Hitachi Data Systems



Hitachi Data Systems Foundations — Modular

THI0515

Courseware Version 5.1

Notice: This document is for informational purposes only, and does not set forth any warranty, express or implied, concerning any equipment or service offered or to be offered by Hitachi Data Systems. This document describes some capabilities that are conditioned on a maintenance contract with Hitachi Data Systems being in effect, and that may be configuration-dependent, and features that may not be currently available. Contact your local Hitachi Data Systems sales office for information on feature and product availability.

Hitachi Data Systems sells and licenses its products subject to certain terms and conditions, including limited warranties. To see a copy of these terms and conditions prior to purchase or license, please call your local sales representative to obtain a printed copy. If you purchase or license the product, you are deemed to have accepted these terms and conditions.

THE INFORMATION CONTAINED IN THIS MANUAL IS DISTRIBUTED ON AN "AS IS" BASIS WITHOUT WARRANTY OF ANY KIND, INCLUDING WITHOUT LIMITATION, ANY IMPLIED WARRANTIES OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE OR NONINFRINGEMENT. IN NO EVENT WILL HDS BE LIABLE TO THE END USER OR ANY THIRD PARTY FOR ANY LOSS OR DAMAGE, DIRECT OR INDIRECT, FROM THE USE OF THIS MANUAL, INCLUDING, WITHOUT LIMITATION, LOST PROFITS, BUSINESS INTERRUPTION, GOODWILL OR LOST DATA, EVEN IF HDS EXPRESSLY ADVISED OF SUCH LOSS OR DAMAGE.

Hitachi Data Systems is registered with the U.S. Patent and Trademark Office as a trademark and service mark of Hitachi, Ltd. The Hitachi Data Systems logotype is a trademark and service mark of Hitachi, Ltd.

The following terms are trademarks or service marks of Hitachi Data Systems Corporation in the United States and/or other countries:

Hitachi Data Systems Registered Trademarks

Essential NAS Platform Hi-Track ShadowImage TrueCopy Universal Storage Platform

Hitachi Data Systems Trademarks

HiCard HiPass Hi-PER Architecture Hi-Star
Universal Star Network

All other trademarks, trade names, and service marks used herein are the rightful property of their respective owners.

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.

©2010, Hitachi Data Systems Corporation. All Rights Reserved

HDS Academy 1050

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

Modular Storage Systems

- Hitachi Adaptable Modular Storage
- Hitachi Workgroup Modular Storage
- Hitachi Simple Modular Storage
- Hitachi Adaptable Modular Storage 2000 Family

Legacy Products:

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

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

Replication Software

Remote Replication:

- Hitachi Universal Replicator software
- Hitachi TrueCopy® Heterogeneous Remote Replication software bundle
- Hitachi TrueCopy® Remote Replication software bundle (for modular systems)

- Hitachi TrueCopy® Synchronous software
- Hitachi TrueCopy® Asynchronous software
- Hitachi TrueCopy® Extended Distance software

In-System Replication:

- 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 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 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

Other Solutions and Terms

- Hitachi Content Archive Platform
- Hitachi Essential NAS Platform®
- Hitachi High-performance NAS Platform, powered by BlueArc®
- Hi-Star™ crossbar switch architecture
- Hitachi Universal Star Network™ V

Contents

INTRODUCTION XIII
Welcome and Introductions xiii
Course Description xiii
Course Objectives xiv
Course Topics xv
Learning Paths xvi
HDS Academy Is on Twitter and LinkedIn xvii

1. SERVICES ORIENTED STORAGE SOLUTIONS FROM HITACHI DATA SYSTEMS 1-1
Module Objectives 1-1
Customers Speak 1-2
CIO's Perspective 1-3
Business Purpose 1-4
One Platform for All Data 1-6
Virtualization: The Key Enabler 1-7
Architecture Summary 1-8
Solutions Focus 1-9

2. PRODUCT OVERVIEW 2-1
Module Objectives 2-1
Module Topics 2-2
Product Description 2-3
Product Line Positioning 2-4
Features 2-5
Specifications 2-6
Specifications for Regular Expansion Unit 2-7
Specifications for Dense Expansion Unit 2-8
Software and Firmware Offerings 2-9
External Design and Connections 2-10
Certified Platforms 2-11
Highlights 2-12

3. HARDWARE COMPONENTS AND ARCHITECTURE 3-1
Module Objectives 3-1
Architecture — Model 2100 Controller 3-2
Architecture — Model 2300 Controller 3-3
Architecture — Model 2500 Controller 3-4
Hardware Components 3-5
Controller Labeling 3-7
Back End SAS Connections 3-8
Power Supply 3-10
Uninterruptible Power Supply Connections 3-11
Rotary Switch Power Connection Mode Settings 3-12
Model 2500 – FC/iSCSI Host Interface Intermix 3-13
Specifications 3-14
Controller Architecture 3-15
Cache Architecture 3-16
Back End Architecture 3-17
Disk Expansion Tray 3-20
Back End Logical Configuration 3-21
System Management Options Interface 3-22
Management Tool 3-23
System and Data Security 3-24

LUN Formatting....3-25
Supported RAID Levels....3-26
Relative RAID Performance....3-31
Supported RAID Levels....3-32
Sparing Out....3-33
Correction Copy....3-34
Dynamic Sparing....3-35
Copy Back....3-36
No Copy Back....3-37
Sparing....3-38
Hi-Track Monitor....3-39

4. ACTIVE-ACTIVE I/O ARCHITECTURE....4-1
Module Objectives....4-1
Overview....4-2
Internal Transaction....4-3
LU Ownership....4-4
Controller Load Balancing....4-6
Performance Design....4-7
Microcode Updates....4-10
Dual Core Support....4-11
Active-Active Summary....4-13

5. SOFTWARE FEATURE OVERVIEW....5-1
Module Objectives....5-1
Cache Memory Configuration....5-2
Basic Operation of Cache Control....5-3
Cache Partition Manager Feature....5-4
Advantage of Selectable Segment Size – Small I/O....5-5
Advantage of Selectable Segment Size – Large I/O....5-6
Advantage of Global Cache....5-7
Disadvantage of Global Cache....5-8
Advantage of Partitioned Cache....5-9
Advantage of Selectable Stripe Size....5-10
Partitioning Cache....5-11
Cache Partition Manager....5-12
Cache Residency Manager Overview....5-13
Installing Cache Residency Manager Feature....5-14
Functionality....5-15
Conditions that Terminate Cache Residency Manager....5-16
Maximum LUN Size for Cache Residency Manager....5-17
Performance Monitor Feature Overview....5-18
Monitored Statistics....5-19
Performance Monitor Window....5-21
Overview of Modular Volume Migration....5-22
Migration from SAS Drives to SATA Drives....5-23
Migrating Volumes for Performance....5-24
Performance Considerations....5-25
Change LU Capacity — LU Expansion....5-26
Overview of Adding LUs together....5-27
General Conditions and Restrictions....5-28
Change LU Capacity — Grow LU Capacity from Free Space....5-29
Change LU Capacity: LU Shrink....5-30
Manual LU Creation Out of Free Space....5-31
General Specifications....5-32
Online RAID Group expansion....5-33
Example....5-34

Priority Mode 5-36
Remove Expansion 5-37
Hitachi Dynamic Provisioning on AMS2000 5-38
What is Dynamic Provisioning? 5-39
Benefit 1: Easier and Faster Volume Configuration 5-42
Benefit 2: Reducing Volume Initial Cost and Increasing Storage Utilization 5-44
Benefit 3: Simplifying Performance Optimization 5-45
Summary of Hitachi Dynamic Provisioning Benefits 5-46

6. SETUP AND CONFIGURATION OF STORAGE NAVIGATOR MODULAR 2 6-1
Module Objectives 6-1
Features 6-2
Supported Operations 6-4
Architecture 6-5
Default Authentication 6-8
Configure 6-9
Start From Web Brower 6-10
GUI Main Window 6-11
RAID Group Creation 6-12
Logical Units Creation 6-13
LU Configuration 6-14
Host Connectivity Basics 6-15
Host Groups and LUN Mapping 6-16
Host Groups 6-17
Host Group Security (LUN Security) 6-18
Host Group - Options Settings 6-19
World Wide Names 6-20
Host HBA Persistent Binding 6-22
Recorded Product Demonstration 6-22

7. HITACHI DYNAMIC LINK MANAGER 7-1
Module Objectives 7-1
Overview 7-2
Benefits 7-3
Features 7-4
Path Failover and Failback 7-5
Load Balancing 7-6
Load Balancing Algorithms 7-7
Round Robin – Extended Round Robin 7-9
Path State Transitions 7-10
Graphical User Interface 7-11
Operations 7-13
Recorded Product Demonstration 7-14
Module Review 7-14

8. HITACHI REPLICATION SOFTWARE OVERVIEW 8-1
Module Objectives 8-1
Hitachi Replication Products 8-2
ShadowImage Replication Software 8-3
Copy-on-Write Snapshot Software 8-4
TrueCopy Remote Replication Software 8-5
TrueCopy Extended Distance Replication Software 8-6
Hitachi Replication Manager Software 8-7
Tools for Setting up Replication 8-8
Requirements for all replication products 8-9

Replication Operations....8-10
Copy Operations....8-11
Replication Operations....8-12

9. SHADOWIMAGE REPLICATION AND COPY-ON-WRITE SNAPSHOT SOFTWARE....9-1
Module Objectives....9-1
Overview....9-2
Applications....9-4
Benefits....9-5
General Replication Terminology....9-7
Management Resources....9-8
Rules....9-9
Internal ShadowImage Operation....9-10
Operations....9-11
paircreate Command....9-12
pairsplit Command....9-14
pairresync Command....9-16
pairresync Command....9-18
pairsplit -E Command....9-20
Resynchronize Suspended Pair....9-21
pairsplit -S Command....9-22
Volume Grouping....9-23
Pair Status Transitions....9-24
Copy-on-Write Snapshot Software....9-25
Overview....9-25
Components....9-28
Comparison with ShadowImage....9-29
Volume Size....9-30
Operation Scenarios....9-31
Status Transitions....9-36
Recorded Product Demonstration....9-37
Module Review....9-37

10. HITACHI TRUECOPY REMOTE REPLICATION SOFTWARE....10-1
Module Objectives....10-1
Description....10-2
Specifications....10-5
Operations....10-6
Preliminary Steps....10-7
Volume States....10-9
TrueCopy Extended Distance....10-10
Business Needs....10-11
Hardware and Software Components....10-12
Competitive Features....10-13
Write Order Guarantee....10-14
Consistency Groups....10-15
Pair Status Transition....10-16
Supported Configurations....10-17
Recorded Product Demonstration....10-19
Module Review....10-19

11. COMMAND CONTROL INTERFACE OVERVIEW....11-1
Module Objectives....11-1
Overview....11-2
Example — ShadowImage....11-4
Example — TrueCopy....11-6

HORCM Configuration Files 11-7
HORCM_DEV 11-9
Managing One Volume with horcm0.conf 11-10
Managing the Other Volume with horcm1.conf 11-12
Configuration 11-13
Commands 11-14
Module Review 11-15

12. HITACHI STORAGE COMMAND SUITE 12-1
Module Objectives 12-1
Common Software Management Framework 12-2
Vision 12-3
Components 12-4
Graphical User Interface 12-5

13. HITACHI DEVICE MANAGER SOFTWARE 13-1
Module Objectives 13-1
Overview 13-2
Problems Solved 13-5
Unified Interface 13-6
Storage Management 13-7
Architecture 13-8
Components 13-9
Views of the Storage Environment 13-10
Architecture Components 13-12
Operations 13-14
Report Operations 13-15
Command Line Interface 13-17
Storage Navigator Modular 2 versus Device Manager Summary 13-18
Replication Operations 13-19
Pair Volume Maintenance 13-21
Recorded Product Demonstration 13-22
Module Review 13-22

14. HITACHI GLOBAL LINK MANAGER SOFTWARE 14-1
Module Objectives 14-1
Multipathing – Problems and Solutions 14-2
Problems and Solutions 14-3
Configuration Overview 14-4
Features 14-5
Graphical User Interface 14-6
Software Path Management 14-7
SNMP-based Alert Management 14-8
Alert Management 14-9
Host Management 14-10
Recorded Product Demonstration 14-11
Module Review 14-11

15. HITACHI TUNING MANAGER SOFTWARE 15-1
Module Objectives 15-1
Overview 15-2
Product Positioning 15-3
The Challenge 15-5
Scope 15-6
High-Level Architecture 15-7
Components 15-8
Agents for Tuning Manager 15-13

Graphical User Interface 15-15
Module Review 15-16

16. HITACHI REPLICATION MANAGER SOFTWARE....16-1
Module Objectives 16-1
Overview 16-2
Features – Centralized Views 16-3
Features – Single Interface 16-4
Features – Pair Management Functions 16-5
System Configuration 16-6
Architecture in an Open System 16-7
Overview of Prerequisite Software 16-8
Managing Users and Permissions 16-9
Concept of Resource Groups 16-11
Resource Group View 16-12

17. HITACHI TIERED STORAGE MANAGER SOFTWARE 17-1
Module Objectives 17-1
Customer Storage Management Needs 17-2
Overview 17-3
Benefits of Tiered Storage Manager 17-4
Requirements 17-5
Software Options 17-6
Components 17-7
Volumes Used in Tiered Storage Manager 17-8
Concepts 17-9
Storage Domain 17-10
Migration Group 17-11
Storage Tier 17-12
Migration Task 17-13
Operations 17-14
Recorded Demonstration 17-15
Module Review 17-15

18. HITACHI STORAGE CAPACITY REPORTER SOFTWARE 18-1
Course Objectives 18-1
Hitachi Storage Viewer Architecture Platform 18-2
Features 18-3
Key Benefits 18-4
Complete End-To-End Storage Capacity View 18-5
Supported Storage Systems 18-6
Architecture 18-7
End-to-End Storage System Capacity 18-8
Storage System Capacity Planning 18-9
Features 18-10
Review Questions 18-11

19. HITACHI DATA PROTECTION SUITE 19-1
Module Objectives 19-1
What is Data Protection Suite? 19-2
Architecture 19-5
Common Technology Engine 19-6
Data Flow 19-8
Operational Concept – Data Migrator Software 19-10
Data Archiving 19-14
Module Review 19-16

20. HITACHI CONTENT PLATFORM 20-1
Module Objectives 20-1
Email Archiving Challenges 20-2
Document Archiving Storage Challenges 20-3
What an Active Archive Solution Must Deliver 20-4
Archive Object Replication 20-5
Archive Object Backup 20-6
Fixed Content 20-7
Retention Times 20-8
All Roads Lead To Hitachi Content Platform 20-10
Product Lines 20-11
Customer View of the Archive 20-12

21. HITACHI ESSENTIAL NAS PLATFORM OVERVIEW 21-1
Module Objectives 21-1
Hitachi Storage System Connectivity History 21-2
Essential NAS Platform Positioning 21-3
Product Overview – Speeds and Feeds 21-4
Hardware Comparison 21-5
Essential NAS Server 21-6
Platform Introduction 21-7
Data Protection Features 21-8
Rear and Front Views 21-9
Software Components 21-10
Data Storage Overview 21-12
Management 21-13

GLOSSARY

EVALUATING THIS COURSE

Introduction

Welcome and Introductions

- Introductions
  - Name
  - Position
  - Professional skills
  - Expectations from the course

Course Description

- This three-day course provides an overview of the Hitachi Data Systems modular storage products and technology. This includes
  - Modular storage hardware and software
  - Hitachi Storage Command Suite
  - Virtualization
  - Services Oriented Storage Solutions from Hitachi Data Systems
  - Storage Area Management
  - Performance
  - Configuration
  - Business continuity
- This course is part of the **Hitachi Certified Storage Professional 2010** Program and supports the Hitachi Data Systems Storage Foundations - Modular exam (HH0-120).

