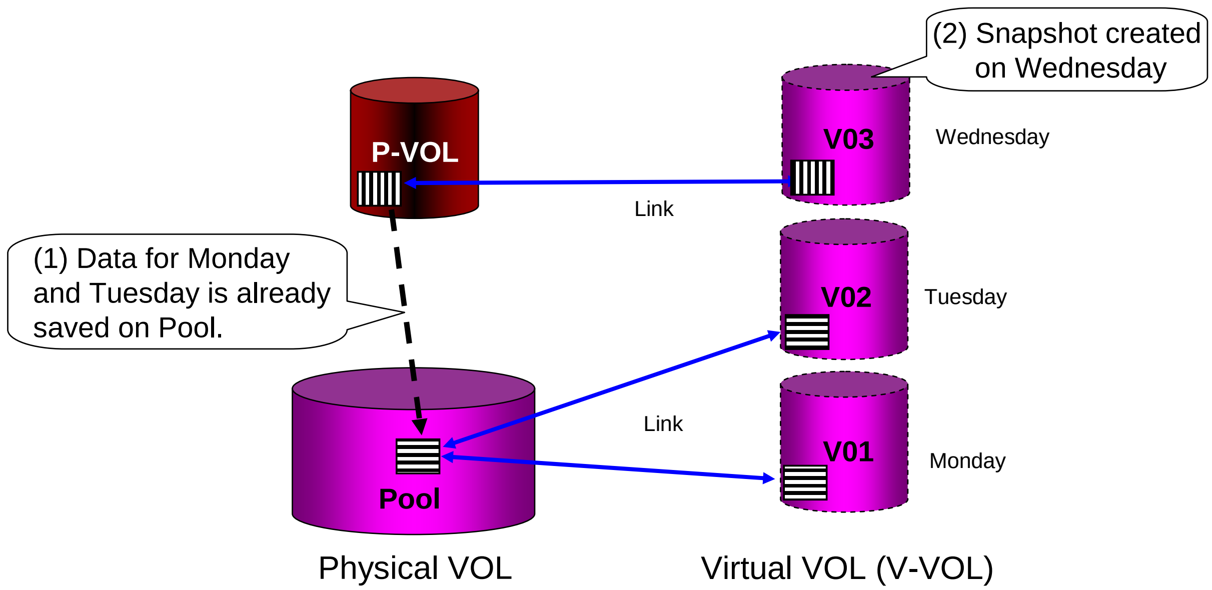

11

On Wednesday, a new Snapshot Image was created. The data block before the write will be taken from the Pool area while the block after the write will be taken from the primary volume (P-VOL).

If there is a request for that block through a V-VOL, the data will be read from the Pool area or from the P-VOL depending on what Snapshot Image is being referenced.

- Write command to the same Data Block again
  - When there is another write to the same area, the data is additionally saved on the Data Pool area (Pool) prior to writing.
    - The V-VOL (that was linked to the P-VOL before the write) views the data saved in the Pool.

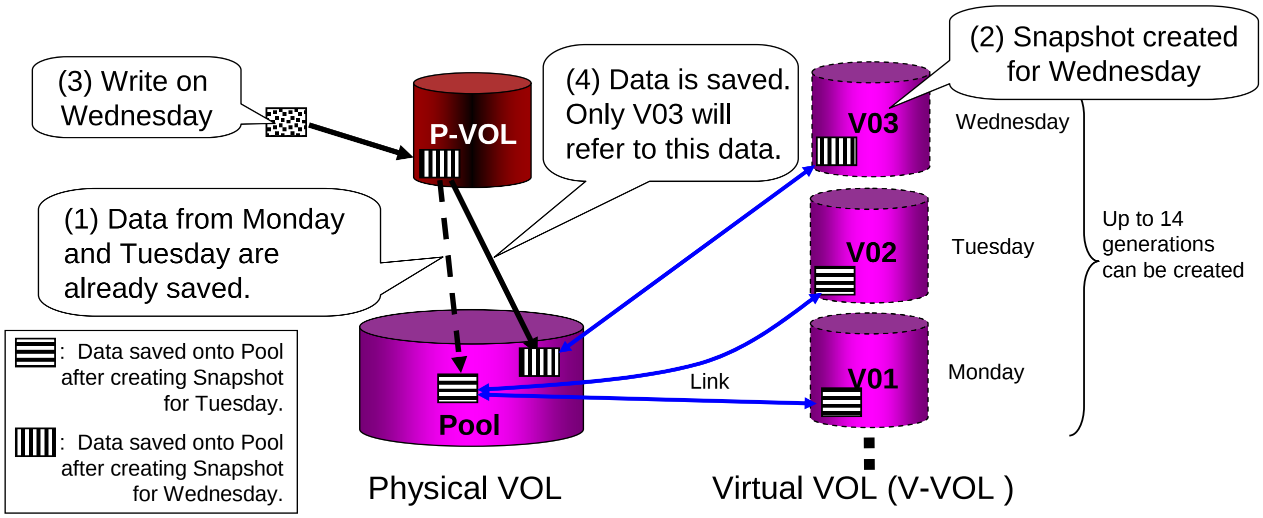

12

An additional write to the same data block on the P-VOL. Before executing the write the block is copied to the Pool area, the pointers that make up the V-VOL are updated. Upon a request for that data block through a V-VOL, the data will be taken from the Pool.

- Restore is possible from any snapshot image (V-VOL)
  - Restores appear to be done instantly to the host
    - Data copying from V-VOL to P-VOL is done in background

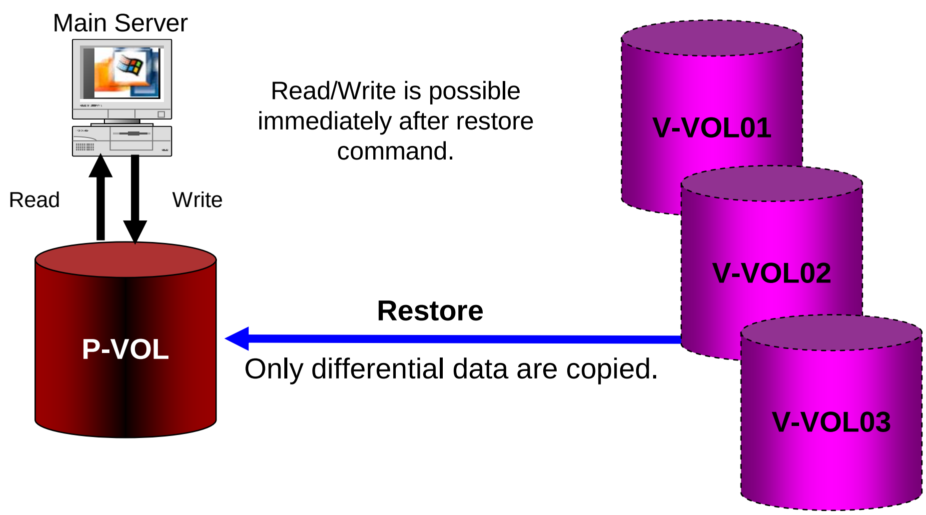

13

Restoration of a primary volume can be done instantly from any V-VOL because it does not involve immediate moving of data from Pool to P-VOL only pointers need to be modified.

The background data will then be copied from Pool to P-VOL.

If the P-VOL became physically damaged, all V-VOLs would be destroyed and a restore is not possible.

Components

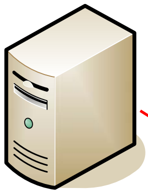

Management Server with SNM2

- Create Pairs
- Split Pairs
- Resynchronize Pairs
- Delete Pairs
- Manages the Pool
- No CMD-DEV necessary
- Communicates via LAN - IP

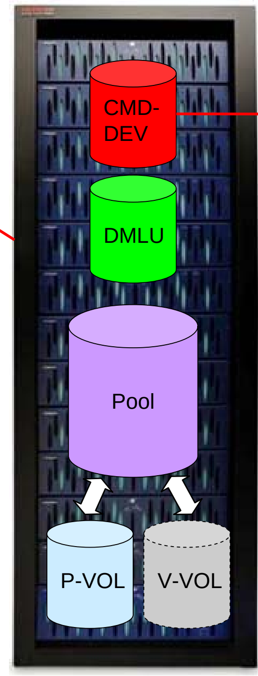

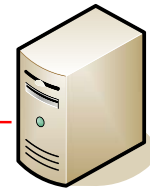

Management Server with CCI

- Create Pairs
- Split Pairs
- Resynchronize Pairs
- Delete Pairs
- CMD-DEV necessary
- Communicates through Fibre Channel

14

Copy-on-Write Snapshot operations involve the P-VOL, Snapshot images (V-VOL), data pool (POOL) in the disk system, the Hitachi Storage Navigator Modular 2 and/or the CCI UNIX and/or Microsoft Windows PC or server.

Volume Grouping

- Perform operations for pairs or a group of pairs

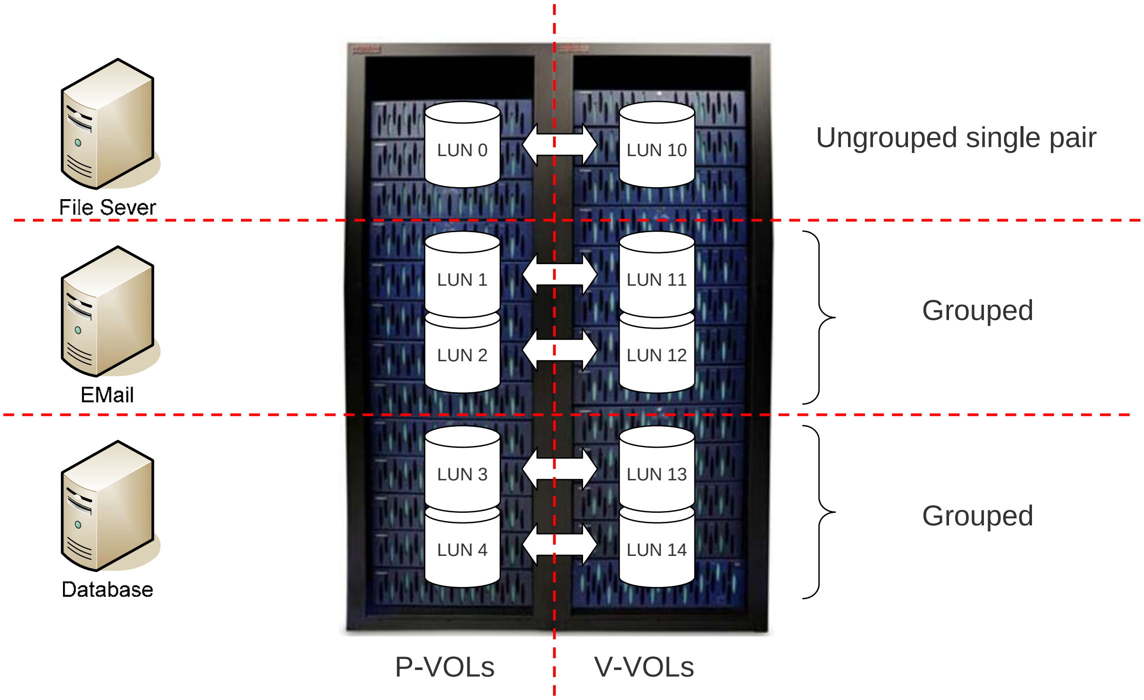

15

You can define or set up Copy-on-Write Snapshot pairs in groups which enable you to issue commands or do operations for single pairs or a group of pairs.

Status Transitions

- Status and commands for Copy-on-Write Snapshot

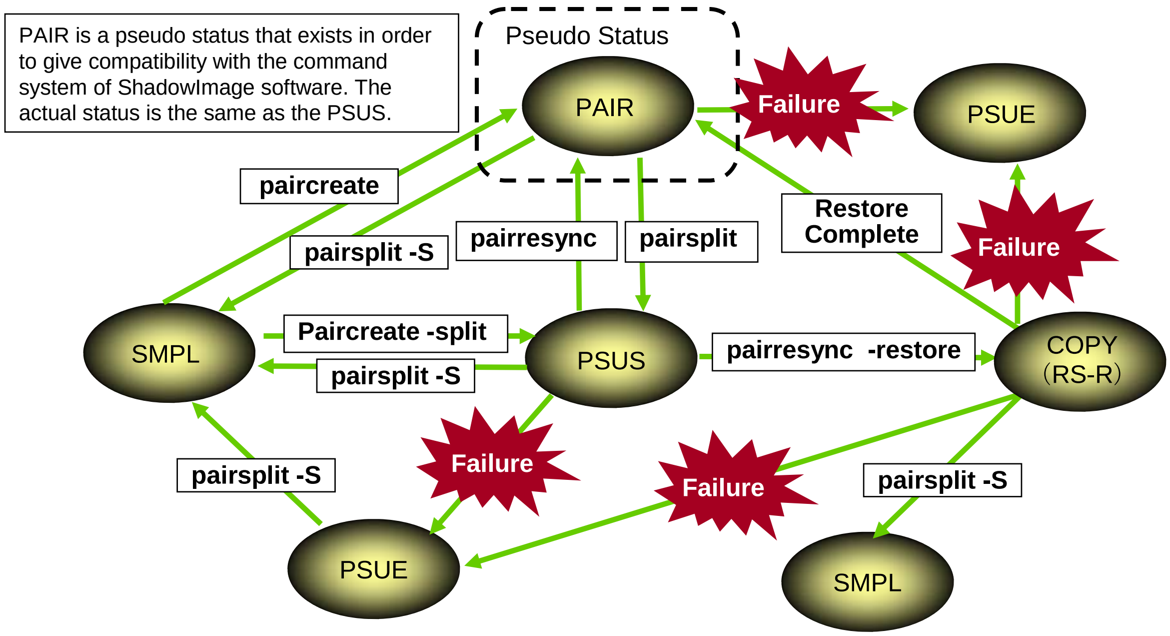

16

If a volume is not assigned to a Copy-on-Write Snapshot pair, its status is SMPL. When a Copy-on-Write Snapshot pair is created by executing the **paircreate –split** or **pairsplit** operation after **paircreate** has been issued, the 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 failure. When the **-S** option is specified by executing **pairsplit**, the pair is split and the pair status changes to SMPL.

Specifications

No	Item	Adaptable Modular Storage 2000 Family
1	User interface of pair operation	Storage Navigator Modular 2 (GUI,CLI) CCI
2	Controller configuration	Dual controllers
3	Maximum number of pairs	Model 2100: **1022** Models 2300 & 2500: **2046**
4	P-VOL: V-VOL	1:32
5	RAID level	RAID-6 , RAID-5, RAID-1, RAID-1+0
6	RAID level of P-VOL to V-VOL	Any to any
7	RAID Group	It is recommended that P-VOL and POOL reside in different RAID Groups.
8	Drive type for a P-VOL/V-VOL	Any
9	V-VOL definition	Do not have to specify P-VOL
10	Split by Consistency Groups	Supported

17

No	Item	Adaptable Modular Storage 2000 Family
11	SNMP Trap	• PSUE • POOL threshold over • POOL over
12	Pair operation and format • Pair operation during format • Format during pair operation	Not supported
13	Combination with ShadowImage pair (See ShadowImage section)	Not supported
14	Combination with another SnapShot pair • V-VOL of a Snapshot pair is P-VOL of another SnapShot pair	Not supported
15	Combination with TrueCopy pair (See TrueCopy section)	Supported (including Swap-resync)
16	Combination with TrueCopy Extended pair (See TrueCopy Extended section)	Supported

18

No	Item	Adaptable Modular Storage 2000 Family
17	Maximum number of POOLs	64/storage system
18	Maximum number of LUs for POOL	128/storage system
19	Maximum Pool Capacity	Model 2100: **6.2**TB Model 2300: **12**TB Model 2500: **24**TB

19

Requirements

- Valid and enabled license
- DM-LU
- CMD-DEV connected to the management server when using CCI to operate Copy-on-Write Snapshot software

20

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

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

Module Review

1. What is Copy-on-Write Snapshot?
2. Why is a Data Pool area required?
3. What is a V-VOL?
4. How is a V-VOL different from a ShadowImage S-VOL?

22

7. Hitachi Copy-on-Write Snapshot Software with Command Control Interface

Module Objectives

- Upon completion of this module, the learner should be able to:
  - Describe the use and components of Hitachi Command Control Interface (CCI) software
  - Set up the CCI for Hitachi Copy-on-Write Snapshot software operations
  - List the CCI commands
  - Describe the web tool monitoring of Copy-on-Write Snapshot

2

Installation

- Hitachi Storage Navigator Modular 2 is required.
- Copy-on-Write Snapshot software license must be installed.
- Command Device and LUNs for P-VOL and Pool must be configured.
- DM-LU must be configured.
- If required, additional snapshot images may be created.

3

CCI is sometimes called RAID Manager.

Overview of CCI

- CCI is the control mechanism for day-to-day handling of the Copy-on-Write Snapshot operation on Hitachi modular storage from a host computer.
- It is used for creating, suspending, and splitting Copy-on-Write Snapshot pairs.
  - These tasks can be automated through the execution of scripts.
- CCI is supported on various platforms.

4

- Components
  - CCI client is loaded onto the OS server.
    - CCI is supported by the following platforms:
      - Sun Solaris, HP-UX, AIX, Linux and Tru64 UNIX, and IRIX
      - Microsoft Windows Server 2000 and 2003 and Microsoft Windows NT
    - The machine on which CCI is installed must have an active Fibre Channel connection to the command devices on the storage system.
  - Copy-on-Write Snapshot key is installed on the storage system.
    - Using Storage Navigator Modular 2
    - Specifying a command device

5

- Components (*continued*)
  - CCI Configuration Files
    - Environmental variables include:
      - /etc/services (UNIX)
      - \WINNT\system32\drivers\etc\services (NT)
    - horcm0.conf and horcm1.conf
      - /etc Directory (UNIX)
      - \HORCM\etc and \WINNT Directory (NT)
  - Storage System Command Device
    - The command device is a host attached LUN that facilitates communication between CCI and storage system

The services file describes the TCP/IP protocol type (udp) and port number

- horcm0.conf describes the environment for the primary volume P-VOL
- horcm1.conf describes the environment for the secondary volume V-VOL

Configuration

- CCI configuration files - /etc/services
  - Environmental variables
    - /etc/services (UNIX)
    - \WINNT\system32\drivers\etc\services (NT)
      - horcm0 11000/udp #HORCM instance 0 port
      - horcm1 11001/udp #HORCM instance 1 port

Any number between 5000 and 65535. Do not duplicate previously assigned ports.

udp - User Datagram Protocol.
A connection-less method of data packet delivery. The application is responsible for data transmission monitoring, error detection and retransmission request.

7

Note: Ensure that the port number is not already assigned to another service.

The words horcm0 and horcm1 are linked to and should be the same as the names defined in the horcm0.conf and horcm1.conf files.

- CCI configuration file - horcm0.conf
  - A text file located in the:
    - /etc Directory (UNIX)
    - \HORCM\etc and \WINNT Directories (NT)
  - Provides a definition of hosts, groups, volumes and command device to the CCI instance
  - Contains four areas that define the CCI environment
    - HORCM_MON
    - HORCM_CMD
    - HORCM_DEV
    - HORCM_INST

8

horcm0.conf describes the operating environment for the primary volume P-VOL. horcm1.conf does the same for the secondary volume V-VOL.

- CCI horcm.conf
  - HORCM_MON – Defines monitoring and communicating parameters. The host IP address, service name, poll rate, and timeout value are all set in this section.
  - HORCM_CMD – Defines the command devices on the Thunder 9500 V Series systems. More than one command device can be used to provide failover functionality should the original command device become inoperable.
  - HORCM_DEV – Defines the device group (names a group of paired logical volumes), the device names (names the paired volume in the device group), the port ID, target ID and LUN are set in the DEV section of the configuration file. Pairs are group and device name dependent.
  - HORCM_INST – Defines the network settings that correspond to the defined group names in the DEV section on the configuration file. The group name, IP address of the other host and the function name are set in the INST section of the configuration.

9

- CCI port conversion – Adaptable Modular Storage (AMS) 2000 family
  - Controller 0
    - 0A = CL1-A
    - 0B = CL1-B
    - 0C = CL1-C
    - 0D = CL1-D
    - 0E = CL1-E
    - 0F = CL1-F
    - 0G = CL1-G
    - 0H = CL1-H
  - Controller 1
    - 1A = CL2-A
    - 1B = CL2-B
    - 1C = CL2-C
    - 1D = CL2-D
    - 1E = CL2-E
    - 1F = CL2-F
    - 1G = CL2-G
    - 1H = CL2-H

10

- CCI horcm0.conf

HORCM_MON

#ip_address	service	poll (ms)	timeout (ms)
SVR1	horcm0	1000	3000

HORCM_CMD

#dev_name
/dev/rdsk/c2t1d1s2

HORCM_DEV

#dev_group	dev_name	port #	target ID	LUN#	MU#
oradb1	disk1	CL1-A	1	2	0
oradb1	disk2	CL1-A	1	2	1
oradb1	disk3	CL1-A	1	2	2

HORCM_INST

#dev_group	ip_address	service
oradb1	SVR1	horcm1

11

HORCM_MON
ip_address: The IP address of the local host.
Service: The port name assigned to the CCI service (registered in the /etc/services). The service parameter defines the CCI instance that runs on the local host.
Poll (10ms): The interval for monitoring paired volumes.
Timeout (10ms): The time-out period of communication with the remote server.

HORCM_CMD
Describes the path to the raw device serving as the command device.

HORCM_DEV
dev_group: Names a group of paired logical volumes. A command is executed for all corresponding volumes according to this group name.
dev_name: Names the paired logical volume within a group.
Port #: Defines the storage system port number of the volume that corresponds with the dev_name volume.
Target ID: Defines the fibre target ID number of the phys. volume on the specified port.
LU #: Defines the fibre logical unit number (LU#) of the physical volume on the specified target ID and port.
MU #: Is a Mirrored Unit number when creating 1 source to multiple targets. 0 = first copy.

HORCM_INST
dev_group: The server name described in dev_group of HORCM_DEV.
ip_address: The network address of the specified remote server.
service: The port name assigned to the HORCM communication path.

- CCI horcm1.conf

HORCM_MON

#ip_address	service	poll (ms)	timeout (ms)
SVR1	horcm1	1000	3000

HORCM_CMD

#dev_name
/dev/rdsk/c2t1d2s2

HORCM_DEV

#dev_group	dev_name	port #	target ID	LUN#	MU#
oradb1	disk1	CL1-A	1	3	0
oradb1	disk2	CL1-A	1	4	0
oradb1	disk3	CL1-A	1	5	0

HORCM_INST

#dev_group	ip_address	service
oradb1	SVR1	horcm0

12

The explanation is exactly the same as for horcm0.conf, only different numbers and names.

- CCI Setup
  - Storage System
    - Identify the source and target (Copy-on-Write Snapshot software) disks
    - Create command devices using Storage Navigator Modular 2
    - Identify command device target ID and LUN ID
    - Create good documentation
  - Production Server (source)
    - Attach/connect logical source disks created on the storage system
    - Create volume groups, logical volumes, and file systems as needed using the source disks
  - Label command device as required by OS host
    - Microsoft Windows NT does not require a signature
    - Sun Solaris, AIX, DEC, HP, and Sequent do not require a label (raw device)

13

- CCI Setup (*continued*)
  - Install CCI software
  - Identify the command device
    - Hardware address looks like:
      - /dev/rdsk/c2t3d5s2 (SUN)
      - /dev/rhdisk5 (AIX)
      - \\.\physicaldrive5(NT)
  - Create configuration files and start the HORC Manager instance

Find a command device in Windows NT using the command, **raidscan –x findcmdev**.

For Sun Solaris, issue the format command to identify the command device. In the output of the format command, search for the “<HITACHI-DF600F-CM>”.

- CCI Tips
  - If the command device in the Horcm0 or Horcm1 file is not correct, HORC Manager will not start. Make sure to specify the correct command device.
  - If the wrong target ID is specified in the HORCM_DEV section, HORC Manager will start without errors. However, the **paircreate** command will produce errors.
  - After starting the HORC instance, verify the target IDs in the HORCM_DEV section by using the following command:
    raidscan – p Port Name
  - Stop and start the instance if changes are made to the configuration files.

15

Execution

- UNIX
  - At the UNIX prompt, enter:
    # cd /HORCM/usr/bin
  - Set the environment variables
    # HORCMINST=0
    # export HORCMINST
    - Defines HORCM Instance 0 as the "Master"
    # HORCC_MRCF=1
    # export HORCC_MRCF
    - Puts the CCI software in ShadowImage mode versus TrueCopy mode. TrueCopy mode is HORCC_MRCF=.
    # ./horcmstart.sh 0 1
    - This starts the HORC Manager instance.

16

- Microsoft Windows
  - At the Microsoft Windows command prompt, enter:
    C:\> cd c:\horcm
  - Set the environment variables
    C:\> set HORCMINST=0
    - Defines HORCM Instance 0 as the "Master"
    C:\> set HORCC_MRCF=1
    - Puts the CCI software in ShadowImage mode versus TrueCopy mode. TrueCopy Mode is HORCC_MRCF=.
    C:\>horcmstart 0 1
    - Starts the HORC Manager instance

17

Commands

- **pairdisplay -g oradb1 –fcx**
  - This command is used to verify the pairing of the volumes and completion of the pair synchronization.
    - -x option is used to display the LDEVs in hex.
    - -fc option is used to display the pair copy completion percentage.
- **pairdisplay** output appears similar to the following:

```
# pairdisplay -g oradb1 -fcx

Group  PairVol(L/R) (Port#,TID,LU-M),Seq#,LDEV#.P/S,Status,Copy%,P-LDEV# W

VG01oradb1(L)    (CL1-A, 0,  0-0)30177   1..P-VOL PAIR,  100     2  -
```

18

On the following pages, you can find examples of commonly used commands. Please refer to the Adaptable Modular Storage 2000 *Copy-on-Write Snapshot Software User's Guide and to the CCI User and Reference Guide* for more details.

- **paircreate -split -g oradb1 –d disk1 –vl**
  - This command starts the snapshot "disk1". Before a P-VOL block is changed, its (old) data is retained in the corresponding V-VOL (pool).
    - The **-vl** option copies data from the source volume to the target volume.
    - Using the **-vr** instead of **-vl** option copies data from the target volumes to the source volumes.
    - Pair Status transition SIMPLEX >>> PSUS.

- **pairsplit -g oradb1 –S**
  - Immediately stops updating the secondary volume and sets the ShadowImage pairs to simplex status via the **-S** option
  - No “Change Data” bitmap is maintained
  - **pairresync** is not possible; must **paircreate**

20

- **pairresync -g oradb1 –d disk1 -restore**
  - psus > copy > pair
  - Copies old data from the snapshot named “disk1” represented by the corresponding V-VOL (pool) back to the primary volume

21

- **pairevtwait -g oradb1 -s pair -t 7200**
  - **-s** option specifies the status transition to monitor.
    - SMPL, COPY, PAIR, PSUS, and PSUE
  - **-t** option specifies the timeout value (in seconds) to wait for the specified status to occur.

22

Commands used in scripts that automate pair manipulation

- **pairmon -g oradb1 -allsnd –nowait**
  - **-allsnd** reports all pair status transition events.
  - **-nowait** specifies the command does not wait if there is no pair transition status to report.

23

Microsoft Windows NT Commands

- **findcmddev**
  - Locates and displays physical drive number of the command device. Does not require CCI to be active.

    C:> raidscan -x findcmddev hdisk0,32

    \\.\PhysicalDrive3

24

Commands used in Windows Server 2000 and 2003 or NT environments.

The **findcmddev** command will locate the command device and output the physical device number, such as \\.\physicaldrive5.

- **sync**
  - Flushes unwritten data from the Microsoft Window NT server memory to the storage system LUNs
    - Must be executed from the server attached to the source disks
    - Does not require CCI to be active

      raidscan -x sync F: G: H: ...

      raidscan -x sync all

      raidscan -x sync drive # (F-R)

- **mount**
  - Allocates the specified logical device to the specified partition using the drive characters.

    raidscan -x mount F: hdisk2 p1 -x mount G: hdisk1 p1

    raidscan -x mount

Drive	FS_name	VOL_name	Device Partition	Port	PathID	Targ	LUN
C:	FAT Null	Harddisk0	Partition1	1	0	0	0
F:	FAT Null	Harddisk1	Partition1	1	0	5	1
G:	NTFS Null	Harddisk2	Partition1	1	0	5	0

26

The **mount** command is a CCI software provided command that serves as the equivalent to a UNIX mount command in a NT environment. NT currently has no such command.

- **unmount**
  - Unmounts the specified logical device and deletes the drive characters.
  - The following example unmounts the F drive and G drive and enables Read/Write access to the secondary volume and displays the mounted devices.

    raidscan -x umount F: -x umount G: -g oradb1 -rw

    raidscan -x mount

Drive	FS_name	VOL_name	Device Partition	Port	PathID	Targ	LUN
C:	FAT Null	Harddisk0	Partition1	1	0	0	0

27

The **umount** command is a CCI software provided command that serves as the equivalent to a UNIX unmount command in a NT environment. NT currently has no such command.

Usage Notes

- Always verify pair configuration prior to **paircreate** command.
  - Verify data flow direction
    - Failure to verify data flow direction can cause data loss.
      - No warning message is issued.
  - If horcm/horcm1.conf changes were applied
- Avoid source and target pairing in the same array group to maximize performance.

28

Web Tool

- Provides information about P-VOLs and pairs

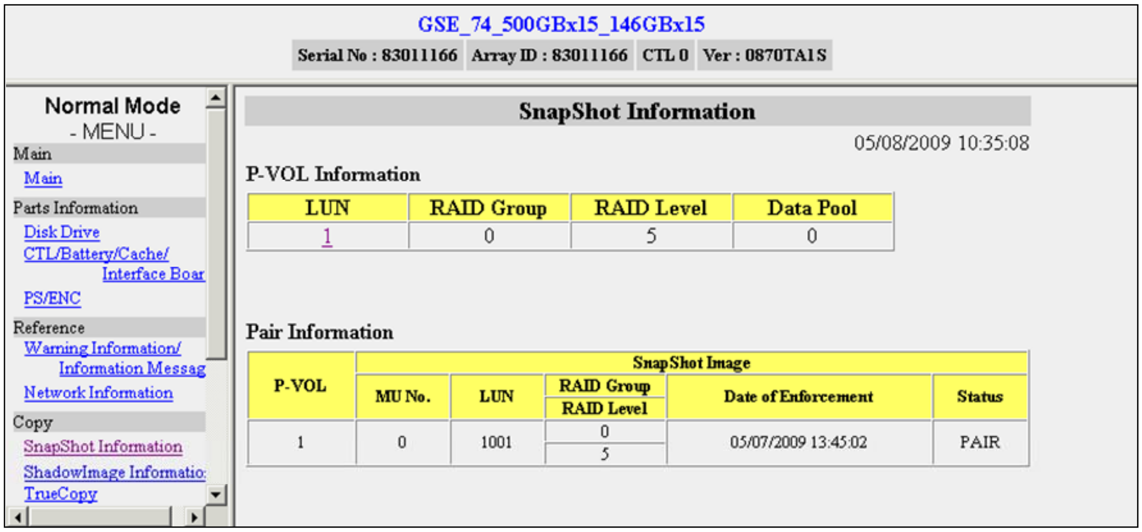

29

The WEB tool provides information about P-VOLs and pairs. This page is available as soon as the product key is installed.

8. Hitachi Copy-on-Write Snapshot Software Storage Navigator Modular 2 Operations

Module Objectives

- Upon completion of this module, the learner should be able to:
  - Describe the Hitachi Storage Navigator Modular 2 GUI and CLI Operations for Hitachi Copy-on-Write Snapshot software

2

License

- A license key for ShadowImage is required and must be enabled.

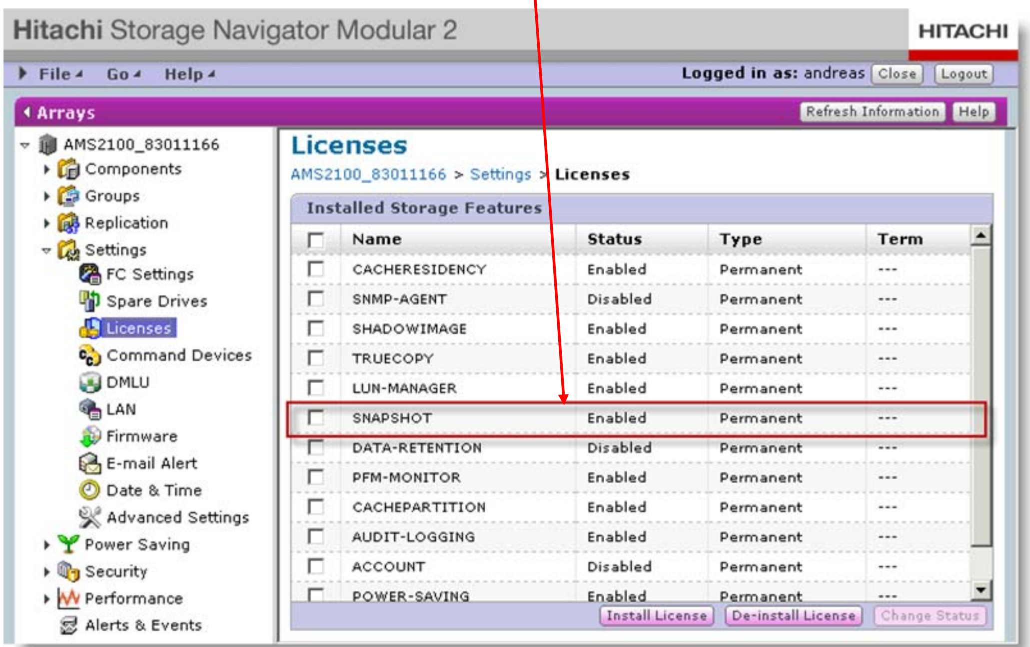

3

Overview of Copy-on-Write Snapshot Software

- Hitachi Storage Navigator Modular 2 (SNM2) GUI and CLI enables you to:
  - Configure and manage data pools
  - Configure and manage Snapshot logical units (V-VOLs)
  - Create Snapshot pairs
  - View the status of Snapshot pairs
  - Split Snapshot pairs
  - Resynchronize Snapshot pairs
  - Restore Snapshot pairs
  - Delete and Release Snapshot pairs

- **No** Command Device, **no** CCI Raid Manager, **no** HORCM script and **no** HORCM instance is necessary when using SNM2 to perform Copy-on-Write Snapshot operations on an Hitachi Adaptable Modular Storage (AMS) 2000 Family storage system.

4

SNM2 CLI – CCI/Raid Manager commands

Action	CLI – Command	CCI - Command
Display pair information	aureplicationlocal -refer	pairdisplay
Create pairs	aureplicationlocal -create	paircreate
Split pairs	aureplicationlocal -split	pairsplit
Resynchronize pairs	aureplicationlocal -resync	pairresync
Restore pairs	aureplicationlocal -restore	pairresync -restore
Delete pairs	aureplicationlocal -simplex	pairsplit -S
Event wait	aureplicationmon -evwait	pairevtwait

5

SNM2 GUI Operations

Storage Navigator Modular 2 Operations

Configure and Manage Data Pools

6

Create a Data Pool with SNM2

1. Select **Data Pools** and click **Create Data Pool.**

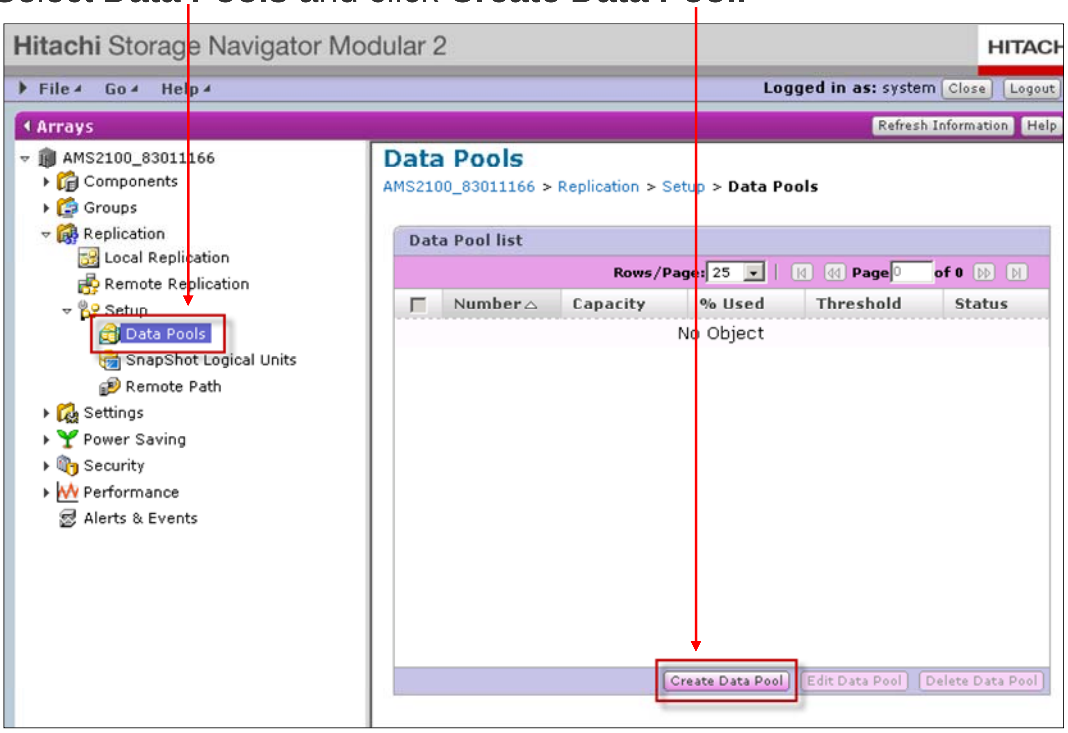

7

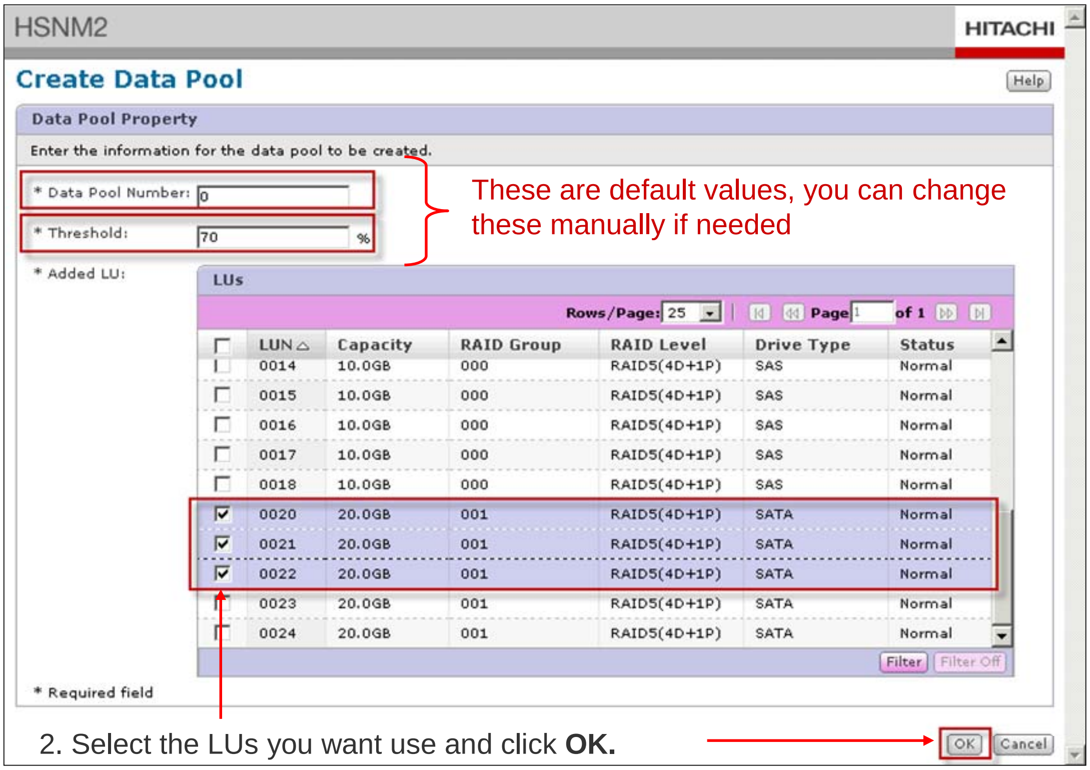

8

Data Pool Number

The Data Pool Number starts at zero and auto increments per default by one. You can specify/change this number manually if needed.

Threshold

The Data Pool usage threshold is set to 70% data usage by default.

You can specify/change this value manually if needed.

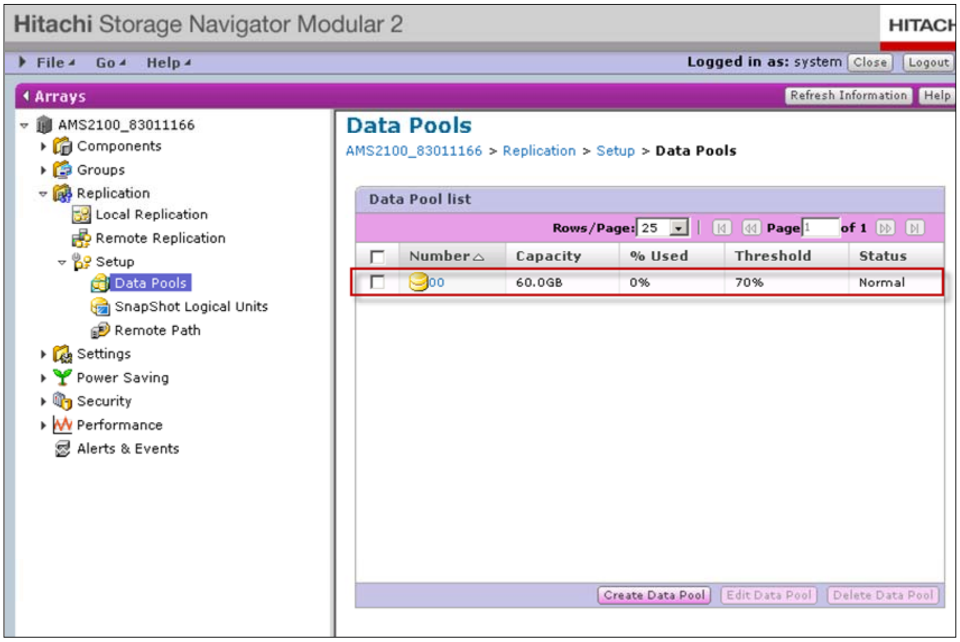

To create a Data Pool with the SNM2 CLI, you can use the CLI command as follows:
aupool -unit 'unit_name' -add -poolno 0 -lu 20 21 22

9

When using SNM 2 CLI to create a Data Pool, the Threshold value will be set to 80% automatically. This can be changed after the Data Pool has been set up. SNM2 CLI does not provide a Data Pool Number automatically.

You will find more information on creating Data Pools using the SNM2 CLI in the *Hitachi Navigator 2 Command Line Interface (CLI) Reference Guide for Replication* on page 1-8

View Data Pool Details with SNM2

- Select **Data Pools** and click on the **Data Pool Number**.

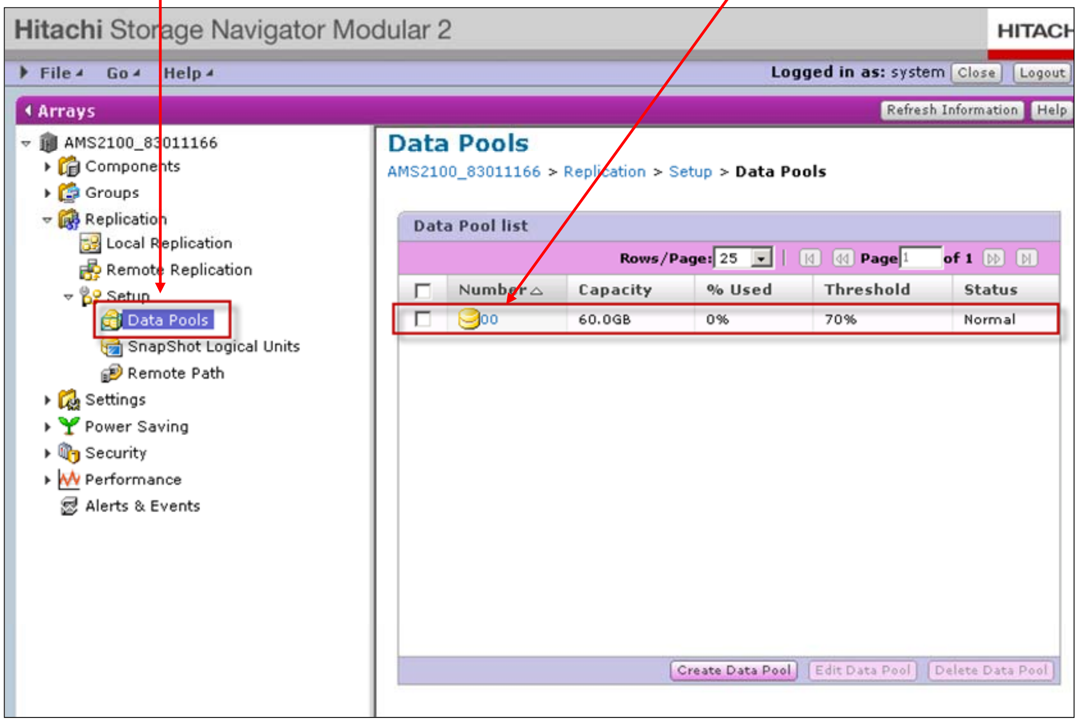

10

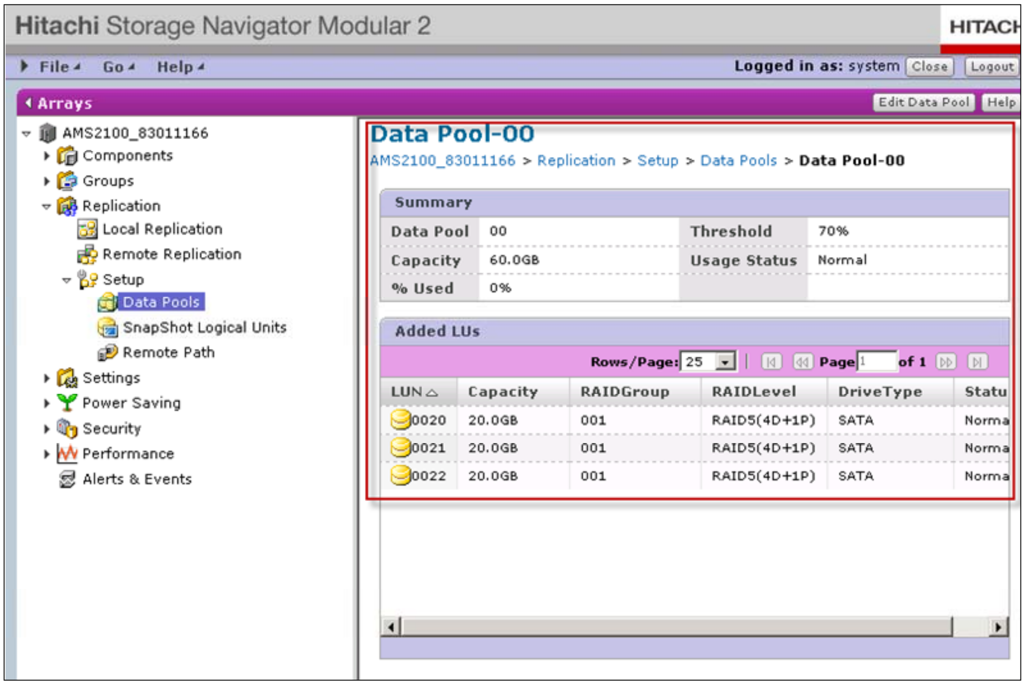

To view the Data Pool details with the SNM2 CLI, you can use the CLI command as follows:
aupool -unit ‘unit_name’ -refer

11

You will find more information on viewing Data Pool details using the SNM2 CLI in the *Hitachi Navigator 2 Command Line Interface (CLI) Reference Guide for Replication* on page 1-8

Expand the Data Pool or Change the Threshold

- Select Data Pools and select the Pool you want change and click Edit Pool.