Course Objectives

- Upon completion of this course, the learner should be able to:
  - Describe the Hitachi Adaptable Modular Storage 2000 family products
  - Describe the essential components and software features within modular storage including models, numbering convention, capacities, types of devices, physical to logical maps and program products
  - Describe the storage performance and data protection strategies for Hitachi Data Systems
  - Describe the Hitachi Resource Manager tool set and what it is used for (Storage Navigator Modular program from volume and machine management)
  - Describe the features, functions, and principles for Hitachi TrueCopy® Remote Replication software bundle
  - Describe the features, functions, and principles for ShadowImage® Replication software
  - Describe the features, functions and principles for Hitachi Copy-on-Write Snapshot software
  - Describe how Hitachi Dynamic Link Manager software and Hitachi Global Link Manager software function

- Upon completion of this course, the learner should be able to (*continued*):
  - Describe the benefits of Services Oriented Storage Solutions
  - Describe the product solutions within Services Oriented Storage Solutions and how they can solve business issues
  - Given a scenario, describe how virtualization is enabled and used to solve business problems
  - Describe how Hitachi Virtual Partition Manager software enables the logical partitioning of Hitachi Adaptable Modular Storage systems
  - Describe how Hitachi Device Manager software functions
  - Describe how Hitachi Tuning Manager software functions
  - Describe how Hitachi Replication Manager software functions
  - Describe how Hitachi Tiered Storage Manager software functions
  - Describe Hitachi Data Systems strategy for content archival
  - Describe Hitachi Essential NAS Platform® benefits to the customer.
  - Describe Hitachi Content Archive Platform
  - Describe the Virtual Tape Library Solutions by Hitachi Data Systems

- The course prepares the learner for **Foundation Certification exam**.
- The course is **not** designed to:
  - Provide hands on labs
  - Provide in-depth technical knowledge on products covered
    - Focus on features, function and benefits
- Other Hitachi Data Systems Academy courses can provide in-depth knowledge and hands-on labs.

Course Topics

Content	Video Demonstrations
1. Services Oriented Storage Solutions from Hitachi Data Systems 2. Product Overview 3. Hardware Components and Architecture 4. Active-Active I/O Architecture 5. Software Feature Overview 6. Setup and Configuration of Storage Navigator Modular 2	Adaptable Modular Storage 2000 family Installation
7. Hitachi Dynamic Link Manager 8. Hitachi Replication Software Overview 9. ShadowImage Replication Software, Copy-on-Write Snapshot Software 10. TrueCopy Remote Replication Software 11. Command Control Interface Overview	Dynamic Link Manager Hitachi In-system Replication Software TrueCopy Extended Distance
12. Hitachi Storage Command Suite 13. Hitachi Device Manager software 14. Hitachi Global Link Manager software 15. Hitachi Tuning Manager software 16. Hitachi Replication Manager 17. Hitachi Tiered Storage Manager 18. Hitachi Storage Capacity Reporter 19. Hitachi Data Protection Suite 20. Hitachi Content Platform 21. Hitachi Essential NAS Platform Overview	Device Manager Software Global Link Manager Software Tuning Manager Software Tiered Storage Manager Software

Learning Paths

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



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.

HDS Academy Is on Twitter and LinkedIn



Follow the HDS Academy on Twitter for regular training updates.

LinkedIn is an online *community* which enables students and instructors to actively participate in online discussions related to Hitachi Data Systems products and training courses.

10

1. Services Oriented Storage Solutions from Hitachi Data Systems

Module Objectives

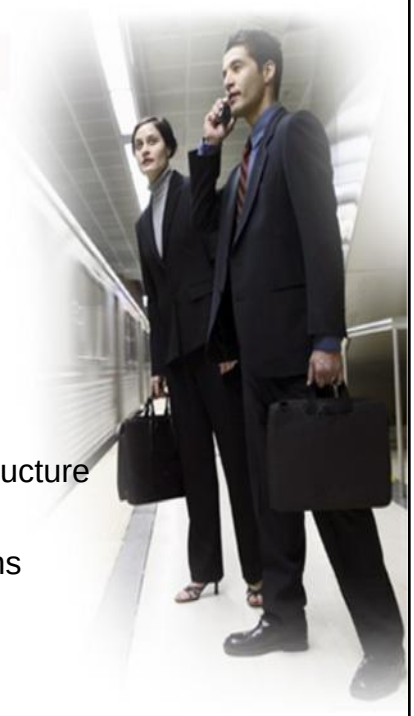
- Upon completion of this module, the learner should be able to:
  - Describe the purpose of Services Oriented Storage Solutions solutions from Hitachi Data Systems
  - Identify the business issues that drive Services Oriented Storage Solutions
  - Describe the purpose and benefits
  - Understand how Hitachi software and hardware products support Services Oriented Storage Solutions

Customers Speak

SNIA
Storage Networking Industry Association

End-users in a recent survey ranked their pain points in the following order, from highest to lowest priority:

- Cost price and total cost of ownership (TCO)
- The challenge of managing storage growth
- The inability to manage storage assets and infrastructure
- The lack of integrated and/or interoperable solutions
- Increasing complexity of storage infrastructure



3

HDS Confidential: For distribution only to authorized parties.

CIO's Perspective

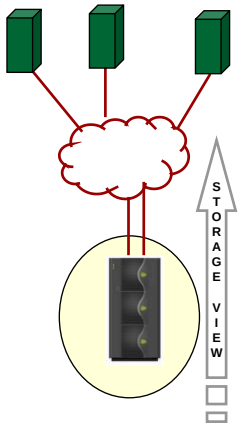
- The Requirement
  - Must align IT with business objectives
  - Must meet service level agreements/objectives
- The Problem
  - Accelerating storage growth
  - Too many islands of storage
  - Paying too much for storage
  - Low utilization of storage
  - Power and cooling challenges
  - Too much budget spent on maintaining old technology
- The Need
  - A services oriented storage approach that provides a common, flexible platform for aligning IT storage resources with business objectives



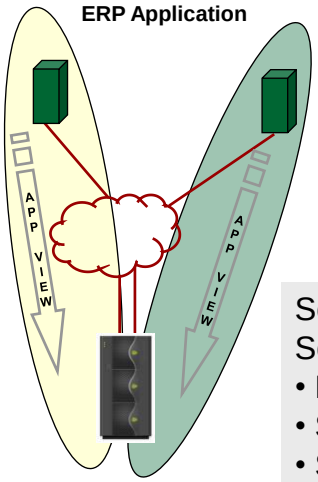
4

Business Purpose

- Enterprises need solutions that manage and optimize their IT infrastructure to meet business requirements – an application-centric approach.



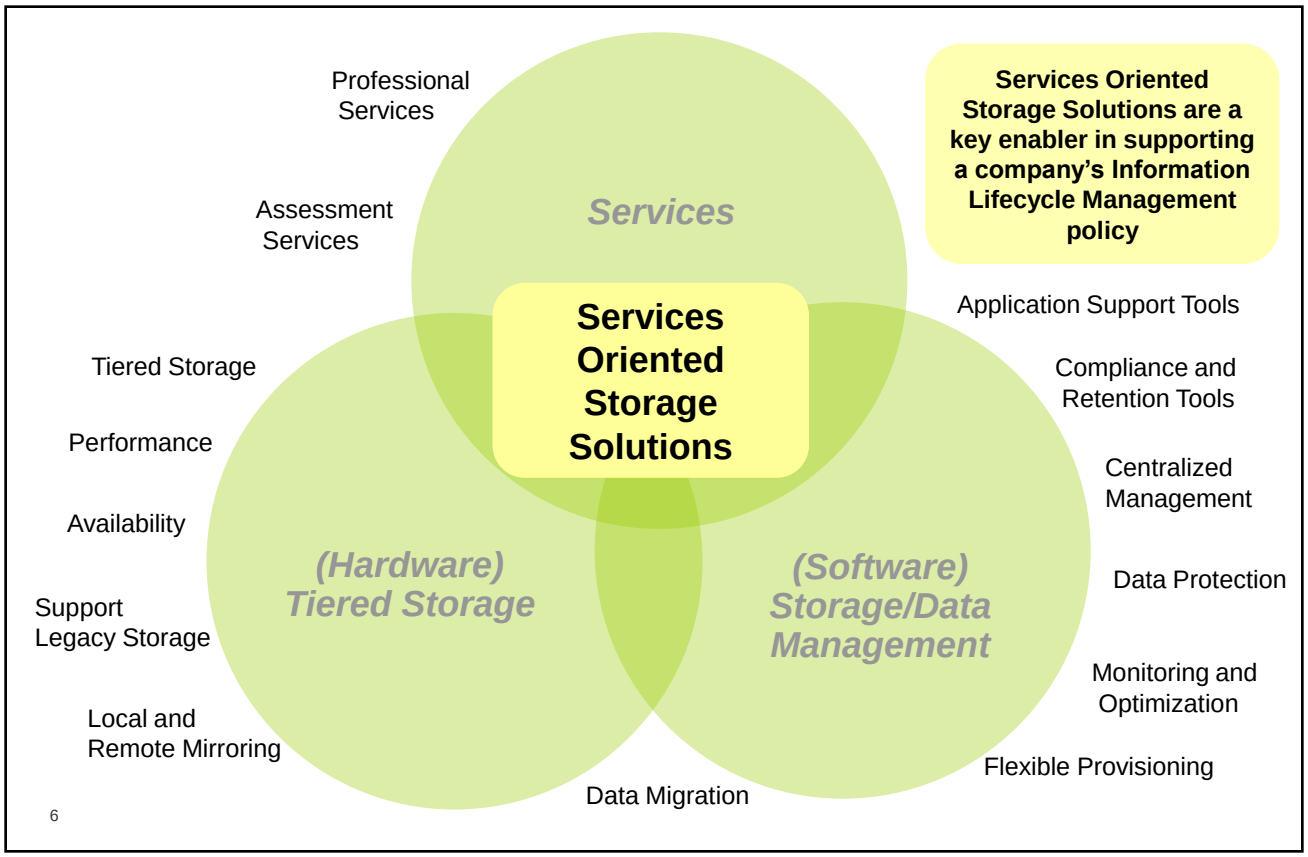
Most storage vendors focus on performance and capacity– often only from the ***storage*** perspective!



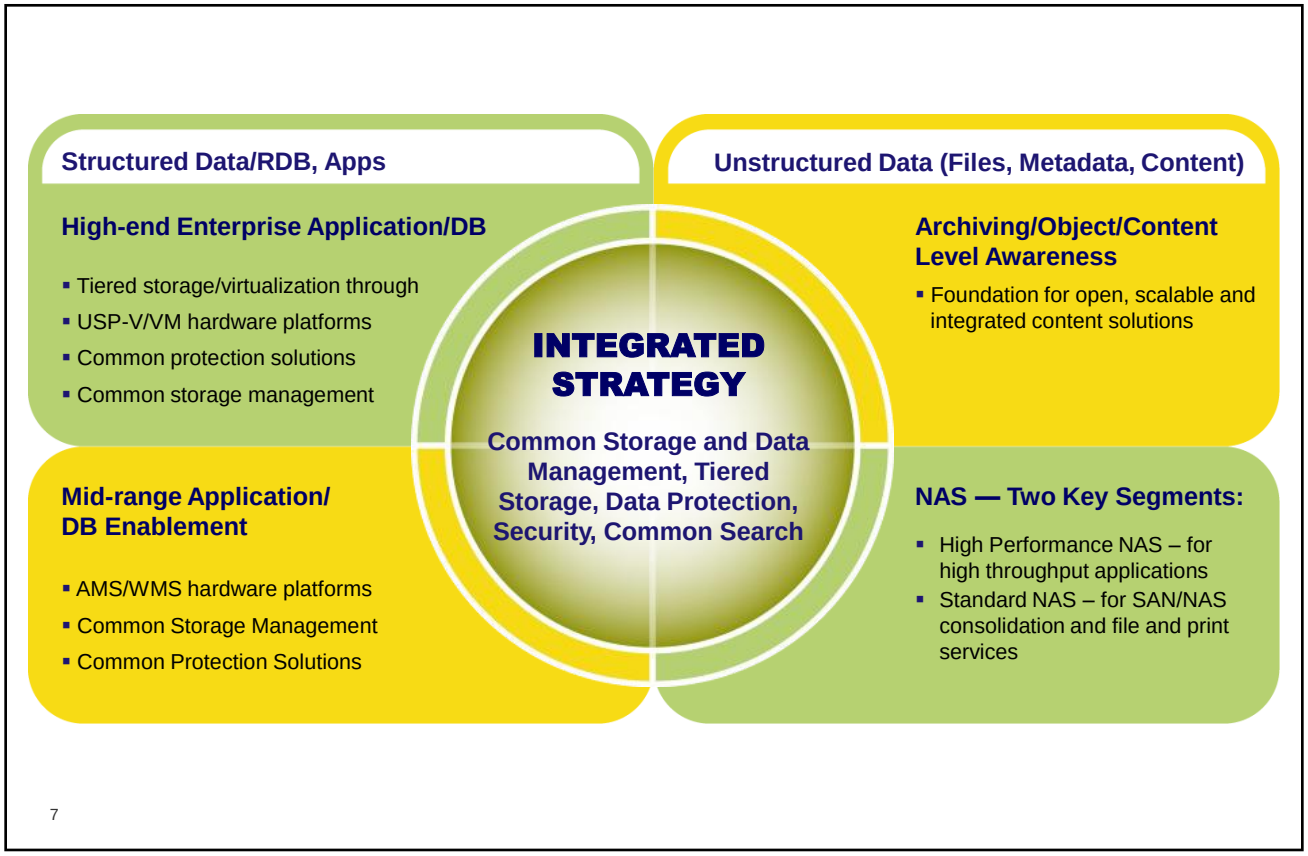
Services Oriented Storage Solutions are comprised of:
- Hardware
- Services
- Software

5

HDS Confidential: For distribution only to authorized parties.