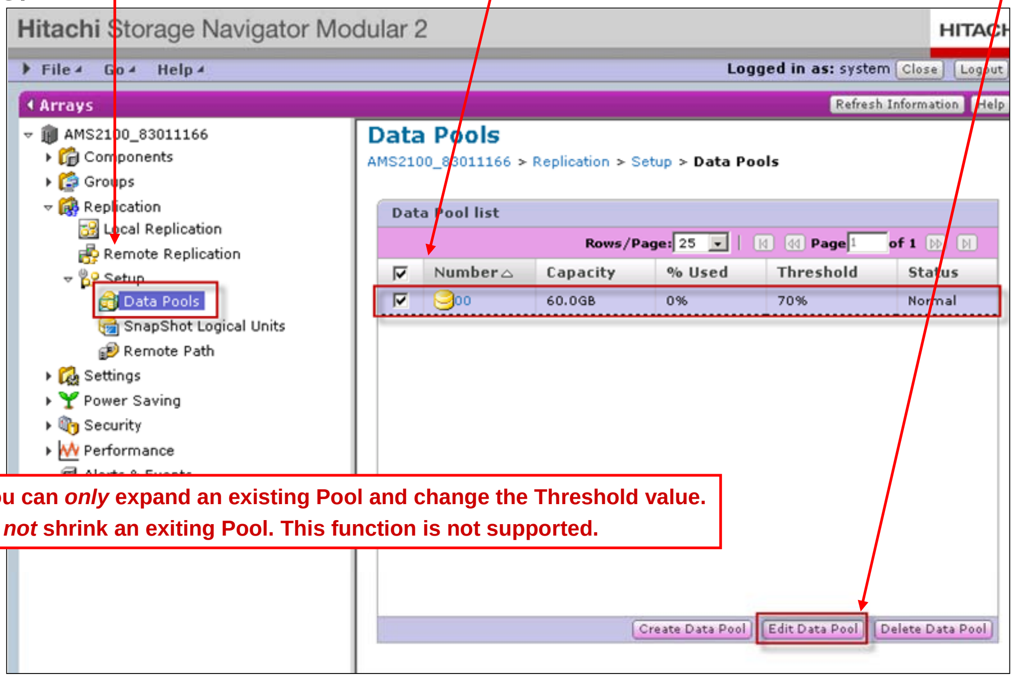

***Note*: You can *only* expand an existing Pool and change the Threshold value. You can *not* shrink an exiting Pool. This function is not supported.**

12

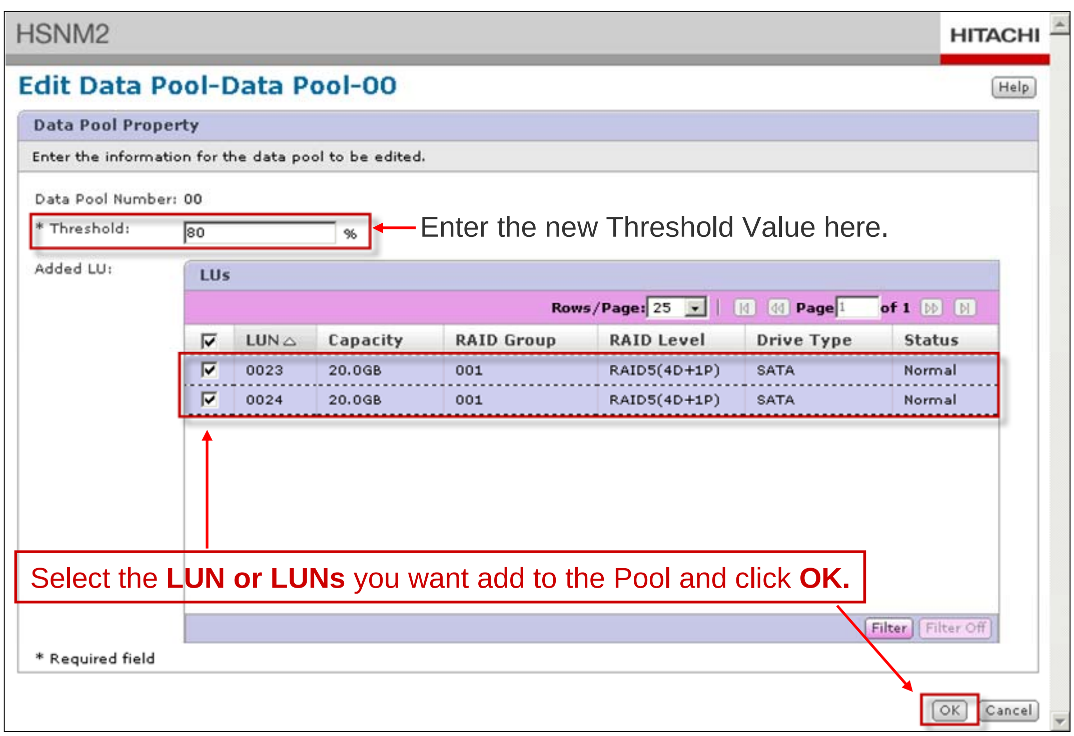

Enter the new Threshold Value here.

Select the **LUN or LUNs** you want add to the Pool and click **OK.**

13

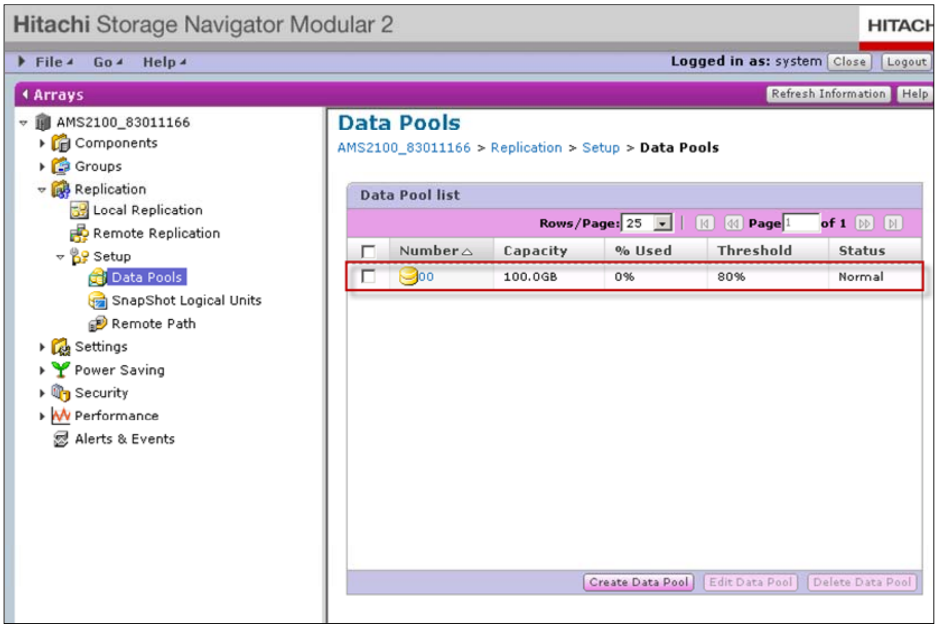

To expand a Data Pool with the SNM2 CLI, you can use the CLI command as follows:

aupool -unit ‘unit_name’ -add -poolno 0 -lu 23 24

To change the Threshold with the SNM2 CLI, you can use the CLI command as follows:

aupool -unit ‘unit_name’ -chg -poolno 0 -thres 50

14

You will find more information on managing Data Pool using the SNM2 CLI in the *Hitachi Navigator 2 Command Line Interface (CLI) Reference Guide for Replication* on page 1-8

Delete a Data Pool

Select **Data Pools**, select the **Pool** you want delete and click **Delete Pool.**

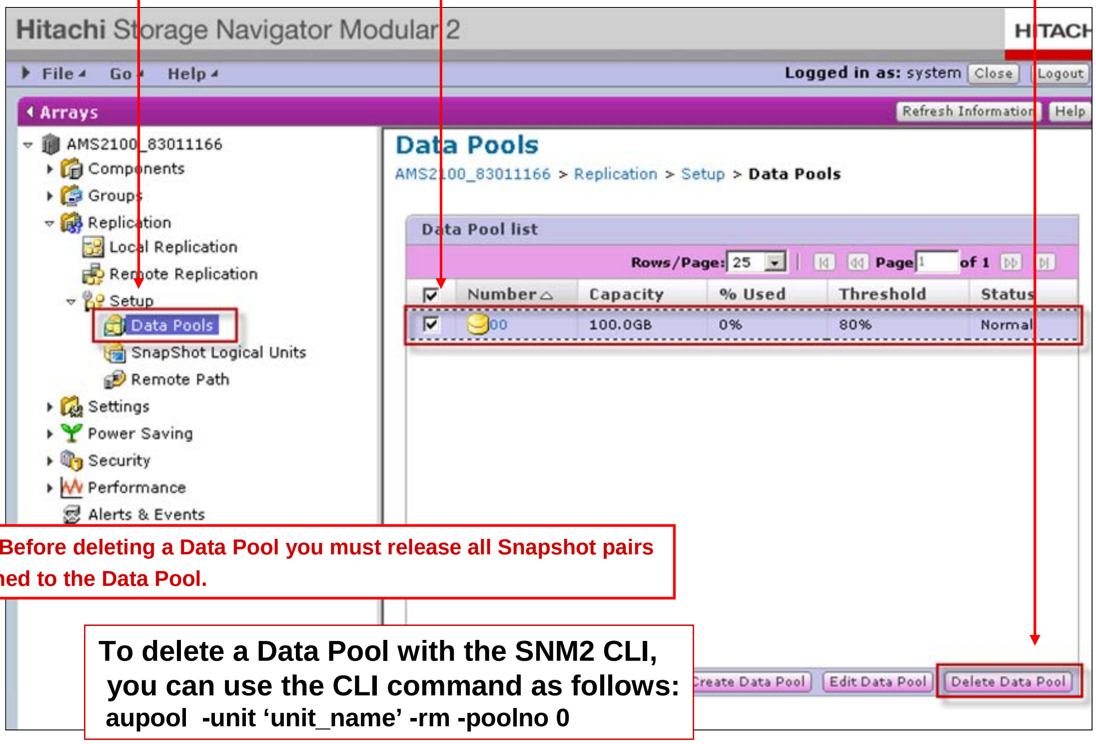

***Note*: Before deleting a Data Pool you must release all Snapshot pairs assigned to the Data Pool.**

To delete a Data Pool with the SNM2 CLI, you can use the CLI command as follows:
aupool -unit 'unit_name' -rm -poolno 0

15

You will find more information on deleting Data Pool details using the SNM2 CLI in the *Hitachi Navigator 2 Command Line Interface (CLI) Reference Guide for Replication* on page 1-8

SNM2 GUI Operations

Storage Navigator Modular 2 Operations

Configure and Manage Snapshot Logical Units (V-VOLs)

16

Create a Snapshot LU (V-VOL)

- Select **SnapShot Logical Unit** and click **Create SnapShot LU**.

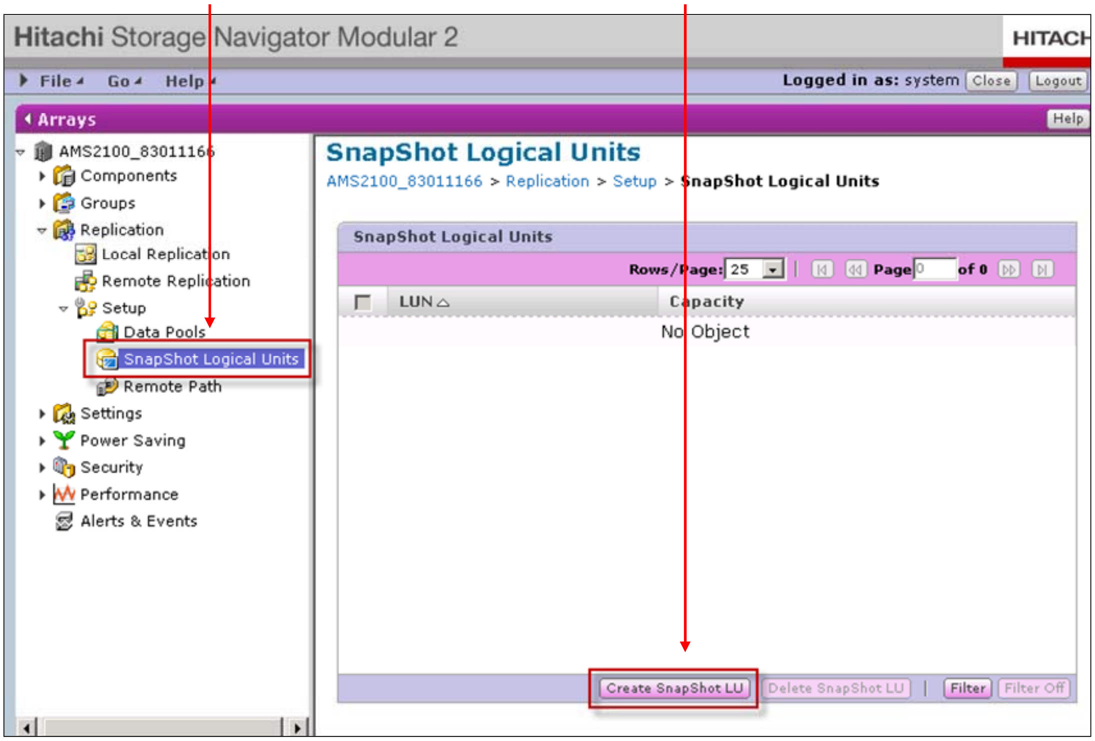

17

Specify the LU Number

Specify the LU Capacity

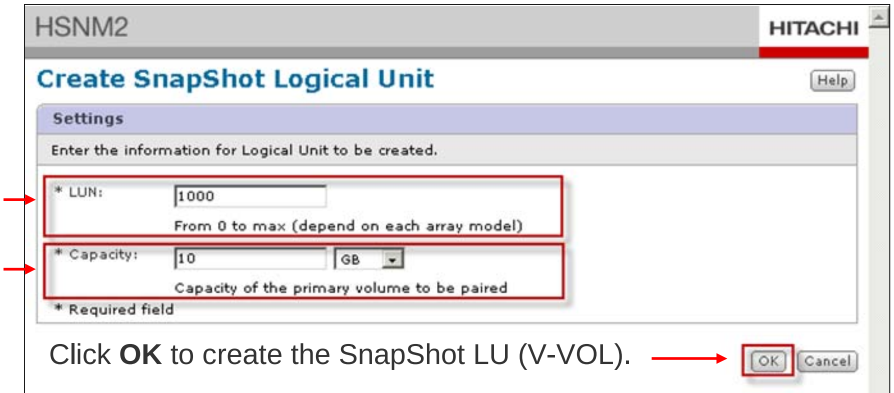

Click **OK** to create the SnapShot LU (V-VOL).

Note: You must create multiple Snapshot LUs to have multiple snapshots from one P-VOL. For example, to create five different snapshots from one P-VOL, you must create five Snapshot LUs of the same size as the P-VOL.

To create a SnapShot Logical unit with SNM2 CLI you can use the CLI command as follows:

aureplicationvvol -unit 'unit_name' -add -lu 1000 -size 10g

18

LUN

Per default SNM2 suggests the next lowest available LU Number and auto increments by one.

Capacity

The capacity must match the capacity of the P-VOL

You will find more information on creating SnapShot Logical Units using the SNM2 CLI in the *Hitachi Navigator 2 Command Line Interface (CLI) Reference Guide for Replication* on page 1-6

Delete a Snapshot LU (V-VOL)

Select **SnapShot Logical Units**, select the **V-VOL** and click **Delete SnapShot LU.**

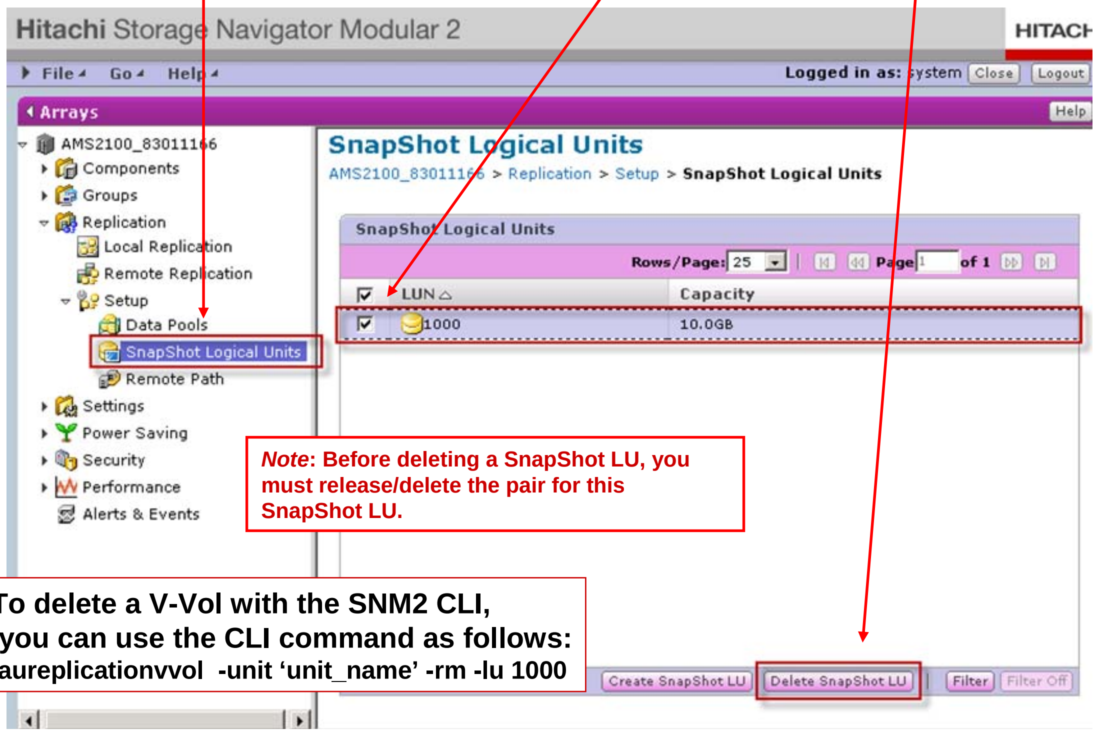

***Note*: Before deleting a SnapShot LU, you must release/delete the pair for this SnapShot LU.**

To delete a V-Vol with the SNM2 CLI, you can use the CLI command as follows:
aureplicationvvol -unit ‘unit_name’ -rm -lu 1000

19

You will find more information on deleting SnapShot Logical Units using the SNM2 CLI in the *Hitachi Navigator 2 Command Line Interface (CLI) Reference Guide for Replication* on page 1-6

SNM2 GUI Snapshot Pair Operations

- Storage Navigator Modular 2 Snapshot Pair Operations
  - Create Snapshot Pairs
  - Reference/View Snapshot Pairs
  - Split Snapshot Pairs
  - Resynchronize Snapshot Pairs
  - Restore Snapshot Pairs
  - Delete/Release Snapshot Pairs

20

Create a SnapShot Pair with SNM2

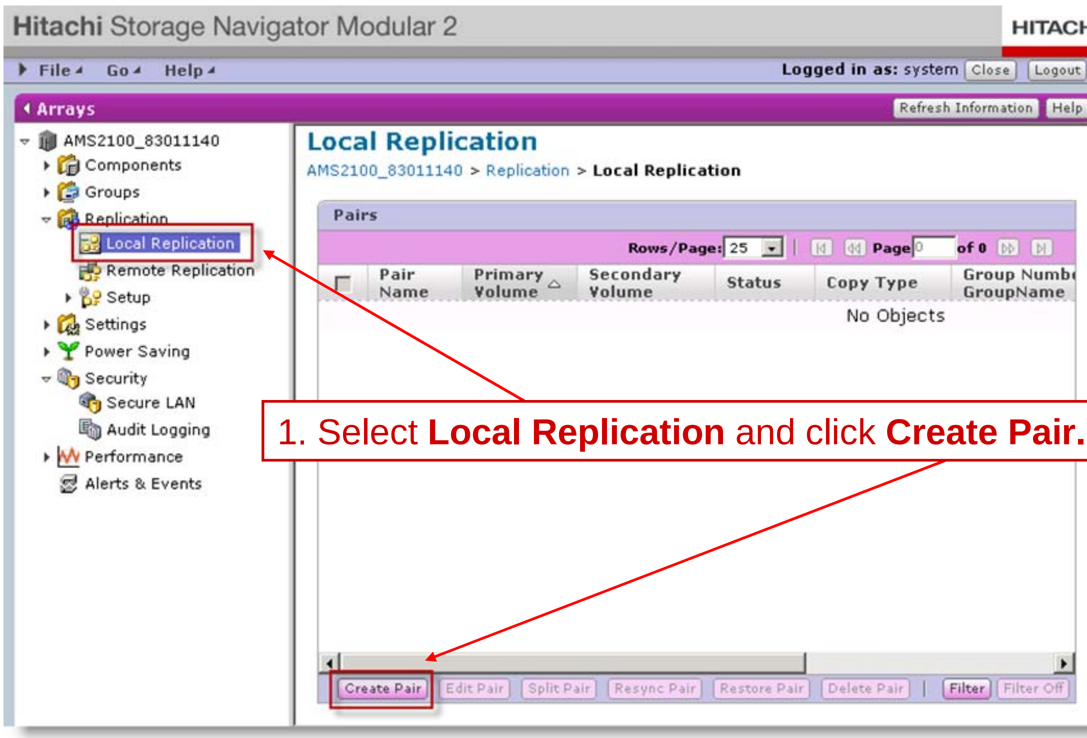

21

2. Select **SnapShot** as **Copy Type**.

22

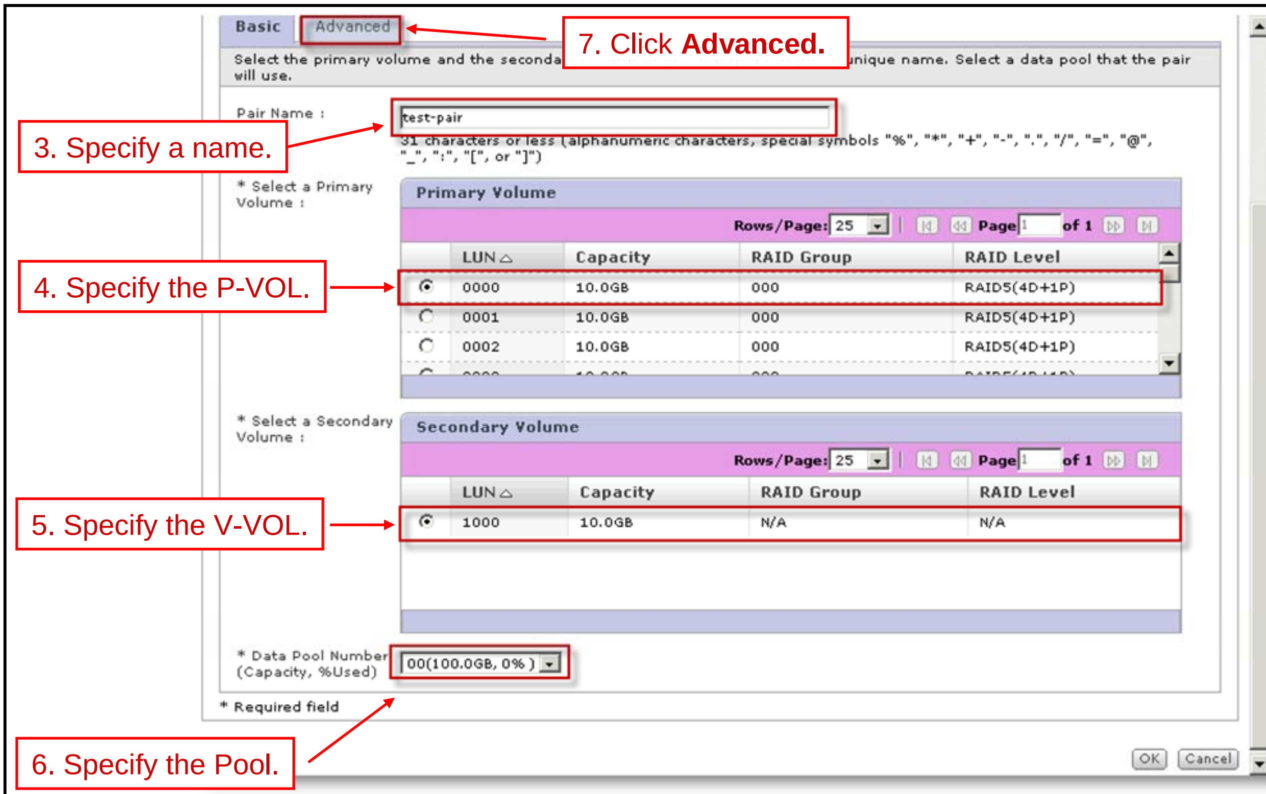

23

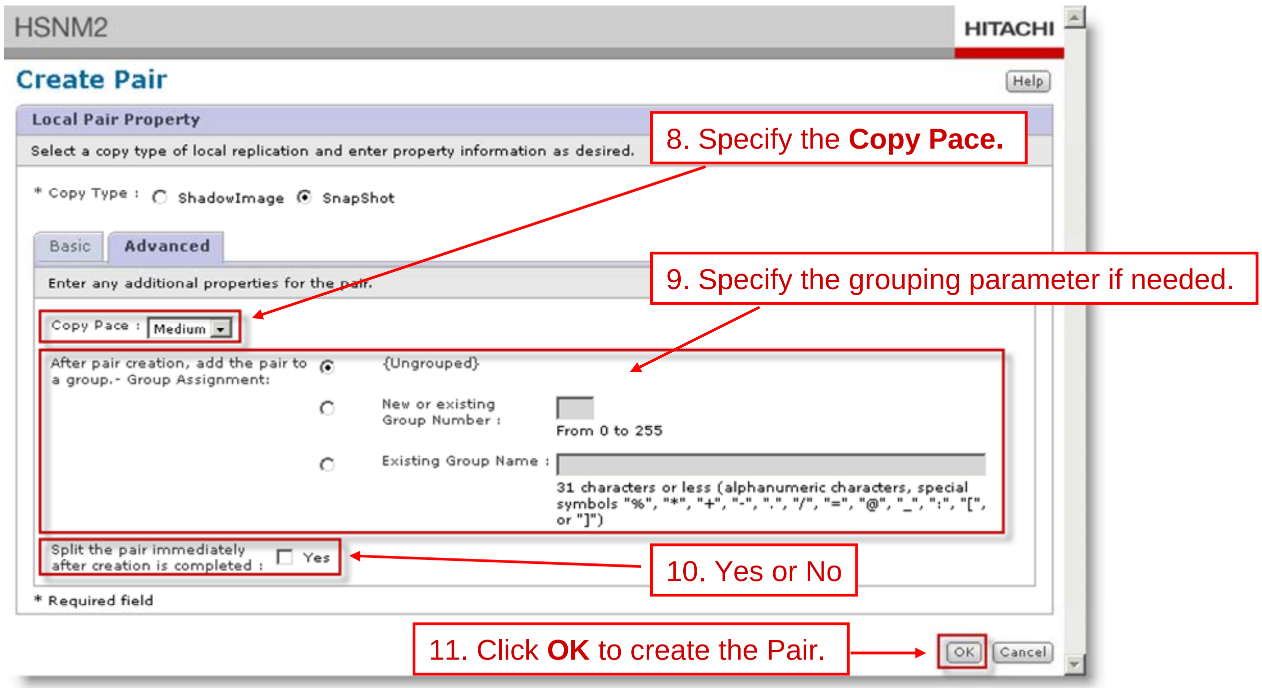

24

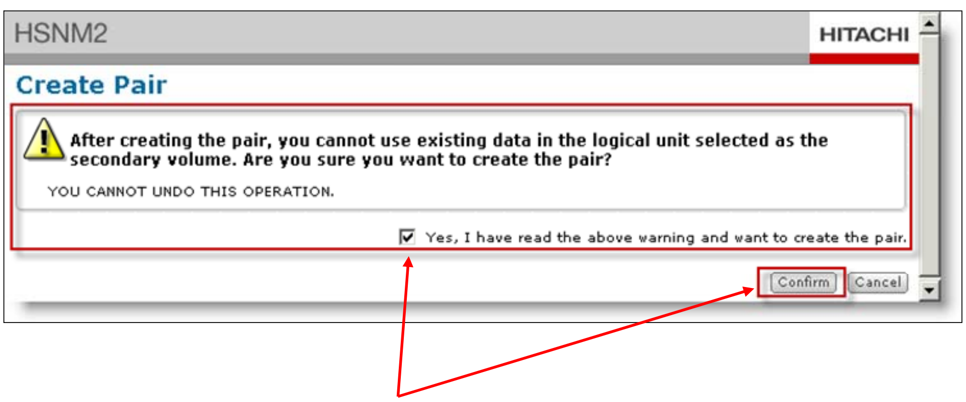

Note: You must confirm the warning message in order to proceed.

25

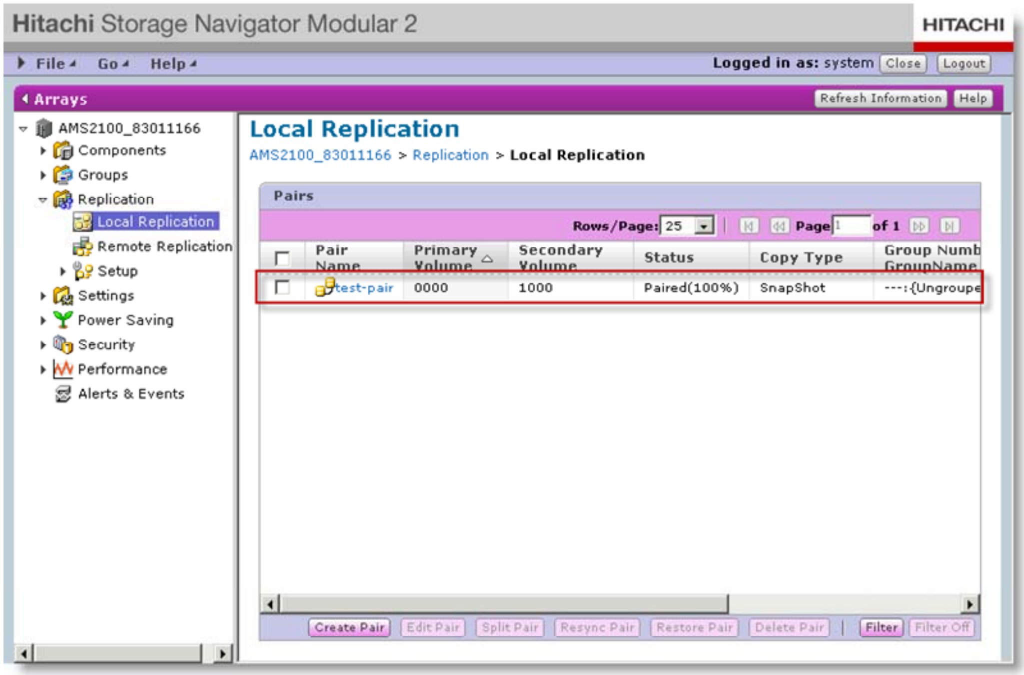

To create the SnapShot pair using SNM2 CLI you can use the CLI command as follows:

aureplicationlocal -unit ‘unit_name’ -create -ss -pvol 0 -svol 1000 -pairname test-pair -localpoolno 0

26

You will find more information on creating SnapShot pairs using the SNM2 CLI in the *Hitachi Navigator 2 Command Line Interface (CLI) Reference Guide for Replication* on page 1-9

Reference/View SnapShot Pairs with SNM2

- Select **Local Replication** and click on the **Pair Name** of the pair you want see the details.

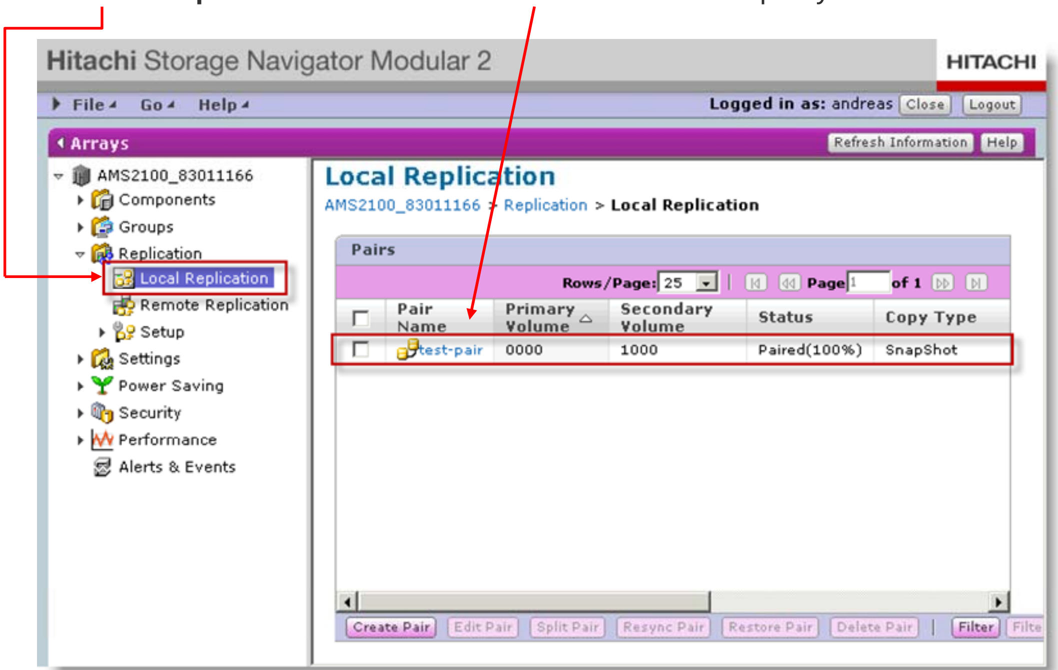

27

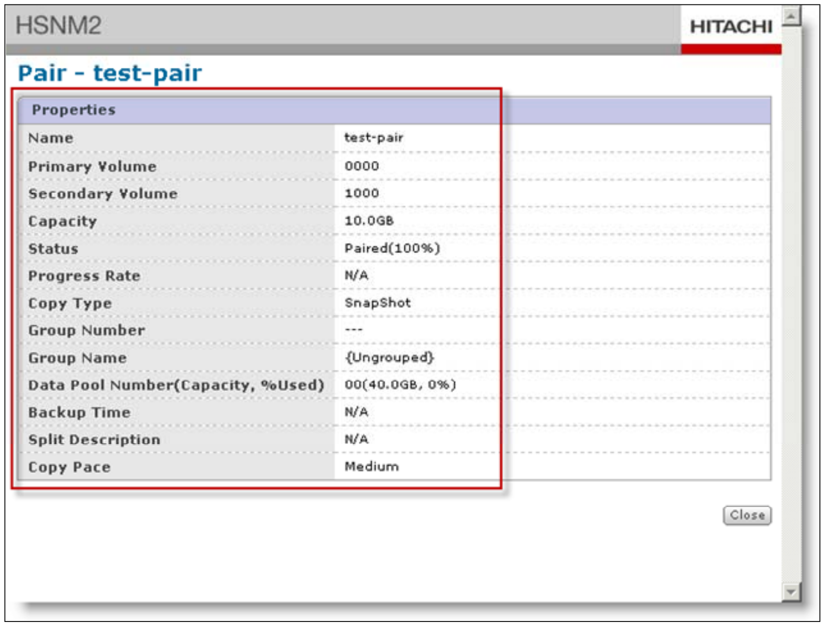

To view the SnapShot pair details using SNM2 CLI you can use the CLI command as follows:

aureplicationlocal -unit ‘unit_name’ -refer -detail -pvol 0 -svol 1000

28

You will find more information on viewing SnapShot pair details using the SNM2 CLI in the *Hitachi Navigator 2 Command Line Interface (CLI) Reference Guide for Replication* on page 1-2

Split SnapShot Pairs with SNM2

1. Select **Local Replication,** select the **Pair/s** you want split and click on **Split Pair** button.

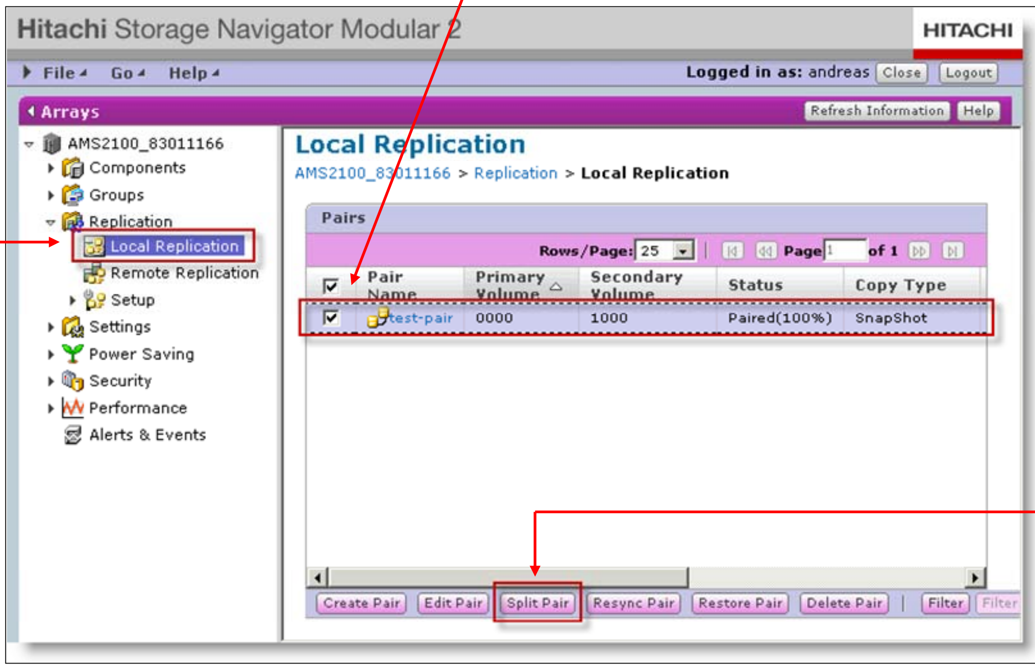

29

2. Optional, you can enter a comment or description, click on the **OK** button to proceed.

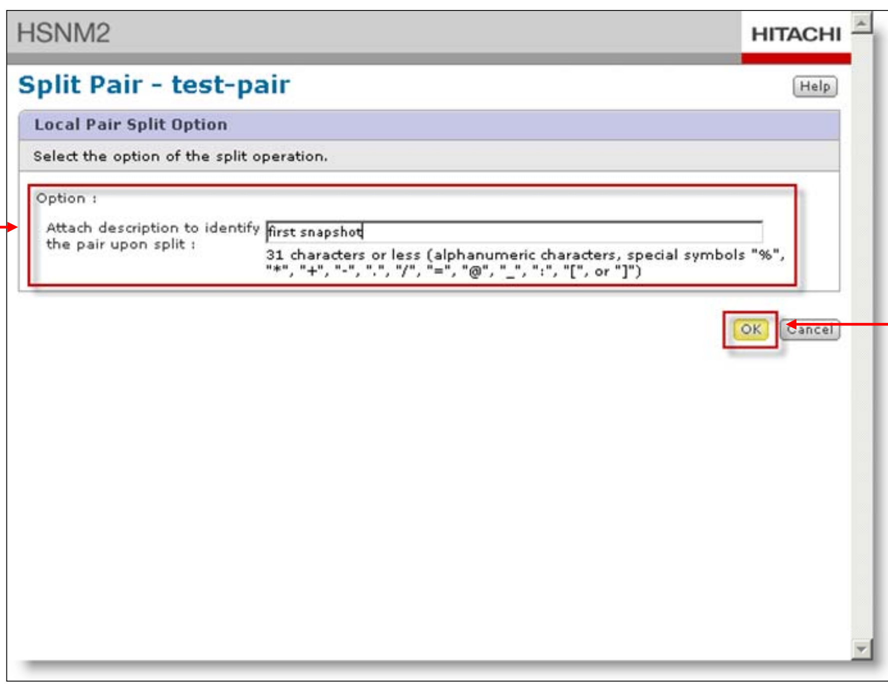

30

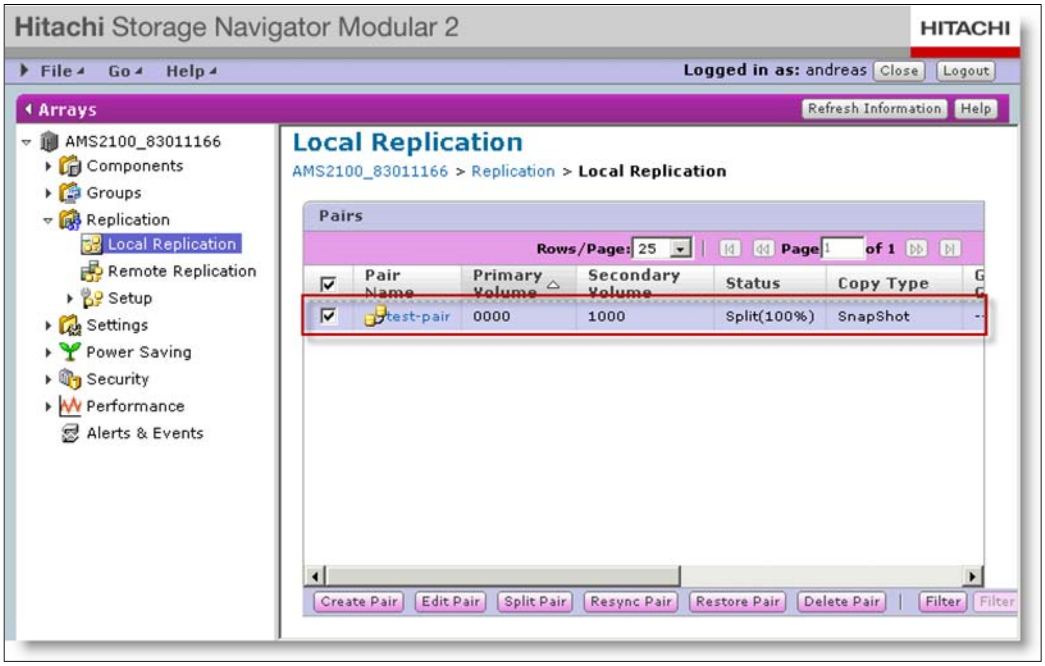

To split the SnapShot pair using SNM2 CLI you can use the CLI command as follows:

aureplicationlocal -unit 'unit_name' -splitt -ss -pvol 0 -svol 1000

31

You will find more information on viewing SnapShot pair details using the SNM2 CLI in the *Hitachi Navigator 2 Command Line Interface (CLI) Reference Guide for Replication* on page 1-12

Resynchronize SnapShot Pairs with SNM2

1. Select **Local Replication,** select the **Pair/s** you want resynchronize and click **Resync Pair.**

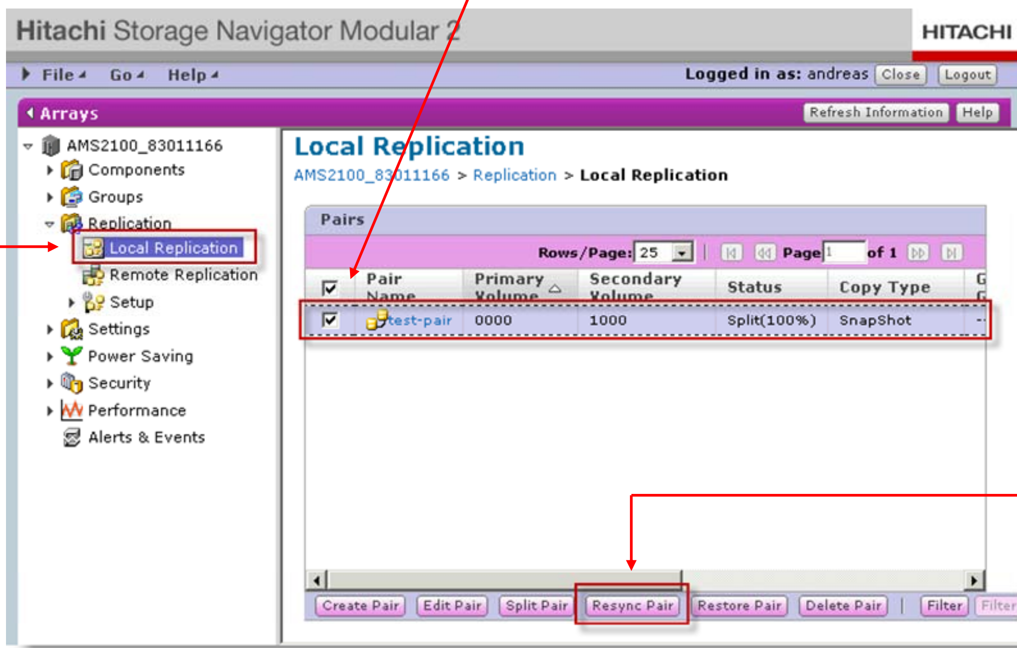

32

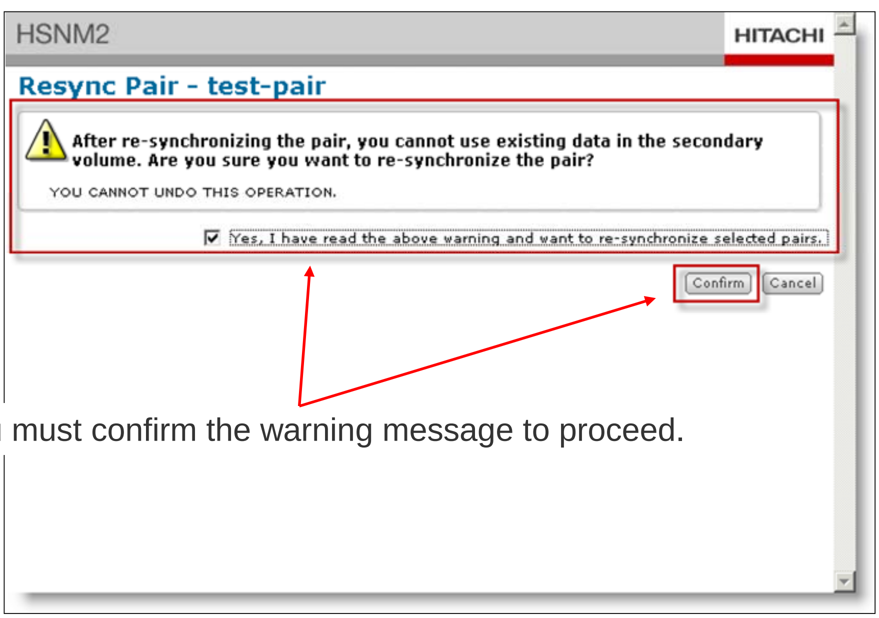

Note: You must confirm the warning message to proceed.