Professional Services
Assessment Services
Services
Services Oriented Storage Solutions are a key enabler in supporting a company's Information Lifecycle Management policy
Application Support Tools
Compliance and Retention Tools
Centralized Management
Data Protection
Monitoring and Optimization
Flexible Provisioning
Data Migration
Services Oriented Storage Solutions
Tiered Storage
Performance
Availability
Support Legacy Storage
Local and Remote Mirroring
(Hardware) Tiered Storage
(Software) Storage/Data Management
6

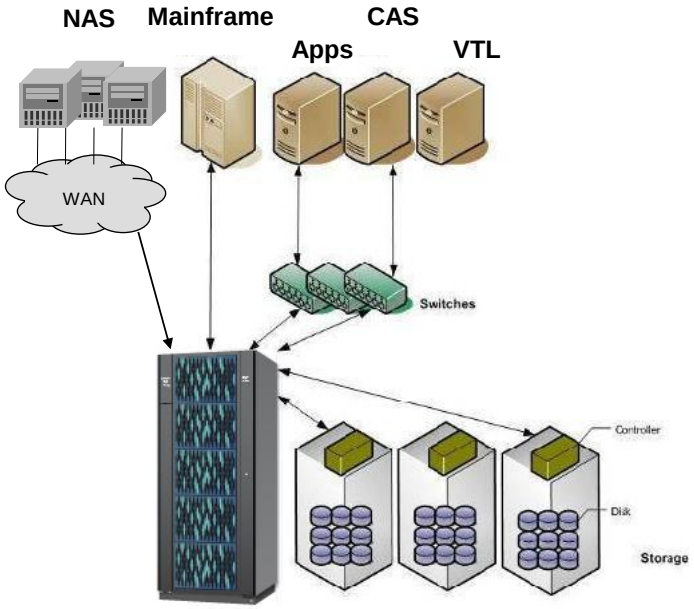
One Platform for All Data



HDS Confidential: For distribution only to authorized parties.

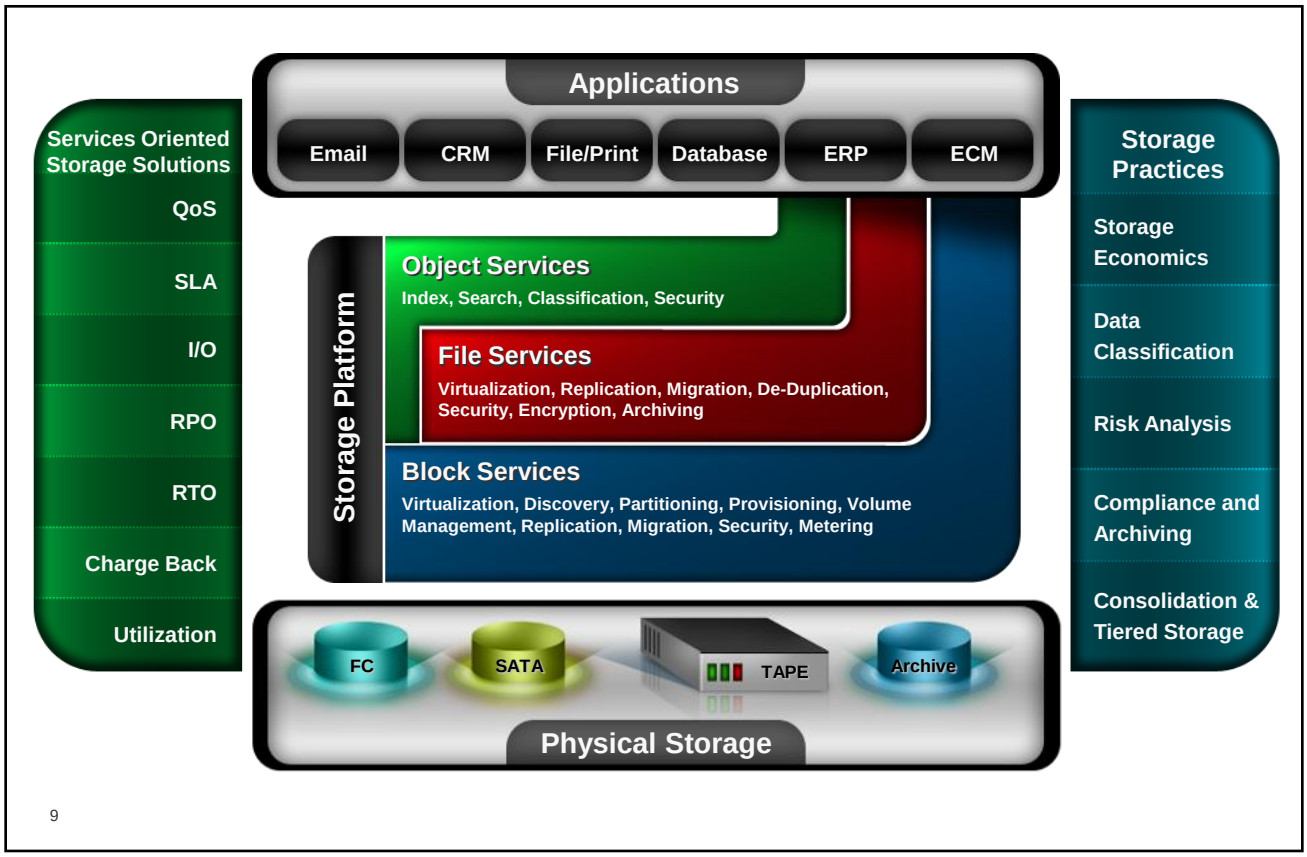
Virtualization: The Key Enabler

- Only Hitachi storage controller based virtualization:
  - Provides common storage and data services for all heterogeneous storage and application storage requirements
  - Adds no complexity or latency
  - Is not limited to SAN connection
  - Preserves end-to-end security
  - Enables seamless integration of future capabilities



8

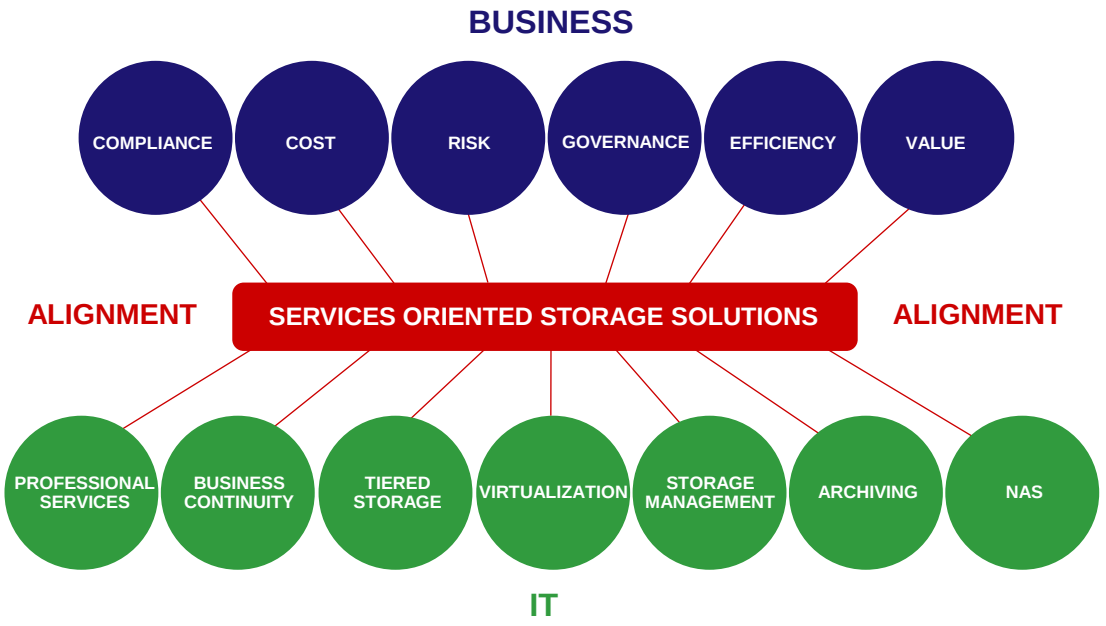
Architecture Summary



HDS Confidential: For distribution only to authorized parties.

Solutions Focus

- A business-centric approach for aligning IT storage resources with constantly changing business requirements, enabled by the Services Oriented Storage platform



10

HDS Confidential: For distribution only to authorized parties.

2. Product Overview

Module Objectives

- Upon completion of this module, the learner should be able to:
  - Present product overview and positioning
  - Discuss the features of:
    - Adaptable Modular Storage 2100
    - Adaptable Modular Storage 2300
    - Adaptable Modular Storage 2500

Hitachi Adaptable Modular Storage (AMS) is the official name. DF800 is an internal name used in the factory and, therefore, it may still show up in various documents and in certain management software.

DF800 consists of Adaptable Modular Storage models:

AMS2100 = DF800S

AMS2300 = DF800M

AMS25000 = DF800H/HE

Workgroup Modular Storage model:

WMS100 = DF700XS

Module Topics

- The Product Overview module covers the following topics:
  - Discuss the features of the Adaptable Modular Storage 2000 family
  - List the models and benefits

3

HDS Confidential: For distribution only to authorized parties.

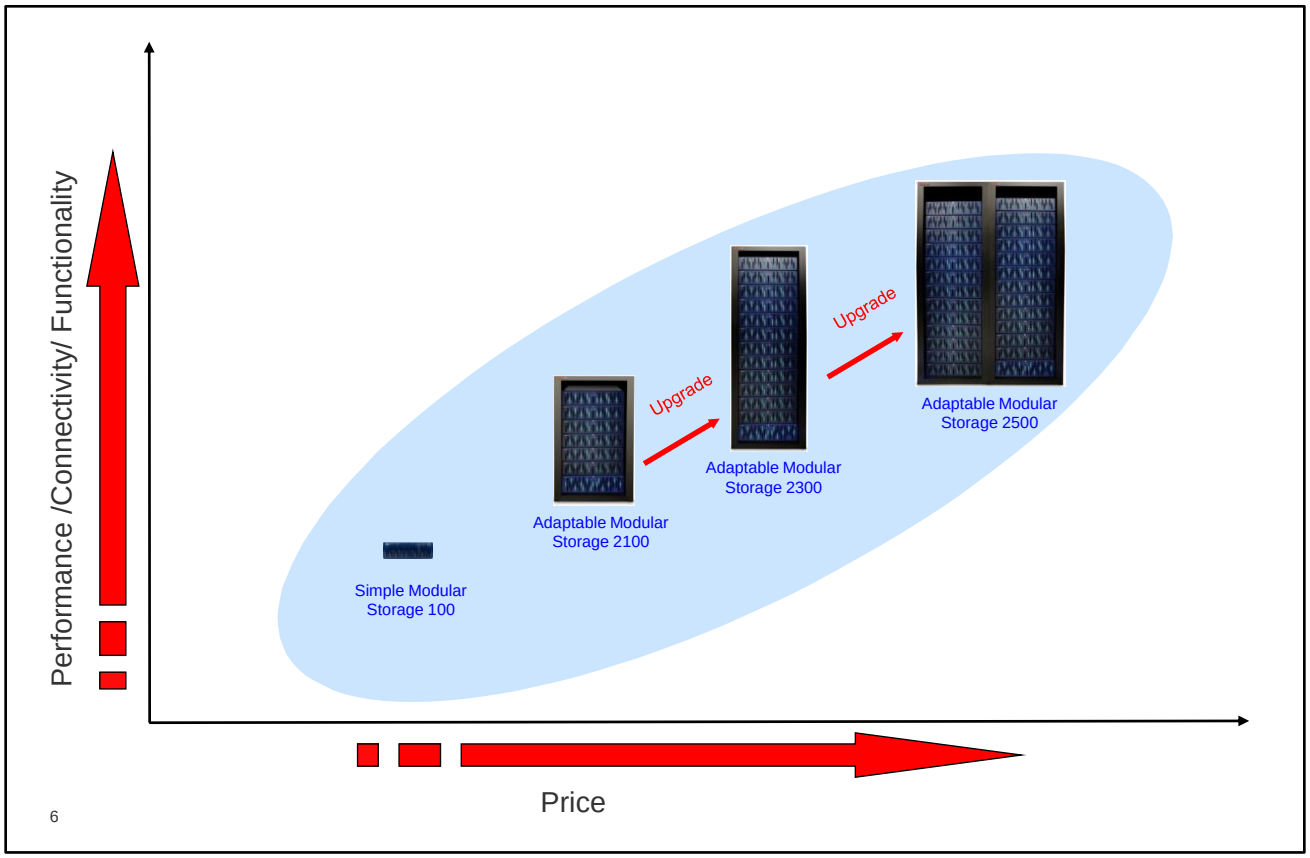
Product Description

- Dense, high-performance modular storage array
- Serial attached SCSI (SAS) back end architecture
  - **SAS** and **SATA II** drives
- Fibre Channel or iSCSI front end host ports
  - Single type of front end interface with models 2100 and 2300
  - Two concurrent types of front end interface with model 2500
- NAS through Fiber Channel "Gateway" offerings
  - Essential NAS
  - High-Performance NAS
- Active-Active Symmetric High Availability Dual Controller functionality
- Straightforward installation and configuration
- Intuitive storage management GUI (Hitachi Storage Navigator Modular 2 program)

4

- Simple maintenance and troubleshooting
  - Improved back end diagnostics due to new serial backend
- Real online firmware upgrade (no path failover required)
- Replication software:
  - Hitachi Copy-on-Write Snapshot software
  - Hitachi ShadowImage Replication software
  - Hitachi TrueCopy Remote Replication software
  - Hitachi TrueCopy Extended Distance software
- Easy data migration to and from previous Adaptable Modular Storage systems using TrueCopy software

Product Line Positioning



Upgrades are data-in-place upgrades. During an upgrade the system is powered down. These are **OFFLINE** procedures.

HDS Confidential: For distribution only to authorized parties.