33

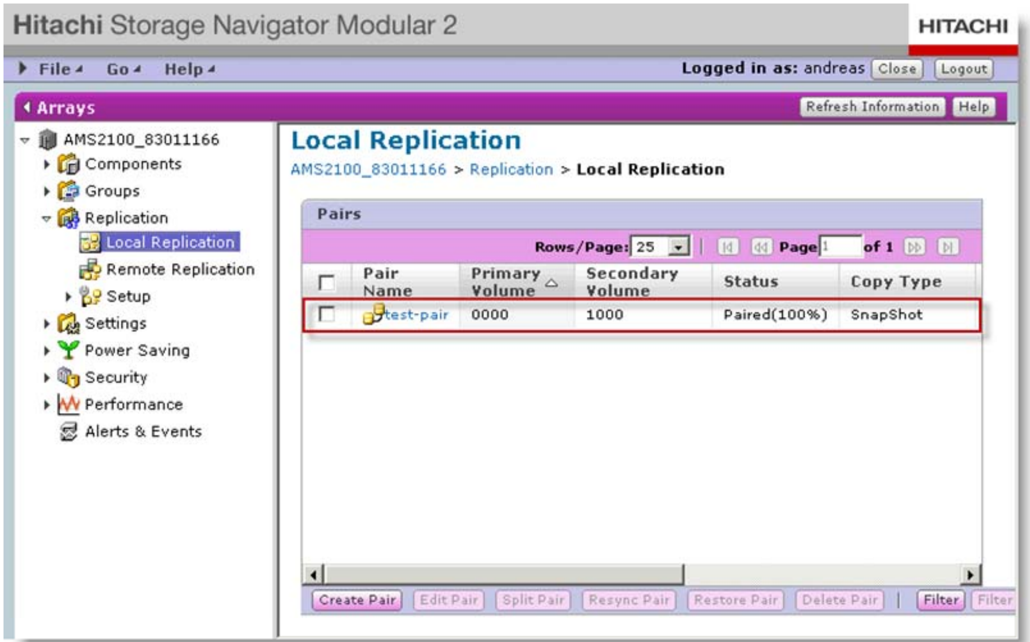

To re-synchronize the SnapShot pair using SNM2 CLI you can use the CLI command as follows:

aureplicationlocal -unit ‘unit_name’ -resync -ss -pvol 0 -svol 1000

34

You will find more information on re-synchronizing SnapShot pairs details using the SNM2 CLI in the *Hitachi Navigator 2 Command Line Interface (CLI) Reference Guide for Replication* on page 1-15

Restore SnapShot Pairs with SNM2

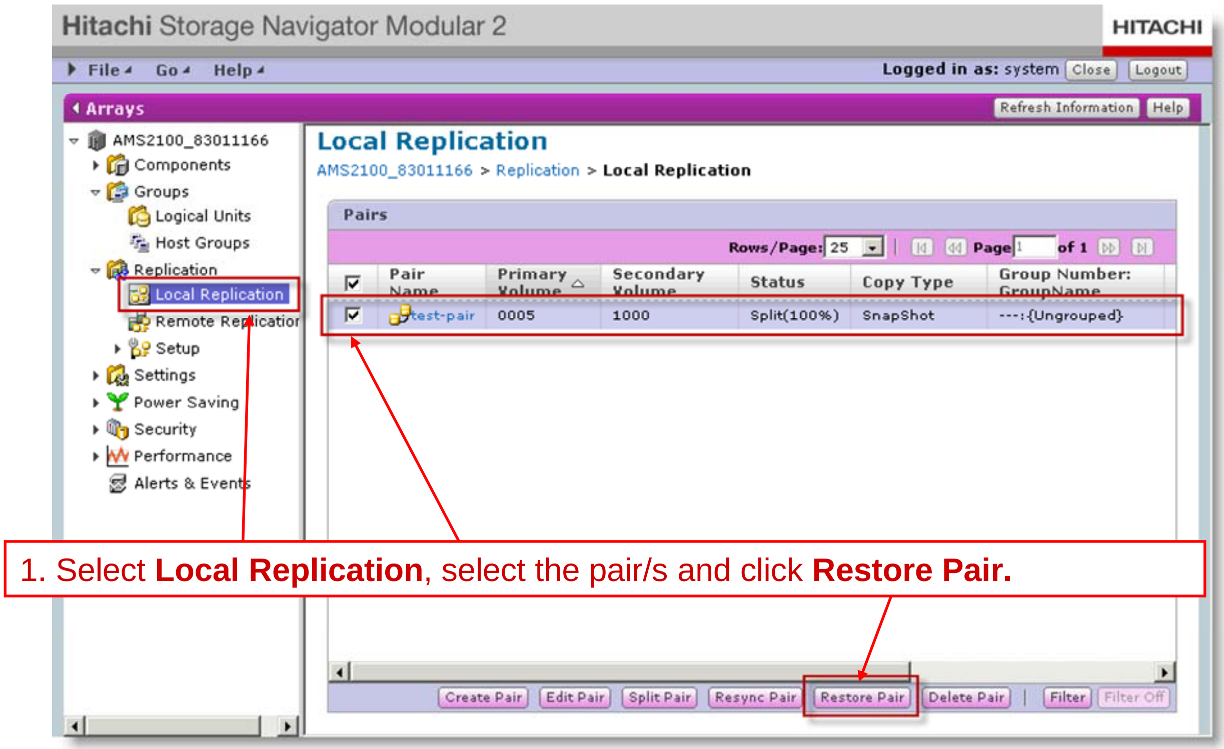

35

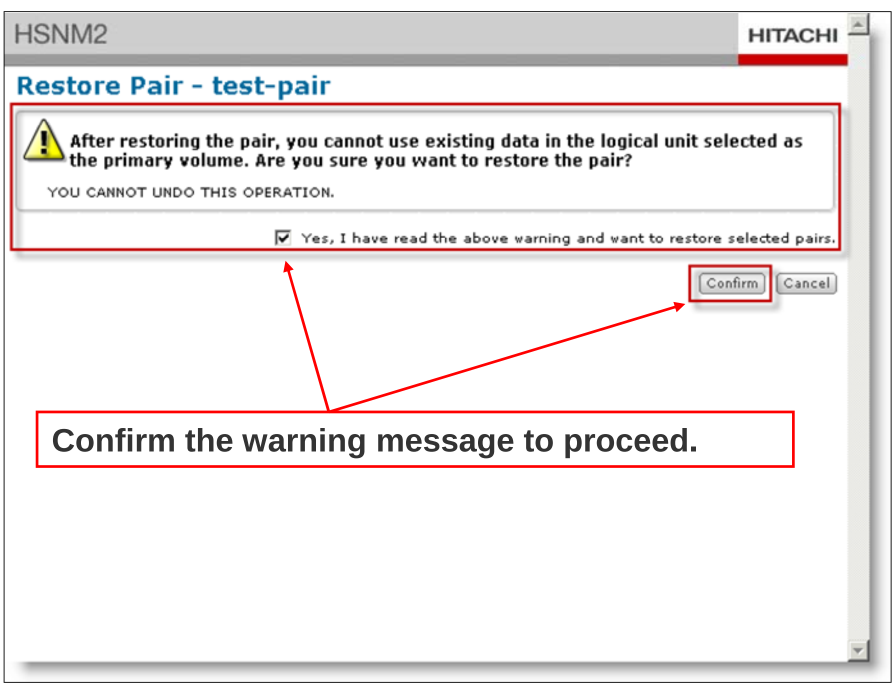

36

Hitachi Data Systems

- The detailed pair information after the restore operation

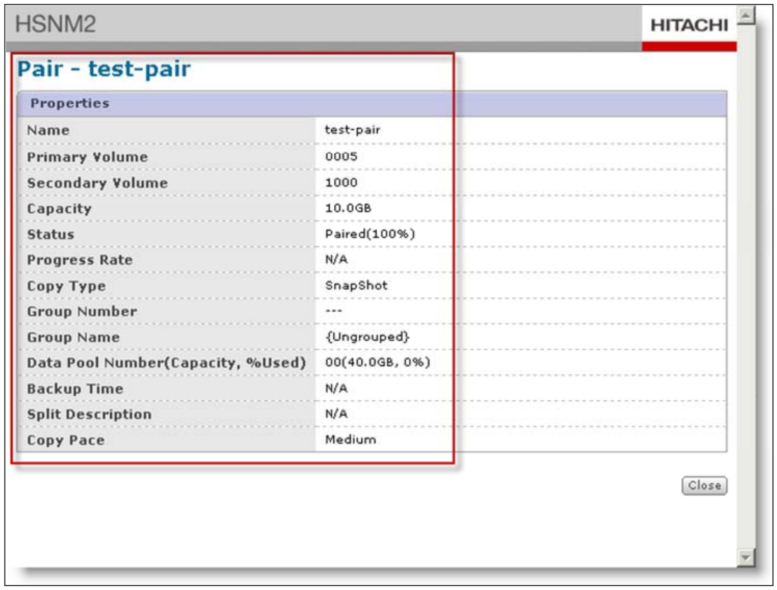

To restore the pair with the SNM2 CLI, you can use the CLI command as follows:
aureplicationlocal -unit ‘unit_name’ -restore -ss -pvol 5 -svol 1000

37

You will find more information on restoring SnapShot pairs using the SNM2 CLI in the *Hitachi Navigator 2 Command Line Interface (CLI) Reference Guide for Replication* on page 1-17

Delete/Release SnapShot Pairs with SNM2

Select **Local Replication**, select the **Pair/s** you want delete and click on the **Delete Pair** button.

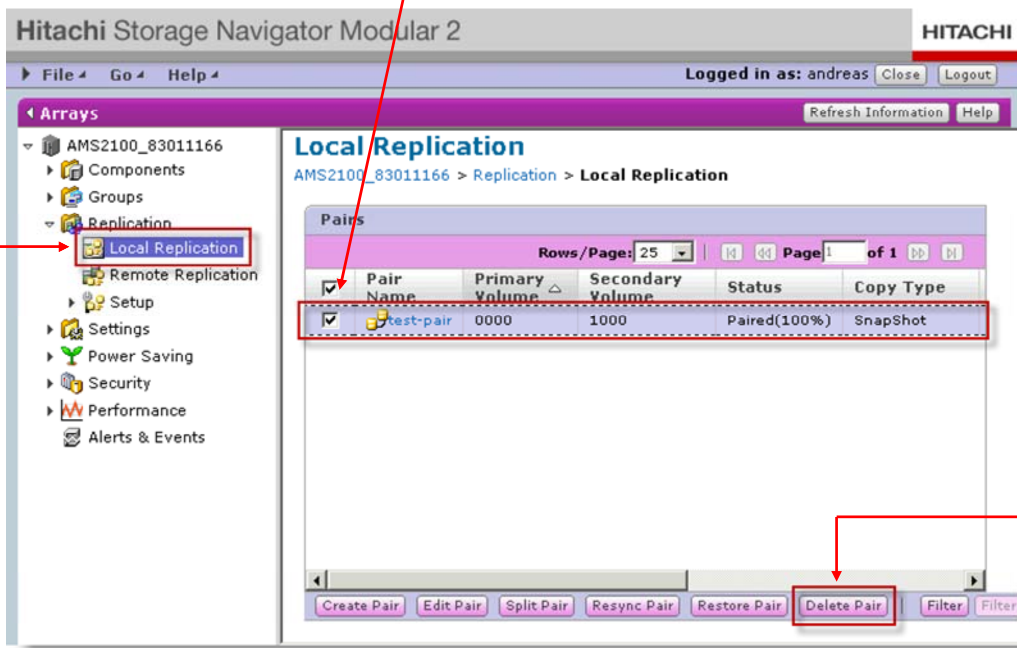

38

Note: You must confirm the warning message to proceed.

39

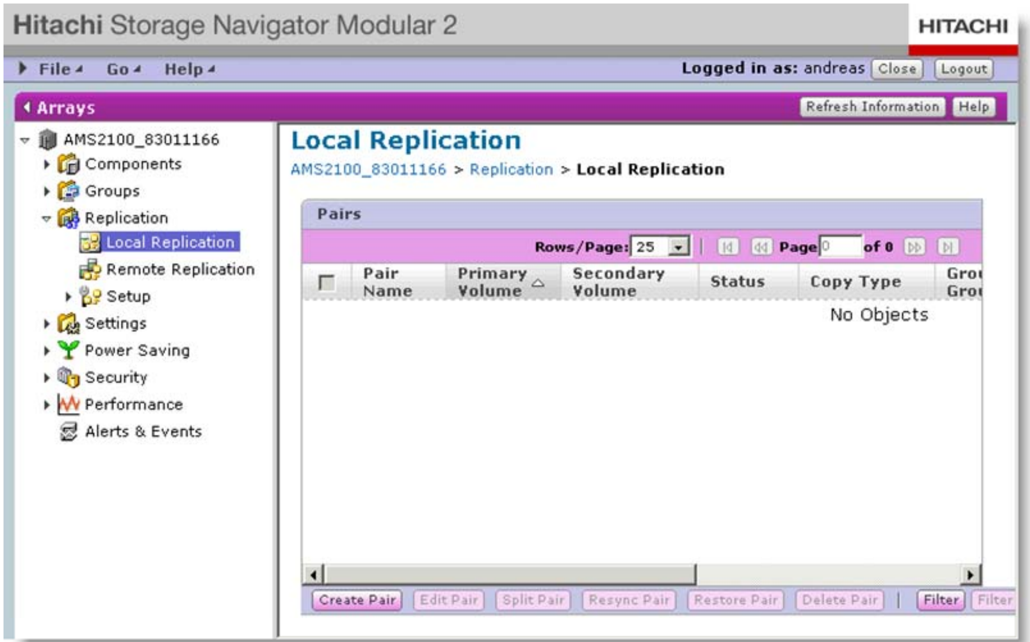

To delete a pair with the SNM2 CLI, you can use the CLI command as follows:

aureplicationlocal -unit 'unit_name' -simplex -ss -pvol 0 -svol 1000

40

You will find more information on deleting ShadowImage pairs using the SNM2 CLI in the *Hitachi Navigator 2 Command Line Interface (CLI) Reference Guide for Replication* on page 1-19

9. Hitachi TrueCopy Remote Replication Software Overview

Module Objectives

- Upon successful completion of this module, you will be able to:
  - Describe the key benefits and features of Hitachi TrueCopy Remote Replication software
  - Define TrueCopy software and associated technologies
  - Describe TrueCopy products
  - List the steps to perform synchronous operations
  - Describe physical and logical paths
  - List volume pair states

2

Disaster Recovery

- Disaster Recovery
  - Protection is key because information is priceless.
  - Business continuity is no longer an option – it is a requirement.

3

- Costs of data unavailability are far reaching.
  - Revenue
    - Billing losses
    - Investment losses
    - Loss of future revenue
    - Additional expenses, for example, shipping
  - Productivity
    - Working time
    - Production delays
  - Damaged Reputation
    - Customers and suppliers
    - Financial markets and banks

4

- Business Continuity includes plans for technology, people and business processes for recovery, which are based around business requirements. Two important considerations are:
  - Requirements for data restoration
  - Recovery time requirements
    - Tape backup?
    - Electronic vaulting?
    - Shared disk?
    - Disk mirroring?
    - Remote disk mirroring?
- Business Impact Assessment (BIA) is the planning phase.

The business continuity team (which includes the business process owner) must translate the business requirements into an overall business continuity plan that includes the technology, people and business processes for recovery.

Two of the most important considerations are:

1. Requirements for data restoration (such as: whether the business requires restoration up to the point of the disaster or can restore from a previous data backup)
2. Recovery time requirements

These considerations determine the technologies and methods used to support the disaster-recovery plan. For example, if the business requires near-continuous recovery of data (with no lost transactions), it will likely use remote mirroring of data and wide-area clusters that enable a “hot” standby application environment. The shorter the recovery time and the less the transaction loss, the higher the cost of the recovery solution. As more and more businesses rely on critical applications to generate revenue (for example, e-commerce and extranets), the requirement for less-than-24-hour and continuous (such as minutes) recovery times is increasing.

A business impact assessment (BIA) is the planning phase that identifies, which business processes, users and applications are critical to the survival of the business. This process feeds into the next planning step for business continuity. Performing a BIA can help identify “costs of data unavailability” (downtime), which can be used to support decisions for various recovery solutions.

- Goals of Disaster Recovery
  - Management of both planned and unplanned situations with minimal or zero disruption to business
  - The ideal scenario when an unplanned event does occur is:
    - Recovery happens almost automatically with no loss of data.
    - Costs of solution and resources are minimal.
    - Impact to the production environment to achieve this is zero.

6

- Recovery Objectives
  - Recovery Time Objective (RTO)
    - Time within which business functions or applications must be restored
  - Recovery Point Objective (RPO)
    - Point in time to which data must be restored to successfully resume processing
  - Cost Involved

7

- RPO and RTO examples

Business Type	RPO	RTO
Greeting Card Manufacturer	Zero	Three days
Online Stock Brokerage	Zero	Seconds
Lottery	Zero	Minutes
ATM Machine	Minutes	Minutes
Semiconductor Fab Plant	Zero	Minutes

8

Disaster Recovery (DR)

- RPO and RTO examples and DR technologies

RTO / RPO	72 hrs	12 hrs	4 hrs	1 hrs	secs
24 hrs	Tape Backup and Vendor DR Site	Tape Backup and Hot Site			
1 hrs			DB Log Shipping Data Vaulting with Hot Site	DB Log Shipping Data Vaulting with Hot Site	
0 hrs		Synchronous Data Mirroring	Synchronous Data Mirroring	Disaster Tolerant Cluster with Disk mirroring	Disaster Tolerant Cluster with bi-directional mirroring

9

Disaster Recovery

- Diversity in data protection requirements

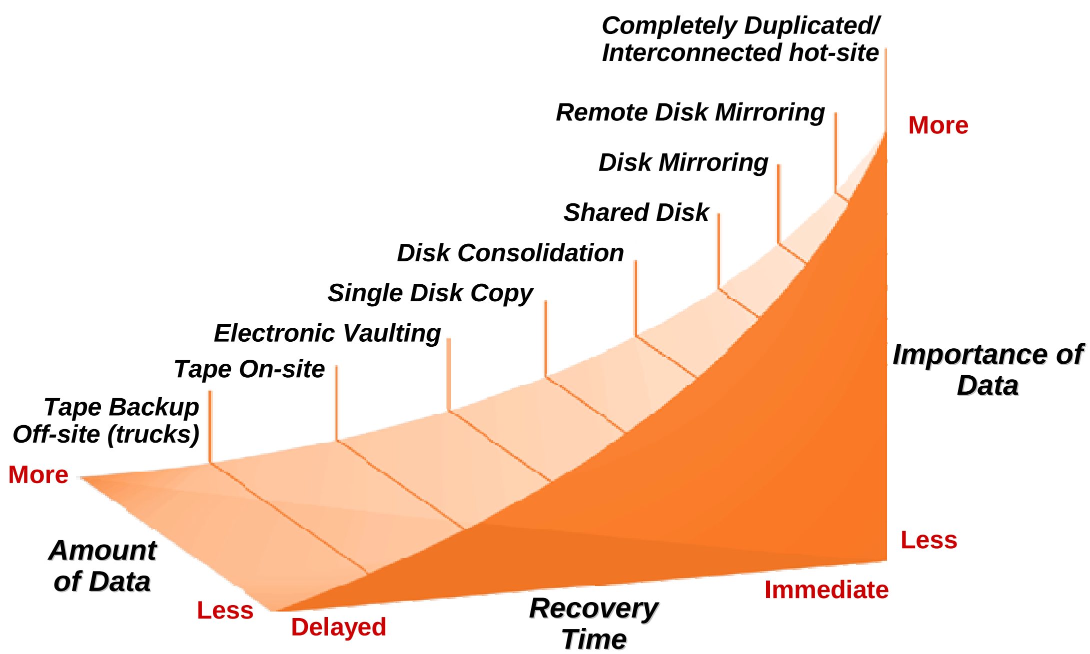

10

All data is not "created equal." It is likely that only a portion of your data is critical to the basic operation of the company. It is key to think through the data protection requirements, for different classes of data. We think it is quite likely that in most scenarios you will have some subset of data that would warrant remote disk mirroring.

- Solution area of cost, performance and distance

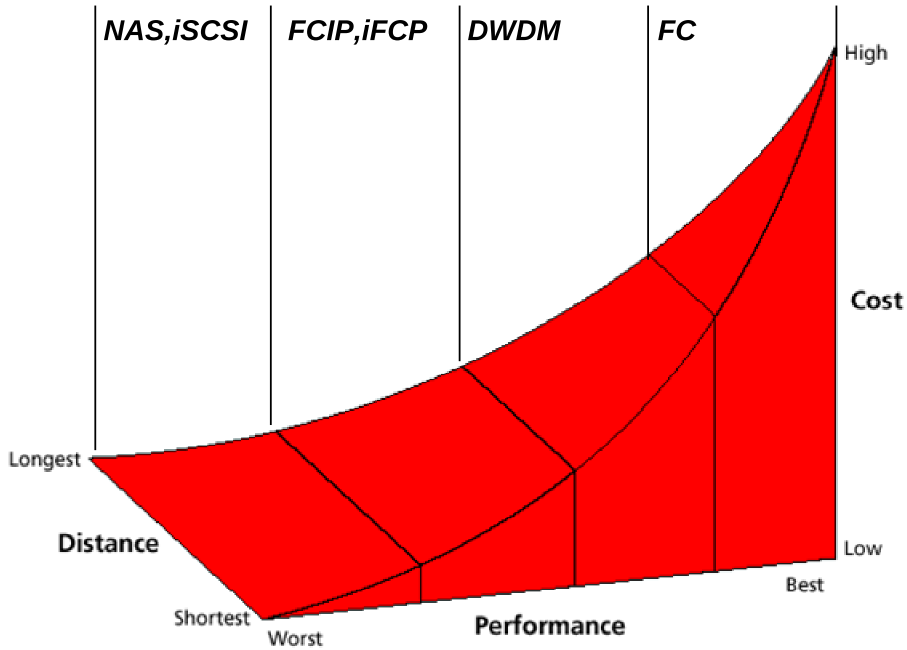

11

Fibre Channel offers the best performance for both short and long distances, but has the highest cost. NAS and iSCSI are at the low end of the cost spectrum and while sacrificing performance for both short and long distances. FCIP, iFCP, and DWDM are solutions, which perform better than NAS and iSCSI but are at a higher cost. However, they are cheaper when comparing them to Fibre Channel.

FCIP: Fibre Channel over IP

iFCP: Internet Fibre Channel Protocol

Remote Replication Terms

- Remote Copy – Refers to the mirroring of data, typically in real time.
- PiT Copy – Point-in-Time Copy refers to a copy of data that that is taken at a specific point in time.
- Data Duplication – Software duplicates data, as in remote copy or PiT snapshots..
- Data Migration – Software migrates data from one storage device to another.

12

Remote Copy: The mirroring of data, typically in real time, to provide an I/O-consistent remote copy of that data. The purpose of remote copy is to protect the data in the event of a business interruption at the customer's production location.

PiT Copy: Point-in-Time Copy refers to a copy of data that that is taken at a specific point in time. Ideally, this copy should be I/O consistent. PiT copies are used in many ways, including backups and checkpoints. More recently, PiT copies have been used in architected disaster recovery solutions.

Data Duplication: This software duplicates data, as in remote copy or PiT snapshots. Data duplication differs from data migration, in that with data duplication, at the end of the process there are two copies of data and with data migration, there is only one.

Data Migration: This software migrates data from one storage device to another. Data migration differs from data duplication in that at the end of the process there is only one copy of data. The purpose of data migration is to reduce operational complexity and costs for storage subsystem upgrades or equipment refurbishment.

Description

- TrueCopy Synchronous Software (at LU level)
  - Purpose:
    - Data migration
    - Data replication
    - Disaster recovery (DR)
  - Benefits:
    - Reliable, seamless management
    - High performance, cost effective solution

13

Data migration is typically a one time occurring event.

Data replication occurs repeatedly at a set schedule.

- TrueCopy Synchronous Software

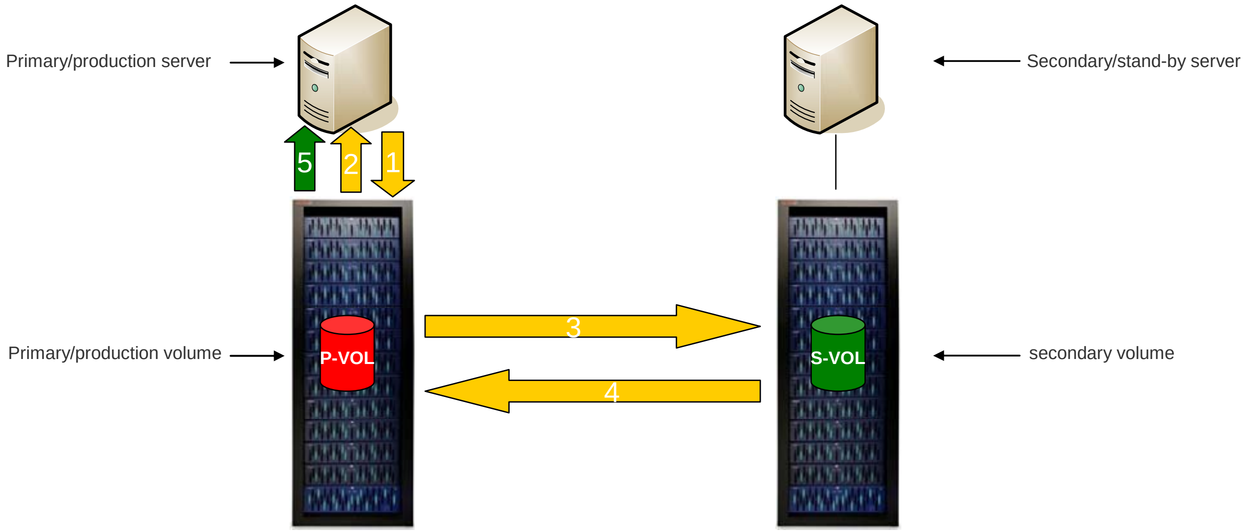

1 – Write to primary
2 – Disconnect from path – path free
3 – Write to secondary
4 – Write complete on secondary
5 – I/O Completion – Application Posted

14

Advantages

- Smallest exposure to "data loss"
- Relatively simple implementation
- Fast disaster-recovery

Disadvantages

- Primary host performance impact
- Short distances only
- Data integrity risk

- Remote Copy Operation Types

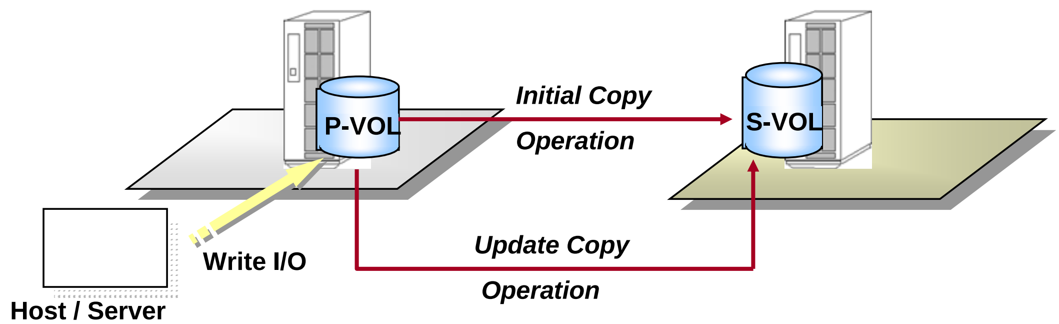

- Initial Copy Operations
  - Syncs P-VOL and S-VOL independently of host I/O
  - Add a pair (paircreate) resume and split pair
- Update Copy Operations
  - Host issues write I/O to P-VOL
  - Sync Mode: P-VOL and S-VOL kept in sync

15

There are two types of copy operations:

- Initial Copy
- Update Copy

Initial Copy: All the content from the primary volume is copied to the secondary volume. As soon as the initial copy starts, the secondary volume becomes unavailable for any type of host I/O.

Update Copy: There are two types of Update Copy. If the updates made to the primary volume is updated to the secondary volume at the same time.

The update copy makes the difference. If the update copy is done synchronously, it is called TrueCopy Synchronous, and if the update copy is done asynchronously, it is called an TrueCopy Asynchronous.

Specifications

- TrueCopy specifications for Adaptable Modular Storage
  - Volume pair control using SNM2 GUI and CLI, or CCI
  - P-VOL : S-VOL = 1 : 1
  - Supports the LUN Expansion (LUSE)
  - Two LUs, which compose a TrueCopy volume pair must be of the same capacity in blocks

16

Volume Grouping

- Perform operations for pairs or a group of pairs

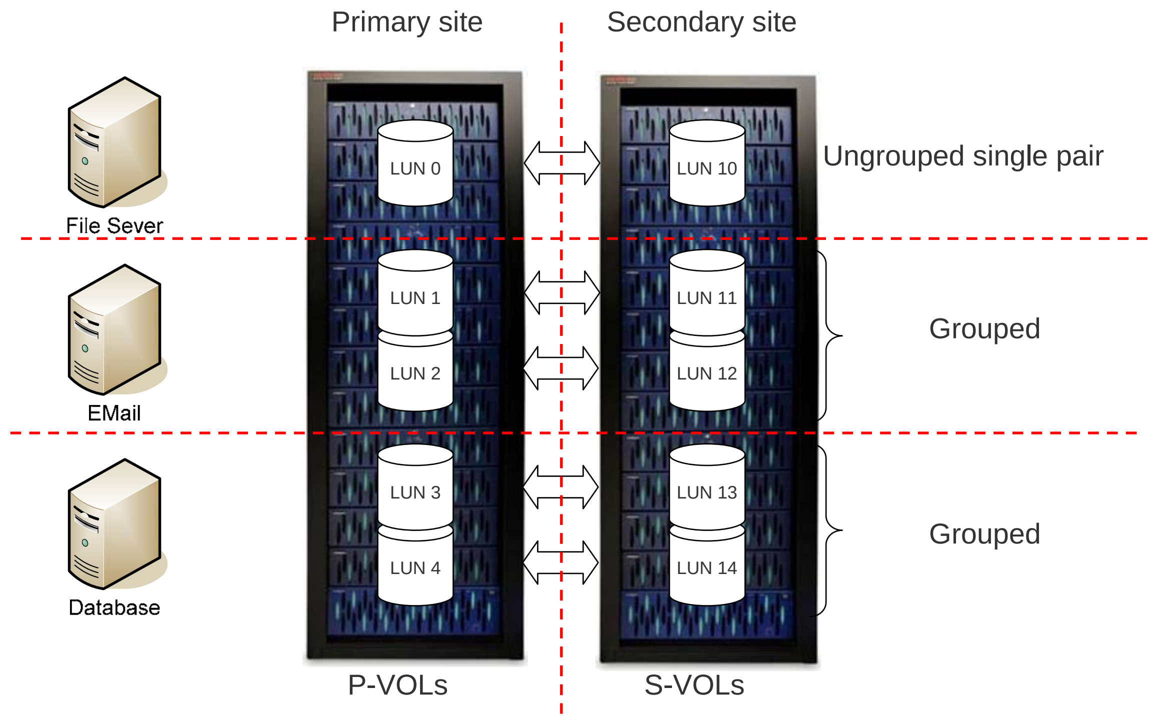

17

You can define or set up TrueCopy pairs in groups which enables you to issue commands or do operations for single pairs or a group of pairs.

Specifications

No	Item	Adaptable Modular Storage 2000 Family
1	User interface of Pair Operation	Storage Navigator Modular 2 (GUI, CLI) CCI
2	Controller configuration	Dual controllers
3	Maximum number of pairs	Model 2100: 2046 Models 2300 & 2500: 4094
4	Local system: Remote system	1:1
5	P-VOL: S-VOL	1:1
6	RAID level	RAID-6, RAID-5, RAID-1, RAID-1+0
7	RAID level of P-VOL to S-VOL	Any to any
8	Drive type for a P-VOL/S-VOL	Any
9	LU size	The size of P-VOL and S-VOL must be the same.

18

No	Item	Adaptable Modular Storage 2000 Family
10	Fibre Channel link	Supported
11	iSCSI link	Supported
12	Remote Path Bandwidth	At least 1.5Mbps per path
13	AMS 2000 family storage to previous AMS modular storage system TrueCopy software	Supported
14	AMS 2000 family storage to Universal Storage Platform TrueCopy software	Not supported
15	Maximum number of CTGs	256 per storage system
16	Swap Takeover	Supported
17	SNMP Trap	- PSUE - Path blockade
18	Pair operation and format • Pair operation during format • Format during pair operation	Not supported

Cross-generation functions revert to previous Adaptable Modular Storage specifications.

No	Item	Adaptable Modular Storage 2000 Family
19	Combination with ShadowImage pair • P-VOL of a ShadowImage pair is P-VOL of a TrueCopy pair • P-VOL of a ShadowImage pair is S-VOL of a TrueCopy pair • S-VOL of a ShadowImage pair is P-VOL of a TrueCopy pair	Supported (including Swap-resync)
20	Combination with SnapShot pair • P-VOL of a SnapShot pair is P-VOL of a TrueCopy pair • P-VOL of a SnapShot pair is S-VOL of a TrueCopy pair • V-VOL of a SnapShot pair is P-VOL of a TrueCopy pair	Supported (including Swap-resync)
21	Concurrent use with TrueCopy Extended	Not supported

20

Operations

- Steps to perform TrueCopy Synchronous operations on an Adaptable Modular Storage 2000 family system
  - Preliminary
    - Information gathering
    - Map S-VOLs to host LUNs at the remote site
  - TrueCopy
    - Install the necessary license key with SNM2
    - Create and configure command devices
    - Create and configure DM-LU
  - Select ports for the remote copy paths.
  - Configure the remote copy paths with SNM2.
  - Install and configure CCI to use a command device on both the local and remote sites.
  - Set up pairs.

21

Raid Manger CCI is optional. You can use SNM2 and the GUI. Command Devices are only necessary for Raid Manger CCI.

Preliminary Steps

- Check the licenses on the local and the remote storage system.
- You must set up and configure:
  - DM-LU
  - Command Devices
  - Remote copy paths

  on the local *and* the remote system.

22

- Detailed Preliminary Steps
  - Find the following information from the customer
    - Identify the master and the remote systems
    - Use SNM2 to identify the system serial numbers
  - Use SNM2
    - At Local Site – record the P-VOL port, LUN, host group and size
    - At Remote Site – record the S-VOL port, LUN, host group and size
    - Make sure the P-VOL and S-VOL are exactly the same size

23

- Gathering Information

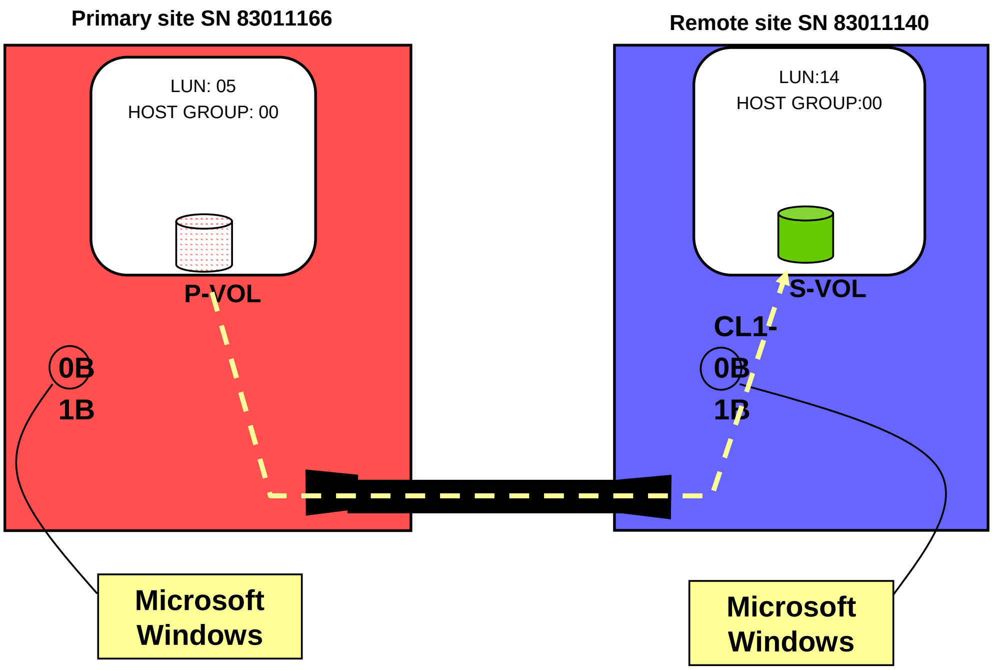

24

- The license for TrueCopy must be installed and enabled.

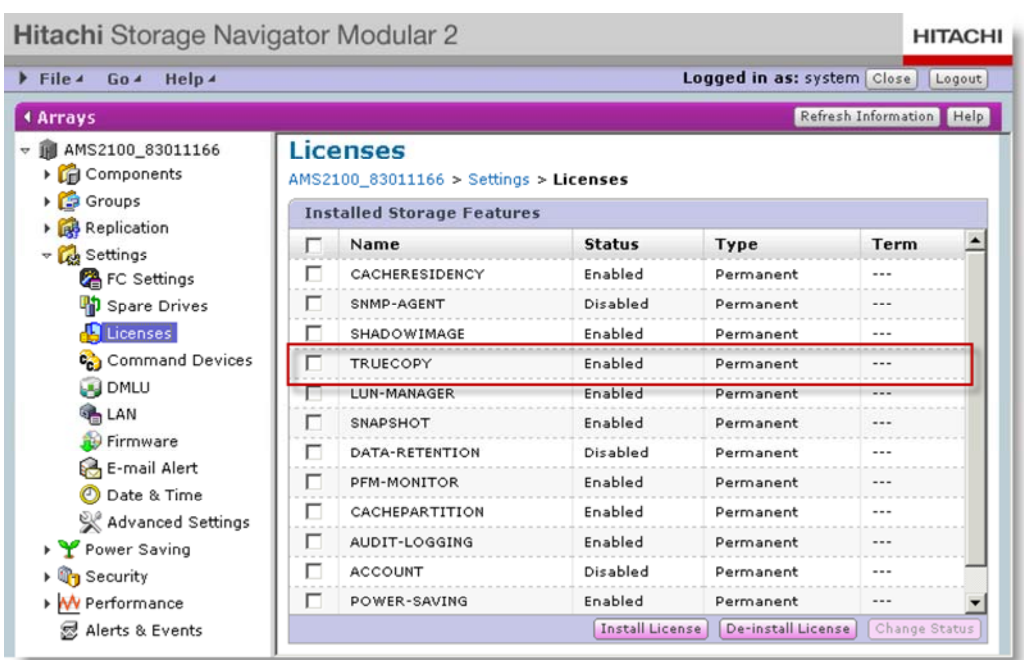

Name	Status	Type	Term
CACHERESIDENCY	Enabled	Permanent	---
SNMP-AGENT	Disabled	Permanent	---
SHADOWIMAGE	Enabled	Permanent	---
TRUECOPY	Enabled	Permanent	---
LUN-MANAGER	Enabled	Permanent	---
SNAPSHOT	Enabled	Permanent	---
DATA-RETENTION	Disabled	Permanent	---
PFM-MONITOR	Enabled	Permanent	---
CACHEPARTITION	Enabled	Permanent	---
AUDIT-LOGGING	Enabled	Permanent	---
ACCOUNT	Disabled	Permanent	---
POWER-SAVING	Enabled	Permanent	---

25

Configure Remote Paths

1. Select **Remote Path** and **Create Path**.

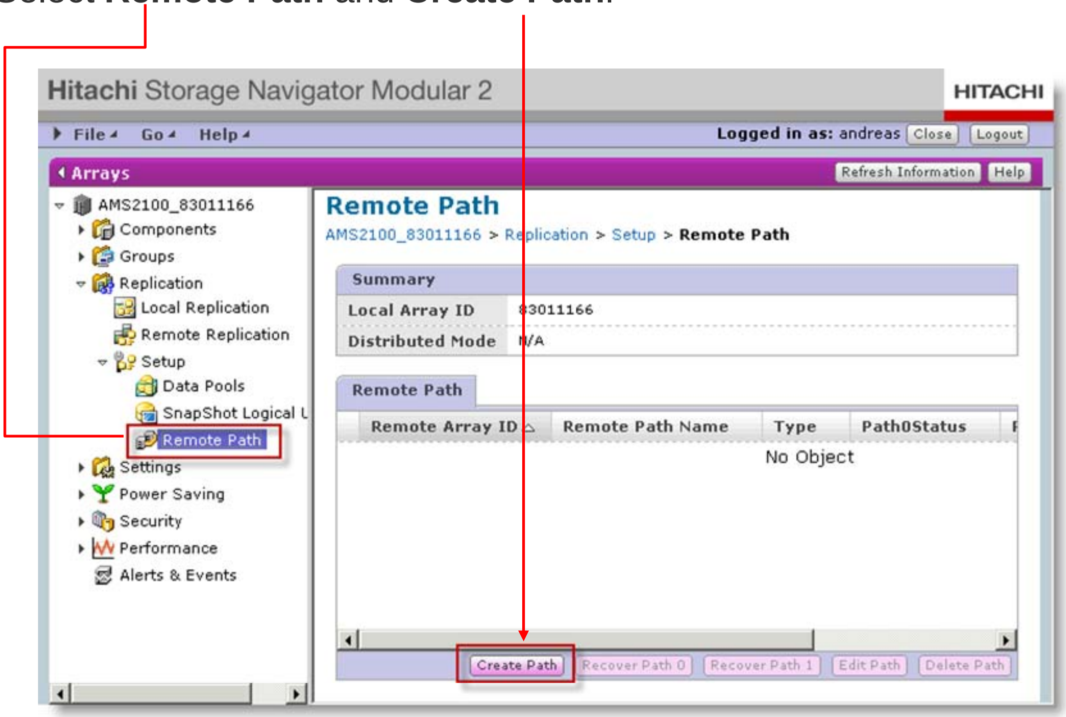

26

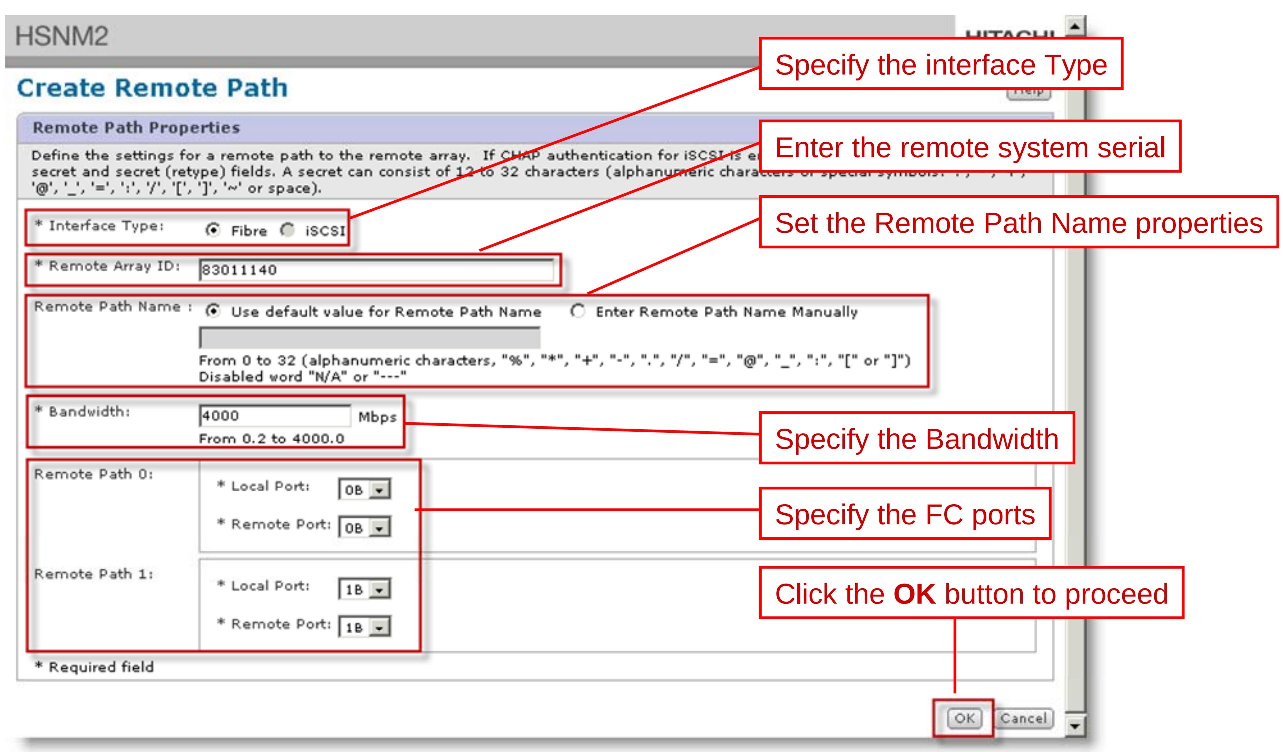

27

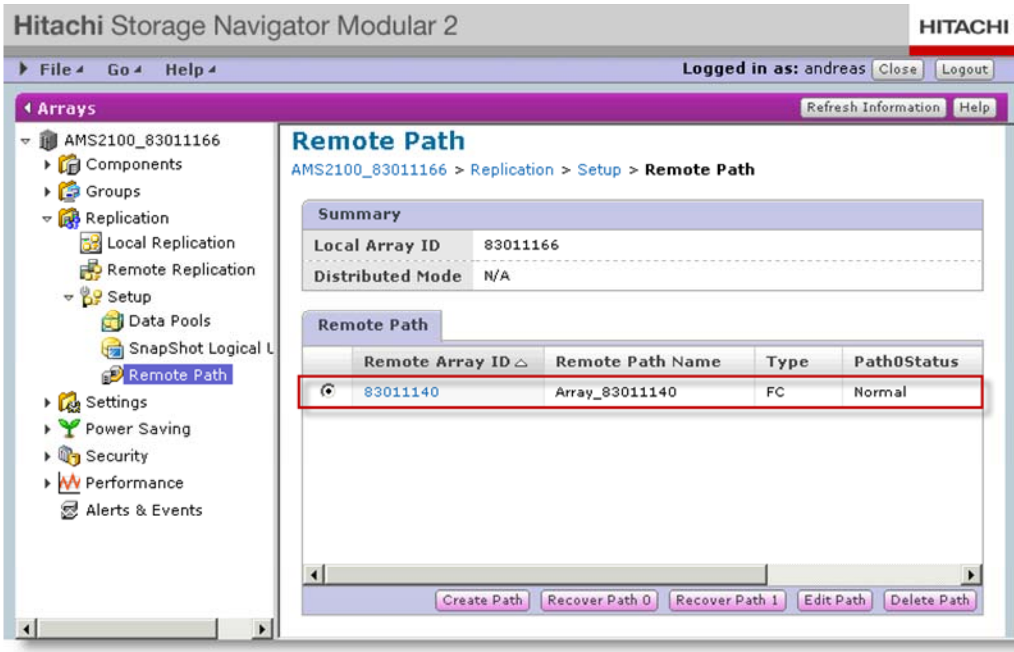

To setup the remote path using SNM2 CLI you can use the CLI command as follows:

aurmtpath -unit „unit_name“ -set -remote “remote_array_ID” -path0 0B 0B -path1 1B 1B -band 40000

28

When using the SNM2 CLI to setup the remote path, the bandwidth gets specified in 0.1 Mbit units.

In the GUI you use 1Mbit units. That’s why you specify 4000 in the GUI and in the CLI 40000 to set up a path with the same bandwidth.

You will find more information on managing remote copy paths using the SNM2 CLI in the *Hitachi Navigator 2 Command Line Interface (CLI) Reference Guide for Replication* on page 2-6

Port Configuration

- Fibre Channel Port Configuration
  - Both sending and receiving ports can function as target or initiator.
  - Can use the same physical link to do remote copy in both the directions.

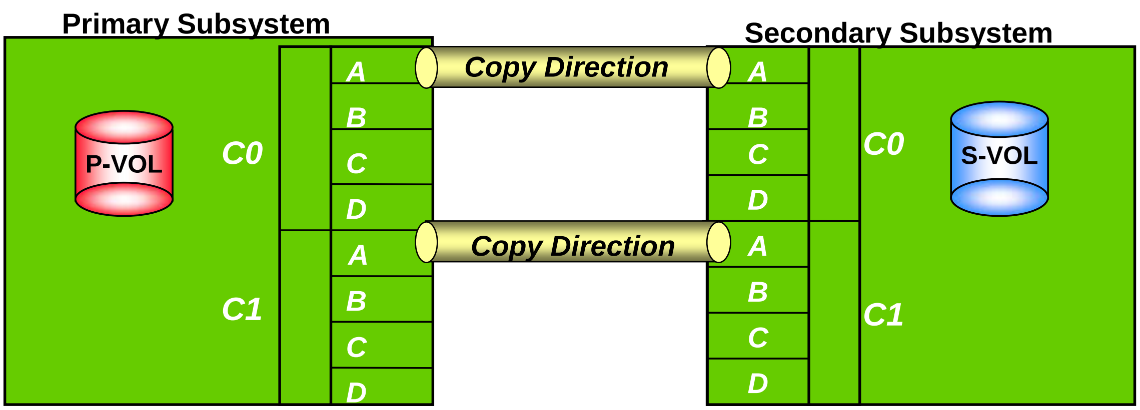

29

Volume States

- TrueCopy volume pairs have five states
  - SMPL
  - COPY
  - PAIR
  - PSUS
  - PSUE

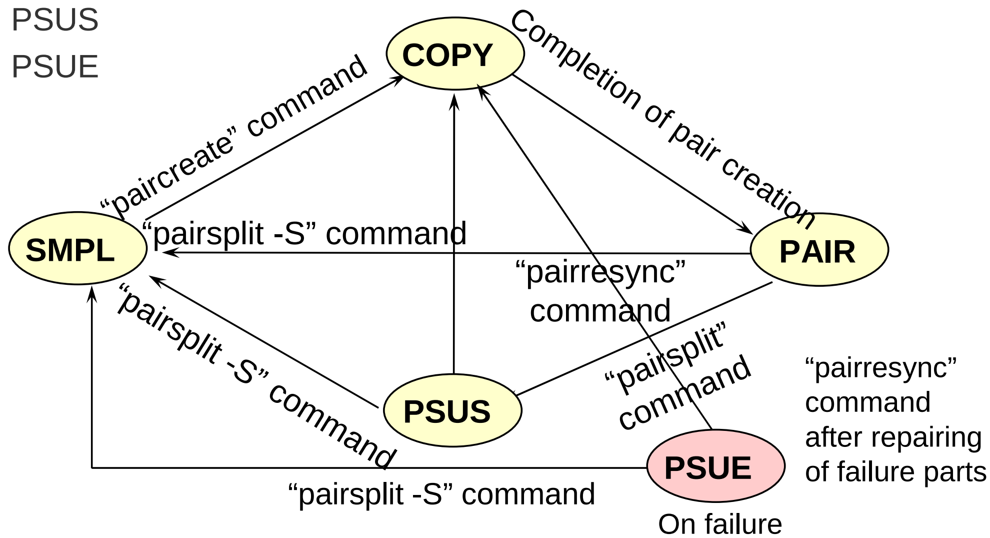

30

TrueCopy volume pairs have typical five states. These states are used to manage health of TrueCopy volume pairs.

- **SMPL (simplex)** – This volume is not currently assigned to a TrueCopy volume pair.
- **COPY** – The initial copy operation for this pair is in progress. This pair is not yet synchronized. During this status, the P-VOL has read and write access, and S-VOL has read only status.
- **PAIR** – This volume is synchronized. The updates to the P-VOL are duplicated to the S-VOL.
- **PSUS** (Pair suspended-Split) – This pair is not synchronized, because the user has split this pair (Pairsplit-r), or because the user had deleted this pair from RCU (Pairsplit-s).
  - When you split a pair from the MCU, the MCU changes the status of the P-VOL and S-VOL to PSUS. When you split a pair from the RCU, the RCU changes the status of the S-VOL to PSUS. The MCU detects this (if path status is normal) and changes P-VOL status to PSUS.
  - When you delete a pair from the RCU, the RCU changes the S-VOL status to SMPL. The MCU detects this (if the path status is normal) and changes the P-

VOL status to PSUS. You must delete the pair from the MCU in order to change the P-VOL status to SMPL.

- **PSUE** (Pair suspend-error) – This pair is not synchronized, because the MCU or RCU has suspended this pair due to an error condition.
    - For TrueCopy synchronous pairs, if the MCU cannot keep the pair synchronized for any reason, the MCU changes the status of the P-VOL and S-VOL (if possible) to PSUE.
    - For TrueCopy asynchronous pairs, if the MCU detects a TrueCopy asynchronous suspension condition, the MCU changes the P-VOL and S-VOL status (if possible) to PSUE.
    - For TrueCopy asynchronous pairs, if the RCU detects a TrueCopy asynchronous suspension condition, the RCU changes the S-VOL status to PSUE, and the MCU detects this (if the path status is normal) and changes the P-VOL status to PSUE.
- **PDUB** – This TrueCopy pair consists of LUSE volumes (for example, OPEN3*n), and an individual LDEV within this TrueCopy LUSE pair has failed due to some error condition. The status of the TrueCopy LUSE volume is PAIR or COPY, and the status of one or more LDEV pairs is PSUE or SMPL.

Module Review

1. What are the critical disaster recovery requirements for business continuity?
2. What are the TrueCopy Synchronous specifications for modular storage systems?
3. Describe the steps to perform TrueCopy synchronous operations.
4. What is the difference between PSUS and PSUE volume status?

32

10. Hitachi TrueCopy Remote Replication Configurations Using CCI

Module Objectives

- Upon successful completion of this module, you should be able to:
  - Describe Hitachi TrueCopy® Synchronous software and Hitachi TrueCopy® Extended Distance software (asynchronous) configurations
  - List conditions that optimize TrueCopy operations and system performance
  - Describe Hitachi Command Control Interface (CCI) operations
  - Identify CCI commands

2

Example Configurations

- TrueCopy Synchronous Software Configurations

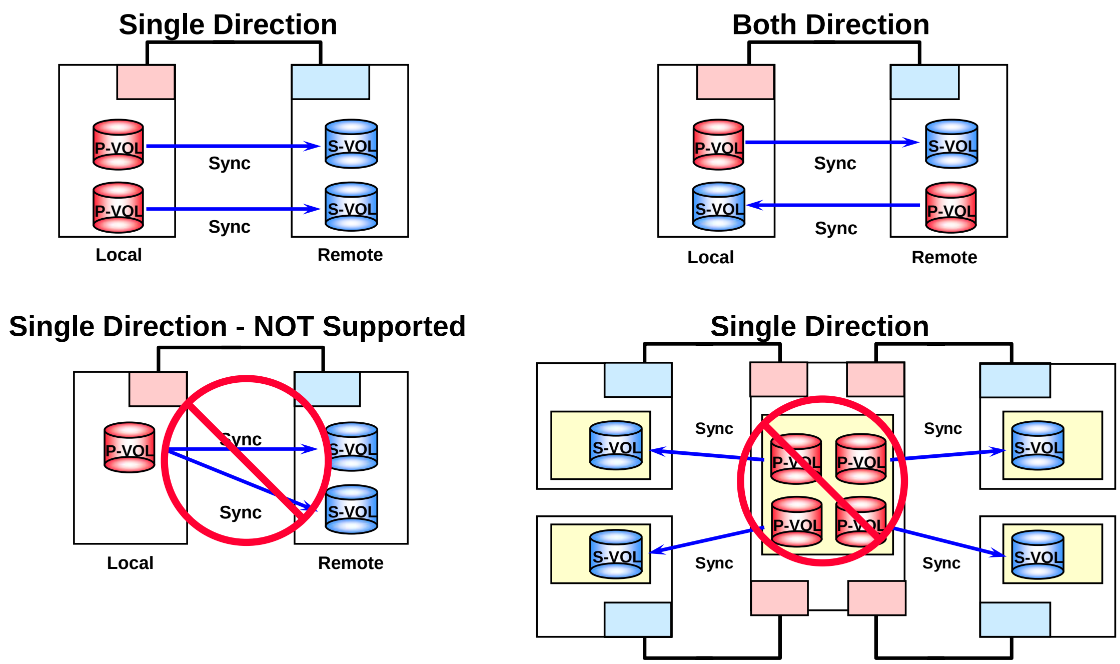

3

- TrueCopy software, TrueCopy Extended software and Hitachi ShadowImage Replication software configurations

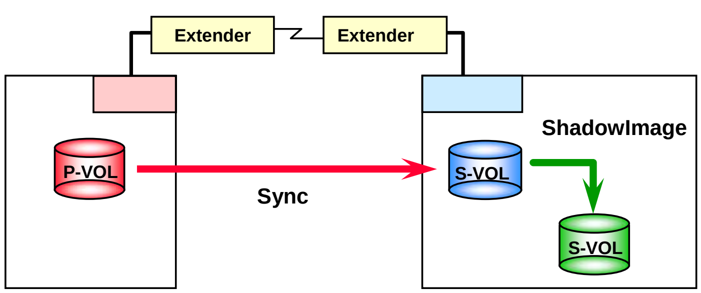

– This configuration allows ShadowImage Replication software to be used to provide multiple backup copies of a single TrueCopy P-VOL at a remote site.

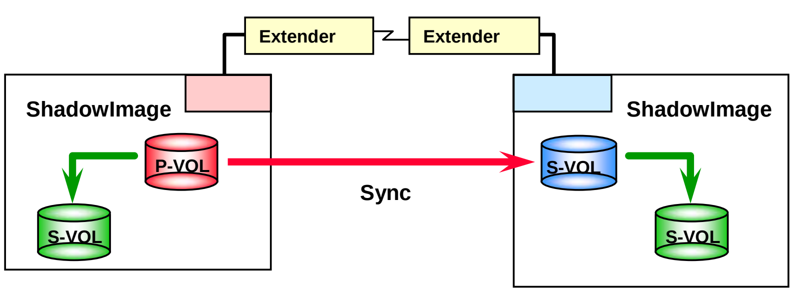

– This configuration allows ShadowImage Replication software to be used to provide multiple backup copies of a single TrueCopy P-VOL at local and remote sites.

4

Optimizing Operations and System Performance

- Write-intensive Workloads
  - Write-intensive workloads such as database logging volumes can have a significant impact on storage system I/O response times.
    - A recommendation to lessen impact would be to minimize queuing by spreading write-intensive data across several volumes.
- Large Block Size
  - Workloads with large write block sizes such as DB deferred writes, can affect performance.
    - A recommendation to lessen performance impact would be to spread workloads with large write block sizes across several volumes.

5

- Sequential write operations
  - TrueCopy operations can have a negative impact on workloads with a high percentage of sequential write operations, such as batch processing operations (for example, dump/restore and sort operations).
  - A recommendation to lessen the negative impact would be to avoid performing restore operations to volumes, which belong to TrueCopy pairs. Instead, restore data to a scratch volume and then create the TrueCopy volume pair.

6

- Link latency (0-100Kms)
  - The formula for link latency is:
    l = ((250 * 2) + (L * 5)) * 10-6
    where L is the fiber optic cable length

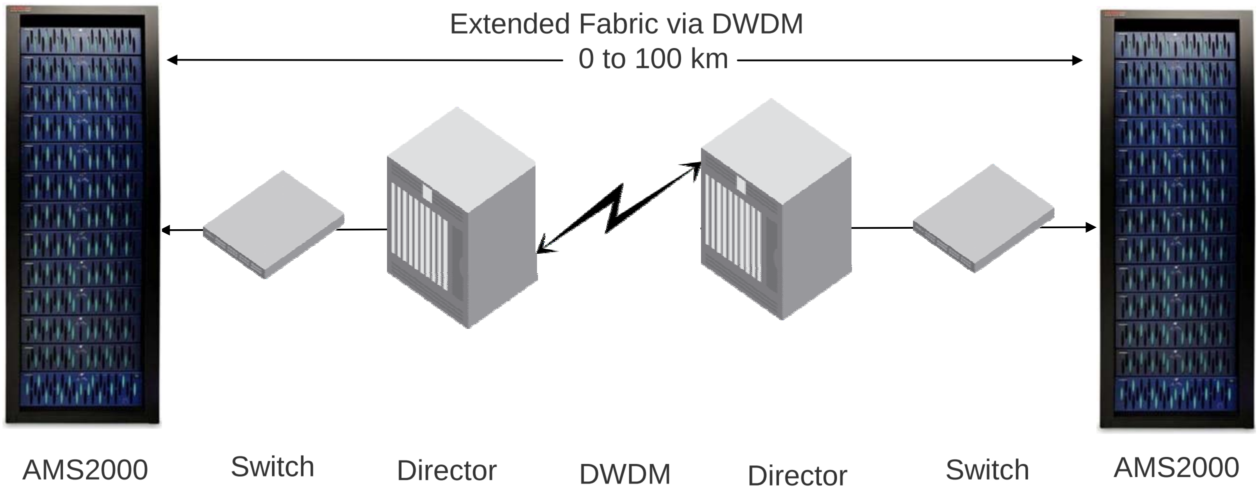

7

The scope of this analysis is to understand the performance impact of increasing the length of the link that interconnects two cascaded switches from ~0 to 100 km. The performance impact of this solution comes from the additional latency due to the link length and extended SAN technology. In addition, protocol implementations – SCSI and Fibre Channel – may also add synchronization delays.

The latency introduced by a fiber optic cable is five nanoseconds per meter (or 5µsec/km). Technologies, such as Dense Wave Division Multiplexer (DWDM) converters[1], add 250 nanoseconds of latency (this is equivalent to the latency introduced by 50 meters of optical fiber). A formula for computing the additional latency introduced by the link is:

l = ((250 * 2) + (L * 5)) * 10-6 (10 to the power of 6)

where:

I is the latency in milliseconds

250 is the latency introduced by the Extended SAN Technology in nanoseconds (DWDM)

5 is the latency of a meter of fiber optic Cable in nanoseconds

L is the fiber optic Cable length in meters

[1] Fibre Channel repeaters should add the same latency as DWDM. Extended long-wave GBICs do not add latency.

Basic Performance

- Configuration

<Host>
CPU : Xeon2.4G
Cache : 2GB
HBA : QLA2342 x1

<AMS2100>
Micro program : 20041109.02_V97
SAS Drive : 146GB 15000rpm
RAID level : RAID5(4D+1P)
LU size : 8GB/LU
1LU/1RG

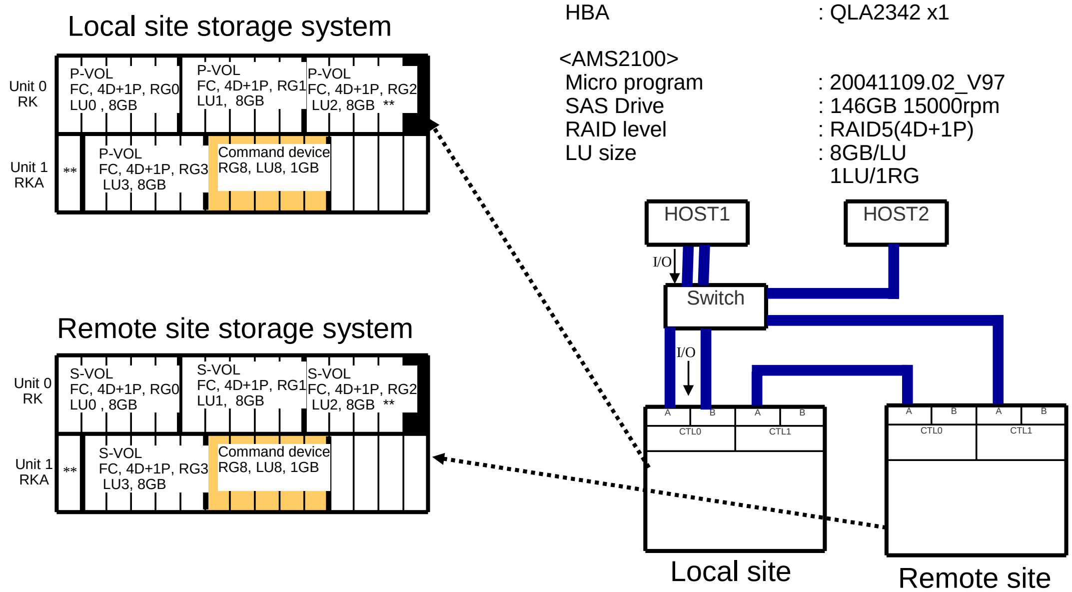

8

This is just a sample and does not reflect customer’s environments.

The results and achieved performance in a real customer environment may vary.

- Response Performance:

Parameters P-VOL receives host I/O: Random Read(70%) Write(30%), 4Kbyte, tag32/LU

RAID level: 5(4D+1P)
Drive type: P-VOL FC, S-VOL FC

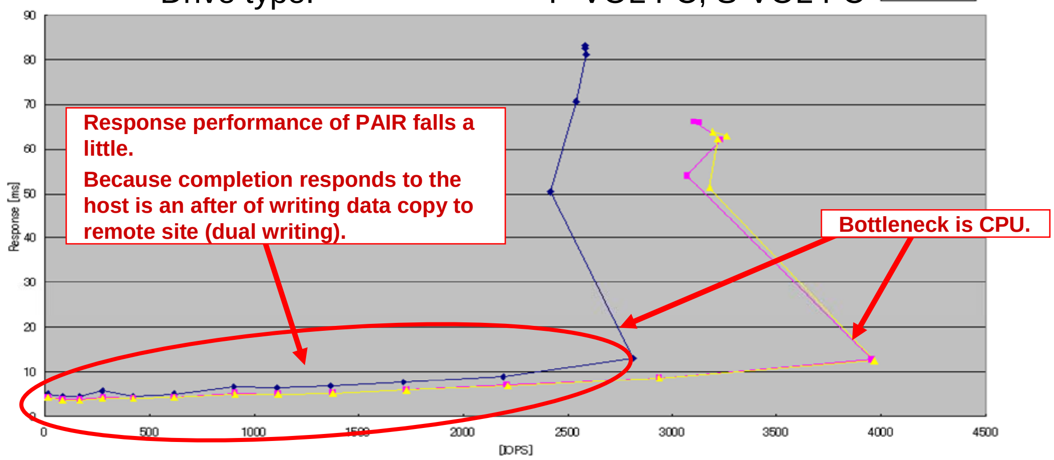

IOPS performances of each status is good up to 2700, because of the delay below 20 milliseconds.

9

This is just a sample and does not reflect customer's environments.

The results and achieved performance in a real customer environment may vary.

- Parameters: P-VOL receives host I/O (random)
  - RAID level: 5 (4D+1P)
  - Drive type: P-VOL FC, S-VOL FC
- Performance falls as the ratio of writes becomes larger

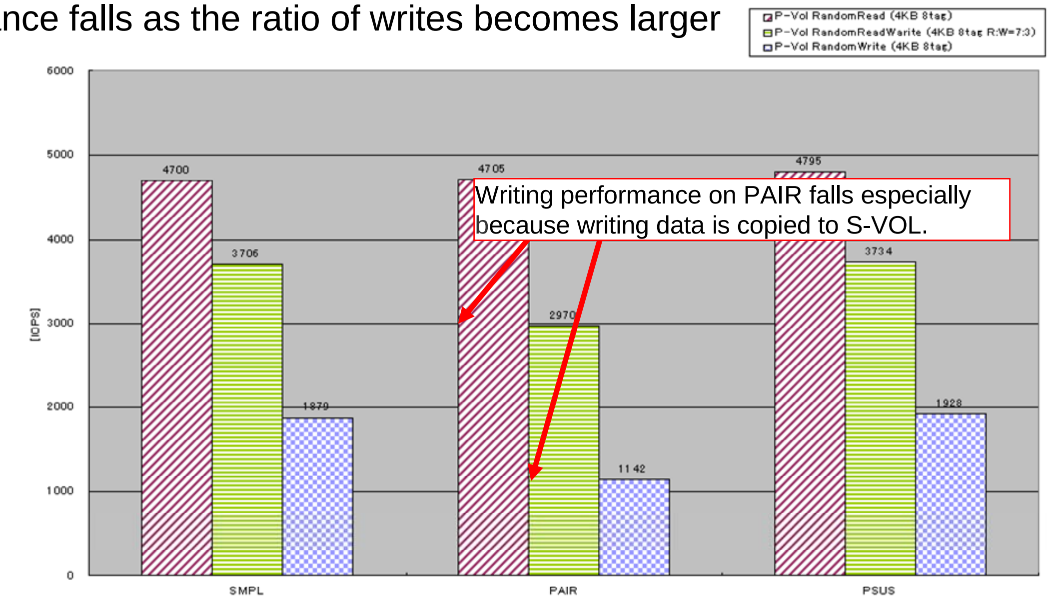

10

This is just a sample and does not reflect customer's environments.

The results and achieved performance in a real customer environment may vary.

- Parameters: P-VOL receives host I/O (sequential)
  - RAID level: 5 (4D+1P)
  - Drive type: P-VOL FC, S-VOL FC
- Performance falls as the ratio of writes becomes larger

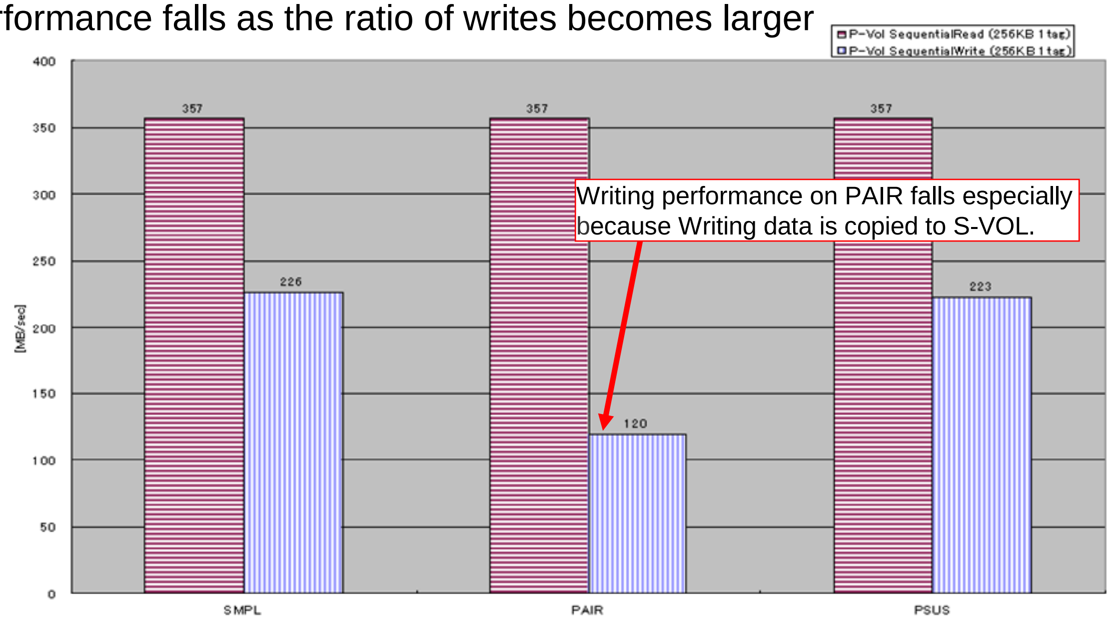

11

This is just a sample and does not reflect customer's environments.

The results and achieved performance in a real customer environment may vary.

- Interswitch Link Performance
  - Maximum number of 2KB frames that can be transmitted per second is:
    - If (N*0.02 < (2*l + 0.02)), X = [N / (2*l + 0.02)] * 103
    - If (N*0.02 >= (2*l + 0.02)), X = 50,000
      - Where ‘X’ is the maximum number of frames per second
      - ‘N’ is the buffer credit value
      - ‘l’ is the link latency in milliseconds

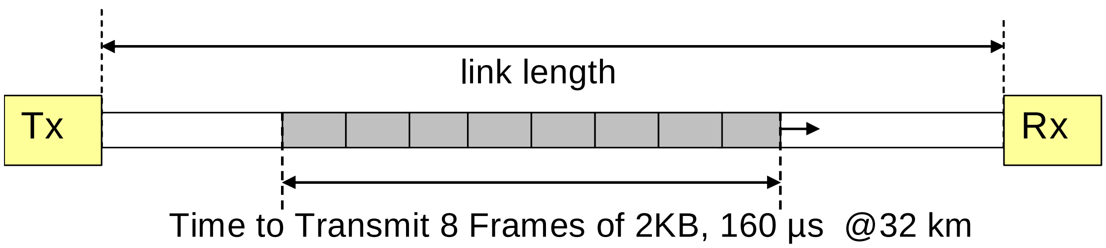

12

The buffer credits available for an (inter-switch link) E_Port (and the latency of the link) define the performance limit of the interswitch link. The buffer credits represent the maximum number of outstanding frames that can be transmitted by a port. Therefore, if the round trip delay of a frame is greater than the time necessary to transmit the maximum number of outstanding frames, it is not possible to keep the “pipe full” and the port will block waiting for the acknowledgement of the first transmitted frame.

Overview of CCI

- Optional feature
- Installed on a host with Fibre Channel connection to the command devices
- Manages the TrueCopy software
- Requirements
  - TrueCopy software license is installed and enabled.
  - Remote paths are configured using Storage Navigator Modular 2 (SNM2).
  - Command Devices are configured using SNM2
- Interfaces with high availability software.
- Supports various TrueCopy commands.
- Requires at least one LU configured as the command device on each storage system and at least one LU configured as DM-LU on each storage system
- Issues an error message when an error occurs in a volume pair

13

- CCI is installed on the host

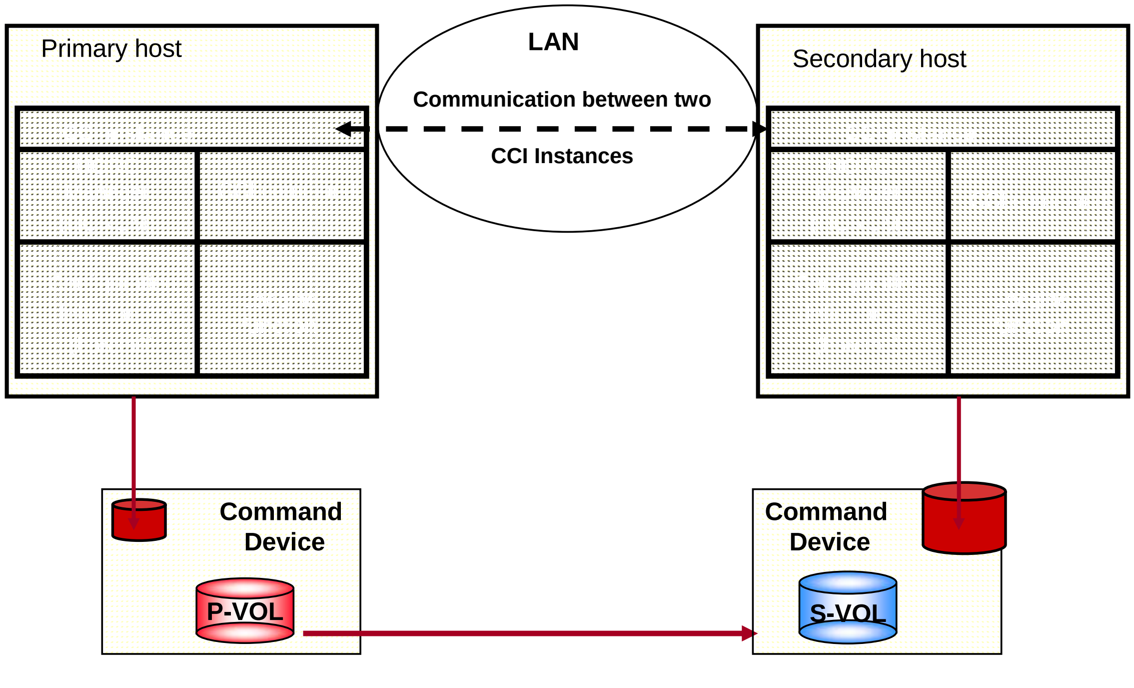

The host, which has CCI installed upon it, is also known as the management server and is connected to the storage system using Fibre Channel to allow in-band management; they also communicate to each other through a regular LAN connection.

The only LUN that the management server needs to see for pair management is the command device. An instance, (referred to as the HORCM instance) runs on the server and is used to issue CCI commands to the command device. The command device accepts commands from the instance and reports results to the HORCM instance.

In the above configuration, the two management servers at both sites are connected to their respective storage systems.

- Case 1
  - 2 hosts and 2 subsystems

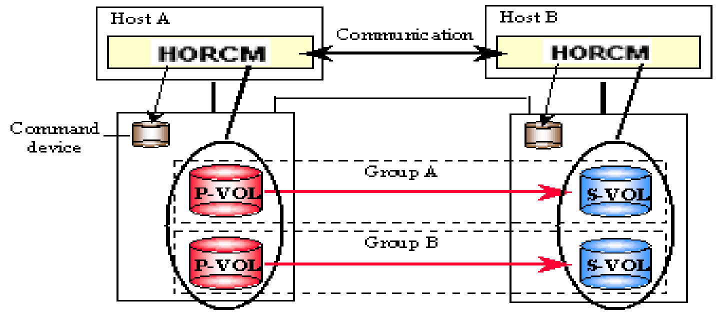

- Case 2
  - 3 hosts and 2 subsystems

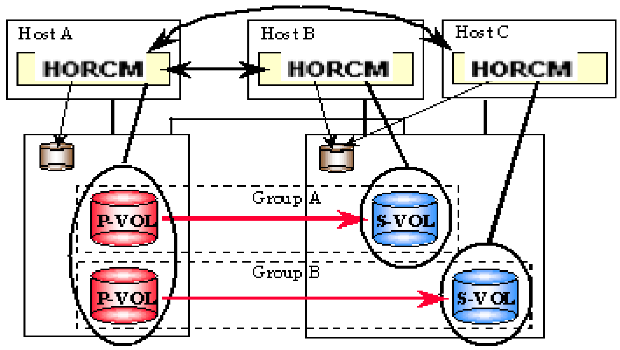

15

Case 2 – There are two hosts running one instance, each on the remote site. The primary site's instance communicates with two secondary site instances to manage multiple pairs independently. Note that there is no LAN connection between the two secondary hosts.

- Case 3
  - 2 hosts and 2 subsystems, 3 instances

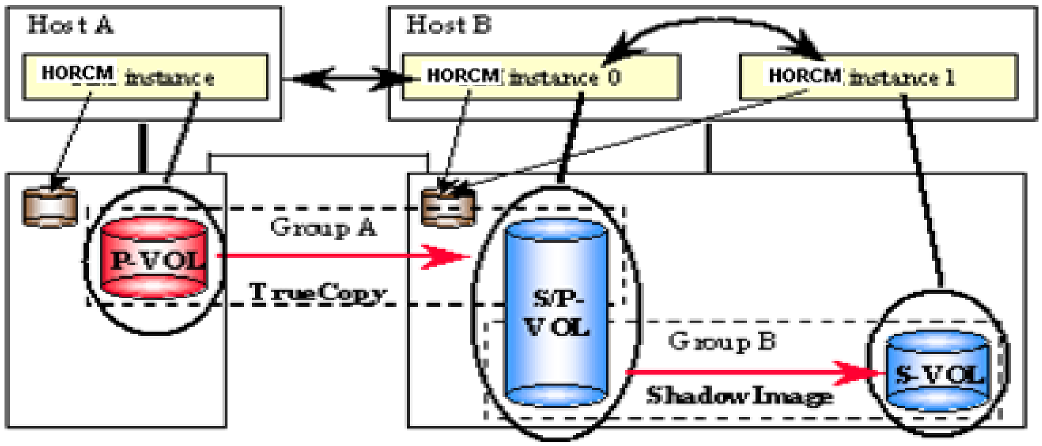

16

Case 3: Host A and B manage TrueCopy pairs between primary site A and secondary site B. Host B and C manage TrueCopy pairs between primary site B and secondary site C.

Case 4: Host A instance and host B instance manage TrueCopy pairs whereas Host B instance and second instance manage ShadowImage pairs.

- CCI configuration file at primary site host

This is LAB A and has instance horcm0 running on it

This is LAB B and has instance horcm1 running on it.

```
HORCM_MON
#ip_address   service   poll(10ms)   timeout(10ms)
LABA          horcm0    1000         3000

HORCM_CMD
#dev_name
\\.\PHYSICALDRIVE3        /dev/rdsk/c0t0d1s2
\\.\CMD-77010047-2-CLI-C              #Windows 2000 and 20003

HORCM_DEV
#dev_group   dev_name      port#    TargetID   LU#
VG01         ntdev1        CL1-B    12         1
VG01         ntdev2        CL1-B    12         3

HORCM_INST
#dev_group   ip_address    service
VG01         LABB          horcm1
```

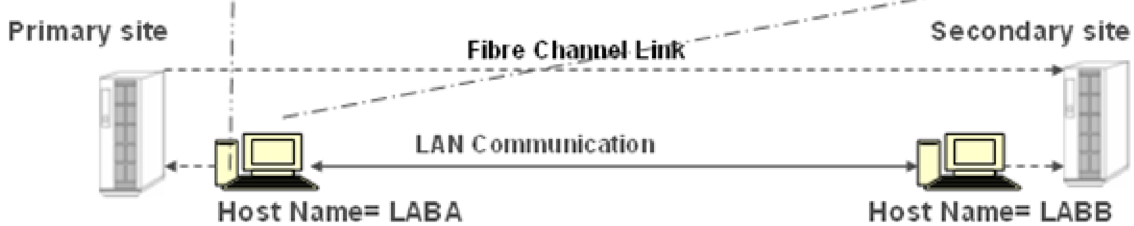

17

- CCI configuration file at secondary site host

This is LABB and has instance horcm1 running on it.

This is LABA and has instance horcm1 running on it.

```
HORCM_MON
#ip_address   service   poll(10ms)   timeout(10ms)
LABB          horcm1    1000         3000
HORCM_CMD
#dev_name
\\.\PHYSICALDRIVE3
\\.\CMD-77010047-2-CLI-C         #Windows 2000 and 20003
HORCM_DEV
#dev_group   dev_name   port#    TargetID   LU#
VG01         ntdev1     CL1-B    12         2
VG01         ntdev2     CL1-B    12         4
HORCM_INST
#dev_group   ip_address   service
VG01         LABA                  horcm0
```

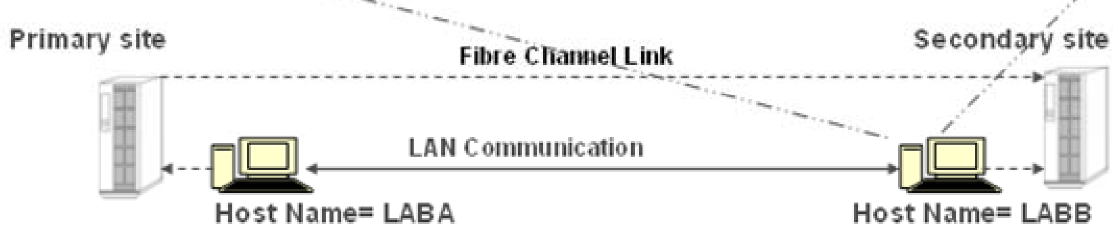

18

Hitachi Data Systems

- Important Pointers for CCI
  - Find out the absolute LUN number for the HORCM configuration file
  - Make sure that you have the right IP address and right HORCM instance name in the HORCM configuration file
  - Watch out for **extra spaces or special characters** in the configuration file It is recommended to use an plain text editor like **Notepad** or similar. |
    Do not use MS Word or WordPad.
  - Always keep in mind that you are editing the services file for the HORCM instances
  - Make sure that you set the right environmental variables
  - Make sure that when using CCI for TrueCopy software, that the HORCC_MRCF environmental variable is **not set**.
    (It is used only when performing Hitachi ShadowImage® In-System Replication software or Hitachi Copy-on-Write Snapshot software operations).

19

CCI uses absolute LUNs to scan a port.

The LUNs in a host group are mapped for the host system so that target ID and LUN, which is indicated by the raidscan command will be different from the target ID and LUN shown by the host system. The target ID and LUN indicated by the raidscan command should be used to create the CCI “config file”.

Commands

CCI Command	Description
horctakover	The host executing horctakover takes ownership of the pair.
paircurchk	Checks the consistency of the data on the secondary volume.
paircreate	Creates a pair.
pairsplit	Splits a pair.
pairresync	Resynchronizes a pair.
pairevtwait	Event waiting command.
pairmon	Monitors a pair and reports changes in the pairs status.
pairvolchk	Checks attributes of volume connected to local or remote hosts.
pairdisplay	Confirms the configuration of a specified pair.
raidscan	Lists the SCSI/fibre port, Target ID, LUN number and LDEV status.
raidar	Reports the I/O activity of a specified LDEV.
raidqry	Confirms the connection of the Thunder 9500 V Series systems/Adaptable Modular Storage and the open system host.
horcctl	Displays the internal trace control parameters.
horcmstart sh	A shell script that starts CCI.
horcmshutdown	Stops the CCI shell script.

20

- Microsoft Windows NT, 2000, and 2003 Commands

CCI Command	Description
findcmddev	Searches for the command device
drivescan	Displays the relationship between the hard disk number and the physical drive
portscan	Displays the physical device on a designated port
sync	Flashes the remaining unwritten data to the physical drive
mount	Mounts the specified device
umount	Unmounts the specified device
setenv	Sets the environmental variables
usetenv	Deletes the environmental variables
env	Displays the environmental variables
sleepetenv	Sets the sleep time for a specified environmental variable

21

- **paircreate** command
  - -g Specifies group name from HORCM.CONF
    - All pairs in group created unless restricted by another switch
  - -d Specifies volume name from HORCM.CONF
    - Restricts operation to one TID/LUN
    - Additional options to improve security
  - -f Specifies fence level - data, status, never
  - -vl Specifies that 'local' device is primary
  - -vr Specifies that 'remote' device is primary
  - -c Specifies number (1 to 15) of tracks to be copied during initial copy operation
  - -nocopy Pair is created without copying primary to secondary (dangerous!)
  - - m <mode> Mode=cyl / Mode=trk

22

paircreate generates a new volume pair from two unpaired volumes. It can be used to create either a paired logical volume or a group of paired volumes. It allows you to specify the direction (local or remote) of the pair (-vl or –vr).

Paircreate example:

C:\HORCM\etc> paircreate -g vg01 -d labadb1 -vl -f never

- **pairdisplay** command
  - -g Specifies group name from HORCM.CONF
    - All volumes in the group checked
  - -d Specifies volume name from HORCM.CONF
    - Restricts operation to one TID/LUN
  - -c Checks pair path and displays illegal pair connections
  - -l Displays volume pair status of ‘local’ host
  - -fc Shows % pair synchronization
  - -fm Displays the bitmap mode (Cyl / Trk)

23

The **pairdisplay** command displays the pair status allowing you to verify completion of pair operations (e.g., **paircreate**, **pairresync**) and can be used to confirm the configuration of the pair connection path (the physical link of paired volumes and servers).

The **pairdisplay** command can be used for a paired volume or a group of paired volumes.

- **pairdisplay** examples

```
C:\HORCM\etc>pairdisplay -g grp1 -d vol1 -fx
Group   PairVol(L/R) (Port#,TID,LU),Seq#,LDEV#.P/S,Status,Fence, Seq#,P-LDEV# M
grp1    vol1(L)    (CL1-A , 1,  0)50003  200..SMPL  ---- ------,-----  ----  -
grp1    vol1(R)    (CL1-C , 0,  0)50001  204..SMPL  ---- ------,-----  ----  -

C:\HORCM\etc>pairdisplay -g grp1 -d vol1 -fx -l
Group   PairVol(L/R) (Port#,TID,LU),Seq#,LDEV#.P/S,Status,Fence, Seq#,P-LDEV# M
grp1    vol1(L)    (CL1-A , 1,  0)50003  200..SMPL  ---- ------,-----  ----  -
```

```
C:\HORCM\etc>pairdisplay -g grp1 -fx
Group   PairVol(L/R) (Port#,TID,LU),Seq#,LDEV#.P/S,Status,Fence, Seq#,P-LDEV# M
grp1    vol1(L)    (CL1-A , 1,  0)50003  200..SMPL  ---- ------,-----  ----  -
grp1    vol1(R)    (CL1-C , 0,  0)50001  204..SMPL  ---- ------,-----  ----  -
grp1    vol2(L)    (CL1-A , 1,  1)50003  201..SMPL  ---- ------,-----  ----  -
grp1    vol2(R)    (CL1-C , 0,  1)50001  205..SMPL  ---- ------,-----  ----  -
grp1    vol3(L)    (CL1-A , 1,  2)50003  100..SMPL  ---- ------,-----  ----  -
grp1    vol3(R)    (CL1-C , 0,  2)50001  104..SMPL  ---- ------,-----  ----  -
grp1    vol4(L)    (CL1-A , 1,  3)50003  101..SMPL  ---- ------,-----  ----  -
grp1    vol4(R)    (CL1-C , 0,  3)50001  105..SMPL  ---- ------,-----  ----  -
```

24

- **pairdisplay** examples

```
C:\HORCM\etc>paircreate -g grp1 -vl -f data

C:\HORCM\etc>pairdisplay -g grp1 -fxc
Group   PairVol(L/R) (Port#,TID,LU),Seq#,LDEV#.P/S,Status,Fence,   % ,P-LDEV# M
grp1    vol1(L)    (CL1-A , 1,  0)50003  200..P-VOL COPY DATA  ,    6   204  -
grp1    vol1(R)    (CL1-C , 0,  0)50001  204..S-VOL COPY DATA  ,-----   200  -
grp1    vol2(L)    (CL1-A , 1,  1)50003  201..P-VOL COPY DATA  ,    6   205  -
grp1    vol2(R)    (CL1-C , 0,  1)50001  205..S-VOL COPY DATA  ,-----   201  -
grp1    vol3(L)    (CL1-A , 1,  2)50003  100..P-VOL COPY DATA  ,    1   104  -
grp1    vol3(R)    (CL1-C , 0,  2)50001  104..S-VOL COPY DATA  ,-----   100  -
grp1    vol4(L)    (CL1-A , 1,  3)50003  101..P-VOL COPY DATA  ,    2   105  -
grp1    vol4(R)    (CL1-C , 0,  3)50001  105..S-VOL COPY DATA  ,-----   101  -
```

```
C:\HORCM\etc>pairdisplay -g grp1 -fxc
Group   PairVol(L/R) (Port#,TID,LU),Seq#,LDEV#.P/S,Status,Fence,   % ,P-LDEV# M
grp1    vol1(L)    (CL1-A , 1,  0)50003  200..P-VOL PAIR DATA  ,  100   204  -
grp1    vol1(R)    (CL1-C , 0,  0)50001  204..S-VOL PAIR DATA  ,  100   200  -
grp1    vol2(L)    (CL1-A , 1,  1)50003  201..P-VOL PAIR DATA  ,  100   205  -
grp1    vol2(R)    (CL1-C , 0,  1)50001  205..S-VOL PAIR DATA  ,  100   201  -
grp1    vol3(L)    (CL1-A , 1,  2)50003  100..P-VOL PAIR DATA  ,  100   104  -
grp1    vol3(R)    (CL1-C , 0,  2)50001  104..S-VOL PAIR DATA  ,  100   100  -
grp1    vol4(L)    (CL1-A , 1,  3)50003  101..P-VOL PAIR DATA  ,  100   105  -
grp1    vol4(R)    (CL1-C , 0,  3)50001  105..S-VOL PAIR DATA  ,  100   101  -
```

25

- **pairsplit** command
  - -g Specifies group name from HORCM.CONF
    - All pairs in group split unless restricted by another switch
  - -d Specifies volume name from HORCM.CONF
    - Restricts operation to one TID/LUN
    - Additional options to improve security
  - -r Specifies split secondary is *read only*
  - -rw Suspended secondary read and write allowed
  - -S Places primary and secondary into simplex
  - -R Forces secondary volume into simplex
  - -l Enables pairsplit (from “local” host) when remote host is down

26

The **pairsplit** commands stops updates to the secondary volume of a pair and can either maintain (status = PSUS) or delete (status = SMPL) the pairing status of the volumes. It can be applied to a paired logical volume or a group of paired volumes. It also allows read access or read/write access to the secondary volume, depending on the selected options. Acceptance of write requests to the primary volume depends on the fence level of the pair (data, status, never, or async). If the **-S** option (simplex) is used; the volume pair is deleted, and returned to the simplex state. To synchronize the volumes, the **pairsplit** command must be issued after write I/Os to the paired volume have completed.

- **pairresync** command
  - -g Specifies group name from HORCM.CONF
    - All pairs in group resynchronised unless restricted by another switch
  - -d Specifies volume name from HORCM.CONF
    - Restricts operation to one TID/LUN
    - Additional options to improve security
  - -c Specifies number (1 to 15) of tracks to be copied during initial copy operation
  - -l Enables pairresync (from “local” host) when remote host is down

27

The **pairresync** command re-establishes a split pair and then restarts the update copy operations to the secondary volume. The normal direction of resynchronization is from the primary volume to the secondary volume. The primary volume remains accessible during **pairresync** while the secondary volume becomes write-disabled when the **pairresync** command is issued. If no data was written to the secondary volume while the pair was split, the differential data on the primary volume is copied. If data was written to the secondary volume, the differential data on the primary volume and secondary volume is copied.

- **pairsplit** and **pairresync** examples

```
C:\HORCM\etc>pairsplit -g grp1 -d vol1 -rw

C:\HORCM\etc>pairdisplay -g grp1 -d vol1 -fxc
Group   PairVol(L/R) (Port#,TID,LU),Seq#,LDEV#.P/S,Status,Fence,   % ,P-LDEV# M
grp1    vol1(L)     (CL1-A , 1,  0)50003  200..P-VOL PSUS DATA  ,   98   204  W
grp1    vol1(R)     (CL1-C , 0,  0)50001  204..S-VOL SSUS DATA  ,  100   200  -
```

```
C:\HORCM\etc>pairresync -g grp1 -d vol1

C:\HORCM\etc>pairdisplay -g grp1 -d vol1 -fxc
Group   PairVol(L/R) (Port#,TID,LU),Seq#,LDEV#.P/S,Status,Fence,   % ,P-LDEV# M
grp1    vol1(L)     (CL1-A , 1,  0)50003  200..P-VOL COPY DATA  ,   84   204  -
grp1    vol1(R)     (CL1-C , 0,  0)50001  204..S-VOL COPY DATA  ,-----   200  -
```

```
C:\HORCM\etc>pairdisplay -g grp1 -d vol1 -fxc
Group   PairVol(L/R) (Port#,TID,LU),Seq#,LDEV#.P/S,Status,Fence,   % ,P-LDEV# M
grp1    vol1(L)     (CL1-A , 1,  0)50003  200..P-VOL PAIR DATA  ,  100   204  -
grp1    vol1(R)     (CL1-C , 0,  0)50001  204..S-VOL PAIR DATA  ,  100   200  -
```

28

- **pairevtwait** command

  pairevtwait { -h . -q . -z . -g <group> . -d <pair Vol> . -d[g] <raw_device> [MU#] . -FHORC | -FMRCF [MU#] | -d[g] <seq#> <LDEV#> [MU#] . -s <status> ... . -t <timeout>[interval] . -nowait . -l . -nomsg }

- The **pairevtwait** command is used to wait for completion of pair creation and pair resynchronization and to check the status.
  - It waits (“sleeps”) until the paired volume status becomes identical to a specified status and then completes. The pair event waiting command waits until the specified status is established, and terminates abnormally if an abnormal status is detected. When the event waiting command with the **`-nowait`** option is issued for a group, the status is returned if the status of each volume in the group is identical.