Features

- All models provide:
  - High-speed response
  - Continuous data availability
  - Scalable connectivity
  - Expandable capacity
- Competitive Features and Functionality
  - Microsoft® environments such as Virtual Disk Service (VDS) and Microsoft Volume Shadow Copy Service (VSS) provider
  - Full Microsoft Windows® Server 2008 support
  - Native Multipath I/O (MPIO, MPxIO, etc.) support
    - Hitachi Dynamic Link Manager can still be used.
  - Functional enhancements to the ShadowImage software features to enable competitive VSS behavior in Microsoft environments
  - LUN Shrink/Grow feature and RAID Group Expansion feature
  - Hitachi Dynamic Provisioning
  - **60TB** LUN support

7

Specifications

Models	2100	2300	2500
Specs	• Dual Controller, 4 GB cache on each controller, 15 drives internal, Symmetric A/A, dual battery, dual redundant power supplies	• Dual Controller, 8 GB cache on each controller, 15 drives internal, Symmetric A/A, dual battery, dual redundant power supplies	• Dual Controller, 16 GB cache on each controller, 0 internal drives, Symmetric A/A, dual battery, dual redundant power supplies
Host Interface Options	• 4 Fibre Channel (FC) auto-sensing 1/2/4 Gbps • 4 iSCSI 1000Base-T copper Ethernet	• 8 Fibre Channel (FC) auto-sensing 1/2/4/8 Gbps • 4 iSCSI 1000Base-T copper Ethernet	• 16 Fibre Channel (FC) auto-sensing 1/2/4/8 Gbps • 8 iSCSI 1000Base-T copper Ethernet
Drive Interface	• 16 Serial Attached SCSI (SAS) • 4x4 wide link, 3 Gbps switched	• 16 Serial Attached SCSI (SAS) • 4x4 wide links, 3 Gbps switched	• 32 Serial Attached SCSI (SAS) • 8x4 wide links, 3 Gbps switched

8

The maximum cache configuration uses 4GB cache memory modules. 2GB and 4GB cache modules are available.

HDS Confidential: For distribution only to authorized parties.

Specifications for Regular Expansion Unit

Models	2100	2300	2500
RAID Levels	RAID-1, 1+0, 5, 6 (SAS & SATAII drives), RAID-0 (SAS drives only)		
Max # of RAID Groups	50	75	100
Max # of Spare Drives	15	30	30
Max # of LUs	2048	4096	4096
Max LU size	60TB		
Supported Drives	146GB/15K, 300GB/15K, 400GB/10K, 450GB/15K SAS,600GB/15K SAS 500GB/7200, 1TB/7200 , 2TB/7200 SATAII		
Upgrades	Model 2100 to 2300, and 2300 to 2500 model via Controller, Data In Place		
Expansion Unit Disk Trays (Optional based on capacity)	15 HDD/Tray (SAS/SATAII Intermix) Up to 7 trays (120 Drives total)	15 HDD/Tray (SAS/SATAII Intermix) Up to 15 trays (240 Drives total)	15 HDD/Tray (SAS/SATAII intermix) no HDDs in the controllers Up to 32 trays (480 Drives total)
Maximum Capacity	313TB	472TB	945TB

9

Specifications for Dense Expansion Unit

Models	2100	2300	2500
RAID Levels	RAID-1, 1+0, 5, 6		
Max # of RAID Groups	50	75	100
Max # of Spare Drives	15	30	30
Max # of LUs	2048	4096	4096
Max LU size	60TB		
Supported Drives	**1TB/7200, 2TB/7200 SATA II**		
Upgrades	Model 2100 to 2300, and 2300 to 2500 model via Controller, Data In Place		
Dense Expansion Unit Disks (No Intermix) (Optional based on capacity)	Min. 4 – max. 48 HDD/Tray Max. **63** Drives total (1 Tray)	Min. 4 – max. 48 HDD/Tray Max. **207** Drives total (4 Trays)	Min. 4 – max. 48 HDD/Tray Max. **480** Drives total (10 Trays)

10

HDS Confidential: For distribution only to authorized parties.

Software and Firmware Offerings

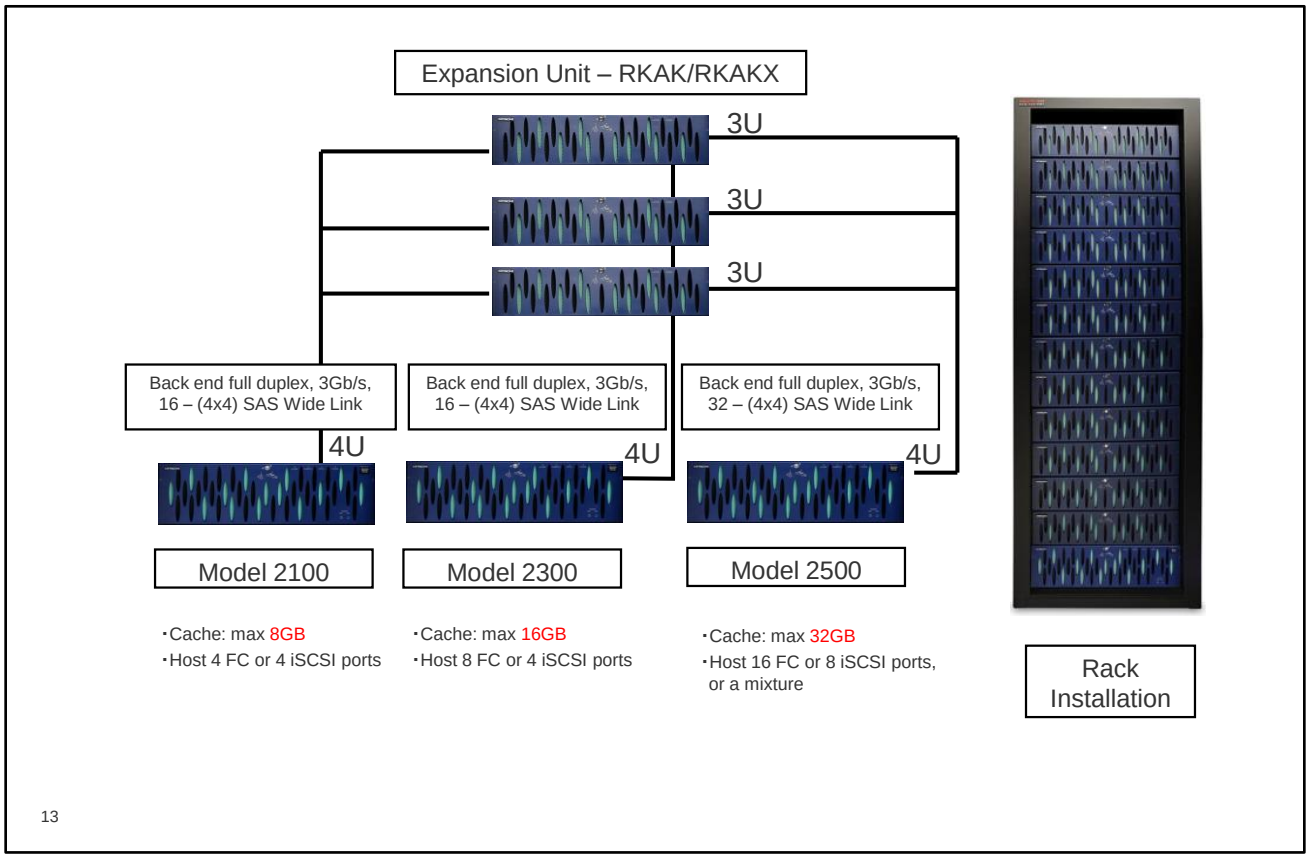
- Storage Management Software
  - Hitachi Storage Navigator Modular 2 (GUI and/or CLI)
  - Hitachi Storage Command Suite
- Bundled Storage Functions (Basic Operating System—BOS M)
  - Account Authentication
  - Audit Logging
  - LUN Manager
  - LU Grow/Shrink
  - Online RAID Group Expansion
  - Cache Residency Manager
  - Cache Partition Manager
  - Modular Volume Migration
  - SNMP Agent Support Function
  - Performance Monitor
  - Hitachi Dynamic Provisioning

11

Optional Storage Features	2000 family Product Line
ShadowImage (Clone)	1 Primary: 8 Secondary, 2048 Max, SNM2 GUI and CLI, CCI RAID Manager management available
Copy-on-Write Snapshot	1 Primary: 32 Snaps, 2048 Max, SNM2 GUI and CLI, CCI RAID Manager management available
TrueCopy (Sync Remote Mirroring)	1 Primary: 1 Secondary, SNM2 GUI and CLI, CCI RAID Manager management available
TrueCopy Extended (Async Remote Mirroring)	1 Primary: 1 Secondary, SNM2 GUI and CLI, CCI RAID Manager management available
Data Retention Utility (DRU)	Yes
Power Savings Feature (Spin Down RAID Groups)	Yes

Note: When using the SNM2 GUI or CLI, there is no HORCM file or Command Device required in order to manage the replication tasks or the Power Savings Feature.

External Design and Connections



The SAS wide link cable provides four 3Gb/sec SAS links. AMS 2100 and AMS 2300 systems have 16 SAS links and the AMS 2500 system has 32 SAS links.

HDS Confidential: For distribution only to authorized parties.

Certified Platforms

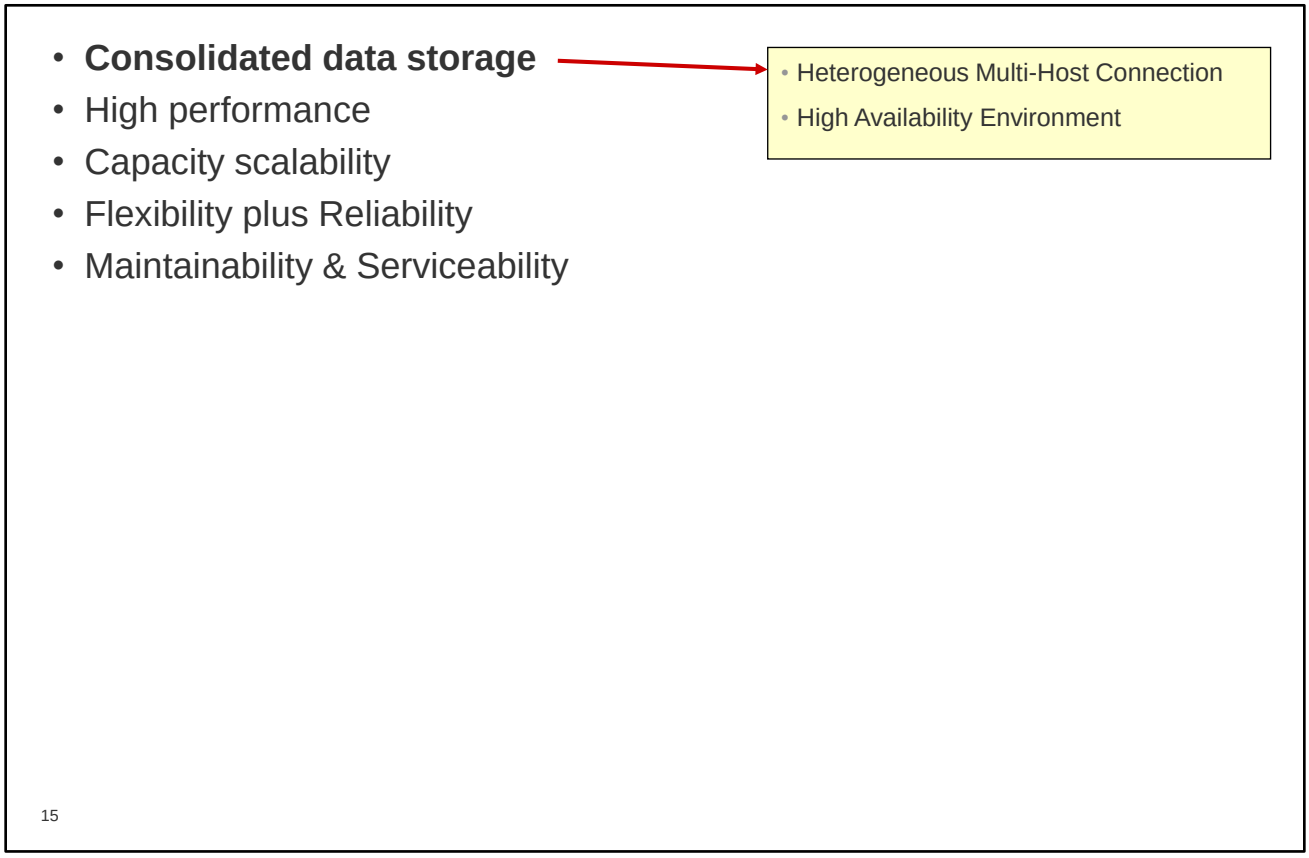
- True Multi-platform Intermix
  - Sun Solaris (Sparc and x64)
  - Microsoft Windows Server 2003 and 2008
  - HP-UX
  - Linux (RedHat, SuSE, Asianux, Oracle EL, RedFlag Linux)
  - IBM AIX
  - Novell Netware
  - Apple Mac OS X
  - HP OpenVMS
  - HP Tru64
  - VMware

14

Full open systems intermix is possible and implemented by configuring **Host Groups**.

The **Host Domain Group** concept allows for setting operating system or application software specific parameters as well as access security for Host Group individually. A Host Group is a group of one or more LUNs and exists behind a Fibre Channel host port in the Model 2000 family Modular Storage system.