29

When the **-nowait** option is specified, then Normal termination:

1. The status is SMPL.
2. The status is COPY or RCPY.
3. The status is PAIR.
4. The status is PSUS.
5. The status is PSUE.

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

Abnormal termination: other than 0 to 127, refer to the execution logs for error details.

When the **-nowait** option is not specified:

Normal termination: 0.

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

Abnormal termination: other than 0 to 127, refer to the execution logs for error details.

- **pairevtwait** options
  - -h: Displays Help/Usage and version information.
  - -q: Terminates the interactive mode and exits this command.
  - -z or -zx: Makes the **pairevtwait** command enter the interactive mode. The -zx option guards performing of the HORCM in the interactive mode. When this option detects a HORCM shut down, interactive mode terminates.
  - -g <group>: Specifies a group name defined in the configuration definition file. This option must always be specified. The command is executed for the specified group unless the -d <pair Vol> option is specified.
  - -d <pair Vol>: Specifies a paired logical volume name defined in the configuration definition file. When this option is specified, the command is executed for the specified paired logical volumes.

30

- **pairevtwait** options (*continued*)
  - -d[g] <raw_device> [MU#]:

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

    Forcibly specifies a cascading TrueCopy Synchronous/Asynchronous software volume for specified pair logical volumes on ShadowImage Replication software environment. If the **-l** option is specified, this option tests status of a cascading TrueCopy volume on a local host (near site). If no **-l** option is specified, this option tests status of a cascading TrueCopy software on a remote host (far site). The target TrueCopy volume must be P-VOL or SMPL.

31

- **pairevtwait** options (*continued*)

  -FMRCF [MU#] or –FBC [MU#]:

  Forcibly specifies a cascading ShadowImage volume for specified pair logical volumes on TrueCopy environment. If the **-l** option is specified, this option tests status of a cascading ShadowImage volume on a local host (near site). If no **-l** option is specified, this option tests status of a cascading ShadowImage volume on a remote host (far site). The target ShadowImage volume must be P-VOL or SMPL.

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

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

32

- **pairevtwait** options (*continued*)

-s <status>:	Specifies the waiting status, which is "smpl", "copy/rcpy", "pair", "psus", or "psue/pdub". If two or more statuses are specified following **-s**, waiting is done according to the logical OR of the specified statuses. This option is valid when the **-nowait** option is not specified.
-t <timeout> [interval]:	Specifies the interval of monitoring a status specified using the **-s** option and the time-out period in units of 1 sec. Unless [interval] is specified, the default value is used. This option is valid when the **-nowait** option is not specified.
-nowait:	When this option is specified, the pair status at that time is reported without waiting. The pair status is set as a returned value for this command. When this option is specified, the **-t** and -s options are not needed.

33

- **pairevtwait** options (*continued*)
  - -l: When this command cannot utilize a remote host for host down, this option executes this command by a local host only. The target volume of a local host must be SMPL or P-VOL. (ShadowImage volumes are able to specify from SVOL.)
  - -nomsg: Suppresses messages to be displayed when this command is executed. It is used to execute a command from a user program. This option must be specified at the beginning of a command argument. The command execution log is not affected by this option.

34

- **pairmon** Command and Example

-d — Selects the default report mode. In the default mode, if there is pair status transition information to be reported, one event is reported and the event is reset.

-allsnd: — Reports all events if there is pair status transition information.

-resevt: — Reports events if there is pair status transition information, and then resets all events.

-nowait: — When this option is specified, the command does not wait when there is no pair status transition information.

-s — Specifies the pair status transition to be reported: smpl, copy (includes rcpy), pair, psus, psue. If two or more statuses are specified following **-s**, masking is done according to the logical OR of the specified statuses.

#pairmon –allsnd -nowait

Group	Pair vol	Port	targ#	lun#	LDEV#	Oldstat	code->	Newstat	code
orabd	oradb1	CL1-A	1	5	145	SMPL	0x00->	COPY	0x01
orabd	oradb2	CL1-A	1	6	146	SMPL	0x02->	PSUS	0x04

35

The pairmon command, which is connected to the HORCM daemon, obtains the pair status transition of each volume pair and reports it. If the pair status changes (due to an error or a user-specified command), the pairmon command issues a message. The pair status transition events exist in the HORCM pair state transfer queue. The **–resevt** option (reset event) deletes one/all events from the HORCM pair state transfer queue. If reset event is not specified, the pair state transfer queue is maintained. If the **–s** option is not specified, pairmon displays all events for which it receives information from HORCM. If the **–s** option is specified, only the specified status transitions are displayed.

- **pairvolchk** Command and Example

-g	Specifies group name from HORCM.CONF. All pairs in group resynchronised unless restricted by another switch.
-d	Specifies volume name from HORCM.CONF.
-c	Checks the conformability of the paired volumes of the local and remote hosts and reports the volume attribute of the remote host.
-ss	Used to acquire the attribute of a volume and the pair status of a volume. If this option is not specified, the volume attribute is reported.

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

36

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

When the -ss option is not specified:

Normal termination:

1: The volume attribute is SMPL.

2: The volume attribute is P-VOL.

3: The volume attribute is S-VOL.

When the -ss option is specified:

Normal termination: 11: The status is SMPL.

TrueCopy Sync/ 22: The status is PVOL_COPY or PVOL_RCPY.

ShadowImage 23: The status is PVOL_PAIR.

24: The status is PVOL_PSUS.

25: The status is PVOL_PSUE.

26: The status is PVOL_PDUB (TrueCopy and LUSE volume only).

29: The status is PVOL_INCSTG (inconsistent status in group). Not returned.

32: The status is SVOL_COPY or SVOL_RCPY.

33: The status is SVOL_PAIR.

34: The status is SVOL_PSUS.

35: The status is SVOL_PSUE.

36: The status is SVOL_PDUB (TrueCopy and LUSE volume only).

39: The status is SVOL_INCSTG (inconsistent status in group). Not returned.

To identify TrueCopy Asynchronous, the pairvolchk command returns a value which is 20 more than the TrueCopy

Sync status code also indicates PFUL and PFUS states to identify sidefile status of TrueCopy Asynchronous.

TrueCopy Async 42: The status is PVOL_COPY.

43: The status is PVOL_PAIR.

44: The status is PVOL_PSUS.

45: The status is PVOL_PSUE.

46: The status is PVOL_PDUB. (TrueCopy and LUSE volume only)

47: The status is PVOL_PFUL.

48: The status is PVOL_PFUS.

52: The status is SVOL_COPY or SVOL_RCPY.

53: The status is SVOL_PAIR.

54: The status is SVOL_PSUS.

55: The status is SVOL_PSUE.

56: The status is SVOL_PDUB. (TrueCopy and LUSE volume only)

57: The status is SVOL_PFUL.

58: The status is SVOL_PFUS.

- **paircurchk** command and example
  - Checks the currency of the TrueCopy Synchronous/ Asynchronous software secondary volumes by evaluating the data consistency based on pair status and fence level

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

```
#paircurchk –g oradb
Group  Pair vol  Port   targ#  lun#  LDEV#  Volstatus  Status  Fence   To be…
orabd  oradb1    CL1-A  1      5     145    S-VOL      PAIR    NEVER   Analyzed
orabd  oradb2    CL1-A  1      6     146    S-VOL      PSUS    STATUS  Suspected
```

37

The CCI paircurchk command checks the currency of the TrueCopy secondary volumes by evaluating the data consistency based on pair status and fence level. This table specifies the data consistency for each possible state of a TrueCopy volume. A paired volume or group can be specified as the target of the paircurchk command. The paircurchk command assumes that the target is an S-VOL. If the paircurchk command is specified for a group, the data consistency of each volume in the group is checked, and all inconsistent volumes are found in the execution log file and displayed.

- **horctakeover** command
  - Is a scripted command for executing several TrueCopy Synchronous/Asynchronous software operations
  - Checks the specified volume's or group's attributes (**paircurchk**), decides the takeover function based on the attributes, executes the chosen takeover function and returns the result
  - The four TrueCopy takeover functions designed for HA software operation:
    - Takeover-switch
    - Swaptakeover
    - PVOL-takeover
    - SVOL-takeover

38

A paired volume or a group can be specified as the target of the TrueCopy takeover command. If SVOL-takeover is specified for a group, the data consistency check is executed for all volumes in the group and all inconsistent volumes are found in the execution log file and displayed (same as **paircurchk** command). The takeover command allows swapping of the primary and secondary volumes, so that if the primary or secondary volume is switched due to a server error or package transfer, duplex operations can be continued using the reversed volumes. When control is handed over to the current node, swapping the volumes again eliminates the need to copy them. The takeover command also allows the secondary volume to be separated for disaster recovery operations.

- **horctakeover** flowchart

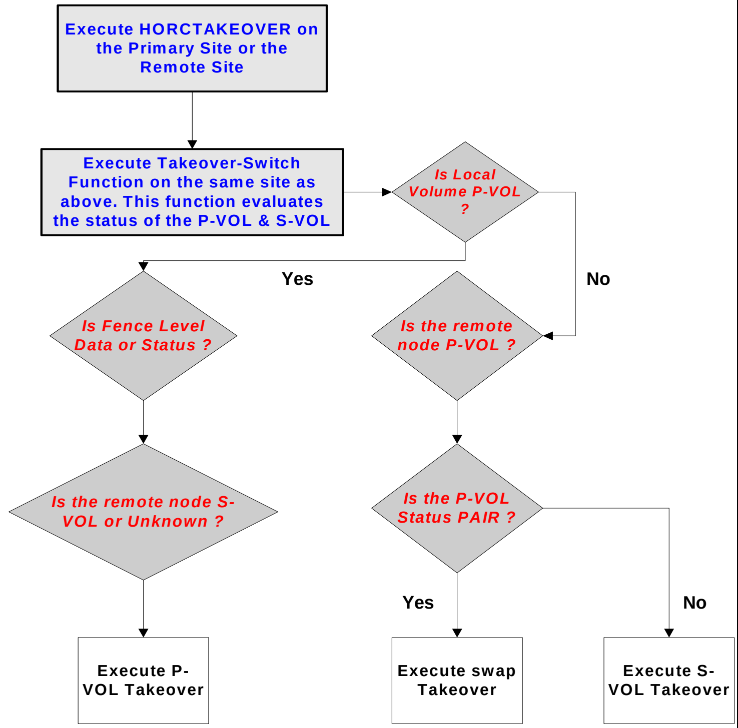

39

- TakeOver Switch

Local Node (Takeover Node)		Remote Node		Takeover Action
Volume Attribute	Fence Level And Status	Volume Attribute	P-VOL Status	
SMPL	-	SMPL	-	NG
		P-VOL	-	Nop-Takeover
		S-VOL	-	Volumes not conform
		Unknown	-	NG
P-VOL (Primary)	Fence = Data or Status Status = PSUE or PDUB	SMPL	-	NG
		P-VOL	-	Volumes not conform
		S-VOL	-	PVOL-Takeover
		Unknown Status (eg. LAN down)	-	PVOL-Takeover
	Fence = Never Status = others	SMPL	-	NG
		P-VOL	-	Volumes not conform
		S-VOL	-	Nop-Takeover
		Unknown Status (eg. LAN down)	-	Nop-Takeover
S-VOL (secondary)	Status – SSWS After SVOL-SSUS-takeover	Any	-	Nop-Takeover
	Others	SMPL	PAIR or PFUL	Volumes not conform
		P-VOL	Other	Swap-Takeover
		S-VOL	-	Volumes not conform
		Unknown	-	SVOL-Takeover

40

The control scripts activated by the HA software are used the same way by all nodes of a cluster; they do not discriminate between primary and secondary volumes. The takeover command, when activated by a control script, checks the combination of attributes of the local and remote volumes and determines the proper takeover action.

The table above lists the volume attributes and specifies the TrueCopy takeover action for each combination of attributes.

- **swaptakeover**
  - The **swaptakeover** function is used by an HA control script when a package is manually moved to an alternate data center while all hardware is operational.
  - This function internally executes the following commands to swap the primary and secondary volumes:
    - Execute suspend for swapping for the local volume (S-VOL). If this step fails, swaptakeover is disabled and an error is returned.
    - Execute resync for swapping of switch to the primary volume that the local volume (S-VOL) is swapped as the NEW_PVOL and resynchronizes the NEW_SVOL based on the NEW_PVOL.

41

When the P-VOL status of the remote node is PAIR and the S-VOL data is consistent, it is possible to swap the primary and secondary volumes. Swap-takeover can be specified for a paired volume or a group.

As for copy tracks, if the remote host is known, the command will use the value of P-VOL specified at paircreate time. If the remote host is unknown, the command will use the default number of tracks (three). If this step fails, swap-takeover returns at S-VOL-SSUS-takeover, and the local volume (S-VOL) is maintained in SSUS(PSUS) state which allows and keeps track of write I/Os using a bitmap for the S-VOL. This special state is displayed as SSWS using the -fc option of the pairdisplay command.

The swap-takeover function does not use SMPL or No Copy mode for swapping to guarantee mirror consistence, and this is included as a function of S-VOL takeover.

For TrueCopy Asynchronous: The CCI software on the S-VOL side will issue a suspend for swapping to the S-VOL side subsystem. Non-transmitted data which remains in the FIFO queue (side file) of the P-VOL will be copied to the S-VOL, and a resync for swapping operation will be performed (after the copy process). The swap operation is required to copy non-transmitted P-VOL data within a given timeout value (specified by the - t <timeout> option).

- **swaptakeover** scenario

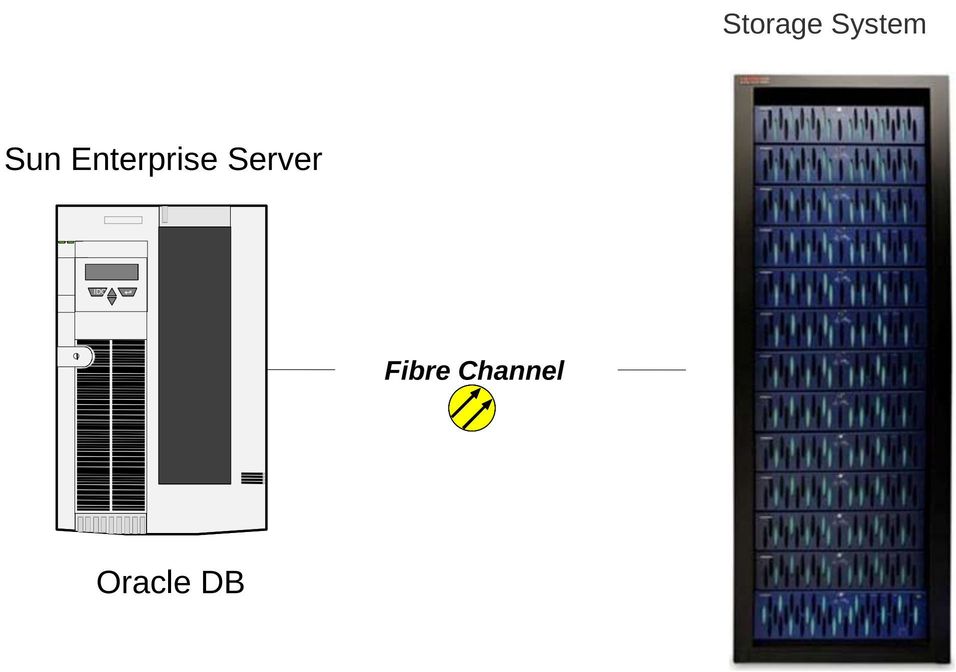

42

A customer is running a Oracle DB on a Sun Enterprise Server.
The REDO logs and data files are located on an AMS2000 family storage system.
The Oracle DB instance must be moved to another newly built datacenter.
At the remote datacenter the following tasks must be performed:

- Install another Sun Enterprise Server and run Oracle on this new server.
  Create the necessary LUNs on the remote AMS2000 storage system for the REDO logs and the data files.
  Install Sun Cluster Services and create a two node Sun Cluster
- Install CCI Software on both the local Sun Enterprise Server and the remote Sun Enterprise Server and ensure that network connectivity is present between the two.
  Configure two HORCM instances one running on the local site and one on the remote site.
- Issue a **paircreate** command and copy data from the P-VOL to the S-VOL and ensure that the status is PAIR.
- Embed the **horctakeover** command within the cluster failover scripts on both nodes.
- Failover the local Sun Enterprise Server to the remote Sun Enterprise Server and ensure that the **horctakeover** command runs on the remote node.

The takeover switch function now runs on the remote node and evaluates the status of the local and remote volumes.

Since the local volume is S-VOL and the remote volume is P-VOL and the status is PAIR the SWAPTAKEOVER function is selected.

The S-VOL is suspended for swapping with the P-VOL (SSWS).

Resync for swapping is executed so that the S-VOL becomes the P-VOL. The new P-VOL is resynchronized with the new S-VOL.

This entire operation is executed while the Oracle DB instance is in production without any Oracle DB downtime.

- PVOL takeover
  - This function allows the takeover node to use the primary volume, that is, reading and writing are enabled. The assumption is that the remote node with the secondary volume cannot be used.
  - The PVOL takeover function executes the following two commands:
    - PVOL-PSUE takeover: Changes the primary volume to the suspend (PSUE, PSUS) state which enables write I/Os to all primary volumes of the group. The action of the PVOL-PSUE- takeover causes PSUE and/or PSUS to be intermingled in the group.
    - PVOL-SMPL takeover: Changes the primary volume to the simplex (SMPL) state. First, P-VOL takeover executes PVOL-PSUE takeover further than PVOL-SMPL takeover. If the PVOL-PSUE takeover function fails, the PVOL-SMPL takeover function is executed.

43

The P-VOL takeover function releases the pair state as a group, since that maintains the consistency of the secondary volume at having accepted horctakeover command when the primary volume is fenced (“data or status” and “PSUE or PDUB” state, “PSUE or PDUB” volume are contained in the group). P-VOL takeover can be specified for a paired volume or a group.

- PVOL takeover scenario

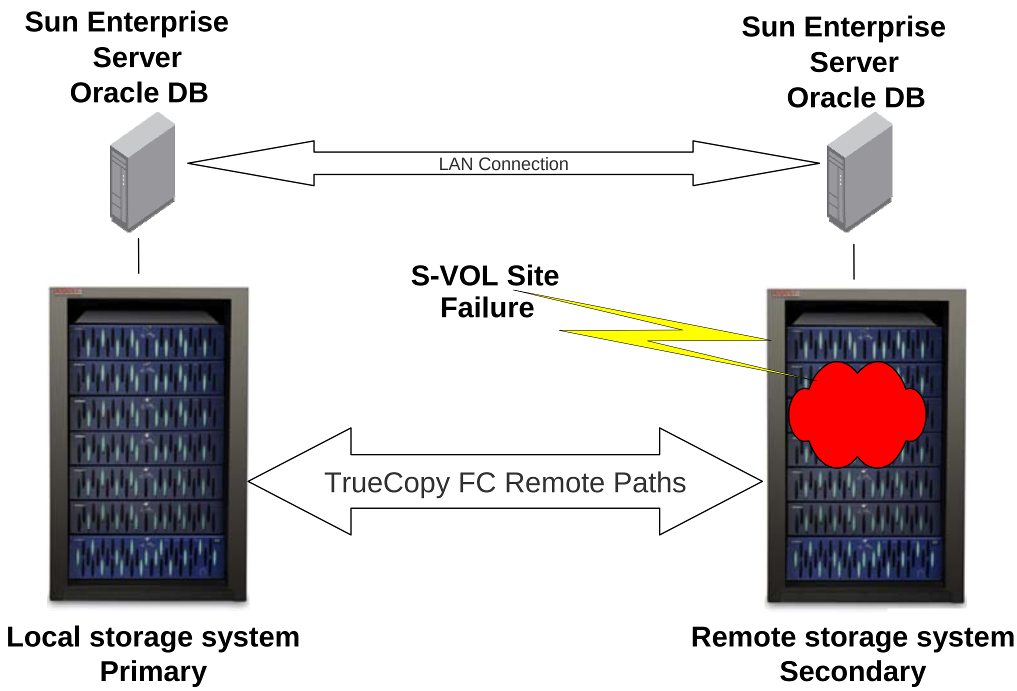

44

1. A customer runs an Oracle DB on a Sun Enterprise server at a production data center.
   The Data Files and REDO logs are on Veritas Volumes located on a AMS2000 storage system.
   The customer has a "hot site" at a remote data center.
   The fence level is set to *data* for the Oracle LUNs.
   CCI is installed on the two Sun Enterprise servers and horctakeover commands are embedded in the cluster failover scripts.
2. There is a failure at the S-VOL site owing to which the primary system is unable to write to the secondary system.
3. The primary system fences off the Oracle LUNs since Fence Level is set to *data*.
4. A **horctakeover** command is issued at the primary site
5. The takeover switch function determines that the command has been run from the P-VOL site and check for the status of the S-VOL.
   The status of the S-VOL is unknown since the S-VOL site is down. The function decides to execute the P-VOL takeover function
6. The CCI software attempts to switch the status of the P-VOL to PSUE so that reads and writes can be enabled to the P-VOL.
   If this fails the software will switch the status of the P-VOL to SMPL.

- SVOL takeover
  - The SVOL takeover function allows the takeover node to use the secondary volume (except in COPY state) in SSUS (PSUS) state. Reading and writing are enabled. The assumption is that the remote node with the primary volume cannot be used.
  - If the primary and secondary volumes are not consistent, the SVOL takeover function fails.
  - If primary and secondary volumes are consistent, the SVOL takeover function attempts to switch to the primary volume using resync for swapping.
  - SVOL takeover can be specified for a paired volume or a group. If the SVOL takeover is specified for a group, a data consistency check is executed for all volumes in the group, and all inconsistent volumes are displayed.

45

The data consistency of the TrueCopy S-VOL is evaluated by its pair status and fence level. If successful, the SVOL takeover function returns swap takeover as the return value of the horctakeover command. If not successful, the SVOL takeover function returns SVOL-SSUS-takeover as the return value of the **horctakeover** command. In case of a host failure, Swap takeover is returned. In case of a Fibre Channel or P-VOL site failure, SVOL-SSUS-takeover is returned.

- SVOL takeover scenario

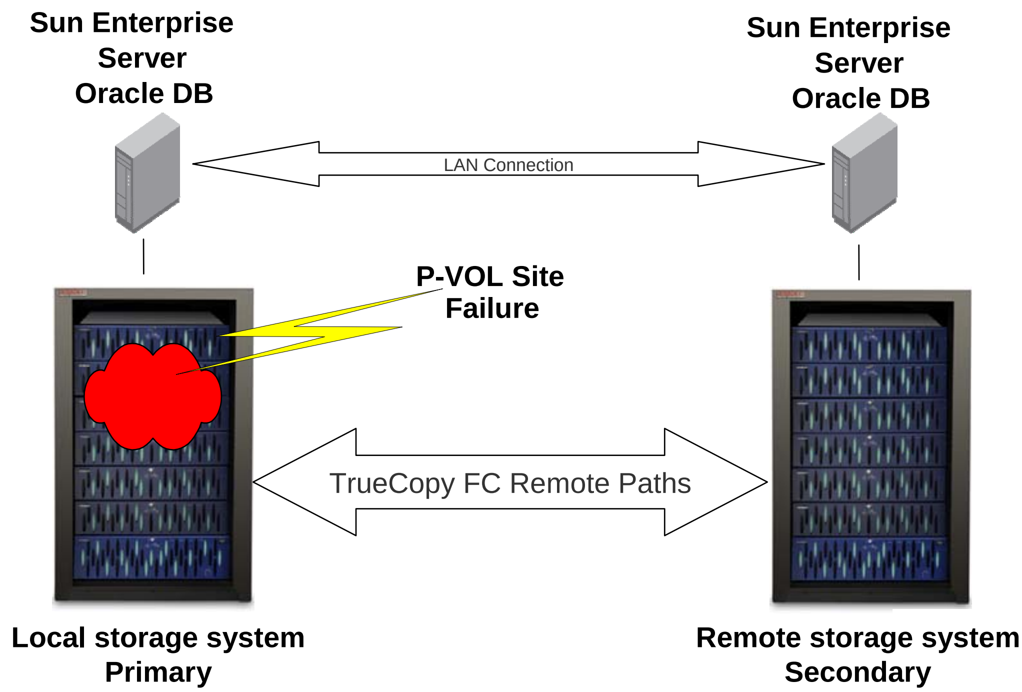

46

1. A customer runs Oracle DB on a Sun Enterprise server at a production data center.
   The Data Files and REDO logs are on Veritas Volumes located on a AMS2000 storage system.
   The customer has a “hot site” at a remote data center.
   The fence level is set to *data* for the Oracle LUNs.
   CCI is installed on the two Sun Enterprise servers and **horctakeover** commands are embedded in the cluster failover scripts.
2. There is a failure at the primary site which results in the Oracle instance failing over to the remote site.
3. The **horctakeover** command embedded in the failover scripts gets executed at the remote site.
4. The takeover switch function determines that the command has been run at the S-VOL site and tries to determine the status of the P-VOL.
   The status of the P-VOL is unknown since the primary site is down.
5. The takeover function decides to execute S-VOL takeover.
6. The CCI software switches the status of the S-VOL to SSUS
   and the S-VOL can now be mounted read-write for Host I/O and the database can be restarted at the remote site in order to continue operations.

- **horctakeover** example

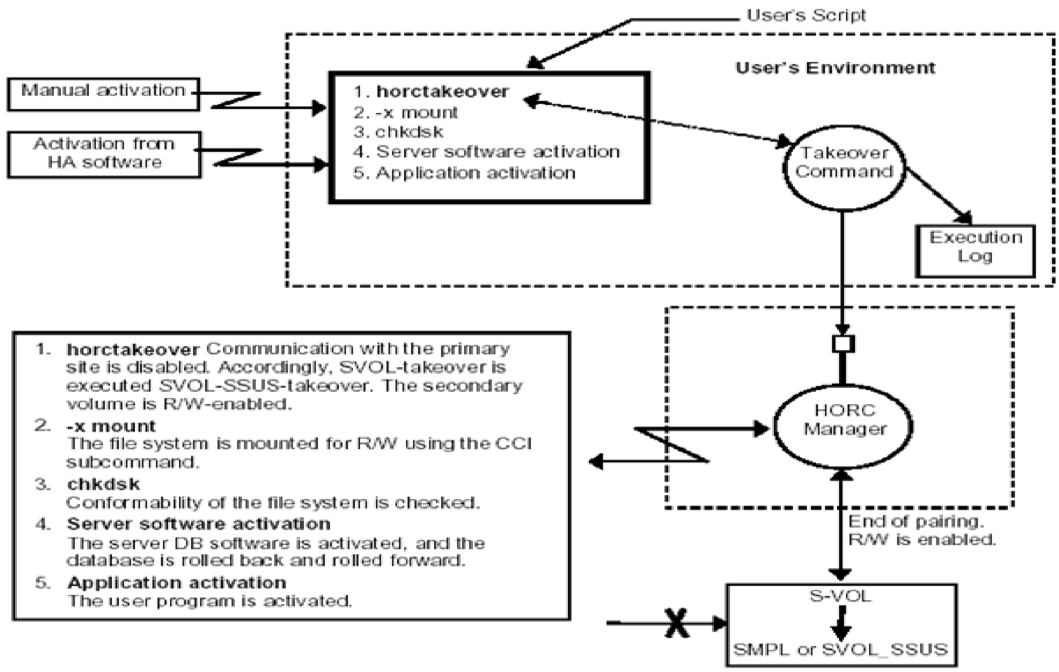

47

The basic TrueCopy commands (takeover, pair creation, pair splitting, pair resynchronization, event waiting) can be combined to enable recovery from a disaster, backup of paired volumes and many other operations (for example, restoration of paired volumes based on the secondary volume, swapping of the paired volumes).

Troubleshooting

- Troubleshooting Information

	INFORMATION TYPE	FILENAME
Problem Starting CCI	HORCM Startup Log	$HORCM_Log/horcm_*HOST*.Log
Problem With CCI Commands	Command Log	$HORCC_Log/horcc_*HOST*.Log
Commands Errors	Error Log	$HORCM_Log/horcmlog_*HOST*/horcm.log

$HORCM_Log Default Directory: /HORCM/Log*n*/Curlog (N=Instance #)

$HORCC_Log Default Directory: /HORCM/Log*n*

48

Module Review

1. What does the **pairevtwait** command specify?
2. What does the **horctakeover** command specify?
3. What is the implementation benefit for using TrueCopy Synchronous and Asynchronous software
4. What is the implementation benefit for using TrueCopy Synchronous, Asynchronous and ShadowImage Replication software together?

50

11. Hitachi TrueCopy Remote Replication with Hitachi Storage Navigator Modular 2

Module Objectives

- Upon successful completion of this module, you should be able to:
  - Describe Hitachi Storage Navigator Modular 2 (SNM2) GUI and CLI operation
  - Identify SNM2 CLI commands

2

TrueCopy Operations with SNM2

- Storage Navigator Modular 2 enables you to:
  - Setup, configure, and restore the TrueCopy Fibre Channel paths
  - Create TrueCopy pairs
  - View pair details and edit the properties of existing pairs
  - Split pairs
  - Resynchronize pairs
  - Swap pairs
  - Delete /Release pairs
- You need not setup Command Devices, CCI or HOCM scripts if SNM2 is sufficient to manage your TrueCopy operations.

3

SNM2 CLI – CCI/Raid Manager commands

Action	CLI – Command	CCI - Command
Display pair information	aureplicationremote -refer	pairdisplay
Create pairs	aureplicationremote -create	paircreate
Split pairs	aureplicationremote -split	pairsplit
Resynchronize pairs	aureplicationremote -resync	pairresync
Swap pairs	aureplicationremote -swaps	pairresync -swaps
Delete pairs	aureplicationremote -simplex	pairsplit -S
Event wait	aureplicationmon -syncwait	pairsynctwait

4

Create a TrueCopy Pair with SNM2

Select **Remote Replication** and click the **Create Pair** button

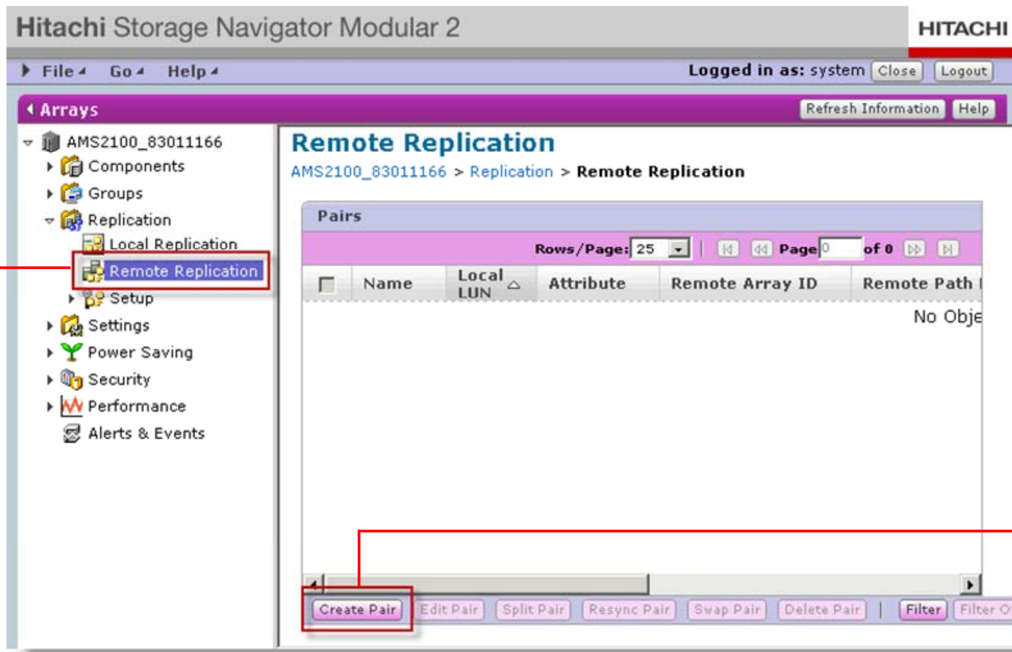

5

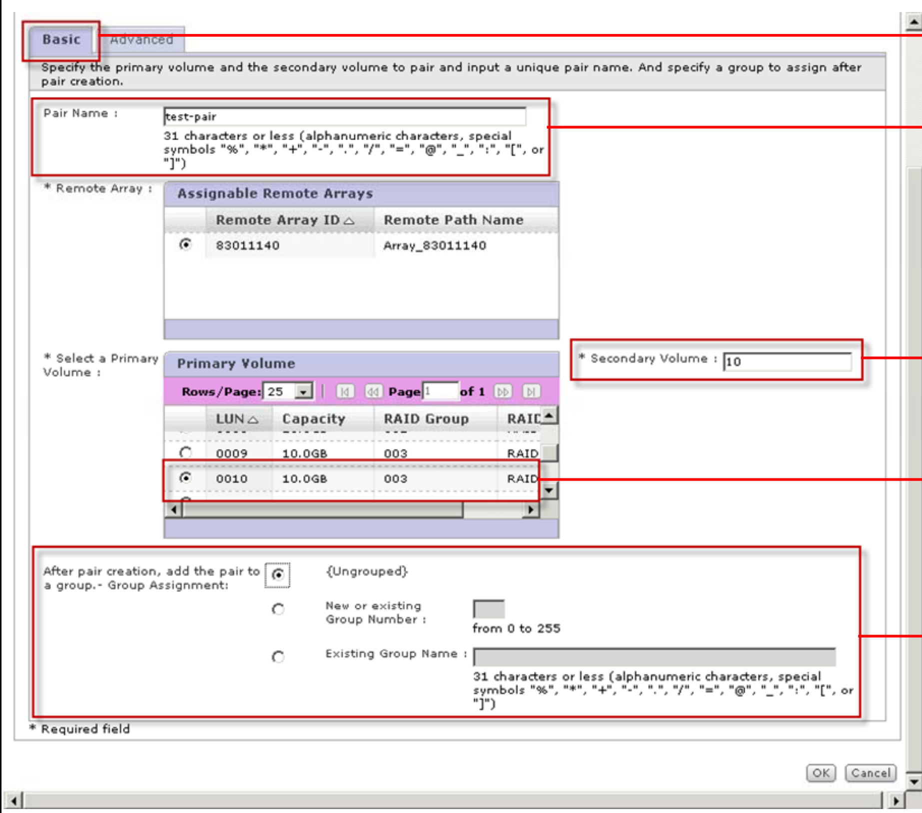

Select **Basic.**

Specify the **Pair Name.**

Specify the S-VOL.

Specify the P-VOL.

Specify the grouping.

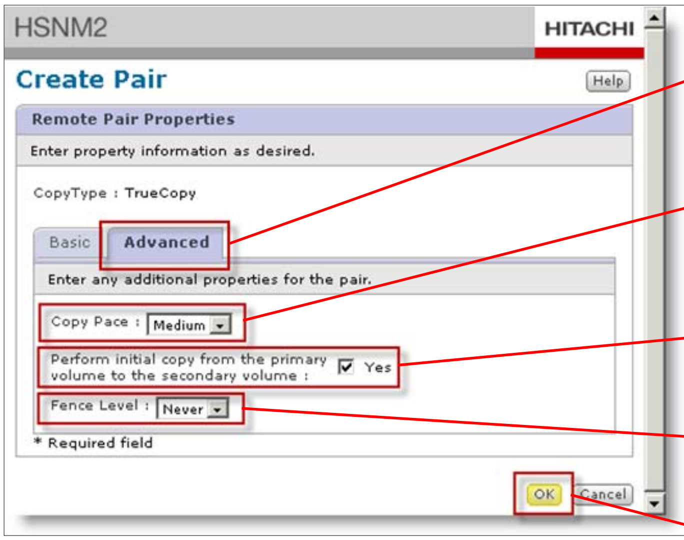

Select **Advanced.**

Specify the **Copy Pace.**
- Slow
- Medium or
- Fast

Specify the initial copy parameter – **Yes** or no.

Set the **Fence Level Never** or **Data.**

Click on **OK** to create the pair.

7

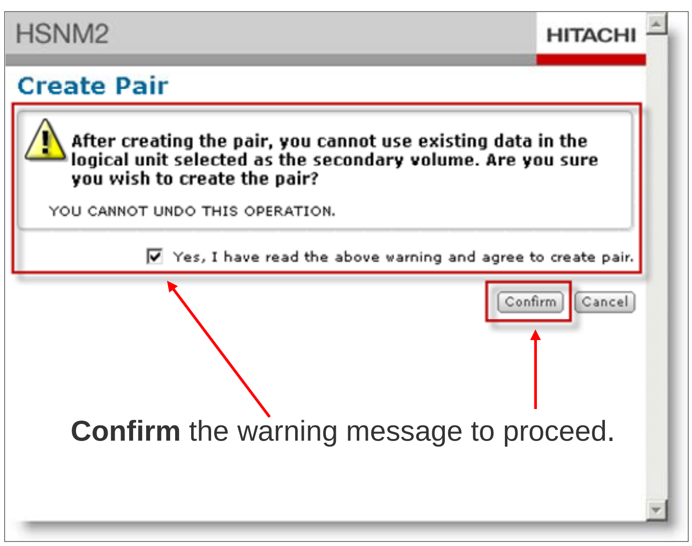

Confirm the warning message to proceed.

8

The local storage system holds the **S-VOL** of a pair.

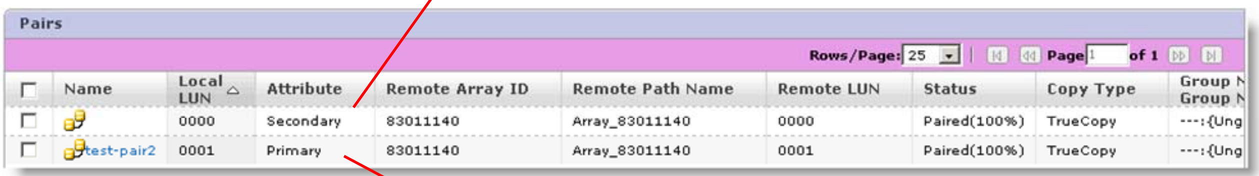

The local storage system holds the **P-VOL** of a pair.

To create the pair using SNM2 CLI you can use the CLI command as follows:

aureplicationremote -unit ‘unit_name’ -create -tc -pvol 0 -svol 0 -pairname pair1 -remote ‘remote_array_id’

9

You will find more information on creating TrueCopy pairs using the SNM2 CLI in the *Hitachi Navigator 2 Command Line Interface (CLI) Reference Guide for Replication* on page 2-16

View TrueCopy Pair Details with SNM2

Select **Remote Replication** and click on the pair **Name.**

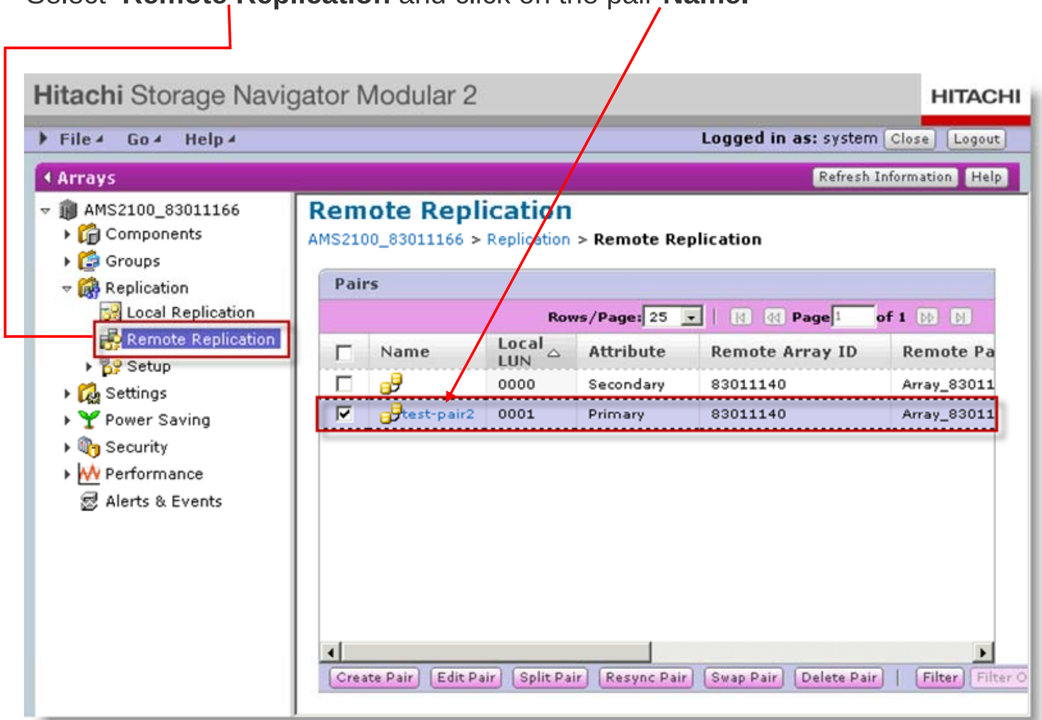

10

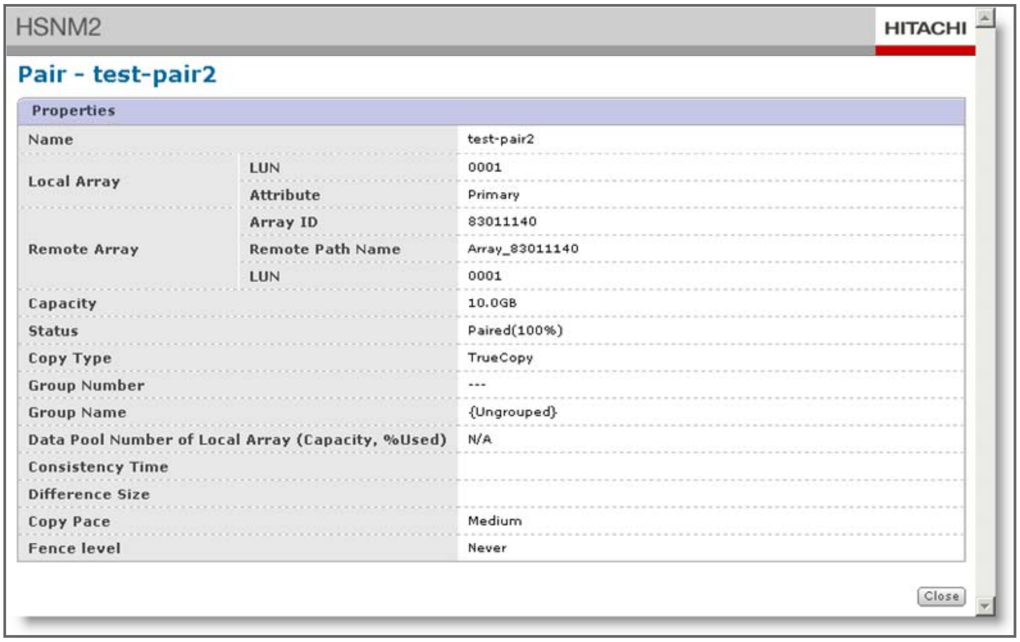

Properties		
Name		test-pair2
Local Array	LUN	0001
	Attribute	Primary
Remote Array	Array ID	83011140
	Remote Path Name	Array_83011140
	LUN	0001
Capacity		10.0GB
Status		Paired(100%)
Copy Type		TrueCopy
Group Number		---
Group Name		{Ungrouped}
Data Pool Number of Local Array (Capacity, %Used)		N/A
Consistency Time		
Difference Size		
Copy Pace		Medium
Fence level		Never

To view the pair details using SNM2 CLI you can use the CLI command as follows:

aureplicationremote -unit ‘unit_name’ -refer

11

You will find more information on creating TrueCopy pairs using the SNM2 CLI in the *Hitachi Navigator 2 Command Line Interface (CLI) Reference Guide for Replication* on page 2-22

Edit TrueCopy Pairs with SNM2

Select **Remote Replication,** Select the **Pair** and click on **Edit Pair**

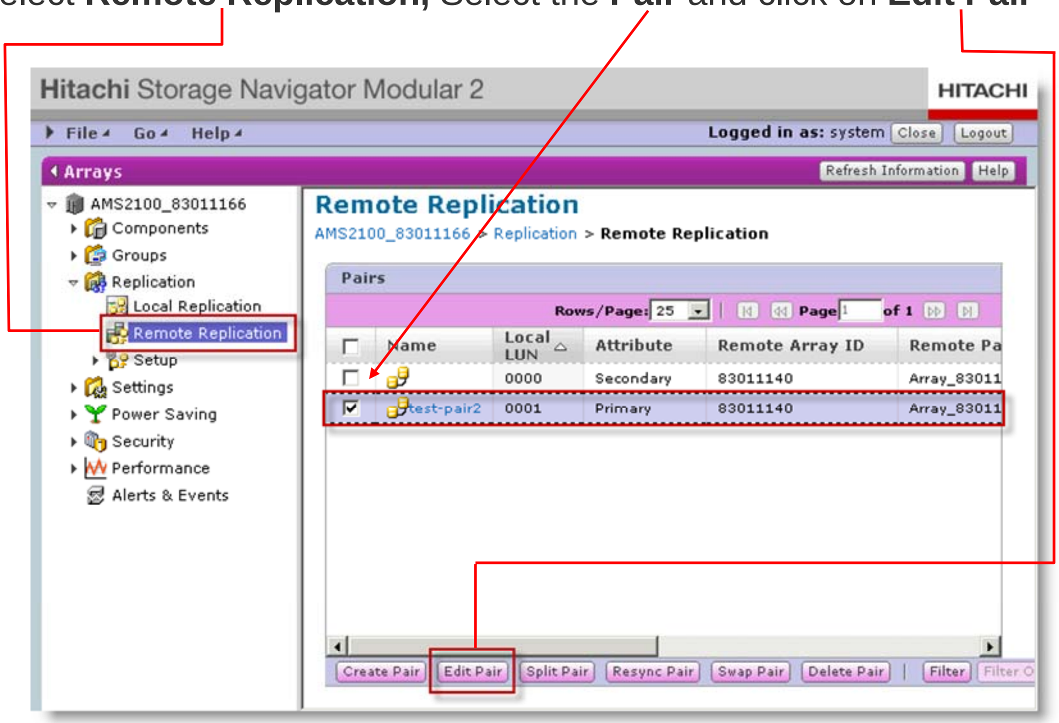

12

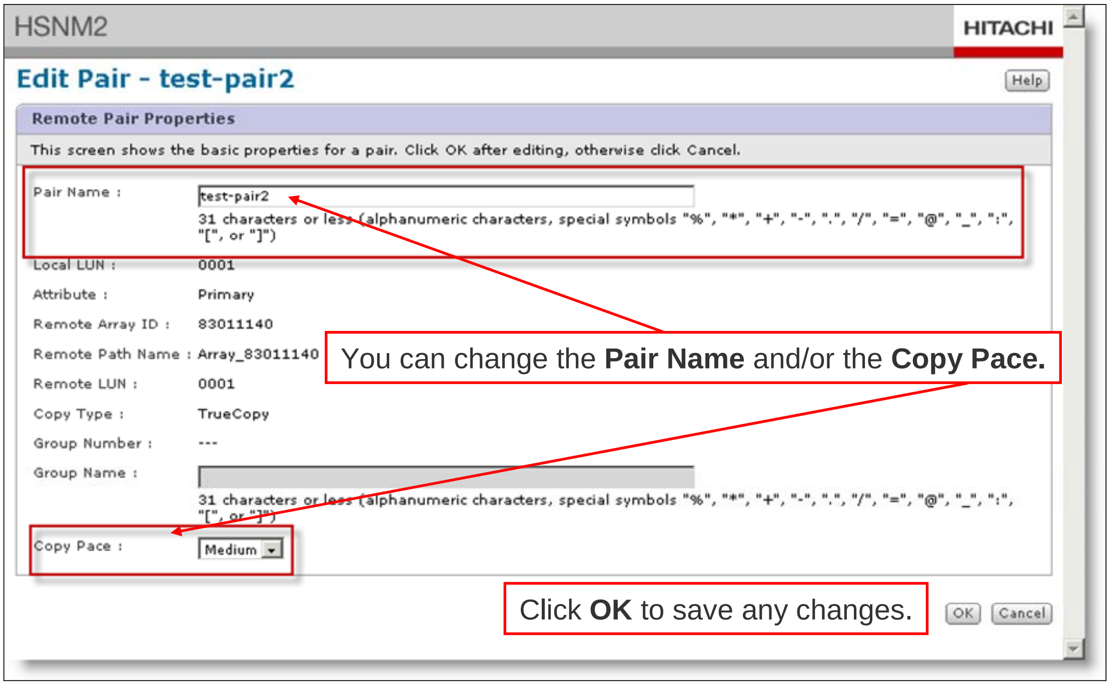

To edit the pair using SNM2 CLI you can use the CLI command as follows:

aureplicationremote -unit 'unit_name' -chg –tc locallun pvol -pvol 0 -svol 0 -remote 'remote_ID' -newpairname Test_Pair2 -pace normal

13

You will find more information on changing/editing TrueCopy pairs using the SNM2 CLI in the *Hitachi Navigator 2 Command Line Interface (CLI) Reference Guide for Replication* on page 2-27

Split TrueCopy Pairs with SNM2

Select **Remote Replication**, select the **Pair** and click on **Split Pair.**

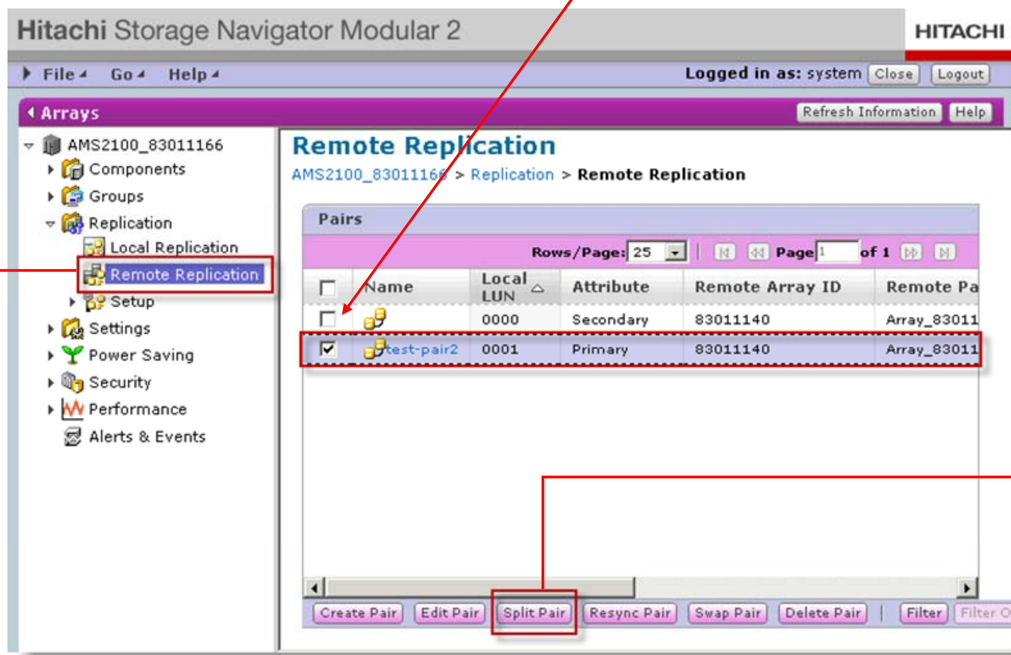

14

Specify the **Accessibility Option** for the **S-VOL** after the split.

- **Read/Write** or
- **Read Only**

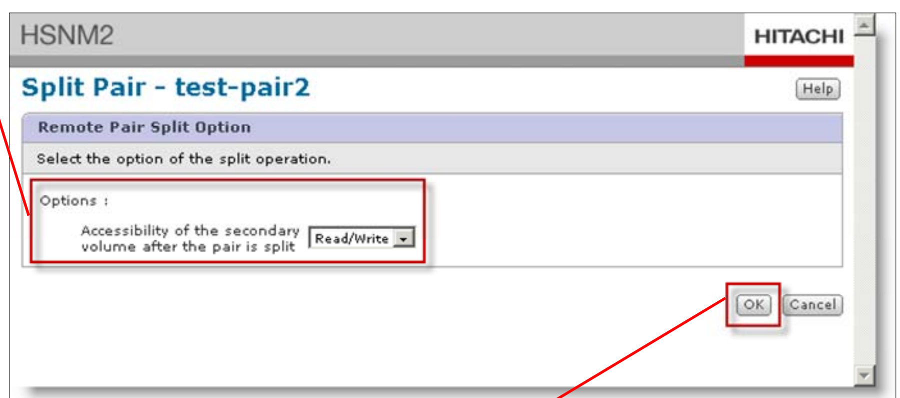

Click **OK** to split the Pair.

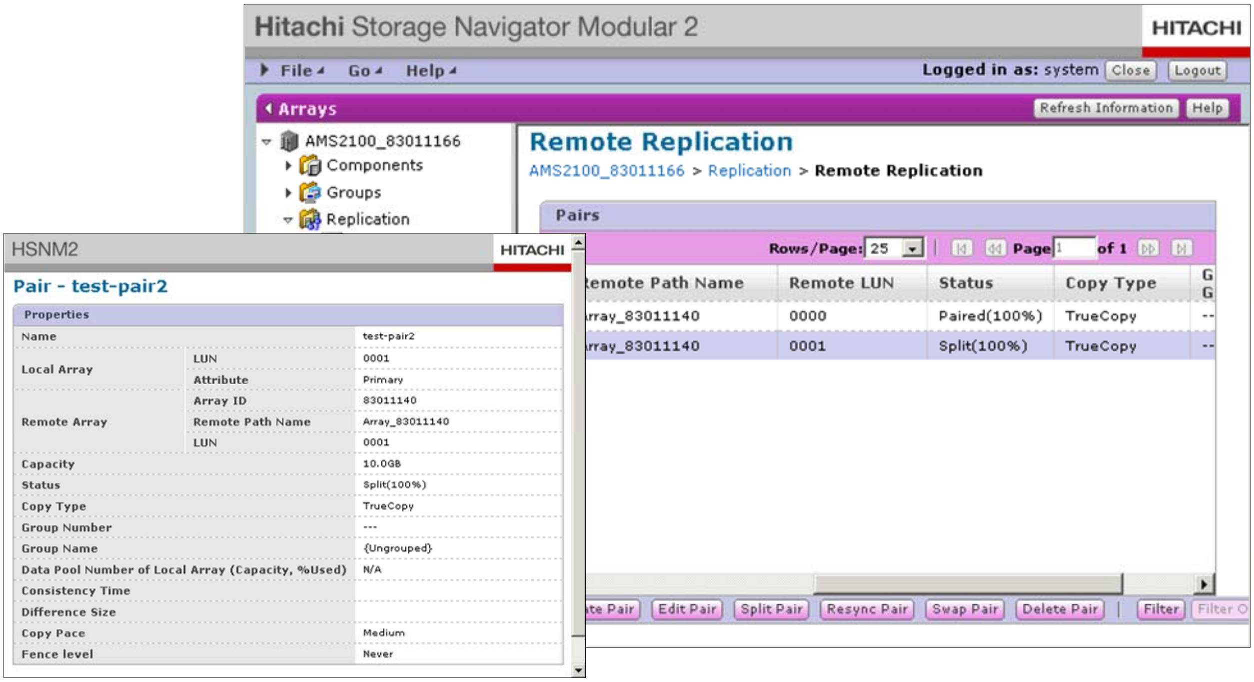

To split the pair using SNM2 CLI you can use the CLI command as follows:

aureplicationremote -unit ‘unit_name’ -split -tc -pvol 1 -svol 1 -remote ‘remote_array_ID’

16

You will find more information on splitting TrueCopy pairs using the SNM2 CLI in the *Hitachi Navigator 2 Command Line Interface (CLI) Reference Guide for Replication* on page 2-19

Resynchronize TrueCopy Pairs with SNM2

- Select **Remote Replication**, select the **Pair** and click on **Resync** Pair.

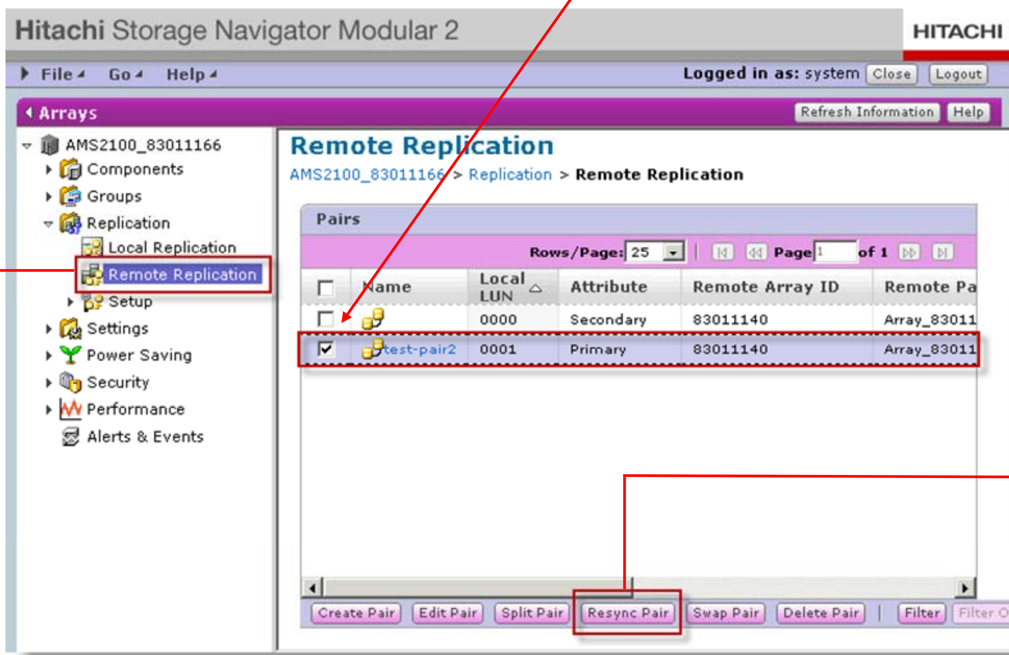

17

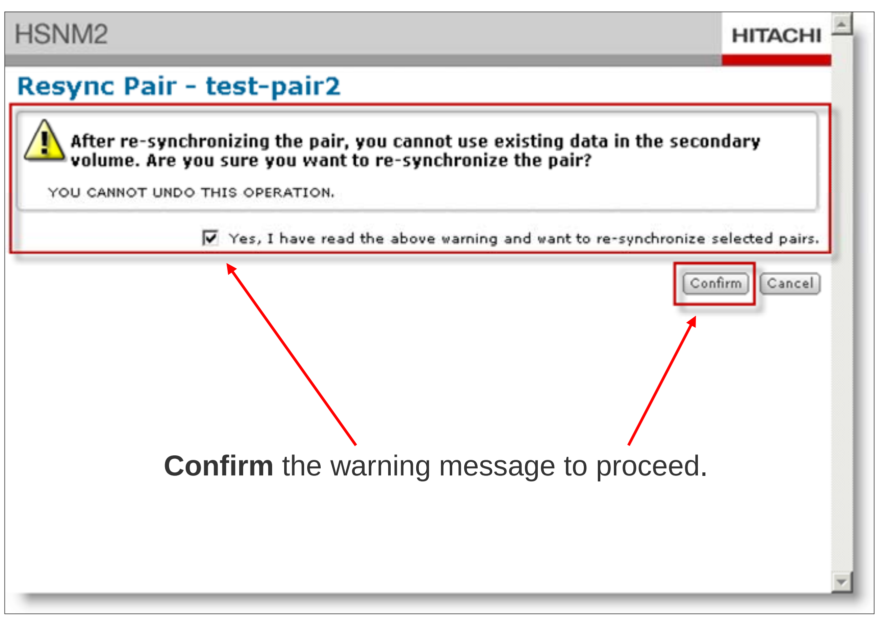

Confirm the warning message to proceed.

18

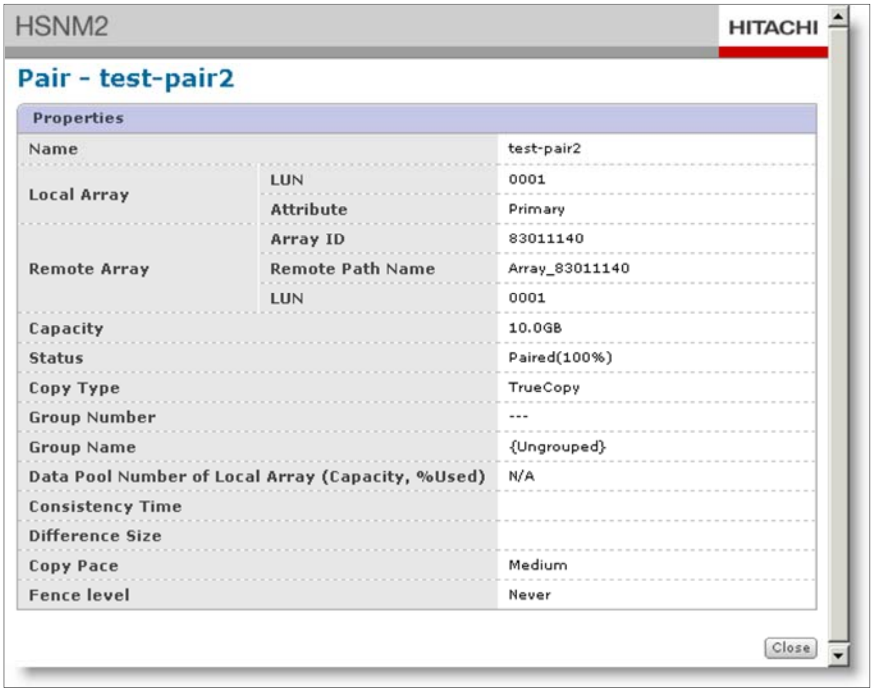

HSNM2 HITACHI

Pair - test-pair2

Properties		
Name		test-pair2
Local Array	LUN	0001
	Attribute	Primary
Remote Array	Array ID	83011140
	Remote Path Name	Array_83011140
	LUN	0001
Capacity		10.0GB
Status		Paired(100%)
Copy Type		TrueCopy
Group Number		---
Group Name		{Ungrouped}
Data Pool Number of Local Array (Capacity, %Used)		N/A
Consistency Time		
Difference Size		
Copy Pace		Medium
Fence level		Never

Close

To re-synchronize the pair using SNM2 CLI you can use the CLI command as follows:
aureplicationremote -unit ‘unit_name’ -resync -tc -pvol 1 -svol 1 -remote ‘remote_array_ID’

19

You will find more information on re-synchronizing TrueCopy pairs using the SNM2 CLI in the *Hitachi Navigator 2 Command Line Interface (CLI) Reference Guide for Replication* on page 2-21

Swap TrueCopy Pairs with SNM2

- Select **Remote Replication,** Select the **Pair** and click on **Swap Pair**

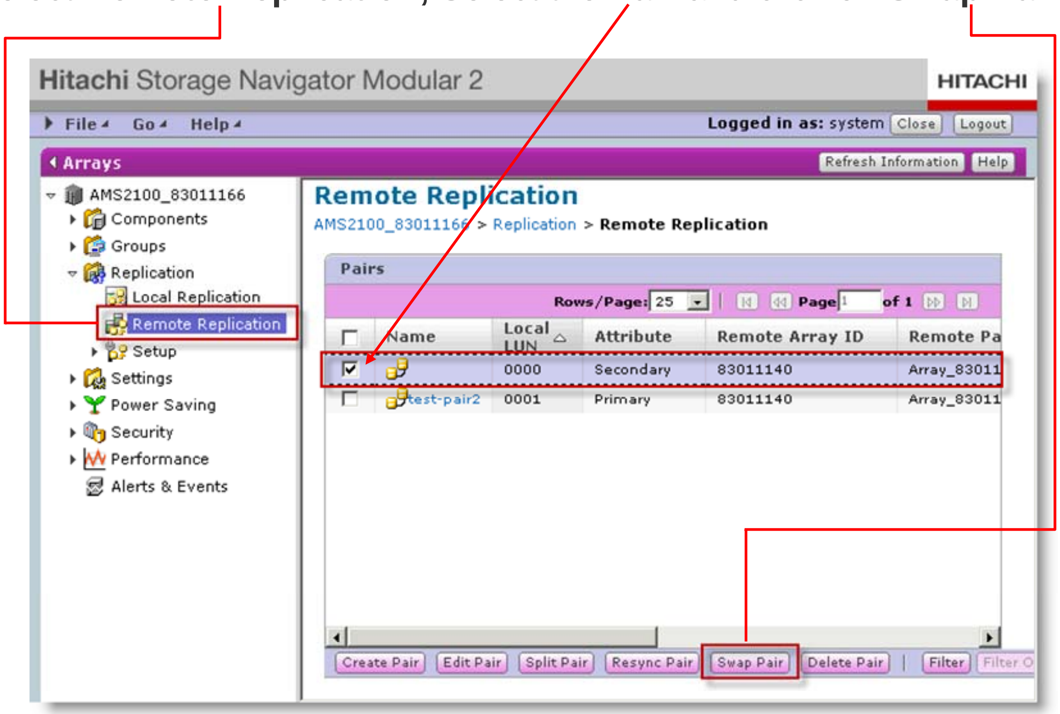

20

Note: You can only swap pairs they have the secondary volumes (S-VOLs) on the current local storage system.

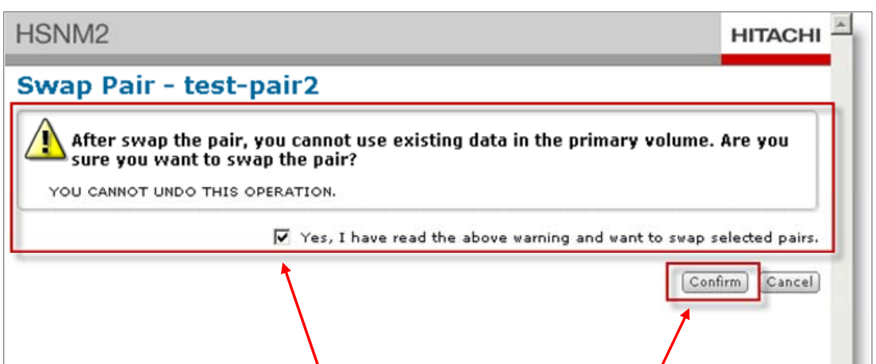

Confirm the warning message to proceed.

- The S-VOL becomes the P-VOL of the swapped pair on the local system.

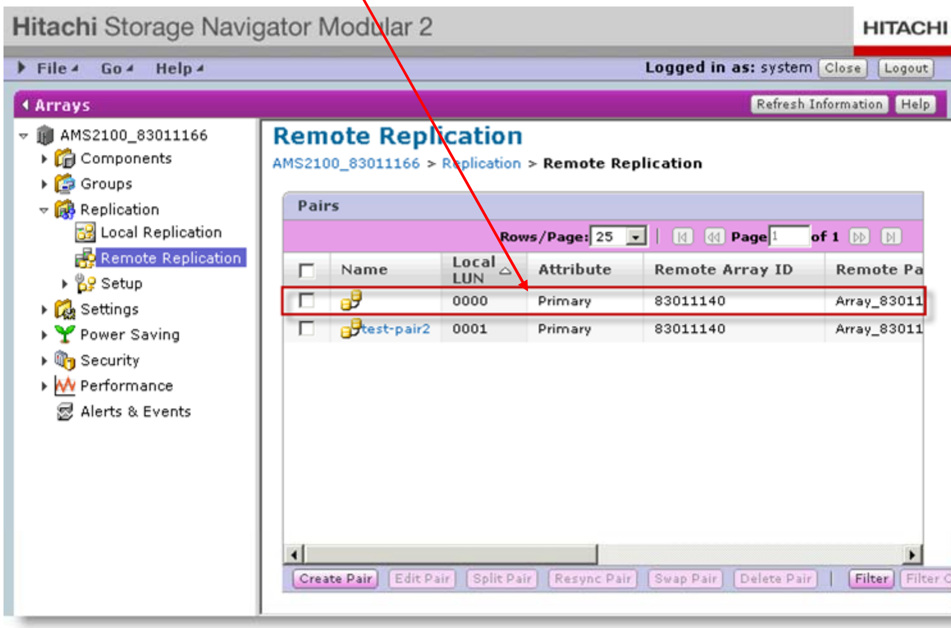

To swap the pair using SNM2 CLI you can use the CLI command as follows:
aureplicationremote -unit 'unit_name' -swaps -tc -svol 1

22

You will find more information on swapping TrueCopy pairs using the SNM2 CLI in the *Hitachi Navigator 2 Command Line Interface (CLI) Reference Guide for Replication* on page 2-23

Delete TrueCopy Pairs with SNM2

- Select Remote Replication, select the Pair and click on Delete Pair.

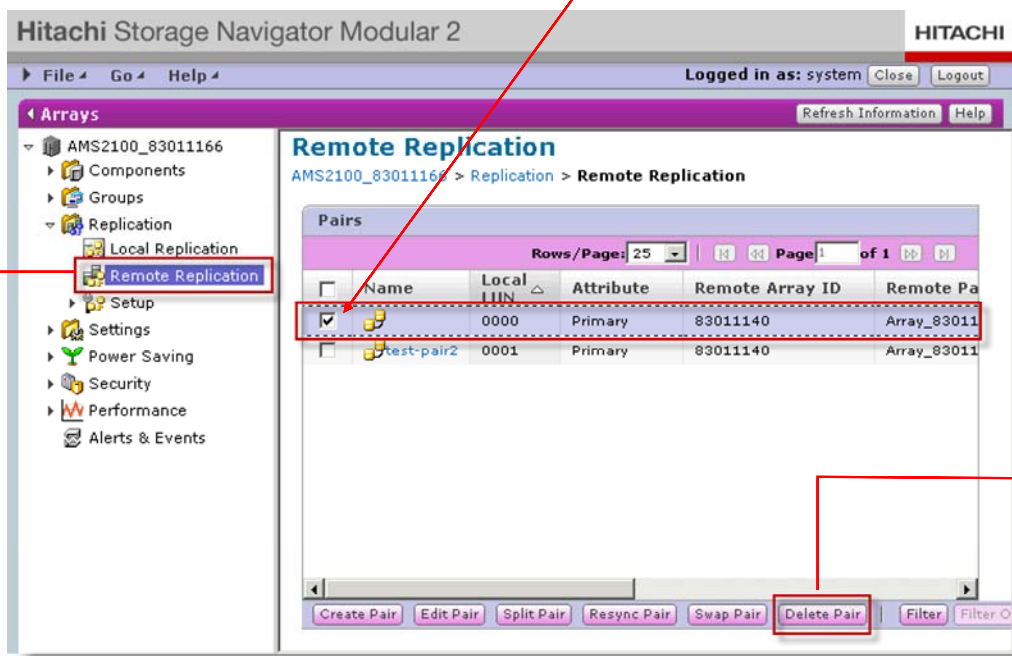

23

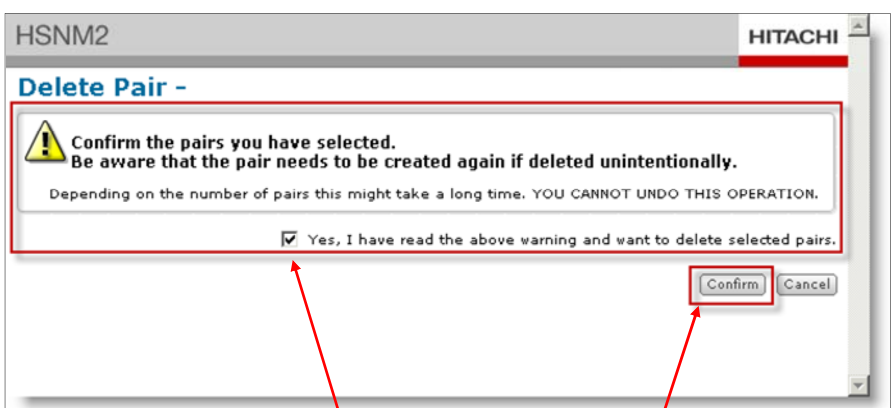

Confirm the warning message to proceed.

24

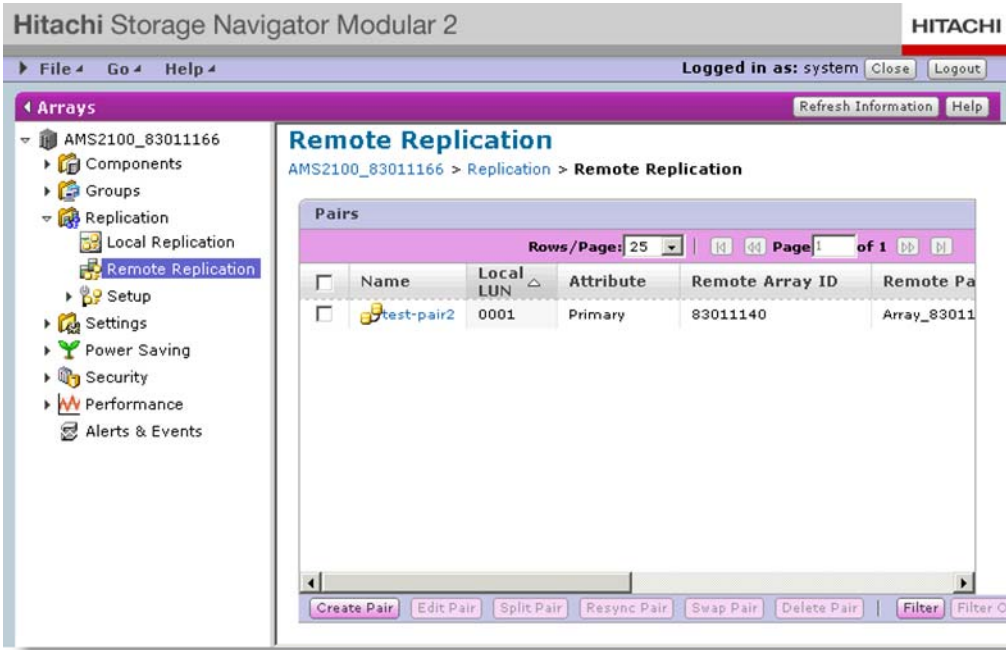

To delete the pair using SNM2 CLI you can use the CLI command as follows:
aureplicationremote -unit ‘unit_name’ -simplex -tc -locallun pvol -pvol 1 -svol 1 -remote ‘remote_ID’

25

You will find more information on deleting/releasing TrueCopy pairs using the SNM2 CLI in the *Hitachi Navigator 2 Command Line Interface (CLI) Reference Guide for Replication* on page 2-25

12. Transport Infrastructure and WAN Connectivity

Module Objectives

- Upon successful completion of this module you should be able to:
  - Describe challenges to transport infrastructure and potential solutions
  - List key vendor products that enable Hitachi TrueCopy® Remote Replication software to meet the needs of wide area networks (WANs)
  - Describe these product key features and benefits
  - Identify successful wide area network configurations
  - List TrueCopy software products and the distances supported
  - Describe wide area networks best practices

2

  Hitachi Data Systems

Transport Infrastructure

- SAN Shortcomings to Date
  - Fibre Channel is relatively new and often seen as immature.
    - Adoption rate is low — estimated at 15-30%.
    - Lack of standards and interoperability cause people to wait.
  - New Fibre Channel infrastructure has been needed for SAN.
    - Net new equipment, infrastructure, software, and practices.
    - Early adopters have found Fibre Channel difficult to manage.
    - QoS (Quality of Service) requires extra effort.
  - Fibre Channel limits of 6 to 75 miles (10 -120km).
  - SAN extension adds complexity and cost.
    - Dense Wave Division Multiplexer (DWDM), FC - IP converters and encapsulation.
  - Not all servers/storage are in a data center.
    - Remote sites have need for pooled data, storage, tape and data management.

3

- IP Storage Networking
  - Uses IP backbone for storage networks instead of the Fibre Channel protocol
  - The benefits over conventional Fibre Channel based SANs include:
    - Unified network management across the enterprise for all devices
    - Intelligent network services, including QoS, traffic management, security, and accounting
    - Highly scalable – LAN to MAN to WAN
    - Interoperability – millions of real world examples
    - Capitalizes on the widespread familiarity, proven technology, and declining cost curve of gigabit Ethernet and IP

MAN is Metro area network.

- Hitachi Data Systems Networks
  - Include Fibre Channel IP technology to provide additional SAN functionality.
  - IP storage targeted at four situations:
    - Provide TrueCopy software replication over IP
    - Stitch together remote, disparate Fibre Channel SANs with IP backbone
    - Build native IP switch fabrics
    - Future-proofing for native IP storage
  - Many IT departments would rather manage and invest in IP rather than Fibre Channel.
  - Fibre Channel has 15-30% penetration and therefore has lots of room for new technologies and solutions.
  - Hitachi Data Systems continues to support new initiatives.

5

- Multi-protocol open approach

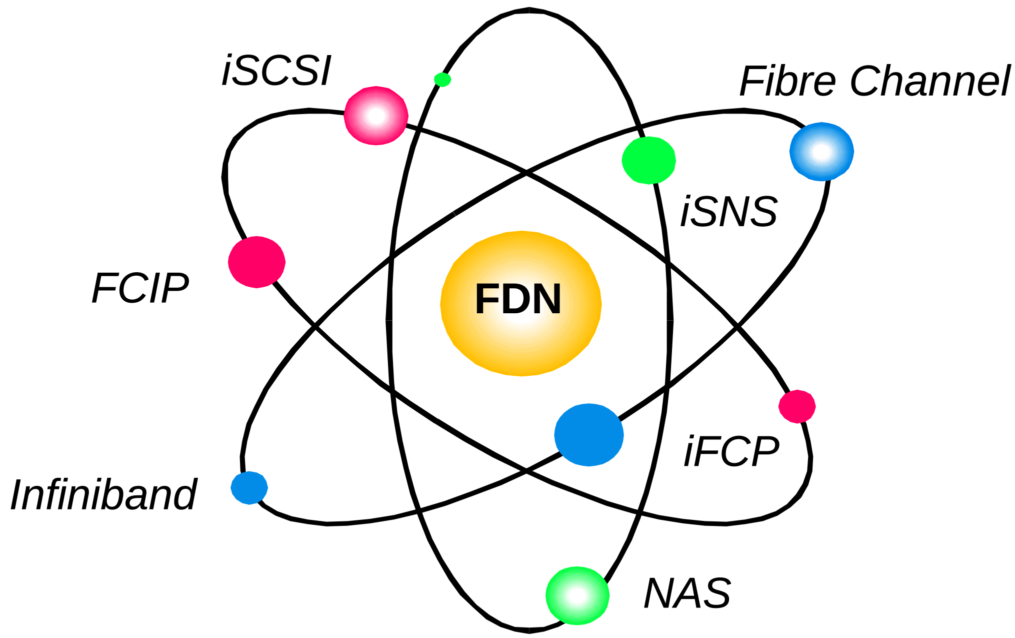

6

Freedom Data Networks

Hitachi Freedom Data Networks (FDN) provides an open architecture that offers organizations the freedom of choice in deploying data access, protection and sharing capabilities across the enterprise. Using multiple technologies and solutions such as SAN and NAS, FDN will build, leverage and augment storage infrastructures providing access to any data from any computer, anytime and anywhere.

- Technology benefits
  - Online asynchronous mirroring, peer-to-peer copy is available as a feature of the storage software
  - Uses existing network technology
  - Shares or separates IP storage networks with traditional IP message networks
  - Uses existing management, capacity planning, and operational coverage for IP networks to extend the storage network
  - iFCP products can connect SANs that are too far apart to allow connection using native Fibre Channel
  - Replicated databases can be provided in far away locations
  - The same IP routing can be managed for local, metro, and wide area networks

7

Protocol Summary

- iSCSI
  - Enables block-level storage accessed from Fibre Channel SANs using IP storage routers or switches
  - Reduces TCO by using existing network and IT support resources
  - Affordably create IP SANs from a number of DAS devices
  - Adaptable Modular Storage 2000 family iSCSI option
- iFCP
  - Large installed base of Fibre Channel devices, with the momentum toward IP storage networking
  - Leverages the high performance and interoperability of the Fibre Channel protocol, while taking advantage of IP networks
  - Environments with a heavy investment in both Fibre Channel SANs and an enterprise-wide IP network backbone
- FCIP
  - Interconnect Fibre Channel SANs to the IP network and to set up connections between SANs, or between Fibre Channel devices and SANs
  - Implementation with customers with Fibre Channel expertise
  - Brocade and Cisco

8

iSCSI

- Enables block-level storage to be accessed from Fibre Channel SANs using IP storage routers or switches, which furthers its applicability as an IP-based storage transport protocol.

iFCP

- iFCP is a gateway-to-gateway protocol where TCP/IP switching and routing components complement and enhance, or replace, the Fibre Channel fabric.
- The primary market drivers for iFCP are the large installed base of Fibre Channel devices, combined with the momentum toward IP storage networking. The emerging iFCP standard leverages the high performance and interoperability of the Fibre Channel protocol, while taking advantage of IP networks. Centralized consolidation of Fibre Channel SANs via iFCP is a consideration for those environments where there is a heavy investment in both Fibre Channel SANs and an enterprise-wide IP network backbone. The driving force behind iFCP is the expansion of IP-based network services to interconnect Fibre Channel devices and SANs.

FCIP

- FCIP uses gateways to interconnect Fibre Channel SANs to the IP network and to set up connections between SANs, or between FC devices and SANs

13. Hitachi TrueCopy® Extended Distance Software

Module Objectives

- Upon completion of this module, the learner should be able to:
  - Identify the benefits of Hitachi TrueCopy® Remote Replication software
  - Explain the basic concepts of TrueCopy® Extended Distance software
  - Identify components, specifications, and rules or limitations
  - List other Hitachi copy products that are compatible and incompatible
  - Identify valid configurations
  - Identify Storage Navigator Modular 2 commands related to TrueCopy Extended Distance software

2

Module Topics

- Need for TrueCopy Extended Distance
- Product Overview
- Internal Overview
- Managing TrueCopy Extended Distance
- Supported System Configurations

3

Business Needs

- Increasing Demands for Business Recoverability
  - Remote data protection
  - Disaster recovery (support for manual operation only)
- Limitations of TrueCopy Synchronous Software
  - Network latency negatively affects host I/O performance.
  - Distance is limited between the two arrays.
    - Maximum theoretical distance extends to 200 kilometers
    - Practical distance is about 30 to 50 kilometers
- Demands for More Recovery Features
  - Process carried out as a background process—does not affect the status of host I/O operations
  - Longer distance disaster recovery and data protection
  - Can be used in lower speed networks
- Acceptable Tradeoffs
  - May lose data when disaster occurred
  - Careful performance design and planning required

4

Network latency may negatively affects host I/O performance as the host is held until acknowledgement is received from the remote system.

Business Objectives

- Disaster Scale
  - What is the assumption for the scale of a disaster or failure?

Type of Disaster	Example	Disaster Scale
Data destruction	File erased by mistake	File
Drive failure	Dual failure on a drive on which RAID5 is being reconfigured	Logical Unit
Subsystem failure	Hardware failure in controller, and more	Subsystem
Site disaster	Fire in building, and more	Site (within some hundreds of meters)
Natural disaster	Earthquake, hurricane and more	Extensive (hundreds of kilometers or more)

- Recovery Point Objective (**RPO**)
  - Up to what point in time should data be recovered in the event of a disaster?
- Recovery Time Objective (**RTO**)
  - How soon should business operations be resumed?
- Cost
  - How much can be budgeted for preparedness measures?

5

- Data Protection
  - Data backups guard against physical destruction and logical destruction
  - Remote backups protect against a site failure
- Disaster Recovery (DR)
  - Work resumes using the copied data
  - Clustering achieves a high RTO

6

Data Protection

- One way to protect data is to periodically create data backups to guard against physical destruction as occurs in dual failures on a drive, and logical destruction that occurs when a file is accidentally erased.
- Backup data can also be created at a remote location as a protection against a site failure or natural disaster. This type of backup is called remote backup.

Disaster Recovery

- In the event of a disaster, work can be resumed using the copied data.
- To achieve a high RTO, clustering configuration is created to be able to continue the work by using the copied data immediately after a failure at the primary site was detected.

Hitachi Recovery Solutions

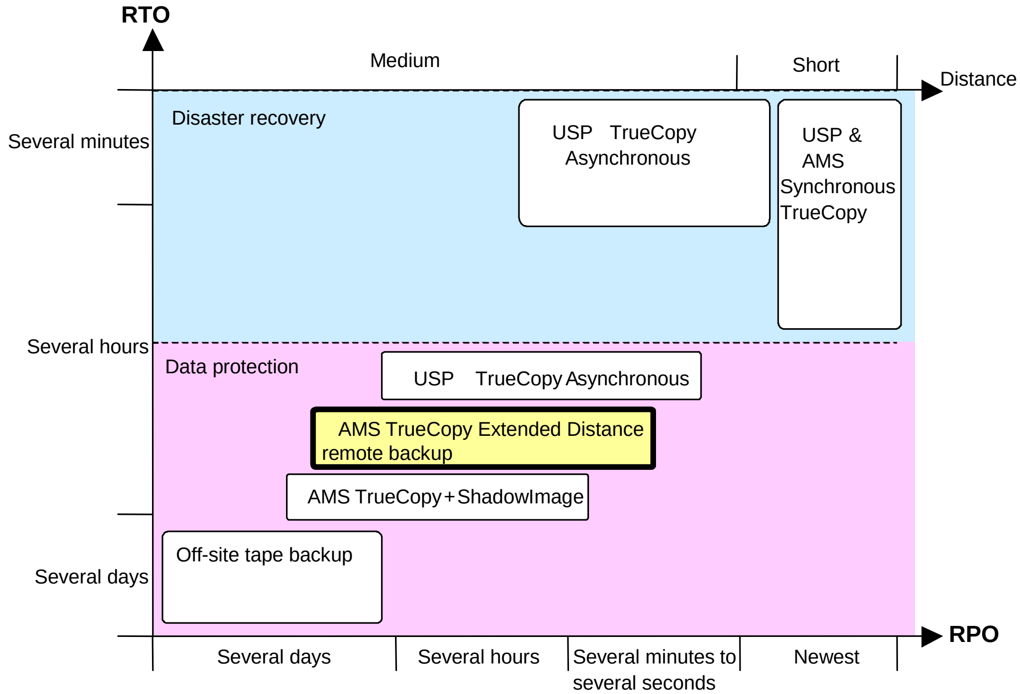

The TrueCopy Extended Distance Software Solution

7

Hitachi ShadowImage Heterogeneous Replication software

Functional Overview

- Because copies are written asynchronously, data on the S-VOL may be older than the data on the primary volume P-VOL.

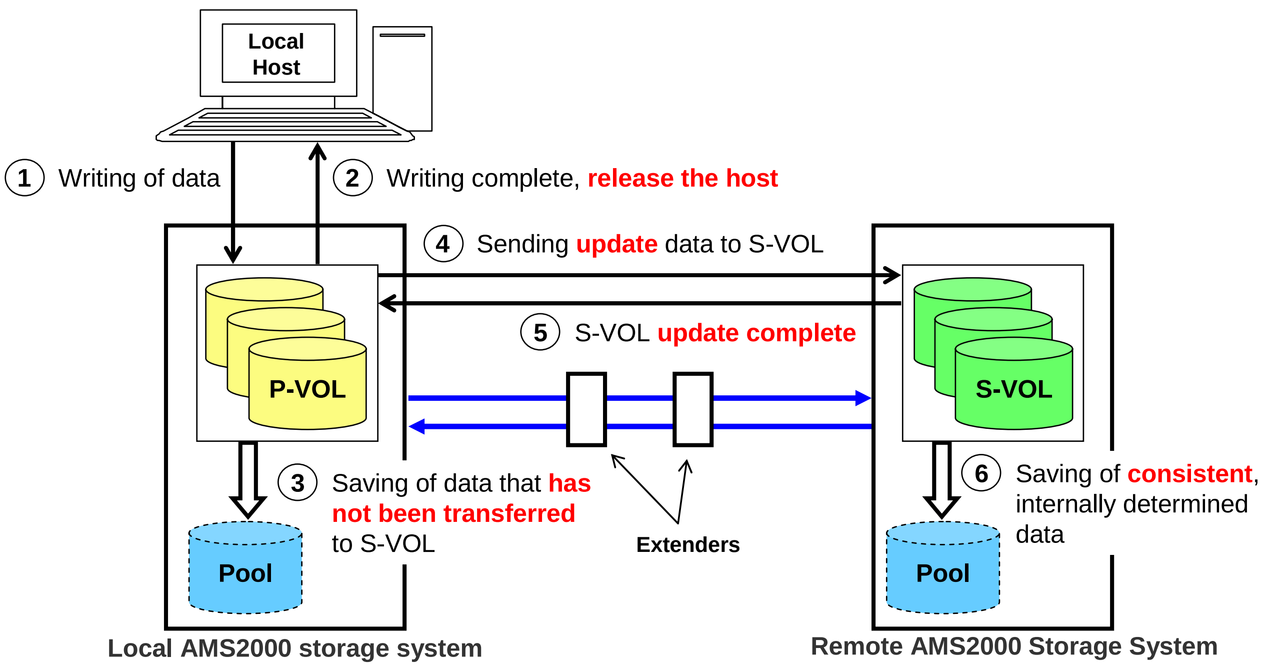

8

Hardware and Software Components

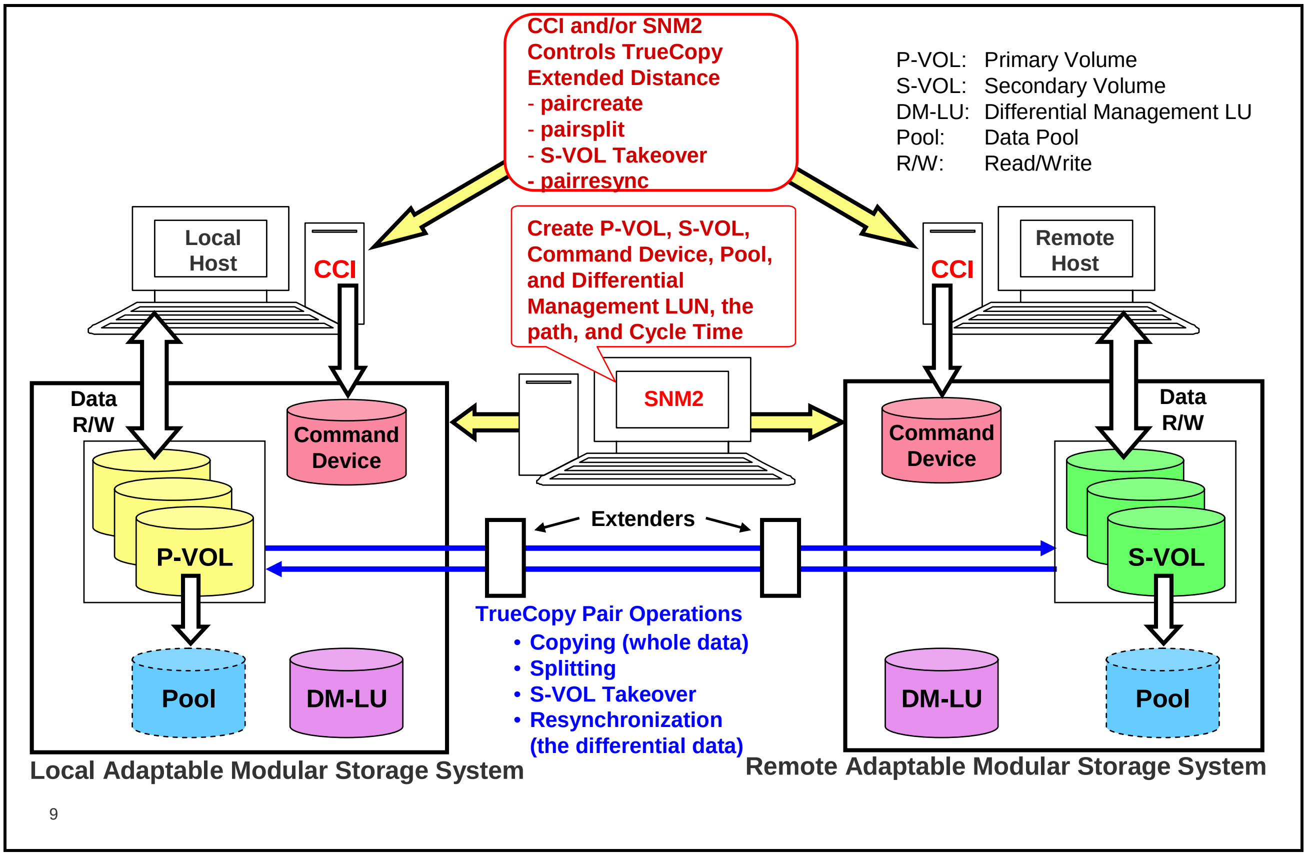

TrueCopy Extended Distance software for a Hitachi Adaptable Modular Storage (AMS) 2000 storage system is a licensed option. The components include:

- Volume Pairs- Primary (P-VOL) is accessible to the production server. The Secondary (S-VOL) can be made accessible to a remote server when in the split state.
- Storage Navigator Modular 2 (SNM2) software is used to create the P-VOL, S-VOL, Pool, Differential Management LU, and the CMD device.
  It is also used to create pairs and then split, resynchronize, delete them and perform a swap.
- Hitachi Command Control Interface (CCI) is used to control TrueCopy Extended Distance software operations from a host using scripts.
- Command Device is used to allow communication between the CCI and the storage system.
- Pool in the local storage system is used to save data that has not been transferred.
- Pool in the remote storage system is used for saving consistent, internally determined data.

If CCI software is not used to manage TrueCopy, Command Devices need not be installed.

Competitive Features

- TrueCopy Extended Distance and Copy-on-Write Snapshot software can be used together. Each requires a separate key.
- Local snapshot creation function (①②).
- Remote snapshot creation (from the local host) command function (③④⑤).
  - Naming function – remote snapshots can be assigned names when created.
  - Time Management function – remote snapshots are time-stamped when created.