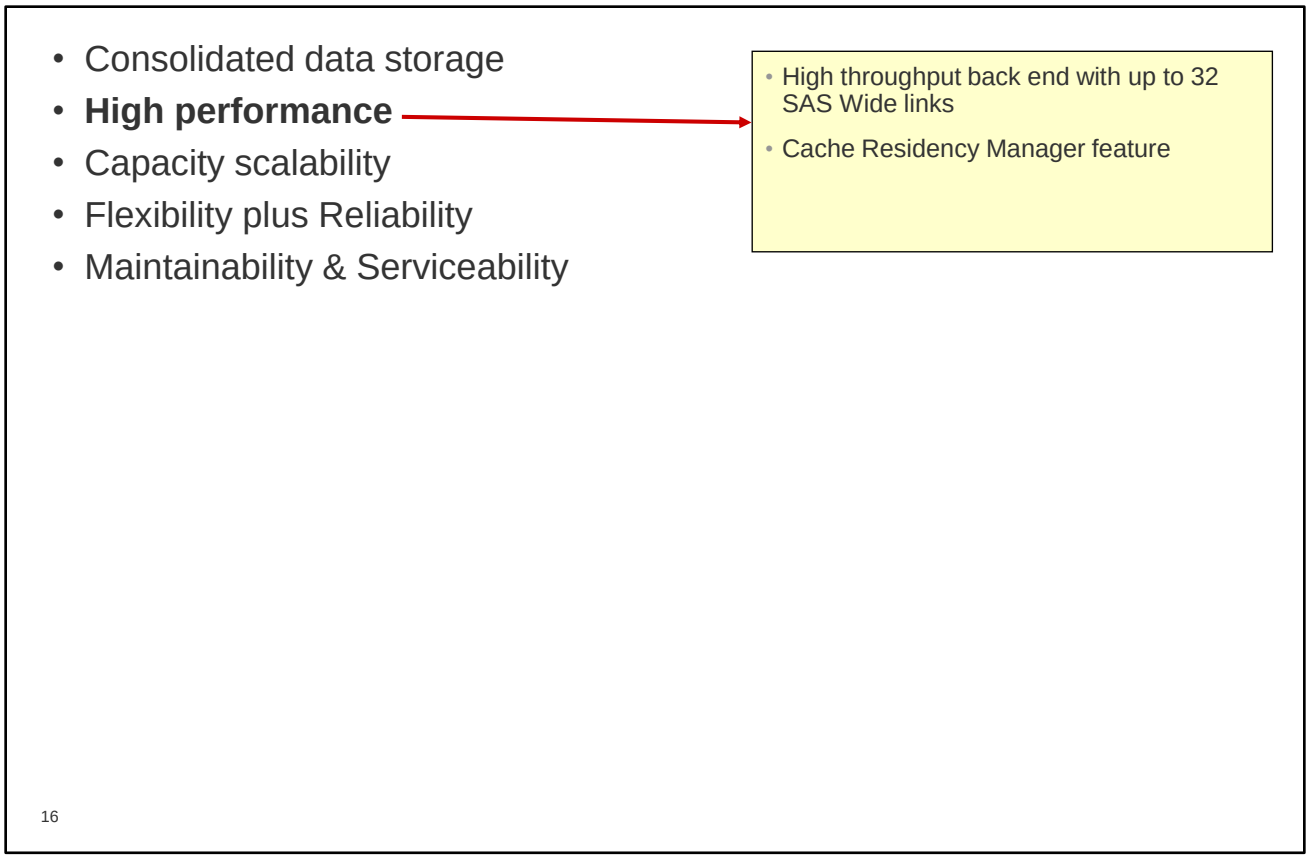
Highlights



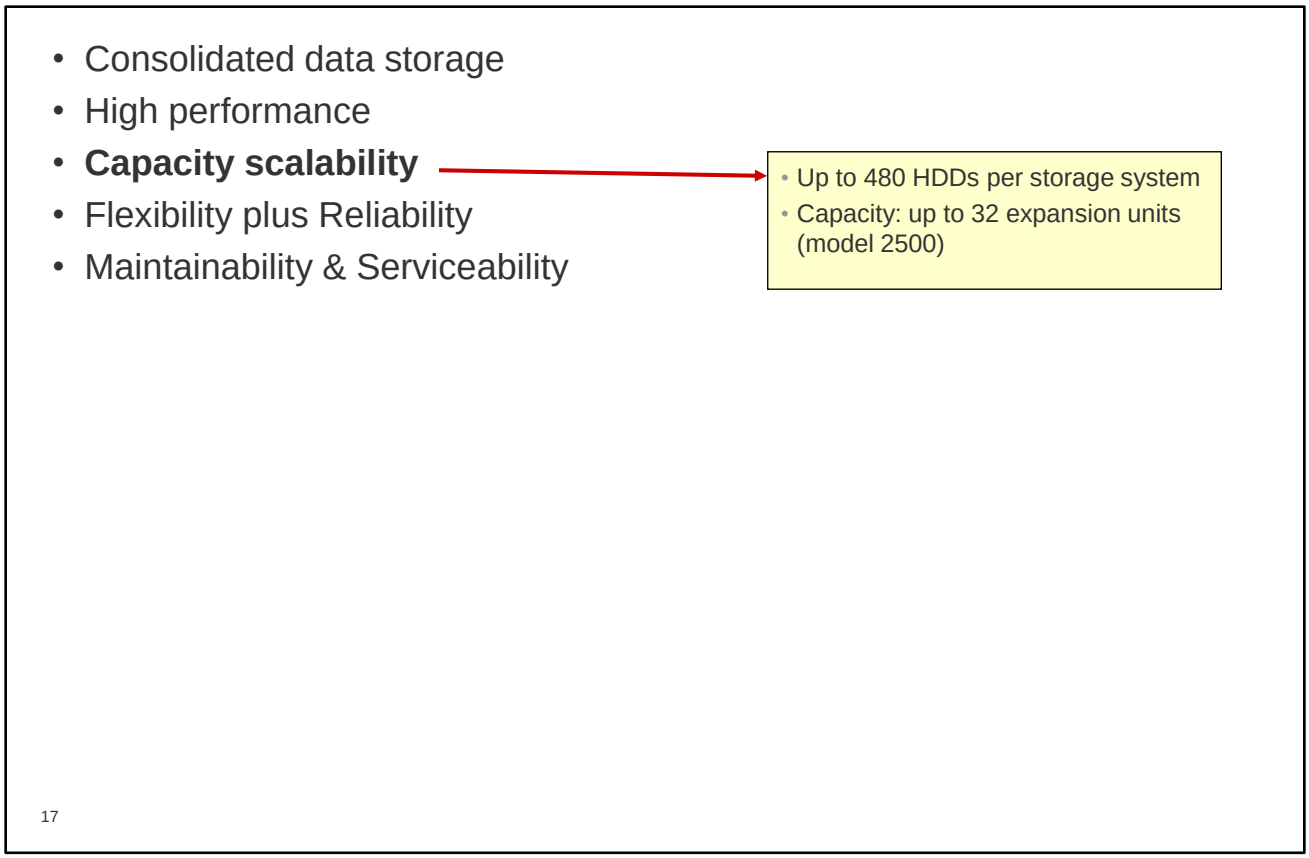
Support for Heterogeneous Multi-Host Connection is accomplished by setting Host-specific parameters for any individual group of LUNs (Host Group).

A High Availability environment (no single point of failure) requires the same level of redundancy in the Hosts, Storage Area Network, and in the Storage Device.

HDS Confidential: For distribution only to authorized parties.



Two back end paths models 2100 and 2300. Also SAS Wide links to drives on the back end.



A system can hold a mix of both high-speed (expensive) HDDs for performance and slower (cheaper) drives for capacity.

- Performance: Used for online transactions, and more.
- Capacity: Used for audio and video streaming, backup's and more.

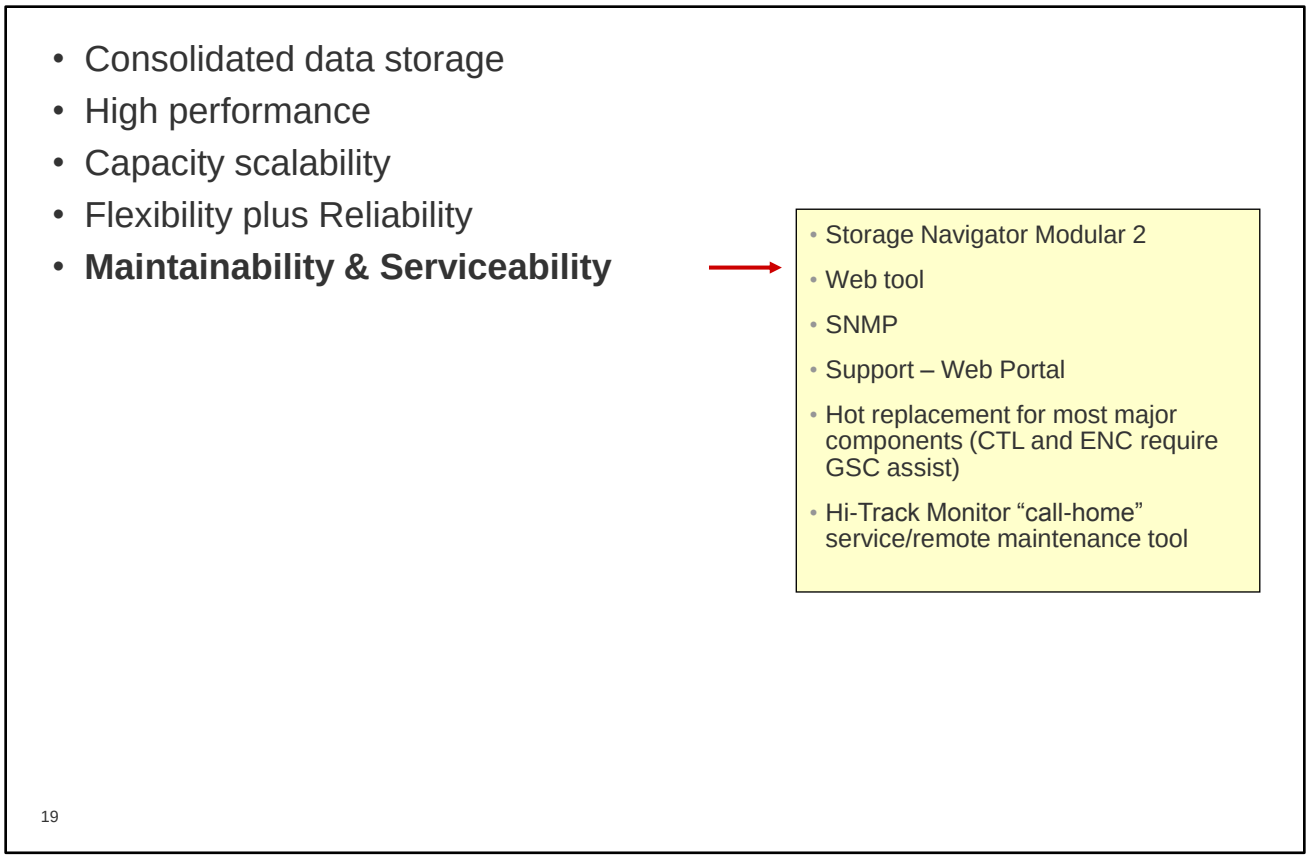
HDS Confidential: For distribution only to authorized parties.

- Consolidated data storage
- High performance
- Capacity scalability
- **Flexibility plus Reliability**
- Maintainability & Serviceability

- Online capacity upgrade
- RAID levels supported 0, 1, 5, 6, 1+0
- Cache Partition Manager
- HDD Roaming
- Up to 30 global spare drives (models 2300, 2500)
- Online Verify and Dynamic Sparing
- LUN Mapping, Host Group Mode, and HG Security
- 8-byte Data Assurance Code

18

Online capacity upgrade: HDDs and expansion units can be added online; Controllers and Cache Memory cannot.



Storage Navigator Modular 2 is shipped with the array. The build center, or CTO, will install and enable feature keys for certain basic software features.

HDS Confidential: For distribution only to authorized parties.

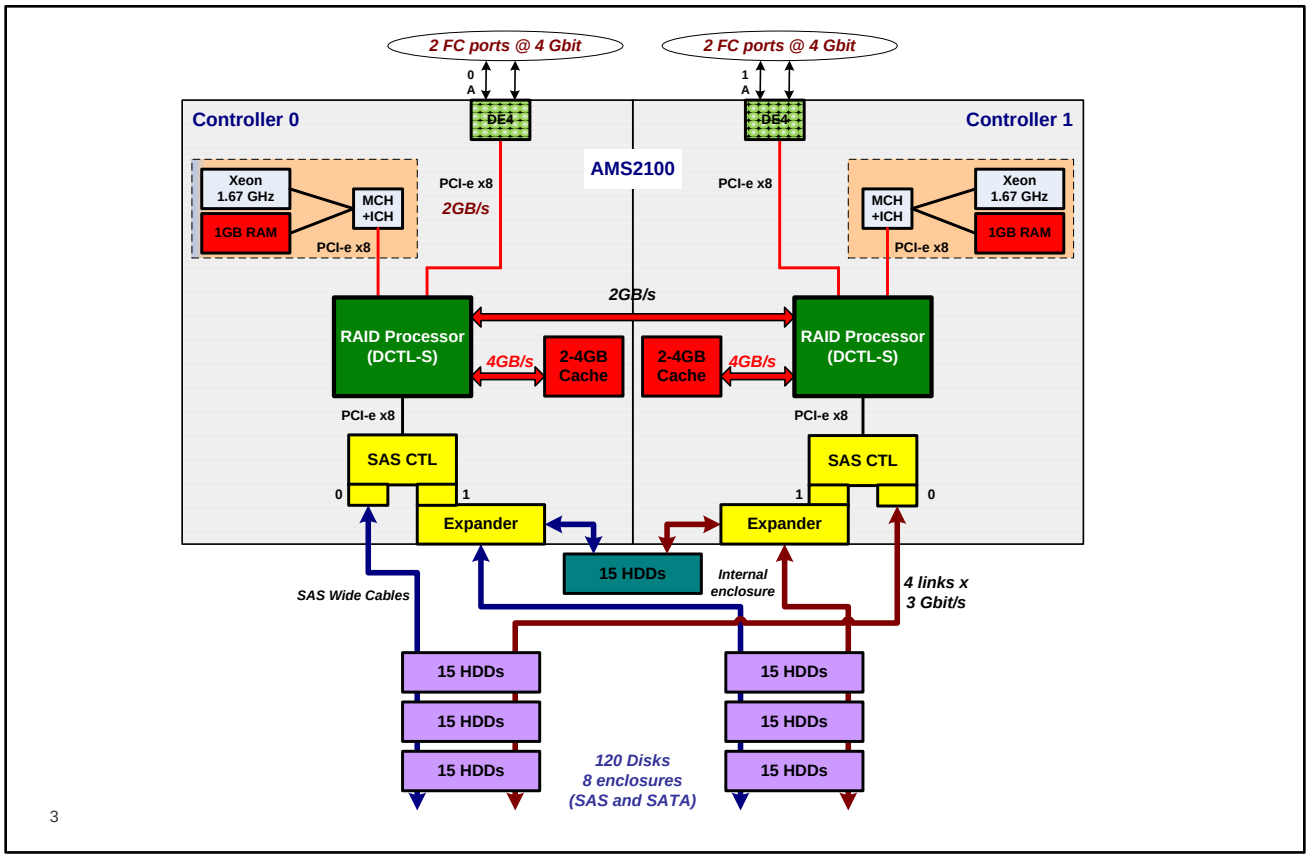
3. Hardware Components and Architecture

Module Objectives

- Upon completion of this module, the learner should be able to:
  - Identify the hardware components of the Adaptable Modular Storage 2000 family that affect installation, configuration, and troubleshooting
  - Explain the architecture that affects installation, configuration, and troubleshooting