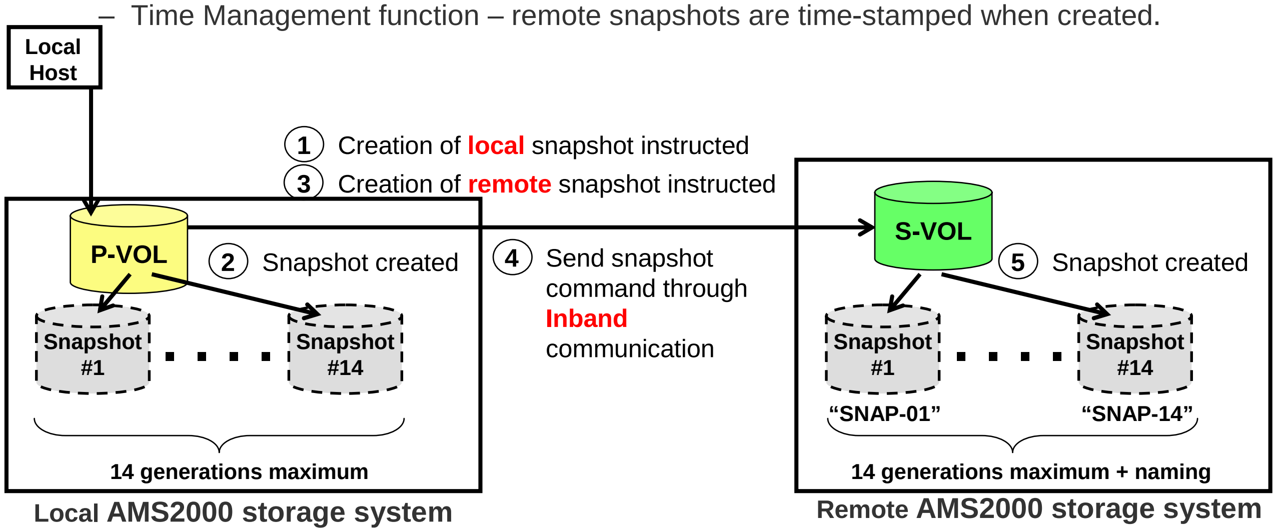

10

Naming function

- The Naming Function adds a human-readable character string to a remote snapshot. Because a snapshot can be identified by, a character string rather than an LU number (Logical Unit Number), a snapshot that includes files to be restored can be easily found from several generations and it reduces the risk of operator error, and more
- The Naming Function also adds day and time string of the server. Example: ***01.12.2006_23:11:30_Backup01***
- At the local host, use the ***pairsplit –mscas*** command to create the remote snapshot.

Time Management function

- The AMS2000 storage system manages at which time a remote snapshot holds P-VOL data. This function simplifies snapshot aging, meaning finding and deleting old snapshots.
- The managed time is the time on the **local** DF700 side. The local and remote controllers have **independent** clocks, respectively. Therefore, when the time management of the remote snap shot is used, it must be set as that there is no big difference in the time of both controllers. It is recommended to use **NTP** (Network Time Protocol) to adjust clocks among the controllers.

Write Order Guarantee

- If a P-VOL failure occurs, I/O operations can be continued by using the S-VOL (the horctakeover command).
- S-VOL data must be updated in the same order that the local host updated the P-VOL.
- Differential data is transmitted in individual cycles, represented as the Time TX cycles in the illustration below. (Cycle Time is discussed later.)

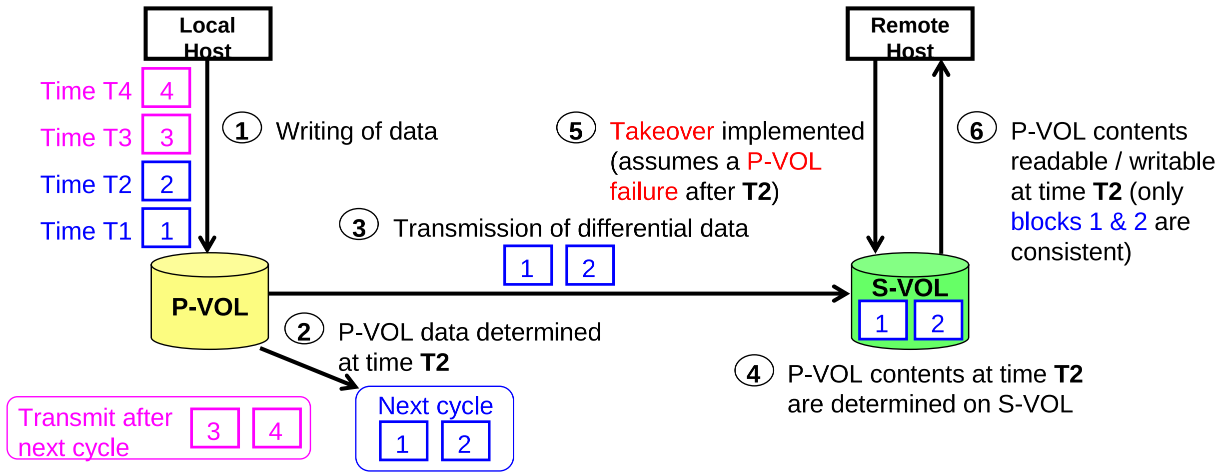

11

Consistency Groups

- TrueCopy Extended Distance software supports Consistency Groups (CTG) which manages multiple logical units as a group and guarantees write order across the logical units.
- At takeover, all internal predetermined data of S-VOLs is restored to the corresponding S-VOLs at the same time, so a host can access consistent data across the S-VOLs.

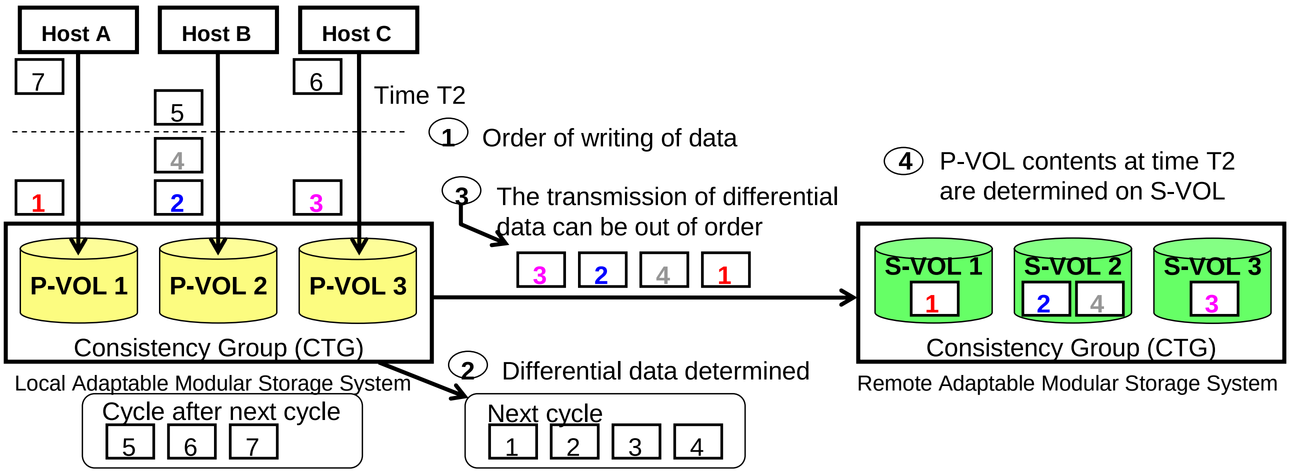

12

Specifications

No	Item	Adaptable Modular Storage 2000 Family
1	User interface of Pair Operation	Storage Navigator Modular 2 (GUI,CLI) CCI (RAID Manager)
2	Controller configuration	Dual controllers
3	Maximum number of pairs	Model 2100: 1024 Models 2300 & 2500: 2046
6	Local system: Remote system	1:1
7	P-VOL: S-VOL	1:1
8	RAID level	RAID-6 , RAID-5, RAID-1, RAID-1+0
9	RAID level of P-VOL to S-VOL	Any to any
10	Drive type for a P-VOL/S-VOL	Any
11	LU size	The size of P-VOL and S-VOL must be the same.
12	FC link	Supported
13	iSCSI link	Supported
14	Remote Path Bandwidth	At least 1.5Mbps per path

13

No	Item	Adaptable Modular Storage 2000 Family
15	Inflow limitation for TrueCopy Extended Distance	Unlimited
16	AMS2000 family storage to previous modular storage TrueCopy Extended Distance	Supported
17	AMS2000 family storage to Universal Storage Platform TrueCopy Extended Distance	Not supported
18	Maximum number of CTGs	16 per storage system
19	Swap Resync	Supported
20	SNMP Trap	• PSUE • POOL threshold over • POOL over • cycle time over • path blockade
21	Indication about cycle information	Local site and Remote site: • Time for determined data • Actual last cycle time

14

Cross-generation functions revert to Adaptable Modular Storage 500 and 1000 specifications. Previous AMS includes models 500 and 1000.

No	Item	Adaptable Modular Storage 2000 Family
22	Pair operation and format • Pair operation during format • Format during pair operation	Not supported
23	Combination with ShadowImage pair • P-VOL of a ShadowImage pair is P-VOL of a TrueCopy Extended Distance pair • P-VOL of a ShadowImage pair is S-VOL of a TrueCopy Extended Distance pair • S-VOL of a ShadowImage pair is P-VOL of a TrueCopy Extended Distance pair	Not supported
24	Combination with SnapShot pair • P-VOL of a SnapShot pair is P-VOL of a TrueCopy Extended Distance pair • P-VOL of a SnapShot pair is S-VOL of a TrueCopy Extended Distance pair	Supported
25	Concurrent use with TrueCopy	Not supported
26	Installation	Reboot is required after installation

15

Pair Status Transition

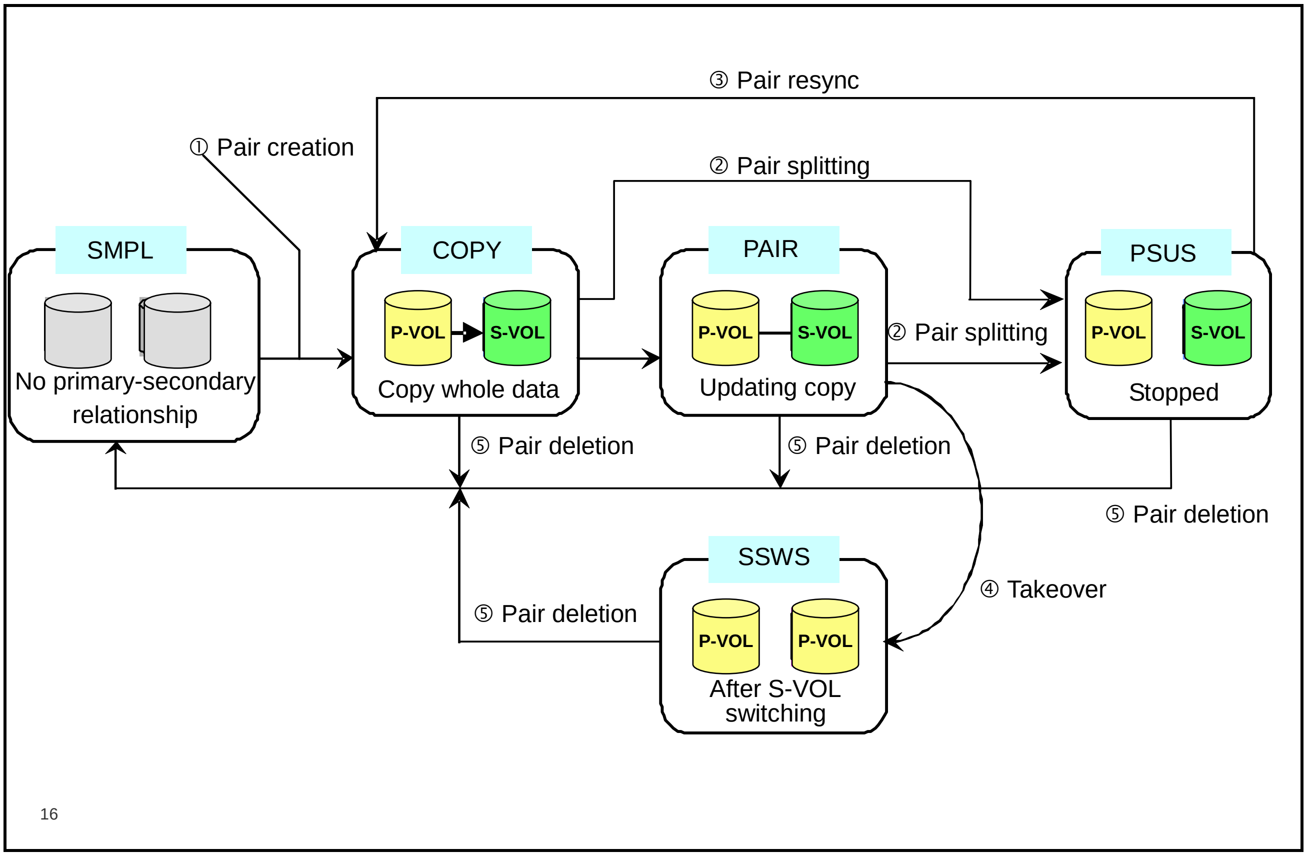

Pair Status Conditions

Pair Status	Description	S-VOL Consistency		Read/Write		Check Tool	
		Pair	CTG	P-VOL	S-VOL	CCI	SNM
SMPL	A TrueCopy Extended Distance pair has not been created	—	—	R/W	R/W	√	√
COPY	Initial copy or recovery copy processing is in progress	x	x	R/W	R/W	√	√
PAIR	Update copy is in progress	OK	OK	R/W	R/W	√	√
PSUS	The pair is suspended, and update copy processing is in progress	OK	x	R/W	R/W or RO (PSUS (N))	√	√
SSWS	Takeover of the S-VOL has been executed	OK	OK	R/W	R/W	√	√
PSUE	Failure stopped Update copy	OK	OK	R/W	R/W	√	√
PFUS	Update copy has stopped by pool capacity overflow	OK	OK	R/W	R/W	√	√
PSUS (N)	Show S-VOL is not consist	x	x	R/W	R/W	√	√
SSWS (R)	S-VOL being internally restored	OK	OK	R/W	R/W	—	√

17

Pair Operations

Pair Operations (1 of 2)				
When to Use	Tasks	Command Name	Contents	Remarks
Configuration and expansion	Pair creation	paircreate	Creates a TrueCopy Extended Distance pair	
	Pair configuration check	pairdisplay	Checks the configuration of the created TrueCopy Extended Distance pair	
	PAIR status waiting	pairevtwait	Waits for transition to the specified status	Wait completion such as initial copy and pair split
Operation and management	Pair status monitoring	pairdisplay	Displays the pair status to confirm whether operation is normal	The -fc option shows the detailed information
	Synchronization monitoring	pairsyncwait	- Checks the time difference between P-VOL & S-VOL - Waits for the completion of copying of specific data written to the P-VOL to the S-VOL	The larger the time difference between the P-VOL and S-VOL, the poorer the RPO performance
	Remote snapshot	pairsplit -mscas	Creates the remote snapshot of the TrueCopy Extended Distance pair P-VOL	P-VOL image when command was issued is taken as a snapshot

18

Pair Operations (2 of 2)				
When to Use	**Tasks**	**Command Name**	**Contents**	**Remarks**
Maintenance	Pair splitting	pairsplit	Temporarily stops update copy for maintenance	
	Pair recovery	pairresync	- Resumes update copy after maintenance or failure	
Disaster recovery	Takeover	horctakeover	Continues operations using S-VOL	- Unsent data at P-VOL is not reflected on S-VOL. - Can be executed only when S-VOL is PAIR, PSUE, PFUS, or PSUS.
System reduction	Pair deletion (Normal)	pairsplit -S	Deletes a TrueCopy Extended Distance pair	
	Pair deletion (Remote Only)	pairsplit -R	Deletes the S-VOL if the S-VOL cannot be correctly deleted	

19

Comparison of TrueCopy Products

Difference Between TrueCopy software and TrueCopy Extended Distance			
	Comparison Item	TrueCopy software Synchronous remote copy	TrueCopy Extended Distance software Asynchronous remote copy
1	Write sequence guarantee when S-VOL is updated	Yes	Each cycle
2	RPO performance	Yes	If an error occurs before the end of a cycle, TrueCopy Extended Distance cannot reflect that data on S-VOL
3	IOPS performance	Yes	Because of pool saving processing (in the local and remote system), I/O performance of TrueCopy Extended Distance is lower than that of TrueCopy software
4	Response performance	Depends on the line delay	Yes
5	Long-distance connection	TrueCopy software is not suitable, because as line delay time increases, host I/O timeouts occur	Yes
6	Pool	Not required	Required
7	Application	Backup, disaster recovery, data transfer	Backup, disaster recovery, data transfer. TrueCopy Extended Distance is not recommended for disaster recovery because it does not support a swap function that replaces P-VOL with S-VOL

20

Comparison with Copy-on-Write Snapshot Software

Comparison with Copy-on-Write Snapshot		
	Snapshot	**TrueCopy Extended Distance**
Threshold Over	Pair statuses change to **PFUS** Pool status also changes to PFUS	Pair statuses **do not change** Pool status changes to PFUS
Pool Over at Local	Pair status changes to **PSUE**	P-VOL with PAIR status changes to **PFUS** P-VOL with COPY status changes to **PSUE**
Pool Over at Remote	Pair status changes to **PSUE**	P-VOL changes to **PSUE** S-VOL with PAIR status changes to **PFUS** S-VOL with COPY status changes to **PSUS(N)**
Data Consistency	V-VOL data is **invalid** after pool over	S-VOL is **consistent** at CTG level after pool over
Recover from PSUE	**Delete and recreate** a pair	**Resync** a pair
Split-Maker	**Not supported**	**Supported** with pairsplit -mscas option
Failures	Pair status changes to **PSUE** V-VOL data is **invalid**	1) Failures at Local: P-VOL changes to **PSUE** S-VOL does not change and is **valid** 2) Failures at Remote: P-VOL change to **PSUE** S-VOL changes to PFUS and is **invalid**
Specifying CTG	paircreate **-m grp** [CTG#]	paircreate **-f async** [CTG#]
# of CTG	128	16

21

TrueCopy Extended Distance and Copy-on-Write Snapshot share same pool so they affect each other's performance.

Internal Functions

- TrueCopy Extended Distance remote replication operations are recurring processes executed in specific time periods (cycles), during which, updates to the P-VOL are written to the S-VOL

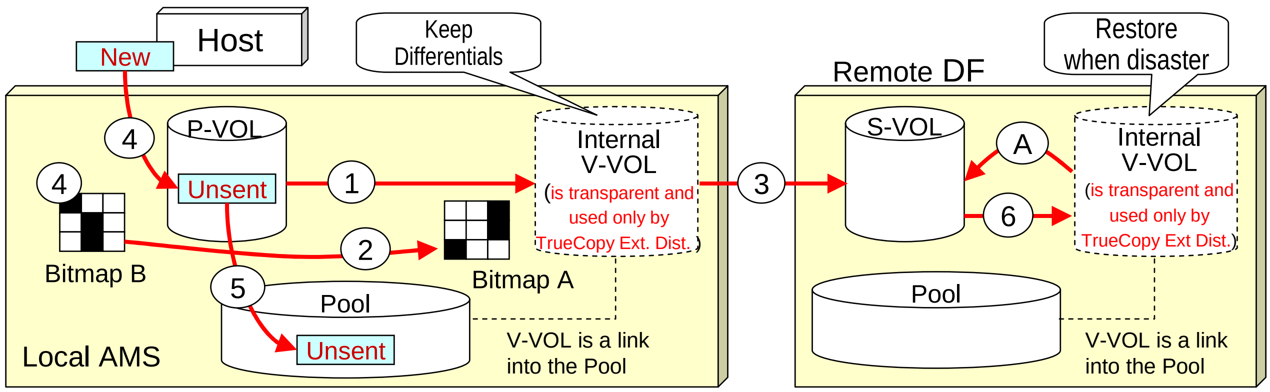

1 Create snapshot of P-VOL (assumes a previous write has already occurred)

2 Copy Bitmap B to A. Clear Bitmap B

3 Send differentials to S-VOL according to Bitmap A

4 New write is cached. Update Bitmap B

5 Copy-on-write to a pool when un-sent differential is overwritten

6 Create snapshot of S-VOL when all differentials have been copied to S-VOL

Note: One cycle shows repetition of 1 to 6

A Restore V-VOL to S-VOL when a disaster occurs, and then use S-VOL to restore P-VOL

22

TrueCopy Extended Distance uses a subset of the Copy-on-Write Snapshot functions for internal control of the V-VOL. It does not require that the Copy-on-Write Snapshot license key be installed.

- Cycle Update Timing Chart

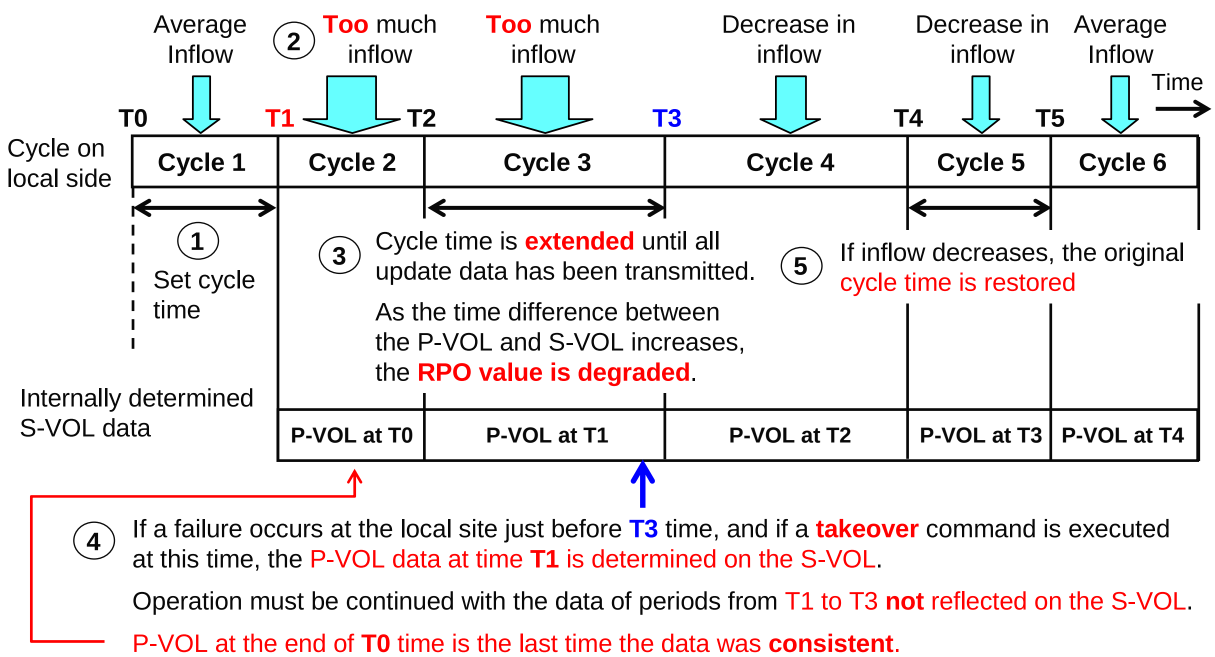

23

Transmitting of the data continues beyond the end of the set cycle time until all differential has been transmitted. It is only in the event of an error that the system has to backup to the cycle when consistent data was determined on S-VOL.

- Use Storage Navigator Modular to set the Cycle Time
- Settings
  - Minimum 30 seconds
  - Default 300 seconds (5 minutes)
  - Maximum 3600 seconds (60 minutes)
- Cycle Time must be larger than or equal to the number of used Consistency Groups (CTG) x 30 seconds

Cycle Time Minimum Values	
Number of CTGs	**Minimum Cycle Time**
1	30 seconds
2	60 seconds
3	90 seconds
. . .	. . .
16	240 seconds

24

- Differential Data Management – Determining differential data
  - What does the Cycle Update process do?
  - Copies differential data during a specific cycle to S-VOL and makes S-VOL data identical to P-VOL data at the beginning of the cycle
  - S-VOL determination
    - S-VOL data becomes identical to P-VOL at the beginning of the cycle
    - The remote system keeps the image as internal pre-determined data which becomes available by executing the takeover command
    - See next slide for timing chart

25

- Determining differential data at S-VOL

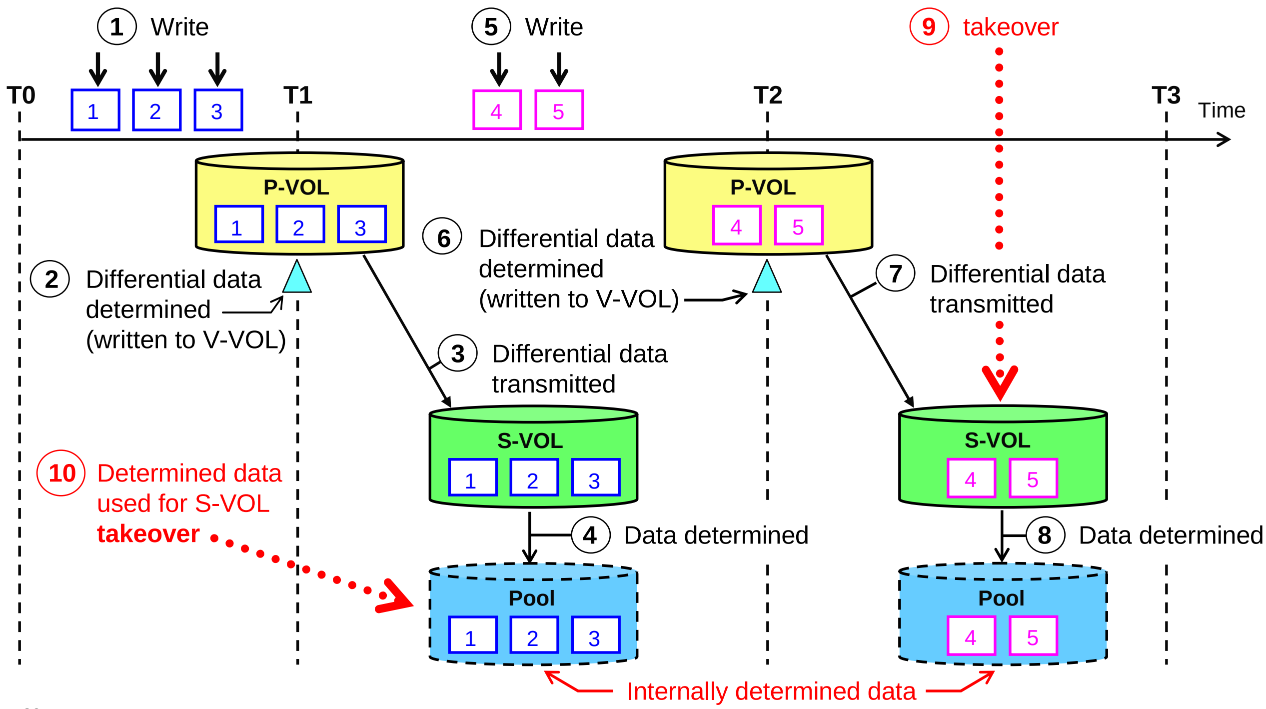

26

- Pool-threshold over and S-VOL recoverability
  - When a pool-over occurs:
    - Snapshot V-VOLs change to **PSUE** (TrueCopy Extended Distance, and Copy-on-Write Snapshot software).
    - TrueCopy Extended Distance keeps internally determined data.
    - During **takeover**, the internally determined data is restored from the pool to the S-VOL.

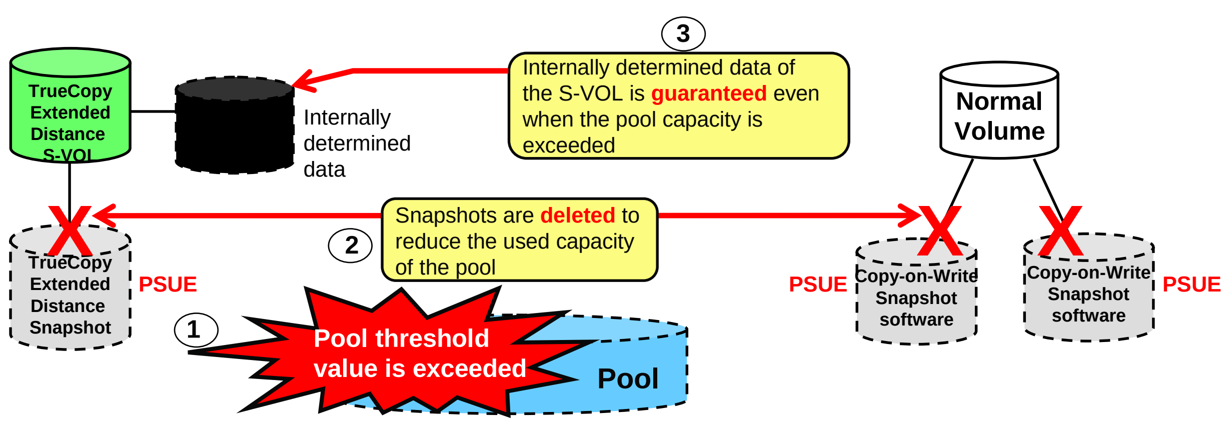

27

Management Software

- Function Chart

Desired Function		SNM2	CCI	Performance Monitor	WebTool
Design	Inflow estimation			✓	
Configuration	Command device creation	✓			
	DM LU creation	✓			
	Logical unit creation	✓			
	Pool creation	✓			
	Remote path setting	✓			
	Cycle time setting	✓			
	TrueCopy Extended Distance pair creation	✓	✓		
	Pair configuration check	✓	✓		
	PAIR status wait	✓	✓		
	Forecasting initial copy completion time	✓	✓		
Operation and Monitoring	Monitoring pair status		✓		
	Monitoring synchronization	✓	✓		
	Monitoring used pool capacity used	✓			
	Collecting performance information	✓		✓	
	Collecting failure information	✓			✓
	Remote snapshot	✓	✓		

28

- Function Chart (*continued*)

Desired Function		SNM	CCI	Performance Monitor	Web Tool
Maintenance	Pair splitting	✓	✓		
	Pair recovery	✓	✓		
	Changing of pool size	✓			
	Recovery after path blocked	✓			✓
Disaster Recovery	takeover/swap	Note 1	Note 1		
System Reduction	Pair deletion	✓	✓		
	Logical unit and pool deletion, and more	✓			

Note 1: Pair swapping and pair re-synchronization after takeover are not supported.

29

Performance Monitor of SNM2

- Performance Monitor of SNM2
  - Design – Measure inflow from the host in order to design system performance
  - Operation and Maintenance
    - Check if the inflow is within the estimation.
    - Measure the outflow to remote system.

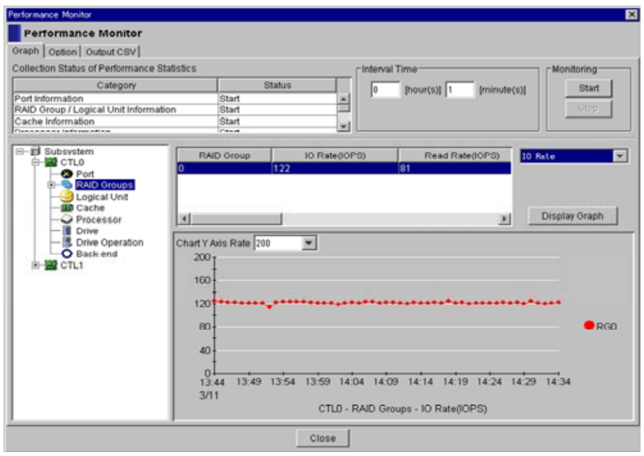

30

Web Tool

- Graphical User Interface
  - Operation and monitoring – Web Tool is used to check failure information (messages).
  - Maintenance – If the remote path between the local storage system and remote storage system is broken, Web Tool is used to recover the path.

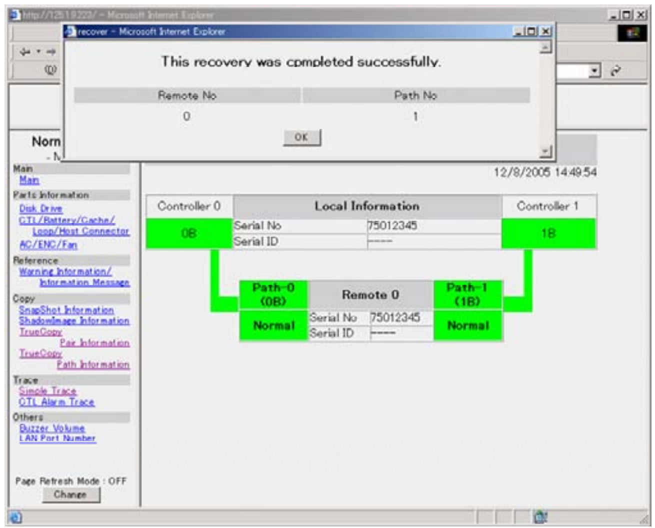

31

Supported Configurations

- Examples of Supported Configurations

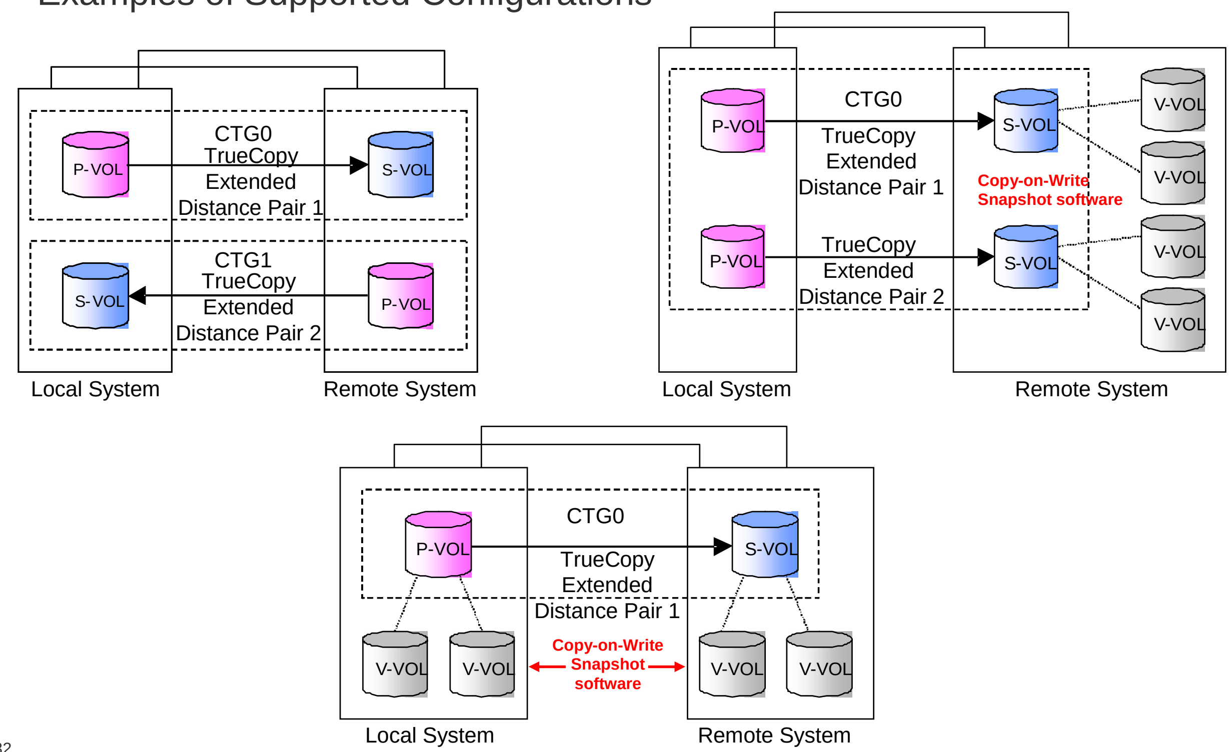

32

- Examples of Unsupported Configurations

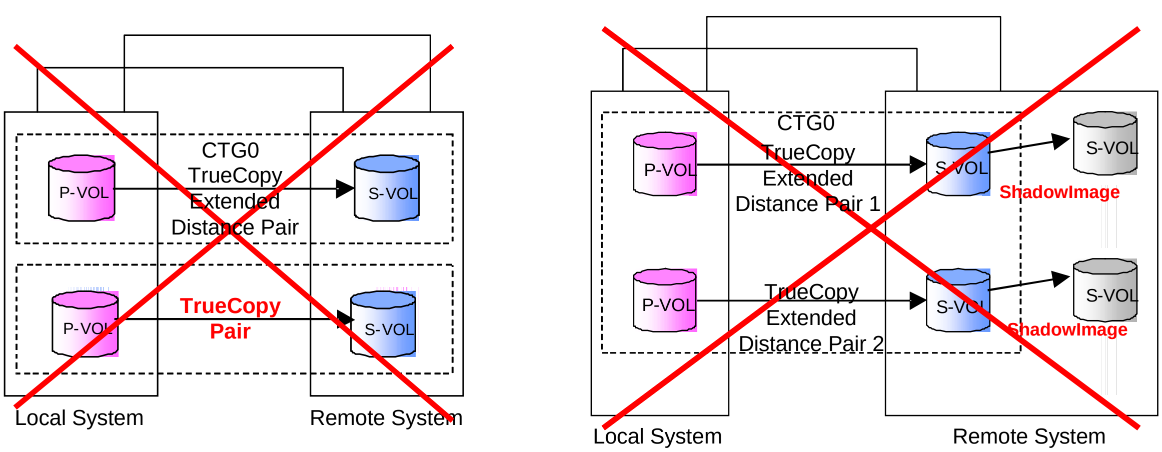

33

- Example of TrueCopy Extended Distance and Copy-on-Write Snapshot

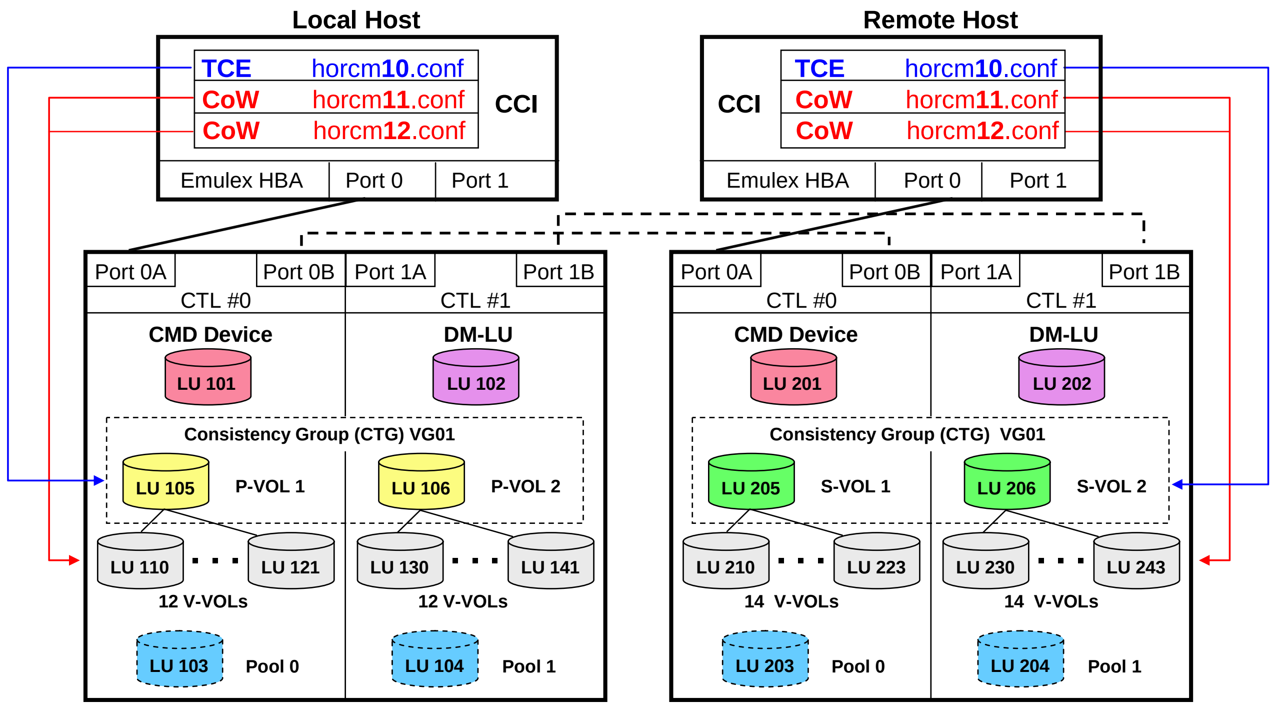

34

In the diagram,

TCE stands for TrueCopy Extended Distance software

CoW stands for Copy-on-Write Snapshot software

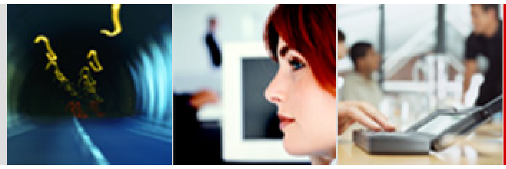

Training Course Glossary

Click a letter below to jump to that letter's terms in the glossary.

A B C D E F G H I J K L M N O P Q R S T U V W X Y Z

—A—

ACC— Action Code. A SIM System Information Message. Will produce an ACC which takes an engineer to the correct fix procedures in the ACC directory in the MM (Maintenance Manual)

ACE (Access Control Entry) — Stores access rights for a single user or group within the Windows security model

ACL (Access Control List)— stores a set of ACEs, so describes the complete set of access rights for a file system object within the Microsoft Windows security model

ACP (Array Control Processor) — Microprocessor mounted on the disk adapter circuit board (DKA) that controls the drives in a specific disk array. Considered part of the back-end, it controls data transfer between cache and the hard drives.

ACP PAIR — Physical disk access control logic. Each ACP consists of two DKA PCBs. To provide 8 loop paths to the real HDDs

Actuator (arm) — read/write heads are attached to a single head actuator, or actuator arm, that moves the heads around the platters

AD — Active Directory

ADC — Accelerated Data Copy

ADP —Adapter

ADS — Active Directory Service

Address— A location of data, usually in main memory or on a disk. A name or token that identifies a network component. In local area networks (LANs), for example, every node has a unique address

AIX — IBM UNIX

AL (Arbitrated Loop) — A network in which nodes contend to send data and only one node at a time is able to send data.

AL-PA — Arbitrated Loop Physical Address

AMS —Adaptable Modular Storage

APID — An ID to identify a command device.

APF (Authorized Program Facility) — In z/OS and OS/390 environments, a facility that permits the identification of programs that are authorized to use restricted functions.

Application Management —The processes that manage the capacity and performance of applications

ARB — Arbitration or "request"

Array Domain—all functions, paths, and disk drives controlled by a single ACP pair. An array domain can contain a variety of LVI and/or LU configurations.

ARRAY UNIT - A group of Hard Disk Drives in one RAID structure. Same as Parity Group

ASIC — Application specific integrated circuit

ASSY — Assembly

Asymmetric virtualization — See Out-of-band virtualization.

Asynchronous— An I/O operation whose initiator does not await its completion before proceeding with other work. Asynchronous I/O operations enable an initiator to have multiple concurrent I/O operations in progress.

ATA — Short for ***A****dvanced* ***T****echnology* ***A****ttachment,* a disk drive implementation that integrates the controller on the disk drive itself, also known as IDE (Integrated Drive Electronics) Advanced Technology Attachment is a standard designed to connect hard and removable disk drives

Authentication — The process of identifying an individual, usually based on a username and password.

Availability — Consistent direct access to information over time

-back to top-

—B—

B4 — A group of 4 HDU boxes that are used to contain 128 HDDs

Backend— In client/server applications, the client part of the program is often called the front-end and the server part is called the back-end. Backup image—Data saved during an archive operation. It includes all the associated files, directories, and catalog information of the backup operation.

BATCTR — Battery Control PCB

BED — Back End Director. Controls the paths to the HDDs

Bind Mode — One of two modes available when using FlashAccess™, in which the FlashAccess™ extents hold read data for specific extents on volumes (see Priority Mode).

BST — Binary Search Tree

BTU— British Thermal Unit

Business Continuity Plan — Describes how an organization will resume partially- or completely interrupted critical functions within a predetermined time after a disruption or a disaster. Sometimes also called a Disaster Recovery Plan.

-back to top-

—C—

CA — Continuous Access software (see HORC)

Cache — Cache Memory**.** Intermediate buffer between the channels and drives. It has a maximum of 64 GB (32 GB x 2 areas) of capacity. It is available and controlled as two areas of cache (cache A and cache B). It is fully battery-backed (48 hours) .

Cache hit rate — When data is found in the cache, it is called a cache hit, and the effectiveness of a cache is judged by its hit rate.

Cache partitioning — Storage management software that allows the virtual partitioning of cache and allocation of it to different applications

CAD — Computer-Aided Design

Capacity — Capacity is the amount of data that a drive can store after formatting. Most data storage companies, including HDS, calculate capacity based on the assumption that 1 megabyte = 1000 kilobytes and 1 gigabyte=1,000 megabytes.

CAPEX - capital expenditure **-** is the cost of developing or providing non-consumable parts for the product or system. For example, the purchase of a photocopier is the CAPEX, and the annual paper and toner cost is the OPEX. (See OPEX).

CAS — Column address strobe is a signal sent to a dynamic random access memory (DRAM) that tells it that an associated address is a column address. CAS- column address strobe sent by the processor to a DRAM circuit to activate a column address.

CCI — Command Control Interface

CE — Customer Engineer

Centralized management —Storage data management, capacity management, access security management, and path management functions accomplished by software.

CentOS— **C**ommunity **Ent**erprise **O**perating **S**ystem

CFW— Cache Fast Write

CHA (Channel Adapter) — Provides the channel interface control functions and internal cache 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.

CH — Channel

CHA — Channel Adapter

CHAP — Challenge-Handshake Authentication Protocol

CHF — Channel Fibre

CHIP (Client-Host Interface Processor) — Microprocessors on the CHA boards that process the channel commands from the hosts and manage host access to cache.

CHK— Check

CHN — **CH**annel adapter **N**AS

CHP — Channel Processor or Channel Path

CHPID — Channel Path Identifier

CH S— Channel SCSI

CHSN — Cache memory Hierarchical Star Network

CHT—Channel tachyon, a Fibre Channel protocol controller

CIFS protocol — common internet file system is a platform-independent file sharing system. A network file system access protocol primarily used by Windows clients to communicate file access requests to Windows servers.

CIM — Common Information ModelCKD (Count-key Data) — A format for encoding data on hard disk drives; typically used in the mainframe environment.

CKPT — Check Point

CL — See Cluster

CLI — Command Line Interface

CLPR (**C**ache **L**ogical **P**a**R**tition) — Cache can be divided into multiple virtual cache memories to lessen I/O contention.

Cluster — A collection of computers that are interconnected (typically at high-speeds) for the purpose of improving reliability, availability, serviceability and/or performance (via load balancing). Often, clustered computers have access to a common pool of storage, and run special software to coordinate the component computers' activities.

CM (Cache Memory Module) — Cache Memory**.** Intermediate buffer between the channels and drives. It has a maximum of 64 GB (32 GB x 2 areas) of capacity. It is available and controlled as two areas of cache (cache A and cache B). It is fully battery-backed (48 hours)

CM PATH (Cache Memory Access Path) — Access Path from the processors of CHA, DKA PCB to Cache Memory.

CMD — Command

CMG — Cache Memory Group

CNAME — Canonical NAME

CPM (Cache Partition Manager) — Allows for partitioning of the cache and assigns a partition to a LU; this enables tuning of the system’s performance.

CNS— Clustered Name Space

Concatenation — A logical joining of two series of data. Usually represented by the symbol “|”. In data communications, two or more data are often concatenated to provide a unique name or reference (e.g., S_ID | X_ID). Volume managers concatenate disk address spaces to present a single larger address spaces.

Connectivity technology — a program or device's ability to link with other programs and devices. Connectivity technology allows programs on a given computer to run routines or access objects on another remote computer

Controller — A device that controls the transfer of data from a computer to a peripheral device (including a storage system) and vice versa.

Controller-based Virtualization — Driven by the physical controller at the hardware microcode level versus at the application software layer and integrates into the infrastructure to allow virtualization across heterogeneous storage and third party products

Corporate governance — Organizational compliance with government-mandated regulations

COW — Copy On Write Snapshot

CPS — Cache Port Slave

CPU — Central Processor Unit

CRM — Customer Relationship Management

CruiseControl — Now called Hitachi Volume Migration software

CSV — Comma Separated Value

CSW (Cache Switch PCB) — The cache switch (CSW) connects the channel adapter or disk adapter to the cache. Each of them is connected to the cache by the Cache Memory Hierarchical Star Net (C-HSN) method. Each cluster is provided with the two CSWs, and each CSW can connect four caches. The CSW switches any of the cache paths to which the channel adapter or disk adapter is to be connected through arbitration.

CU (Control Unit) — The hexadecimal number to which 256 LDEVs may be assigned

CUDG —Control Unit DiaGnostics. Internal system tests.

CV — Custom Volume

CVS (Customizable Volume Size) — software used to create custom volume sizes. Marketed under the name Virtual LVI (VLVI) and Virtual LUN (VLUN)

-back to top-

—D—

DAD (Device Address Domain) — Indicates a site of the same device number automation support function. If several hosts on the same site have the same device number system, they have the same name.

DACL — Discretionary ACL - the part of a security descriptor that stores access rights for users and groups.

DAMP (Disk Array Management Program) — Renamed to Storage Navigator Modular (SNM)

DAS — Direct Attached Storage

DASD—Direct Access Storage Device

Data Blocks — A fixed-size unit of data that is transferred together. For example, the X-modem protocol transfers blocks of 128 bytes. In general, the larger the block size, the faster the data transfer rate.

Data Integrity —Assurance that information will be protected from modification and corruption.

Data Lifecycle Management — An approach to information and storage management. The policies, processes, practices, services and tools used to align the business value of data with the most appropriate and cost-effective storage infrastructure from the time data is created through its final disposition. Data is aligned with business requirements through management policies and service levels associated with performance, availability, recoverability, cost and what ever parameters the organization defines as critical to its operations.

Data Migration— The process of moving data from one storage device to another. In this context, data migration is the same as Hierarchical Storage Management (HSM).

Data Pool— A volume containing differential data only.

Data Striping — Disk array data mapping technique in which fixed-length sequences of virtual disk data addresses are mapped to sequences of member disk addresses in a regular rotating pattern.

Data Transfer Rate (DTR) — The speed at which data can be transferred. Measured in kilobytes per second for a CD-ROM drive, in bits per second for a modem, and in megabytes per second for a hard drive. Also, often called simply *data rate*.

DCR (Dynamic Cache Residency) — see FlashAccess™

DE— Data Exchange Software

Device Management — Processes that configure and manage storage systems

DDL — Database Definition Language

DDNS —Dynamic DNS

DFS — Microsoft Distributed File System

DFW —DASD Fast Write

DIMM—Dual In-line Memory Module

Direct Attached Storage — Storage that is directly attached to the application or file server. No other device on the network can access the stored data

Director class switches — larger switches often used as the core of large switched fabrics

Disaster Recovery Plan (DRP) — A plan that describes how an organization will deal with potential disasters. It may include the precautions taken to either maintain or quickly resume mission-critical functions. Sometimes also referred to as a *Business Continuity Plan*.

Disk Administrator — An administrative tool that displays the actual LU storage configuration

Disk Array — A linked group of one or more physical independent hard disk drives generally used to replace larger, single disk drive systems. The most common disk arrays are in daisy chain configuration or implement RAID (Redundant Array of Independent Disks) technology. A disk array may contain several disk drive trays, and is structured to improve speed and increase protection against loss of data. Disk arrays organize their data storage into Logical Units (LUs), which appear as linear block paces to their clients. A small disk array, with a few disks, might support up to 8 LUs; a large one, with hundreds of disk drives, can support thousands.

DKA (Disk Adapter) — Also called an array control processor (ACP); it provides the control functions for data transfer between drives and cache. The DKA contains DRR (Data Recover and Reconstruct), a parity generator circuit. It supports four fibre channel paths and offers 32 KB of buffer for each fibre channel path.

DKC (Disk Controller Unit) — In a multi-frame configuration, the frame that contains the front end (control and memory components).

DKCMN — Disk Controller Monitor. Monitors temperature and power status throughout the machine

DKF (fibre disk adapter) — Another term for a DKA.DKU (Disk Unit) — In a multi-frame configuration, a frame that contains hard disk units (HDUs).

DLIBs — Distribution Libraries

DLM —Data Lifecycle Management

DMA— Direct Memory Access

DM-LU (Differential Management Logical Unit) — DM-LU is used for saving management information of the copy functions in the cache

DMP — Disk Master Program

DNS — Domain Name System

Domain — A number of related storage array groups. An "ACP Domain" or "Array Domain" means all of the array-groups controlled by the same pair of DKA boards.
OR
— The HDDs managed by one ACP PAIR (also called BED)

DR — Disaster Recovery

DRR (Data Recover and Reconstruct) —Data Parity Generator chip on DKA

DRV — Dynamic Reallocation Volume

DSB — Dynamic Super Block

DSP — Disk Slave Program

DTA —Data adapter and path to cache-switches

DW — Duplex Write

DWL — Duplex Write Line

Dynamic Link Manager — HDS software that ensures that no single path becomes overworked while others remain underused. Dynamic Link Manager does this by providing automatic load balancing, path failover, and recovery capabilities in case of a path failure.

-back to top-

—E—

ECC — Error Checking & Correction

ECC.DDR SDRAM — **E**rror **C**orrection **C**ode **D**ouble **D**ata **R**ate **S**ynchronous **D**ynamic **RA**m **M**emory

ECN — Engineering Change Notice

E-COPY — Serverless or LAN free backup

ENC — Stands for **EN**closure **C**ontroller, the units that connect the controllers in the DF700 with the Fibre Channel disks. They also allow for online extending a system by adding RKAs

ECM— Extended Control Memory

EOF — End Of Field

EPO — Emergency Power Off

ENC — Enclosure

EREP — Error REporting and Printing

ERP — Enterprise Resource Management

ESA — Enterprise Systems Architecture

ESC — Error Source Code

ESCD — ESCON Director

ESCON (Enterprise Systems Connection) — An input/output (I/O) interface for mainframe computer connections to storage devices developed by IBM.

Ethernet — A local area network (LAN) architecture that supports clients and servers and uses twisted pair cables for connectivity.

EVS — Enterprise Virtual Server

ExSA — Extended Serial Adapter

-back to top-

—F—

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

Failback — The restoration of a failed system share of a load to a replacement component. For example, when a failed controller in a redundant configuration is replaced, the devices that were originally controlled by the failed controller are usually failed back to the replacement controller to restore the I/O balance, and to restore failure tolerance. Similarly, when a defective fan or power supply is replaced, its load, previously borne by a redundant component, can be failed back to the replacement part.

Failed over — A mode of operation for failure tolerant systems in which a component has failed and its function has been assumed by a redundant component. A system that protects against single failures operating in failed over mode is not failure tolerant, since failure of the redundant component may render the system unable to function. Some systems (e.g., clusters) are able to tolerate more than one failure; these remain failure tolerant until no redundant component is available to protect against further failures.

Failover — A backup operation that automatically switches to a standby database server or network if the primary system fails, or is temporarily shut down for servicing. Failover is an important fault tolerance function of mission-critical systems that rely on constant accessibility. Failover automatically and transparently to the user redirects requests from the failed or down system to the backup system that mimics the operations of the primary system.

Failure tolerance — The ability of a system to continue to perform its function or at a reduced performance level, when one or more of its components has failed. Failure tolerance in disk subsystems is often achieved by including redundant instances of components whose failure would make the system inoperable, coupled with facilities that allow the redundant components to assume the function of failed ones.

FAIS — Fabric Application Interface Standard

FAL — File Access Library

FAT — File Allocation Table

Fault Tolerat — Describes a computer system or component designed so that, in the event of a component failure, a backup component or procedure can immediately take its place with no loss of service. Fault tolerance can be provided with software, embedded in hardware, or provided by some hybrid combination.

FBA — Fixed-block Architecture. Physical disk sector mapping.

FBA/CKD Conversion — The process of converting open-system data in FBA format to mainframe data in CKD format.

FBA — Fixed Block Architecture

FBUS — Fast I/O Bus

FC — Fibre Channel is a technology for transmitting data between computer devices; a set of standards for a serial I/O bus capable of transferring data between two ports

FC-0 — Lowest layer on fibre channel transport, it represents the physical media.

FC-1 — This layer contains the 8b/10b encoding scheme.

FC-2 — This layer handles framing and protocol, frame format, sequence/exchange management and ordered set usage.

FC-3 — This layer contains common services used by multiple N_Ports in a node.

FC-4 — This layer handles standards and profiles for mapping upper level protocols like SCSI an IP onto the Fibre Channel Protocol.

FCA — Fibre Adapter. Fibre interface card. Controls transmission of fibre packets.

FC-AL — Fibre Channel Arbitrated Loop. A serial data transfer architecture developed by a consortium of computer and mass storage device manufacturers and now being standardized by ANSI. FC-AL was designed for new mass storage devices and other peripheral devices that require very high bandwidth. Using optical fiber to connect devices, FC-AL supports full-duplex data transfer rates of 100MBps. FC-AL is compatible with SCSI for high-performance storage systems.

FC-P2P — Fibre Channel Point-to-Point

FC-SW — Fibre Channel Switched

FCC — Federal Communications Commission

FC — Fibre Channel or Field-Change (microcode update)

FCIP –**F**ibre **C**hannel over **IP**, a network storage technology that combines the features of Fibre Channel and the Internet Protocol (IP) to connect distributed SANs over large distances. FCIP is considered *a tunneling protocol*, as it makes a transparent point-to-point connection between geographically separated SANs over IP networks. FCIP relies on TCP/IP services to establish connectivity between remote SANs over LANs, MANs, or WANs. An advantage of FCIP is that it can use TCP/IP as the transport while keeping Fibre Channel fabric services intact.

FCP — Fibre Channel Protocol

FC RKAJ (Fibre Channel Rack Additional) — Acronym referring to an additional rack unit(s) that houses additional hard drives exceeding the capacity of the core RK unit of the Thunder 9500V/9200 subsystem.

FCU— File Conversion Utility

FD — Floppy Disk

FDR— Fast Dump/Restore

FE — Field Engineer

FED — Channel Front End Directors

Fibre Channel — A serial data transfer architecture developed by a consortium of computer and mass storage device manufacturers and now being standardized by ANSI. The most prominent Fibre Channel standard is Fibre Channel Arbitrated Loop (FC-AL).

FICON (Fiber Connectivity) — A high-speed input/output (I/O) interface for mainframe computer connections to storage devices. As part of IBM's S/390 server, FICON channels increase I/O capacity through the combination of a new architecture and faster physical link rates to make them up to eight times as efficient as ESCON (Enterprise System Connection), IBM's previous fiber optic channel standard.

Flash ACC — Flash access. Placing an entire LUN into cache

FlashAccess — HDS software used to maintain certain types of data in cache to ensure quicker access to that data.

FLGFAN — Front Logic Box Fan Assembly.

FLOGIC Box — Front Logic Box.

FM (Flash Memory) — Each microprocessor has FM. FM is non-volatile memory which contains microcode.

FOP — Fibre Optic Processor or fibre open

FPC — Failure Parts Code or Fibre Channel Protocol Chip

FPGA — Field Programmable Gate Array

Frames — An ordered vector of words that is the basic unit of data transmission in a Fibre Channel network.

Front-end — In client/server applications, the client part of the program is often called the front end and the server part is called the back end.

FS — File System

FSA — File System Module-A

FSB — File System Module-B

FSM — File System Module

FSW (Fibre Channel Interface Switch PCB) — A board that provides the physical interface (cable connectors) between the ACP ports and the disks housed in a given disk drive.

FTP (File Transfer Protocol) — A client-server protocol which allows a user on one computer to transfer files to and from another computer over a TCP/IP network

FWD — Fast Write Differential

-back to top-

—G—

GARD — General Available Restricted Distribution

GB — Gigabyte

GBIC — Gigabit Interface Converter

GID — Group Identifier

GID — **Gr**oup **Id**entifier within the Unix security model

GigE — Giga Bit Ethernet

GLM — Gigabyte Link Module

Global Cache — Cache memory is used on demand by multiple applications, use changes dynamically as required for READ performance between hosts/applications/LUs.

Graph-Track™ — HDS software used to monitor the performance of the Hitachi storage subsystems. Graph-Track™ provides graphical displays, which give information on device usage and system performance.

GUI — Graphical User Interface

-back to top-

—H—

H1F — Essentially the Floor Mounted disk rack (also called Desk Side) equivalent of the RK. (See also: RK, RKA, and H2F).

H2F — Essentially the Floor Mounted disk rack (also called Desk Side) add-on equivalent similar to the RKA. There is a limitation of only one H2F that can be added to the core RK Floor Mounted unit. (See also: RK, RKA, and H1F).

HLU (Host Logical Unit) — A LU that the Operating System and the HDLM recognizes. Each HLU includes the devices that comprise the storage LU

H-LUN — Host Logical Unit Number (See LUN)

HA — High Availability

HBA — Host Bus Adapter—An HBA is an I/O adapter that sits between the host computer's bus and the Fibre Channel loop and manages the transfer of information between the two channels. In order to minimize the impact on host processor performance, the host bus adapter performs many low-level interface functions automatically or with minimal processor involvement.

HDD (Hard Disk Drive) — A spindle of hard disks that make up a hard drive, which is a unit of physical storage within a subsystem.

HD — Hard Disk

HDev (Hidden devices) — Hitachi Tuning Manager Main Console may not display some drive letters in its resource tree, and information such as performance and capacity is not available for such invisible drives. This problem occurs if there is a physical drive with lower PhysicalDrive number assigned that is in “damaged (SCSI Inquiry data cannot be obtained)” or “hidden by HDLM” status.

HDS — Hitachi Data Systems

HDU (Hard Disk Unit) — A number of hard drives (HDDs) grouped together within a subsystem.

HDLM — Hitachi Dynamic Link Manager software

Head — See read/write head

Heterogeneous — The characteristic of containing dissimilar elements. A common use of this word in information technology is to describe a product as able to contain or be part of a *heterogeneous network,"* consisting of different manufacturers' products that can interoperate. Heterogeneous networks are made possible by standards-conforming hardware and software interfaces used in common by different products, thus allowing them to communicate with each other. The Internet itself is an example of a heterogeneous network.

HiRDB — Hitachi Relational Database

HIS — High Speed Interconnect

HiStar — Multiple point-to-point data paths to cache

Hi Track System — Automatic fault reporting system.

HIHSM — Hitachi Internal Hierarchy Storage Management

HMDE — Hitachi Multiplatform Data Exchange

HMRC F — Hitachi Multiple Raid Coupling Feature

HMRS — Hitachi Multiplatform Resource Sharing

HODM — Hitachi Online Data Migration

Homogeneous — Of the same or similar kind

HOMRCF — Hitachi Open Multiple Raid Coupling Feature; Shadow Image, marketing name for HOMRCF

HORC — Hitachi Open Remote Copy — See TrueCopy

HORCM — Hitachi Open Raid Configuration Manager

Host — Also called a server. A Host is basically a central computer that processes end-user applications or requests.

Host LU — See HLU

Host Storage Domains—Allows host pooling at the LUN level and the priority access feature lets administrator set service levels for applications

HP — Hewlett-Packard Company

HPC — High Performance Computing

HRC — Hitachi Remote Copy — See TrueCopy

HSG — Host Security Group

HSM — Hierarchical Storage Management

HSSDC — High Speed Serial Data Connector

HTTP — Hyper Text Transfer Protocol

HTTPS — Hyper Text Transfer Protocol Secure

Hub — A common connection point for devices in a network. Hubs are commonly used to connect segments of a LAN. A hub contains multiple ports. When a packet arrives at one port, it is copied to the other ports so that all segments of the LAN can see all packets. A switching hub actually reads the destination address of each packet and then forwards the packet to the correct port.

HXRC — Hitachi Extended Remote Copy

Hub — Device to which nodes on a multi-point bus or loop are physically connected

-back to top-

—I—

IBR — Incremental Block-level Replication

IBR —Intelligent Block Replication

ID — Identifier

IDR — Incremental Data Replication

iFCP — Short for the *Internet* ***F****ibre* ***C****hannel* ***P****rotocol*, iFCP allows an organization to extend Fibre Channel storage networks over the Internet by using TCP/IP. TCP is responsible for managing congestion control as well as error detection and recovery services. iFCP allows an organization to create an IP SAN fabric that minimizes the Fibre Channel fabric component and maximizes use of the company's TCP/IP infrastructure.

In-band virtualization — Refers to the location of the storage network path, between the application host servers in the storage systems. Provides both control and data along the same connection path. Also called symmetric virtualization.

Interface —The physical and logical arrangement supporting the attachment of any device to a connector or to another device.

Internal bus — Another name for an internal data bus. Also, an expansion bus is often referred to as an internal bus.

Internal data bus — A bus that operates only within the internal circuitry of the CPU, communicating among the internal caches of memory that are part of the CPU chip's design. This bus is typically rather quick and is independent of the rest of the computer's operations.

IID — Stands for Initiator ID. This is used to identify LU whether it is NAS System LU or User LU. If it is 0, that means NAS System LU and if it is 1, then the LU is User LU.

IIS — Internet Information Server

I/O — Input/Output — The term I/O (pronounced "eye-oh") is used to describe any program, operation or device that transfers data to or from a computer and to or from a peripheral device.

IML — Initial Microprogram Load

IP — Internet Protocol

IPL — Initial Program Load

IPSEC — IP security

iSCSI (Internet SCSI) — Pronounced eye skuzzy. Short for Internet SCSI, an IP-based standard for linking data storage devices over a network and transferring data by carrying SCSI commands over IP networks. iSCSI supports a Gigabit Ethernet interface at the physical layer, which allows systems supporting iSCSI interfaces to connect directly to standard Gigabit Ethernet switches and/or IP routers. When an operating system receives a request it generates the SCSI command and then sends an IP packet over an Ethernet connection. At the receiving end, the SCSI commands are separated from the request, and the SCSI commands and data are sent to the SCSI controller and then to the SCSI storage device. iSCSI will also return a response to the request using the same protocol. iSCSI is important to SAN technology because it enables a SAN to be deployed in a LAN, WAN or MAN.

iSER — iSCSI Extensions for RDMA

ISL — Inter-Switch Link

iSNS — Internet Storage Name Service

ISPF — Interactive System Productivity Facility

ISC — Initial shipping condition

ISOE — iSCSI Offload Engine

ISP — Internet service provider

-back to top-

—J—

Java (and Java applications). — Java is a widely accepted, open systems programming language. Hitachi's enterprise software products are all accessed using Java applications. This enables storage administrators to access the Hitachi enterprise software products from any PC or workstation that runs a supported thin-client internet browser application and that has TCP/IP network access to the computer on which the software product runs.

Java VM — Java Virtual Machine

JCL — Job Control Language

JBOD — Just a Bunch of Disks

JRE —Java Runtime Environment

JMP —Jumper. Option setting method

-back to top-

—K—

kVA— Kilovolt Ampere

kW — Kilowatt

-back to top-

—L—

LACP — Link Aggregation Control Protocol

LAG — Link Aggregation Groups

LAN— Local Area Network

LBA (logical block address) — A 28-bit value that maps to a specific cylinder-head-sector address on the disk.

LC (Lucent connector) — Fibre Channel connector that is smaller than a simplex connector (SC)

LCDG—Link Processor Control Diagnostics

LCM— Link Control Module

LCP (Link Control Processor) — Controls the optical links. LCP is located in the LCM.

LCU — Logical Control Unit

LD — Logical Device

LDAP — Lightweight Directory Access Protocol

LDEV (Logical Device) — A set of physical disk partitions (all or portions of one or more disks) that are combined so that the subsystem sees and treats them as a single area of data storage; also called a volume. An LDEV has a specific and unique address within a subsystem. LDEVs become LUNs to an open-systems host.

LDKC — Logical Disk Controller
Manual.

LDM — Logical Disk Manager

LED — Light Emitting Diode

LM — Local Memory

LMODs — Load Modules

LNKLST — Link List

Load balancing — Distributing processing and communications activity evenly across a computer network so that no single device is overwhelmed. Load balancing is especially important for networks where it's difficult to predict the number of requests that will be issued to a server. If one server starts to be swamped, requests are forwarded to another server with more capacity. Load balancing can also refer to the communications channels themselves.

LOC — Locations section of the Maintenance

Logical DKC (LDKC) — An internal architecture extension to the Control Unit addressing scheme that allows more LDEVs to be identified within one Hitachi enterprise storage system. The LDKC is supported only on Universal Storage Platform V/VM class storage systems. As of March 2008, only one LDKC is supported, LDKC 00. Refer to product documentation as Hitachi has announced their intent to expand this capacity in the future.

LPAR — Logical Partition

LRU — Least Recently Used

LU — Logical Unit; Mapping number of an LDEV

LUN (Logical Unit Number) — One or more LDEVs. Used only for open systems. LVI (logical volume image) identifies a similar concept in the mainframe environment.

LUN Manager — HDS software used to map Logical Units (LUNs) to subsystem ports.

LUSE (Logical Unit Size Expansion) — Feature used to create virtual LUs that are up to 36 times larger than the standard OPEN-x LUs.

LVDS — Low Voltage Differential Signal

LVM — Logical Volume Manager

-back to top-

—M—

MAC — Media Access Control (MAC address = a unique identifier attached to most forms of networking equipment.

MIB — Management information base

MMC — Microsoft Management Console

MPIO — multipath I/O

Mapping — Conversion between two data addressing spaces. For example, mapping refers to the conversion between physical disk block addresses and the block addresses of the virtual disks presented to operating environments by control software.

Mb — Megabits

MB — Megabytes

MBUS — Multi-CPU Bus

MC — Multi Cabinet

MCU — Main Disk Control Unit; the local CU of a remote copy pair.

Metadata — In database management systems, data files are the files that store the database information, whereas other files, such as index files and data dictionaries, store administrative information, known as metadata.

MFC — Main Failure Code

MIB — Management Information Base, a database of objects that can be monitored by a network management system. Both SNMP and RMON use standardized MIB formats that allow any SNMP and RMON tools to monitor any device defined by a MIB.

Microcode — The lowest-level instructions that directly control a microprocessor. A single machine-language instruction typically translates into several microcode instructions.

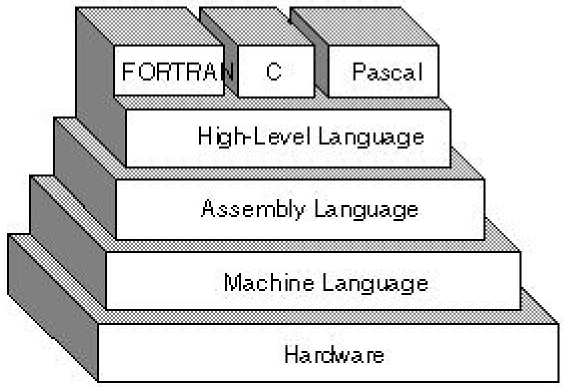

Microprogram — See Microcode

Mirror Cache OFF — Increases cache efficiency over cache data redundancy.

MM — Maintenance manual.

MPA — Micro-processor adapter

MP — Microprocessor

MPU— Microprocessor Unit

Mode— The state or setting of a program or device. The term mode implies a choice -- that you can change the setting and put the system in a different mode.

MSCS — Microsoft Cluster Server

MS/SG — Microsoft Service Guard

MTS — Multi-Tiered Storage

MVS — Multiple Virtual Storage

-back to top-

—N—

NAS (Network Attached Storage) — A disk array connected to a controller that gives access to a LAN Transport. It handles data at the file level.

NAT — Network Address Translation

NAT — Network Address Translation

NDMP — Network Data Management Protocol, is a protocol meant to transport data between NAS devices

NetBIOS — Network Basic Input/Output System

Network — A computer system that allows sharing of resources, such as files and peripheral hardware devices

NFS protocol — Network File System is a protocol which allows a computer to access files over a network as easily as if they were on its local disks.

NIM — Network Interface Module

NIS — Network Information Service (YP)

Node — An addressable entity connected to an I/O bus or network. Used primarily to refer to computers, storage devices, and storage subsystems. The component of a node that connects to the bus or network is a port.

Node name — A Name_Identifier associated with a node.

NTP — Network Time Protocol

NVS — Non Volatile Storage

-back to top-

—O—

OEM — Original Equipment Manufacturer

OFC — Open Fibre Control

OID — Object identifier

OLTP — On-Line Transaction Processing

ONODE — Object node

OPEX – Operational Expenditure – An operating expense, operating expenditure, operational expense, operational expenditure or OPEX is an on-going cost for running a product, business, or system. Its counterpart is a capital expenditure (CAPEX).

Out-of-band virtualization — Refers to systems where the controller is located outside of the SAN data path. Separates control and data on different connection paths. Also called asymmetric virtualization.

ORM— Online Read Margin

OS — Operating System

-back to top-

—P—

Parity — A technique of checking whether data has been lost or written over when it's moved from one place in storage to another or when it's transmitted between computers

Parity Group — Also called an array group, is a group of hard disk drives (HDDs) that form the basic unit of storage in a subsystem. All HDDs in a parity group must have the same physical capacity.

Partitioned cache memory — Separate workloads in a 'storage consolidated' system by dividing cache into individually managed multiple partitions. Then customize the partition to match the I/O characteristics of assigned LUs

PAT — Port Address Translation

PATA — Parallel ATA

Path — Also referred to as a transmission channel, the path between two nodes of a network that a data communication follows. The term can refer to the physical cabling that connects the nodes on a network, the signal that is communicated over the pathway or a sub-channel in a carrier frequency.

Path failover — See Failover

PAV — Parallel Access Volumes

PAWS — Protect Against Wrapped Sequences

PBC — Port By-pass Circuit

PCB — Printed Circuit Board

PCI — Power Control Interface

PCI CON (Power Control Interface Connector Board)

Performance — speed of access or the delivery of information

PD — Product Detail

PDEV— Physical Device

PDM — Primary Data Migrator

PDM — Policy based Data Migration

PGR — Persistent Group Reserve

PK — Package (see PCB)

PI — Product Interval

PIR — Performance Information Report

PiT — Point-in-Time

PL — Platter (Motherboard/Backplane) - the circular disk on which the magnetic data is stored.

Port — In TCP/IP and UDP networks, an endpoint to a logical connection. The port number identifies what type of port it is. For example, port 80 is used for HTTP traffic.

P-P — Point to Point; also P2P

Priority Mode— Also PRIO mode, is one of the modes of FlashAccess™ in which the FlashAccess™ extents hold read and write data for specific extents on volumes (see Bind Mode).

Provisioning — The process of allocating storage resources and assigning storage capacity for an application, usually in the form of server disk drive space, in order to optimize the performance of a storage area network (SAN). Traditionally, this has been done by the SAN administrator, and it can be a tedious process.

In recent years, automated storage provisioning, also called auto-provisioning, programs have become available. These programs can reduce the time required for the storage provisioning process, and can free the administrator from the often distasteful task of performing this chore manually

Protocol — A convention or standard that enables the communication between two computing endpoints. In its simplest form, a protocol can be defined as the rules governing the syntax, semantics, and synchronization of communication. Protocols may be implemented by hardware, software, or a combination of the two. At the lowest level, a protocol defines the behavior of a hardware connection.

PS — Power Supply

PSA — Partition Storage Administrator

PSSC — Perl SiliconServer Control

PSU — Power Supply Unit

PTR — Pointer

P-VOL — Primary Volume

-back to top-

—Q—

QD — Quorum Device

QoS — Quality of Service —In the field of computer networking, the traffic engineering term quality of service (QoS), refers to resource reservation control mechanisms rather than the achieved service quality. Quality of service is the ability to provide different priority to different applications, users, or data flows, or to guarantee a certain level of performance to a data flow.

-back to top-

—R—

R/W — Read/Write

RAID (Redundant Array of Independent Disks, or Redundant Array of Inexpensive Disks) — A group of disks that look like a single volume to the server. RAID improves performance by pulling a single stripe of data from multiple disks, and improves fault-tolerance either through mirroring or parity checking and it is a component of a customer's SLA.

RAID-0 — Striped array with no parity

RAID-1 — Mirrored array & duplexing

RAID-3 — Striped array with typically non-rotating parity, optimized for long, single-threaded transfers

RAID-4 — Striped array with typically non-rotating parity, optimized for short, multi-threaded transfers

RAID-5 — Striped array with typically rotating parity, optimized for short, multithreaded transfers

RAID-6 — Similar to RAID-5, but with dual rotating parity physical disks, tolerating two physical disk failures

RAM — Random Access Memory

RAM DISK — A LUN held entirely in the cache area.

Read/Write Head — Read and write data to the platters, typically there is one head per platter side, and each head is attached to a single actuator shaft

Redundant — Describes computer or network system components, such as fans, hard disk drives, servers, operating systems, switches, and telecommunication links that are installed to back up primary resources in case they fail. A well-known example of a redundant system is the redundant array of independent disks (RAID). Redundancy contributes to the *fault tolerance* of a system.

Reliability —level of assurance that data will not be lost or degraded over time

Resource Manager — Hitachi Resource Manager™ utility package is a software suite that rolls into one package the following four pieces of software:

- Hitachi Graph-Track™ performance monitor feature
- Virtual Logical Volume Image (VLMI) Manager (optimizes capacity utilization),
- Hitachi Cache Residency Manager feature (formerly FlashAccess) (uses cache to speed data reads and writes),
- LUN Manager (reconfiguration of LUNS, or logical unit numbers).

RCHA — RAID Channel Adapter

RC — Reference Code or Remote Control

RCP — Remote Control Processor

RCU — Remote Disk Control Unit

RDMA — Remote Direct Memory Access

Redundancy — Backing up a component to help ensure high availability.

Reliability — An attribute of any commuter component (software, hardware, or a network) that consistently performs according to its specifications.

RID — **R**elative **Id**entifier that uniquely identifies a user or group within a Microsoft Windows domain

RISC — Reduced Instruction Set Computer

RK (Rack) — Acronym referring to the main "Rack" unit, which houses the core operational hardware components of the Thunder 9500V/9200 subsystem. (See also: RKA, H1F, and H2F)

RKA (Rack Additional) — Acronym referring to "Rack Additional", namely additional rack unit(s) which house additional hard drives exceeding the capacity of the core RK unit of the Thunder 9500V/9200 subsystem. (See also: RK, RKA, H1F, and H2F).

RKAJAT — Rack Additional SATA disk tray

RLGFAN — Rear Logic Box Fan Assembly

RLOGIC BOX — Rear Logic Box

RMI (Remote Method Invocation) — A way that a programmer, using the Java programming language and development environment, can write object-oriented programming in which objects on different computers can interact in a distributed network. RMI is the Java version of what is generally known as a RPC (remote procedure call), but with the ability to pass one or more objects along with the request.

RoHS — Restriction of Hazardous Substances (in Electrical and Electronic Equipment)

ROI — Return on Investment

ROM — Read-only memory

Round robin mode — A load balancing technique in which balances power is placed in the DNS server instead of a strictly dedicated machine as other load techniques do. Round robin works on a rotating basis in that one server IP address is handed out, then moves to the back of the list; the next server IP address is handed out, and then it moves to the end of the list; and so on, depending on the number of servers being used. This works in a looping fashion. Round robin DNS is usually used for balancing the load of geographically distributed Web servers.

Router — a computer networking device that forwards data packets toward their destinations, through a process known as routing.

RPO (Recovery Point Option) — point in time that recovered data should match.

RPSFAN — Rear Power Supply Fan Assembly

RS CON — RS232C/RS422 Interface Connector

RSD — Raid Storage Division

R-SIM—Remote Service Information Message

RTO (Recovery Time Option) — length of time that can be tolerated between a disaster and recovery of data.

-back to top-

—S—

SA — Storage Administrator

SAA — Share Access Authentication - the process of restricting a user's rights to a file system object by combining the security descriptors from both the file system object itself and the share to which the user is connected

SACK — **S**equential **Ack**nowledge

SACL — System ACL - the part of a security descriptor that stores system auditing information

SAN (Storage Area Network) — A network linking computing devices to disk or tape arrays and other devices over Fibre Channel. It handles data at the block level.

SANtinel — HDS software that provides LUN security. SANtinel protects data from unauthorized access in SAN environments. It restricts server access by implementing boundaries around predefined zones and is used to map hosts in a host group to the appropriate LUNs.

SARD — System Assurance Registration Document

SAS — SAN Attached Storage, storage elements that connect directly to a storage area network and provide data access services to computer systems.

SAS — (Serial Attached SCSI) disk drive configurations for Hitachi Simple Modular Storage 100 systems

SATA — (Serial ATA) —Serial Advanced Technology Attachment is a new standard for connecting hard drives into computer systems. SATA is based on serial signaling technology, unlike current IDE (Integrated Drive Electronics) hard drives that use parallel signaling.

SC (simplex connector) — Fibre Channel connector that is larger than a Lucent connector (LC).

SC — Single Cabinet

SCM — Supply Chain Management

SCP — Secure Copy

SCSI — Small Computer Systems Interface. A parallel bus architecture and a protocol for transmitting large data blocks up to a distance of 15-25 meters.

Sector - a sub-division of a track of a magnetic disk that stores a fixed amount of data.

Selectable segment size — can be set per partition

Selectable Stripe Size — Increases performance by customizing the disk access size.

Serial Transmission — The transmission of data bits in sequential order over a single line.

Server — A central computer that processes end-user applications or requests, also called a host.

Service-level agreement (SLA**) -** A contract between a network service provider and a customer that specifies, usually in measurable terms, what services the network service provider will furnish. Many Internet service providers (ISP)s provide their customers with an SLA. More recently, IT departments in major enterprises have adopted the idea of writing a service level agreement so that services for their customers (users in other departments within the enterprise) can be measured, justified, and perhaps compared with those of outsourcing network providers.

Some metrics that SLAs may specify include:

- What percentage of the time services will be available
- The number of users that can be served simultaneously
- Specific performance benchmarks to which actual performance will be periodically compared
- The schedule for notification in advance of network changes that may affect users
- Help desk response time for various classes of problems
- Dial-in access availability

- Usage statistics that will be provided.

Service-Level Objective (SLO) - Individual performance metrics are called service-level objectives (SLOs). Although there is no hard and fast rule governing how many SLOs may be included in each SLA, it only makes sense to measure what matters.

Each SLO corresponds to a single performance characteristic relevant to the delivery of an overall service. Some examples of SLOs would include: system availability, help desk incident resolution time, and application response time.

SES — SCSI Enclosure Services

SENC — Is the SATA (Serial ATA) version of the ENC. ENCs and SENCs are complete microprocessor systems on their own and they occasionally require a firmware upgrade.

SFP — Small Form-Factor Pluggable module Host connector — A specification for a new generation of optical modular transceivers. The devices are designed for use with small form factor (SFF) connectors, and offer high speed and physical compactness. They are hot-swappable.

ShadowImage® — HDS software used to duplicate large amounts of data within a subsystem without affecting the service and performance levels or timing out. ShadowImage replicates data with high speed and reduces backup time.

SHSN — Shared memory Hierarchical Star Network

SI — Hitachi ShadowImage® Replication software

SIM RC — Service (or system) Information Message Reference Code

SID — Security Identifier - user or group identifier within the Microsoft Windows security model

SIMM — Single In-line Memory Module

SIM — Storage Interface Module

SIM — Service Information Message; a message reporting an error; contains fix guidance information

SIz — Hitachi ShadowImage® Replication Software

SLA —Service Level Agreement

SLPR (**S**torage administrator **L**ogical **P**a**R**tition) — Storage can be divided among various users to reduce conflicts with usage.

SM (Shared Memory Module) — Stores the shared information about the subsystem and the cache control information (director names). This type of information is used for the exclusive control of the subsystem. Like CACHE, shared memory is controlled as two areas of memory and fully non-volatile (sustained for approximately 7 days).

SM PATH (Shared Memory Access Path) — Access Path from the processors of CHA, DKA PCB to Shared Memory.

SMB/CIFS — Server Message Block Protocol / Common Internet File System

SMC — Shared Memory Control

SM — Shared Memory

SMI-S — Storage Management Initiative Specification

SMP/E (System Modification Program/Extended) — An IBM licensed program used to install software and software changes on z/OS systems.

SMS — Hitachi Simple Modular Storage

SMTP — Simple Mail Transfer Protocol

SMU — System Management Unit

Snapshot Image — A logical duplicated volume (V-VOL) of the primary volume. It is an internal volume intended for restoration

SNIA — Storage Networking Industry Association, an association of producers and consumers of storage networking products, whose goal is to further storage networking technology and applications.

SNMP (Simple Network Management Protocol) — A TCP/IP protocol that was designed for management of networks over TCP/IP, using agents and stations.

SOAP (simple object access protocol) — A way for a program running in one kind of operating system (such as Windows 2000) to communicate with a program in the same or another kind of an operating system (such as Linux) by using the World Wide Web's Hypertext Transfer Protocol (HTTP) and its Extensible Markup Language (XML) as the mechanisms for information exchange.

Socket — In UNIX and some other operating systems, a software object that connects an application to a network protocol. In UNIX, for example, a program can send and receive TCP/IP messages by opening a socket and reading and writing data to and from the socket. This simplifies program development because the programmer need only worry about manipulating the socket and can rely on the operating system to actually transport messages across the network correctly. Note that a socket in this sense is completely soft - it's a software object, not a physical component.

SPAN — Span is a section between two intermediate supports. See Storage pool

Spare — An object reserved for the purpose of substitution for a like object in case of that object's failure.

SPC — SCSI Protocol Controller

SpecSFS — Standard Performance Evaluation Corporation Shared File system

SSB — Sense Byte

SSC — SiliconServer Control

SSH — Secure Shell

SSID — Subsystem Identifier

SSL — Secure Sockets Layer

SSVP — Sub Service Processor; interfaces the SVP to the DKC

Sticky Bit — Extended Unix mode bit that prevents objects from being deleted from a directory by anyone other than the object's owner, the directory's owner or the root user

STR — Storage and Retrieval Systems

Storage pooling — The ability to consolidate and manage storage resources across storage system enclosures where the consolidation of many appears as a single view.

Striping — A RAID technique for writing a file to multiple disks on a block-by-block basis, with or without parity.

Subsystem — Hardware and/or software that performs a specific function within a larger system.

SVC — Supervisor Call Interruption

S-VOL — Secondary Volume

SVP (Service Processor) — A laptop computer mounted on the control frame (DKC) and used for monitoring, maintenance and administration of the subsystem

Symmetric virtualization — See In-band virtualization.

Synchronous— Operations which have a fixed time relationship to each other. Most commonly used to denote I/O operations which occur in time sequence, i.e., a successor operation does not occur until its predecessor is complete.

Switch— A fabric device providing full bandwidth per port and high-speed routing of data via link-level addressing.

Software — Switch

-back to top-

—T—

T.S.C. (Technical Support Center) — A chip developed by HP, and used in various devices. This chip has FC-0 through FC-2 on one chip.

TCA — TrueCopy Asynchronous

TCO — Total Cost of Ownership

TCP/IP — Transmission Control Protocol over Internet Protocol

TCP/UDP — User Datagram Protocol is one of the core protocols of the Internet protocol suite. Using UDP, programs on networked computers can send short messages known as datagrams to one another.

TCS — TrueCopy Synchronous

TCz — Hitachi TrueCopy® Remote Replication software

TDCONV (Trace Dump CONVerter) — Is a software program that is used to convert traces taken on the system into readable text. This information is loaded into a special spreadsheet that allows for further investigation of the data. More in-depth failure analysis.

TGTLIBs— Target Libraries

Target — The system component that receives a SCSI I/O command, an open device that operates at the request of the initiator

THF — Front Thermostat

Thin Provisioning — Thin Provisioning allows space to be easily allocated to servers, on a just-enough and just-in-time basis.

Throughput — The amount of data transferred from one place to another or processed in a specified amount of time. Data transfer rates for disk drives and networks are measured in terms of throughput. Typically, throughputs are measured in kbps, Mbps and Gbps.

THR — Rear Thermostat

TID — Target ID

Tiered storage —A storage strategy that matches data classification to storage metrics. Tiered storage is the assignment of different categories of data to different types of storage media in order to reduce total storage cost. Categories may be based on levels of protection needed, performance requirements, frequency of use, and other considerations. Since assigning data to particular media may be an ongoing and complex activity, some vendors provide software for automatically managing the process based on a company-defined policy.

Tiered Storage Promotion — Moving data between tiers of storage as their availability requirements change

TISC — The Hitachi Data Systems internal Technical Information Service Centre from which microcode, user guides, ECNs, etc. can be downloaded.

TLS — Tape Library System

TLS — Transport Layer Security

TMP — Temporary

TOC — Table Of Contents

TOD — Time Of Day

TOE — TCP Offload Engine

Topology — The shape of a network or how it is laid out. Topologies are either physical or logical.

TPF — Transaction Processing Facility

Transfer Rate — See Data Transfer Rate

Track — Circular segment of a hard disk or other storage media

Trap — A program interrupt, usually an interrupt caused by some exceptional situation in the user program. In most cases, the Operating System performs some action, and then returns control to the program.

TRC — Technical Resource Center

TrueCopy — HDS software that replicates data between subsystems. These systems can be located within a data center or at geographically separated data centers. The 9900V adds the capability of using TrueCopy to make copies in two different locations simultaneously.

TSC — Technical Support Center

TSO/E — Time Sharing Option/Extended

-back to top-

—U—

UFA — UNIX File Attributes

UID — **U**ser **Id**entifier

UID — **U**ser **Id**entifier within the UNIX security model

UPS — Uninterruptible Power Supply — A power supply that includes a battery to maintain power in the event of a power outage.

URz — Hitachi Universal Replicator software

USP — Universal Storage Platform™

USP V — Universal Storage Platform™ V

USP VM — Universal Storage Platform™ VM

-back to top-

—V—

VCS — Veritas Cluster System

VHDL — VHSIC (Very-High-Speed Integrated Circuit) Hardware Description Language

VHSIC — Very-High-Speed Integrated Circuit

ViVol — Virtual Volume (used in the High-performance NAS Platform)

VI — Virtual Interface, a research prototype that is undergoing active development, and the details of the implementation may change considerably. It is an application interface that gives user-level processes direct but protected access to network interface cards. This allows applications to bypass IP processing overheads (copying data, computing checksums, etc.) and system call overheads while still preventing one process from accidentally or maliciously tampering with or reading data being used by another.

VirtLUN —VLL. Customized volume; size chosen by user

Virtualization —The amalgamation of multiple network storage devices into what appears to be a **single** storage unit. Storage virtualization is often used in a SAN, and makes tasks such as archiving, back up, and recovery easier and faster. Storage virtualization is usually implemented via software applications.

VLL — Virtual Logical Volume Image/Logical Unit Number

VLVI — Virtual Logic Volume Image, marketing name for CVS (custom volume size)

VOLID — Volume ID

Volume — A fixed amount of storage on a disk or tape. The term volume is often used as a synonym for the storage medium itself, but it is possible for a single disk to contain more than one volume or for a volume to span more than one disk.

VTOC — Volume Table of Contents

V-VOL — Virtual volume

-back to top-

—W—

WAN —Wide Area Network

WDIR — Working Directory

WDIR — Directory Name Object

WDS — Working Data Set

WFILE — Working File

WFILE — File Object

WFS — Working File Set

WINS — Windows Internet Naming Service

WMS — Hitachi Workgroup Modular Storage system

WTREE — Working Tree

WTREE — Directory Tree Object

WWN (World Wide Name) — A unique identifier for an open-system 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 Hitachi Volume Security software (formerly SANtinel) parameters because it determines whether the open-system host is to be allowed or denied access to a specified LU or a group of LUs.

WWN — World Wide Name — 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.

WWNN — World Wide Node Name — A globally unique 64-bit identifier assigned to each Fibre Channel node process.

WWPN (World Wide Port Name) — A globally unique 64-bit identifier assigned to each Fibre Channel port. Fibre Channel ports' WWPN are permitted to use any of several naming authorities. Fibre Channel specifies a Network Address Authority (NAA) to distinguish between the various name registration authorities that may be used to identify the WWPN.

-back to top-

—X—

XAUI — "X"=10, AUI = Attachment Unit Interface

XFI — Standard interface for connecting 10 Gig Ethernet MAC device to XFP interface

XFP — "X" = 10 Gigabit Small Form Factor Pluggable

XRC — Extended Remote Copy

-back to top-

—Y—

-back to top-

—Z—

Zone — A collection of Fibre Channel Ports that are permitted to communicate with each other via the fabric

Zoning — A method of subdividing a storage area network into disjoint zones, or subsets of nodes on the network. Storage area network nodes outside a zone are invisible to nodes within the zone. Moreover, with switched SANs, traffic within each zone may be physically isolated from traffic outside the zone.

-back to top-

Evaluating this Course

1. Log in to the Hitachi Data Systems Learning Center page at https://learningcenter.hds.com
2. Select the **Learning** tab on the upper-left corner of the Hitachi Data Systems Learning Center page.
3. On the left panel of the **Learning** page, click **Learning History**. The **Learning History** page appears.
4. From the Title column of the **Learning History** table, select the title of the course in which you have enrolled. The **Learning Details** page for the enrolled course appears.
5. Select the **More Details** tab.

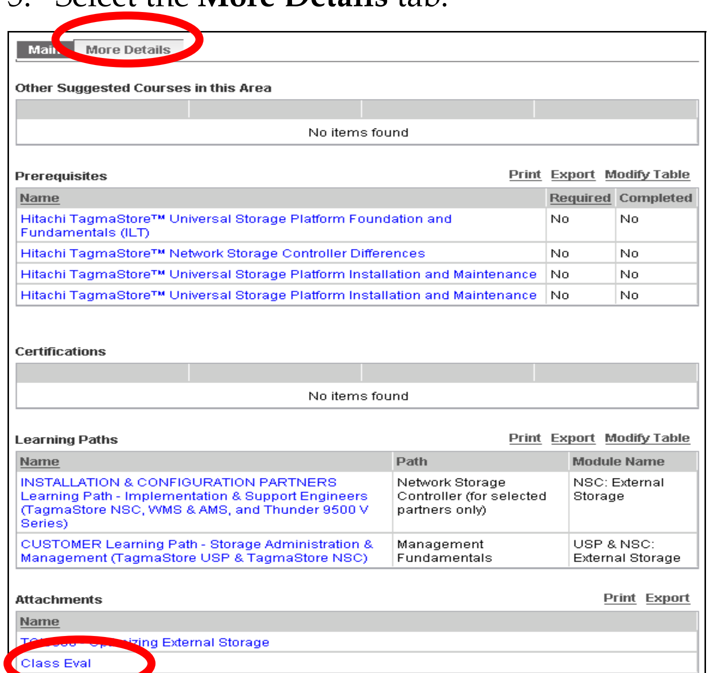

Main | More Details

Other Suggested Courses in this Area

No items found			

Prerequisites — Print Export Modify Table

Name	Required	Completed
Hitachi TagmaStore™ Universal Storage Platform Foundation and Fundamentals (ILT)	No	No
Hitachi TagmaStore™ Network Storage Controller Differences	No	No
Hitachi TagmaStore™ Universal Storage Platform Installation and Maintenance	No	No
Hitachi TagmaStore™ Universal Storage Platform Installation and Maintenance	No	No

Certifications

No items found			

Learning Paths — Print Export Modify Table

Name	Path	Module Name
INSTALLATION & CONFIGURATION PARTNERS Learning Path - Implementation & Support Engineers (TagmaStore NSC, WMS & AMS, and Thunder 9500 V Series)	Network Storage Controller (for selected partners only)	NSC: External Storage
CUSTOMER Learning Path - Storage Administration & Management (TagmaStore USP & TagmaStore NSC)	Management Fundamentals	USP & NSC: External Storage

Attachments — Print Export

Name
T[illegible] - Op[illegible]izing External Storage
Class Eval

6. Under Attachments, click the **Class Eval** link. The **Class Evaluation** form opens.
7. Complete the form and submit.