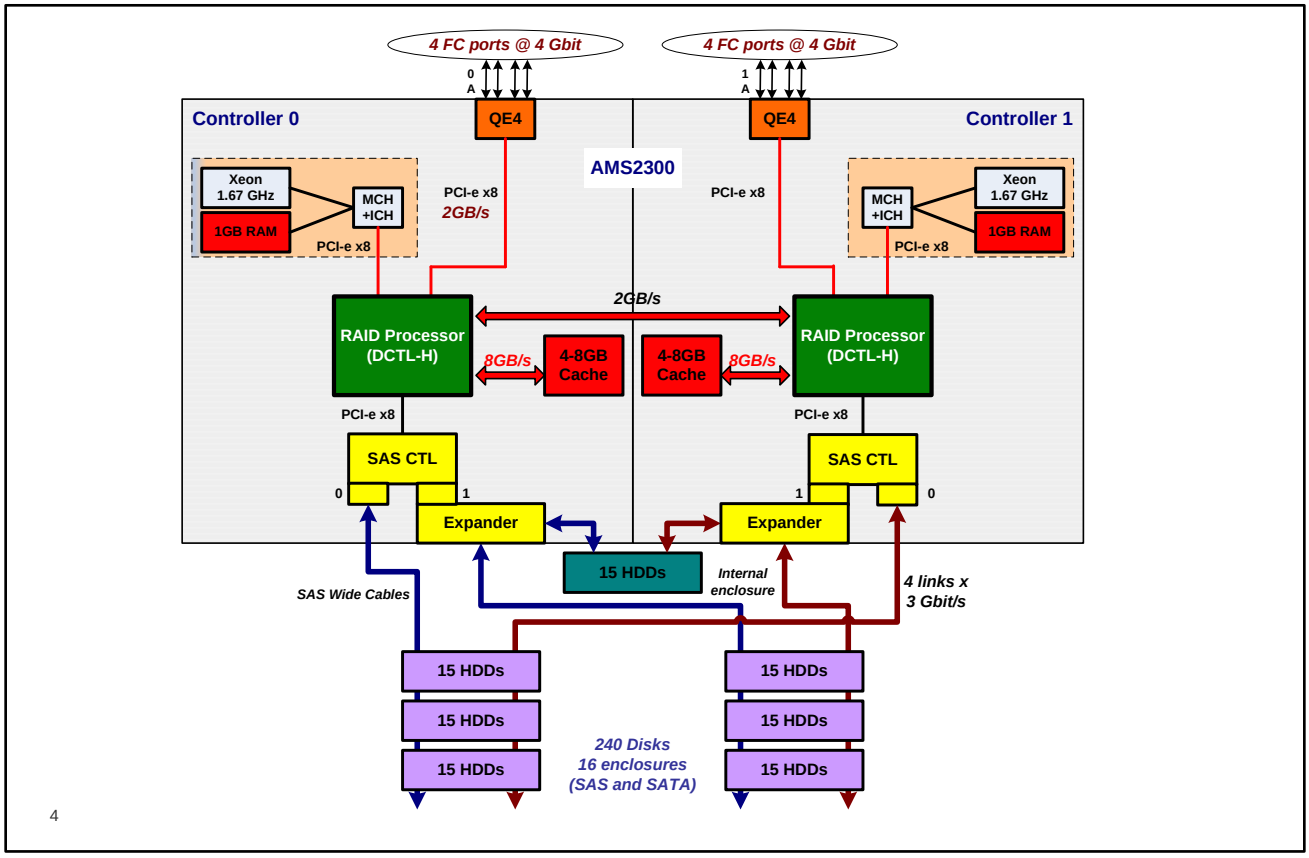
2

Architecture — Model 2100 Controller

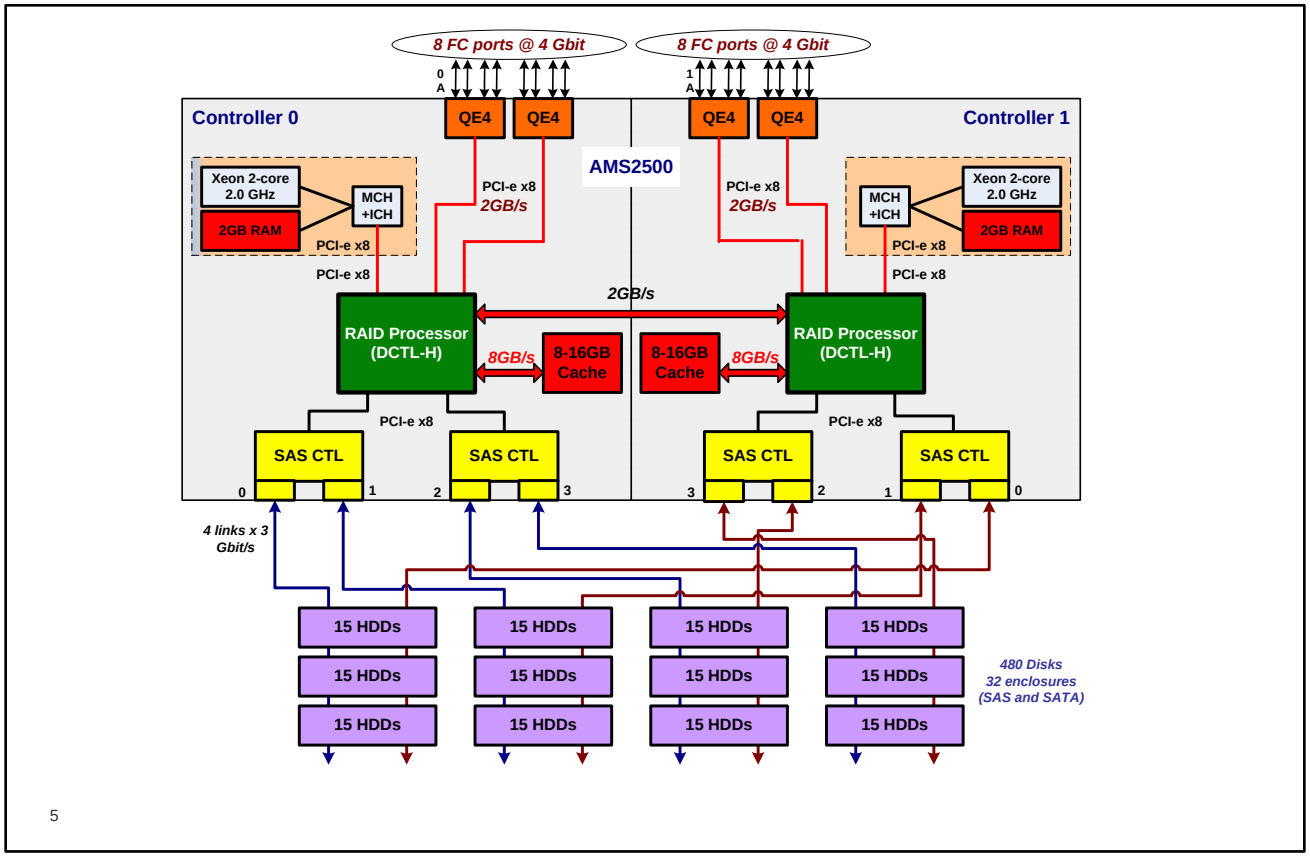


HDS Confidential: For distribution only to authorized parties.

Architecture — Model 2300 Controller



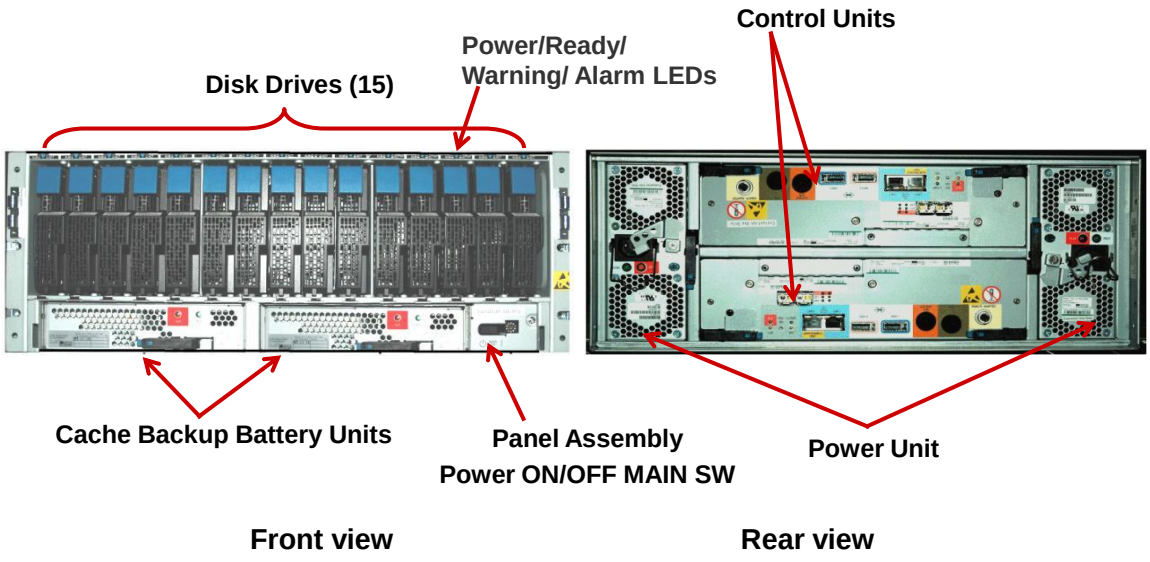
Architecture — Model 2500 Controller



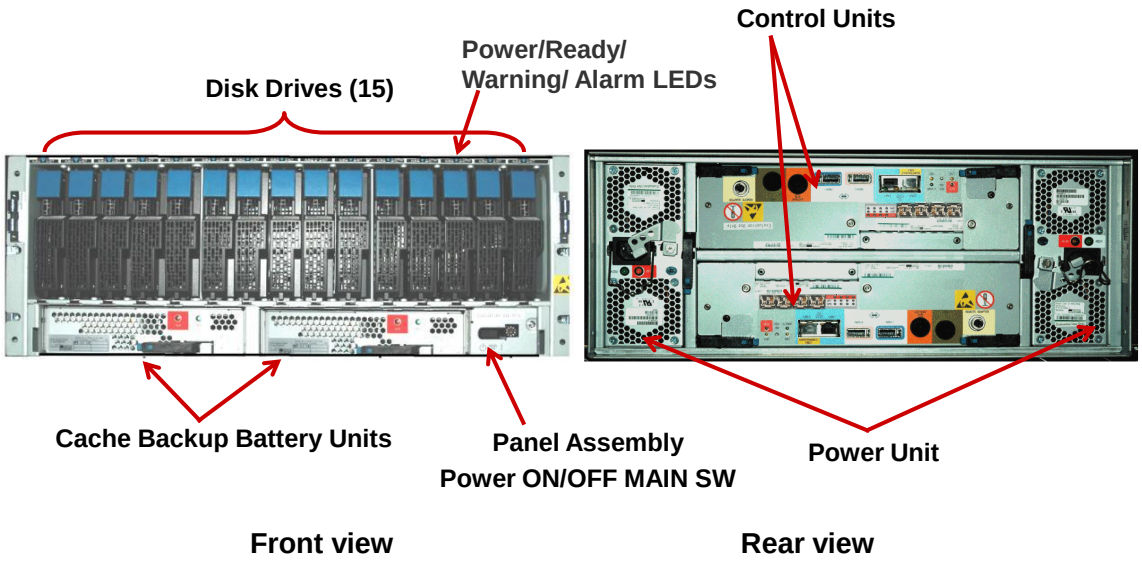
HDS Confidential: For distribution only to authorized parties.

Hardware Components

- Adaptable Modular Storage 2100

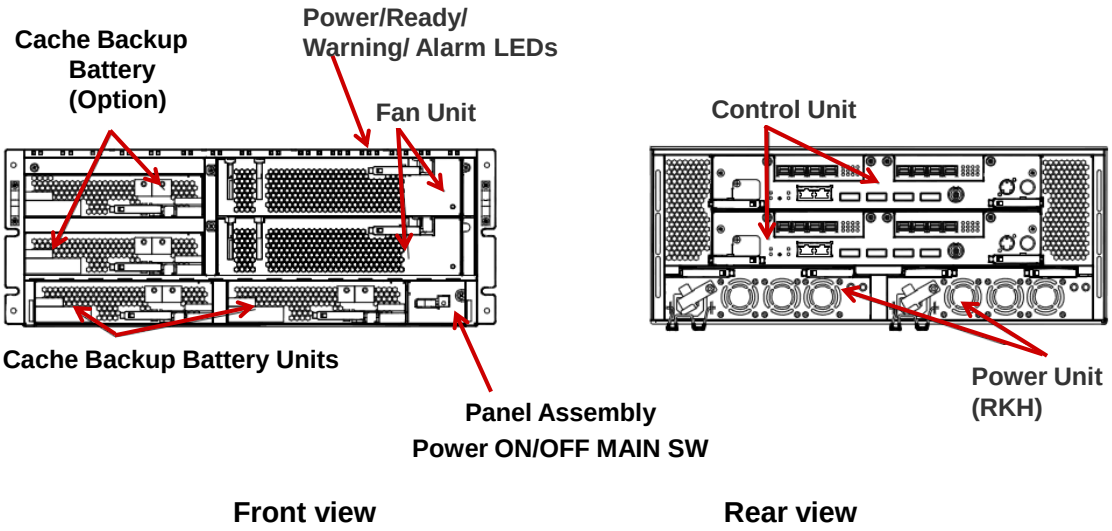


- Adaptable Modular Storage 2300



7

- Adaptable Modular Storage 2500

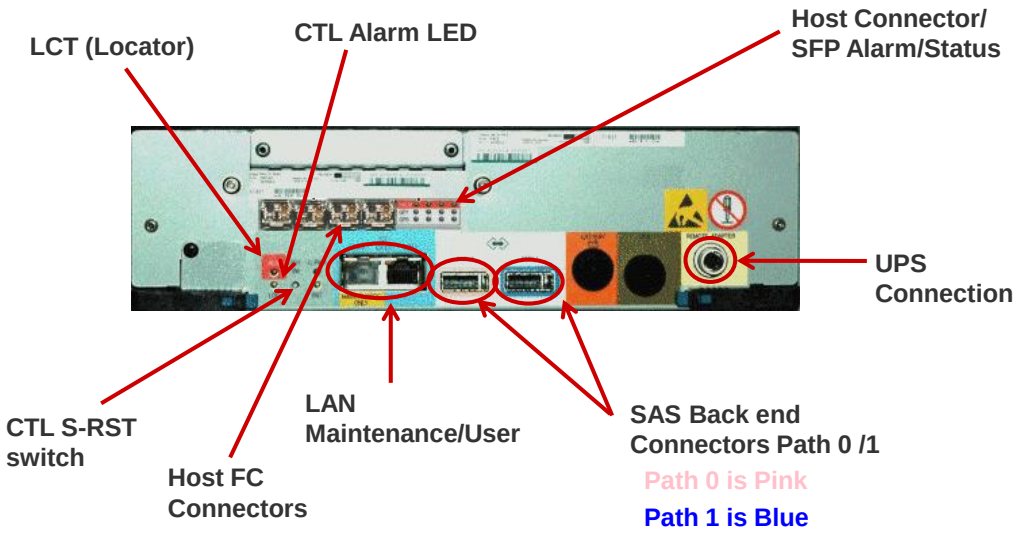


8

HDS Confidential: For distribution only to authorized parties.

Controller Labeling

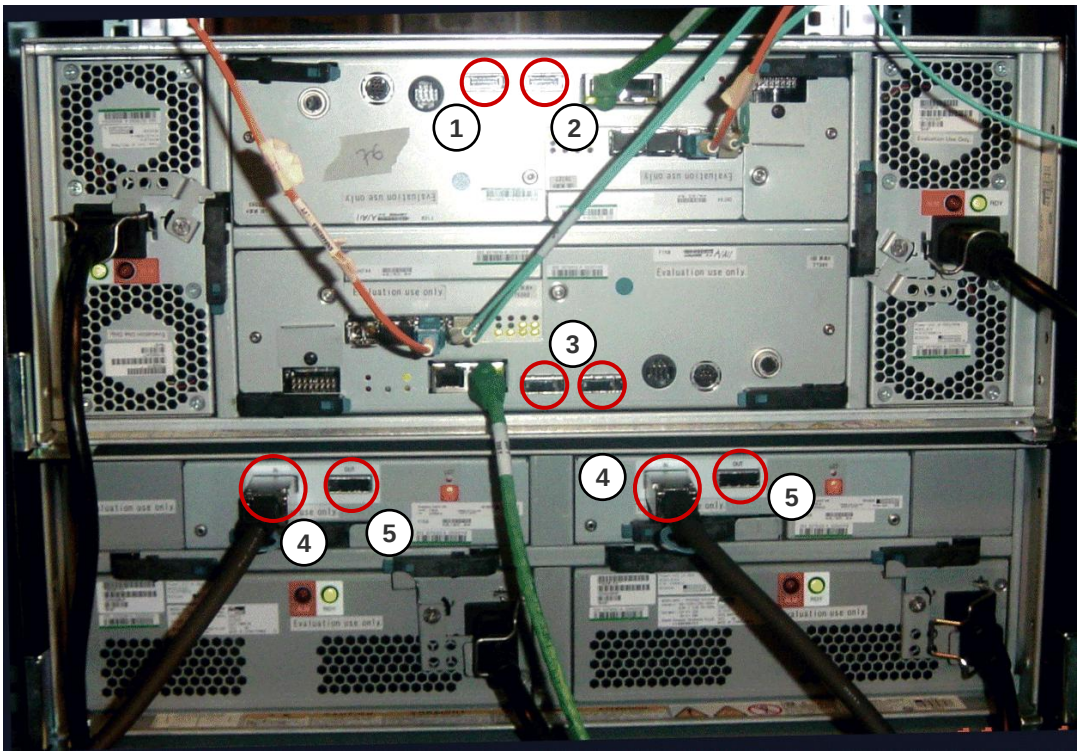
- Easy-to-read symbols
- Color coded components indicate function



9

LCT is a new locator LED. This is also located on the expansion units (front and rear) to identify and locate with Storage Navigator Modular 2.

Back End SAS Connections



10

ENC stands for Enclosure Controller.

This picture shows top unit controller:

- CTL0
- CTL1
- ENC Path0/1 connectors of both controllers

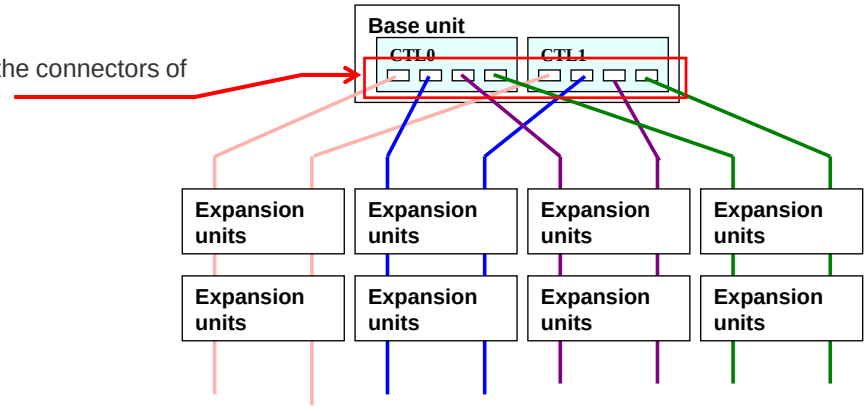
It also shows the bottom expansion unit with ENC connectors:

- IN connector shown plugged in on Path0
- OUT connector

HDS Confidential: For distribution only to authorized parties.

- Expansion Unit Drive Tray Connection
  - New ENC Cables – Simple and easy back end cabling
    - Single keyed cable type for all connections
  - Common HDD tray for SAS and SATA
  - All SAS back end paths active
  - Model 2500 most paths/ENC cables

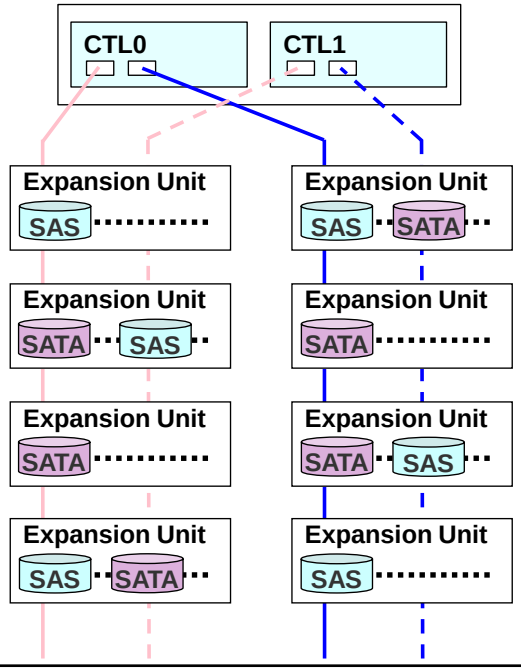
Note: Cables are one color, but the connectors of the controller are color coded.



11

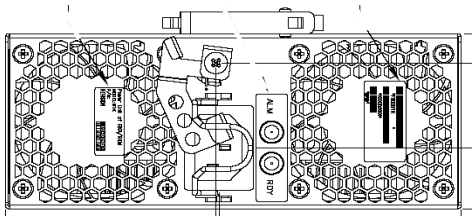
ENC stands for Enclosure Controller

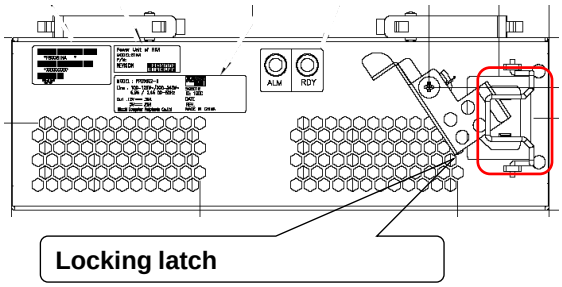
- Typical Cabling Diagram of Model 2100 and 2300



Power Supply

- Power supply does not have a switch.
- To turn power-off, flip power switch
- To completely power off the system, remove power cable..
- Locking latch component added to secure power supply cable.
  - Power supply unit cannot be removed from unit frame while connected.



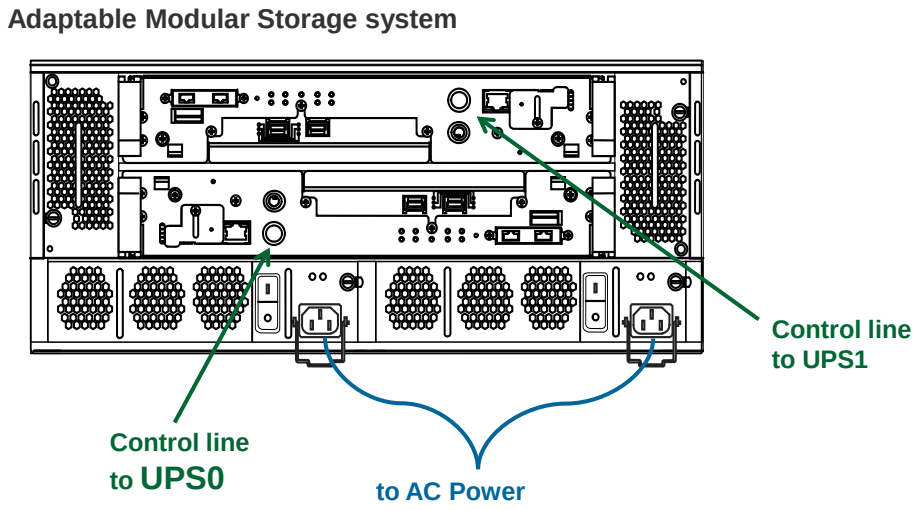


13

HDS Confidential: For distribution only to authorized parties.

Uninterruptible Power Supply Connections

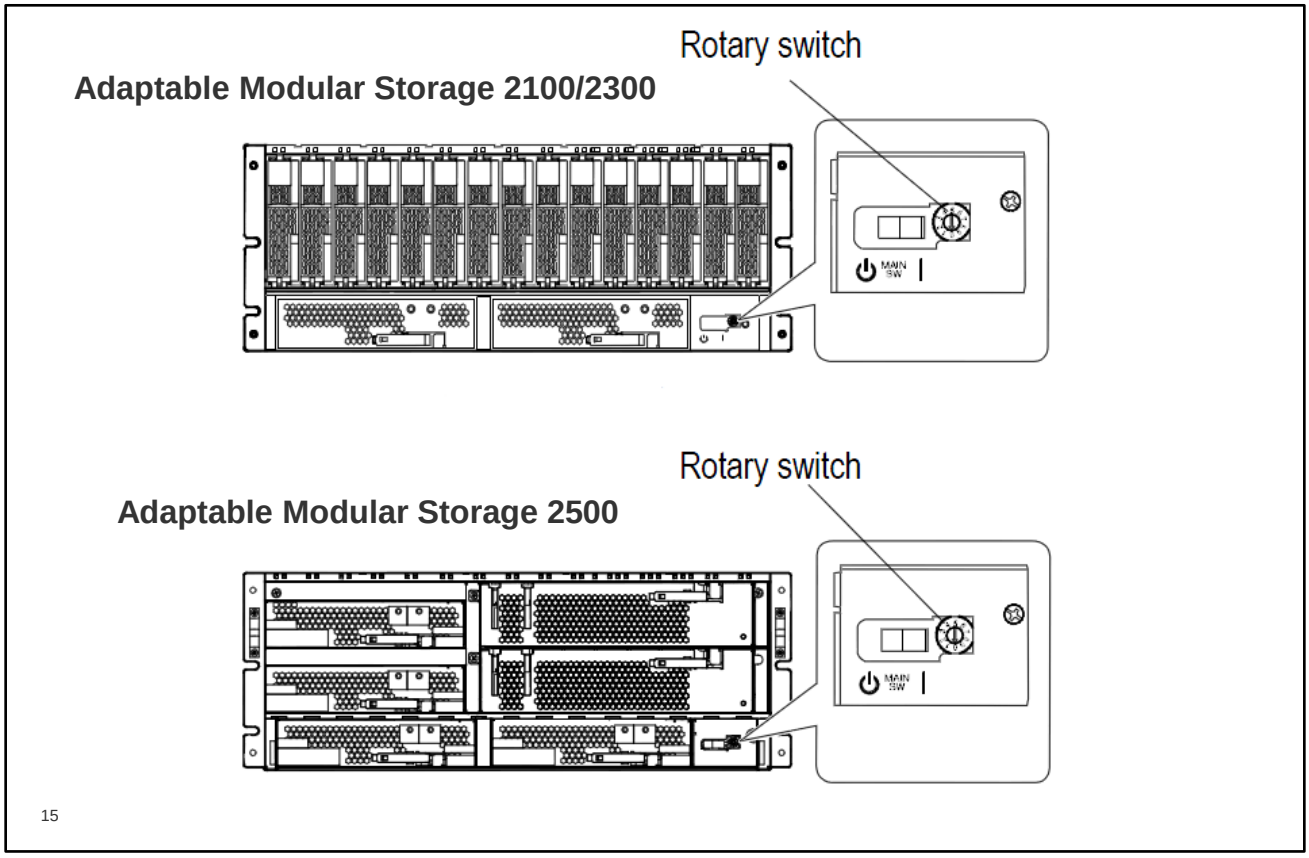
- Uninterruptible power supply (UPS)



14

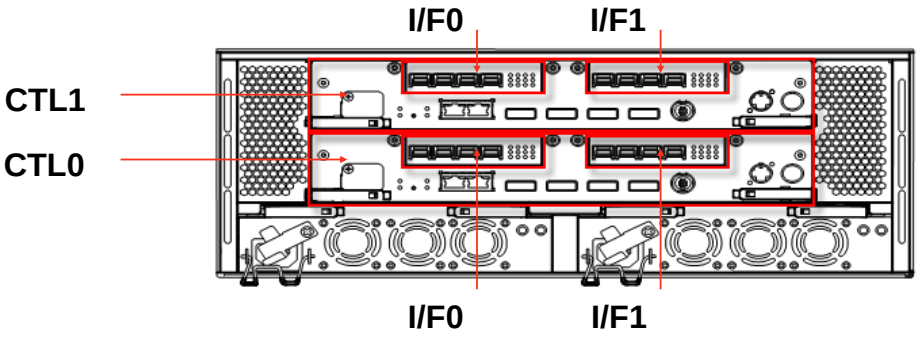
The UPS interface can detect high or low voltages, and more on the Adaptable Modular Storage 2000 family system.

Rotary Switch Power Connection Mode Settings



HDS Confidential: For distribution only to authorized parties.

Model 2500 – FC/iSCSI Host Interface Intermix



- Each interface board on CTL0 and CTL1 must be the same type of interface board or **None**. Otherwise the storage system will not go into the **Ready** status.

  Exception:
  If the interface board type of I/F0 on **CTL0** is **iSCSI**, the interface board I/F0 on **CTL1** must be either **iSCSI** or **None**.

16

Note:

The FC/iSCSI Host Interface intermix on a AMS2500 storage system is supported with the v.6 firmware.

I/F0 and I/F1 in this picture are images for FC Interface board, not for iSCSI Interface board (iSCSI has only 2 ports per Interface board).

Specifications

- Battery Life

		2100	2300	2500
Maximum Cache	2GB DIMMs	4GB per array	8GB per array	16GB per array
	4GB DIMMs	8GB per array	16GB per array	32GB per array
Internal Batteries		2	2	**2 (default)** **2 (optional)**
Internal Battery Backup Time	2GB/4GB DIMMs	72/48 hours	36/24 hours	24/48/96 hours
External Batteries		N.A.	N.A.	1 or 2
External Battery Backup Time (with maximum internal batteries)	One	N.A.	N.A.	90 hours
	Two	N.A.	N.A.	168 hours

17

Note: This is applicable with maximum internal batteries installed, fully loaded, and working properly.

Controller Architecture

- Model **2100** CTL Architecture
  - MPU:IA32 1.67GHz
  - Internal data bus: PCI-express
  - Back end: SAS
  - Fibre Channel interface – 2 Ports
  - iSCSI interface – 2 Ports
- Model **2300** CTL Architecture
  - MPU: IA32 1.67GHz
  - Internal data bus: PCI-express
  - Back end: SAS
  - Fibre Channel interface – 4 Ports
  - iSCSI interface – 2 Ports
- Model **2500** CTL Architecture
  - MPU: IA32 2.0GHz Dual Core
  - Internal data bus: PCI-express
  - Back end: SAS
  - Fibre Channel interface – 8 Ports
  - iSCSI interface – 4 Ports

18

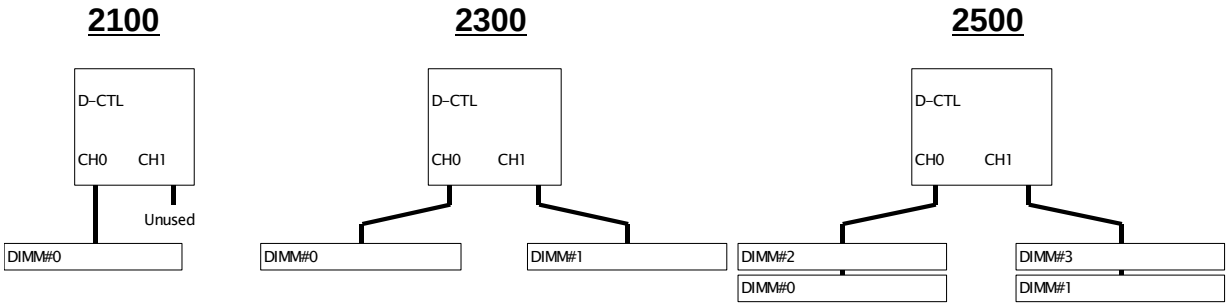
Note:

On models 2100 and 2300 you can not intermix Fibre Channel and iSCSI.

On the 2500 model you can intermix Fibre Channel and iSCSI. In this case you will have 4 Fibre Channel ports per controller and 2 iSCSI ports per controller.

Cache Architecture

- Cache composition in each models controller



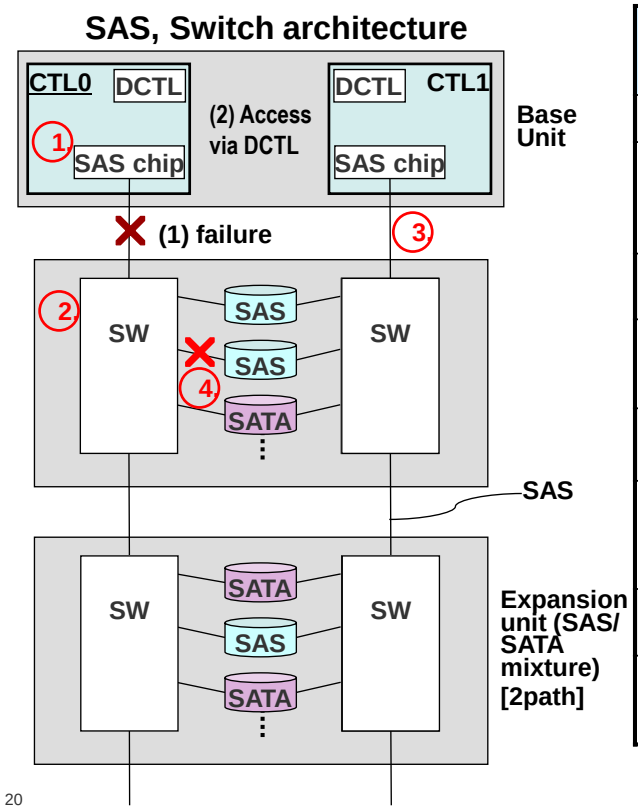
	2100	2300	2500
Maximum Cache per CTL	1	2	4
Additional #	n/a	1	2

19

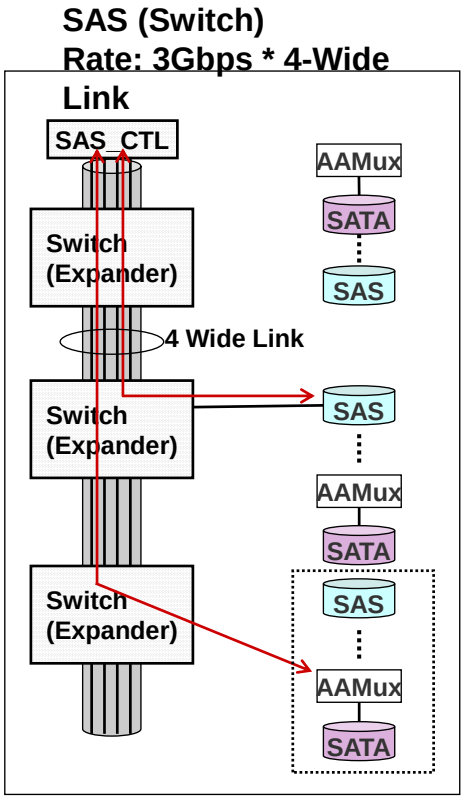
Note: Model 2500 – Cache units must be installed in PAIRS.

HDS Confidential: For distribution only to authorized parties.

Back End Architecture



2000 Advantage	Contents
1. New Topology (Loop –>Switch)	
SAS Switch Architecture	Point to Point connection between SAS protocol chip and each drive. 4-Wide Links are configured in one SAS back end path.
2. Simple and Easy cabling (common expansion unit for SAS/SATA)	
Simple & Easy back end cabling configuration	Back end cabling can be simplified by SAS/SATA coexistence in the same expansion unit chassis.
3. Back end Path failover (through back end path –> DCTL)	
Steady transaction by changing the back end access path	When the back end path failure occurs, the access path failover is executed using DCTL interface.
4. Improvement of Failure diagnostics (Loop –> Switch)	
Improvement of back end failure diagnostics	The failure parts detection is improved by switch device (Expander) functions.



- 4-Wide Links are configured in one SAS back end path.
- Each link is dynamically allocated to any of the disks.

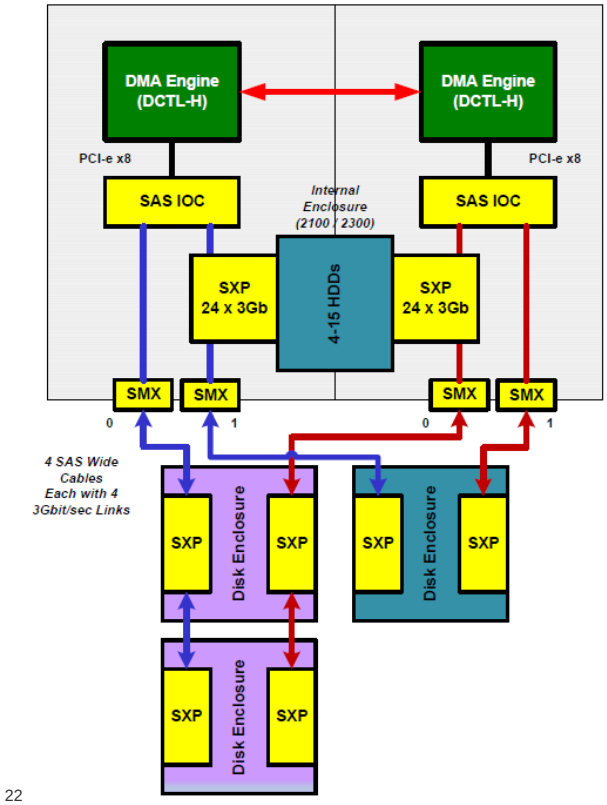
— 12 Gbps; — 3 Gbps

SAS: Serial Attached SCSI

SATA: Serial Advanced Technology Attachment

AAMux: Active-Active Multiplexer

HDS Confidential: For distribution only to authorized parties.



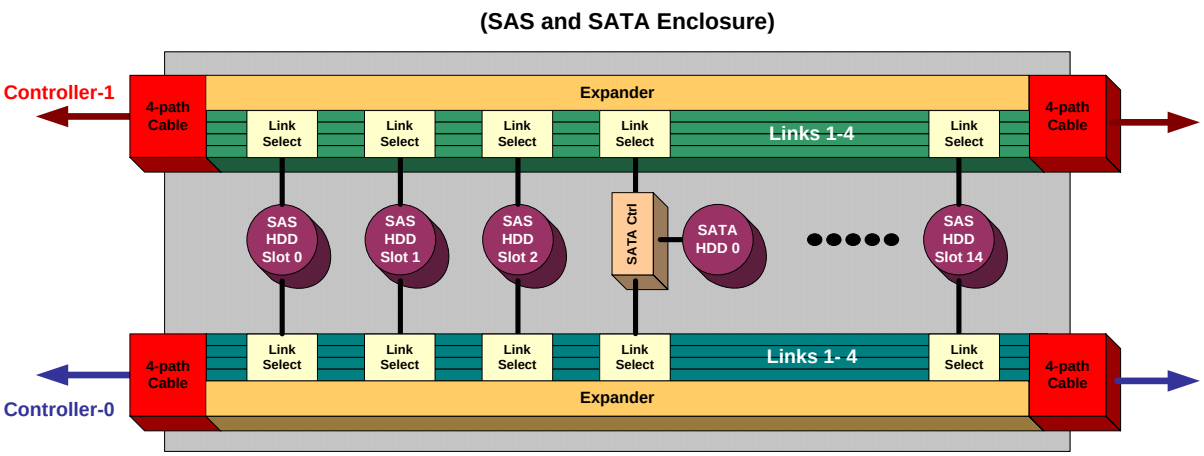
Field Replaceable Unit (FRU)

No	Failure Parts	FRU
1	Controller (including SAS protocol chip, Expander, etc)	Controller (new CTL CRU consists of cache, and host interfaces)
2	ENC (including Expander)	ENC and Cable
3	Cable between chassis	ENC and Cable
4	SAS Drive	SAS Drive
5	SATA Drive (includes AAMux)	SATA Drive (includes AAMux)

WL: Wide Link

AAMux: Active-Active Multiplexer

Disk Expansion Tray



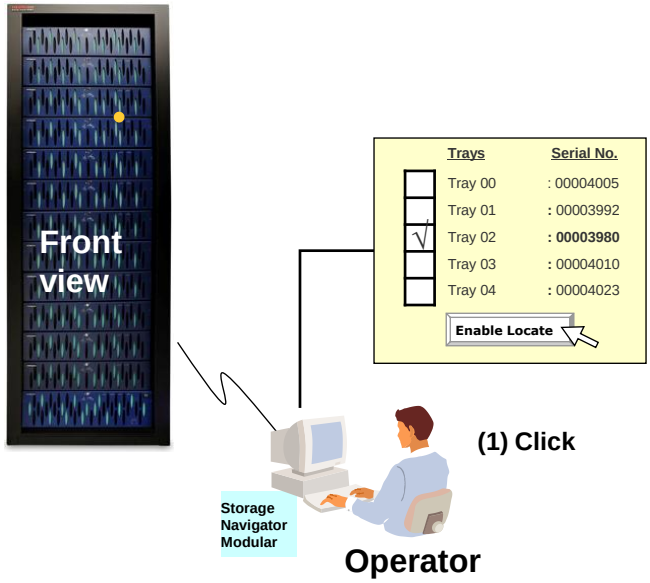
23

An AAmux (SATA Ctrl) chip is installed on every SATA Disk.

HDS Confidential: For distribution only to authorized parties.

Back End Logical Configuration

- Storage Navigator Modular 2 shows the logical configuration and serial number of each expansion unit.
- Users can recognize a target expansion unit by illuminating an LED using Storage Navigator Modular 2 and executing a Check with Serial No. of RKAK.
- Users can recognize a failed component (ENC, PS, etc.) by LED.



24

RKAK refers to the expansion unit.

System Management Options Interface

- Out of band
  - **Hitachi Storage Navigator Modular 2** – CLI and GUI
  - SNMP (trap for failure information only)
  - Maintenance
  - Embedded Web Server for Field Maintenance
- In-band
  - **Hitachi Command Control Interface** (CCI or RAID Manager) for:
  - Hitachi ShadowImage Heterogeneous Replication software
  - Hitachi TrueCopy Heterogeneous Remote Replication software
  - Hitachi Copy-on-Write Snapshot software

25

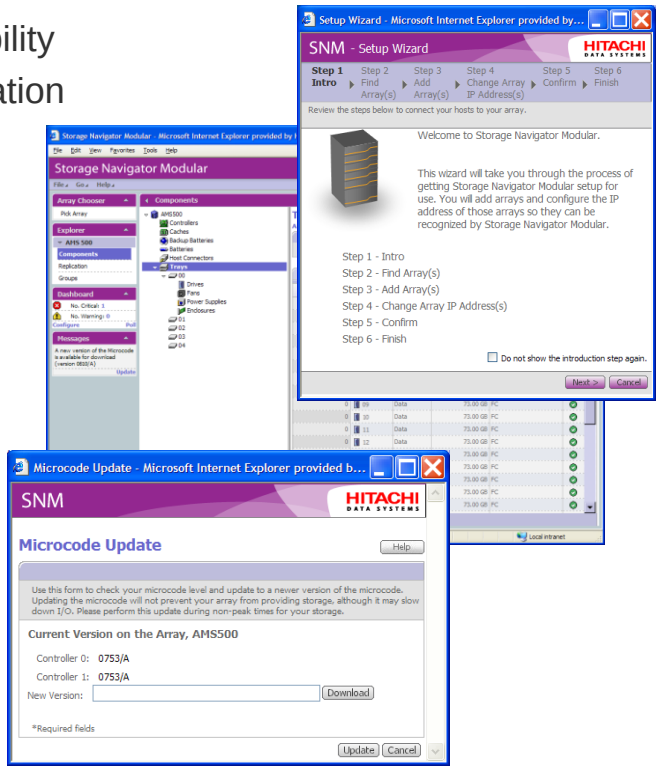
The management interface Storage Navigator Modular 2 is an easy to use GUI utility developed for the Adaptable Modular Storage 2000 and Simple Modular Storage product lines. It is based on HiRDB. It replaces Storage Navigator Modular.

There is also a CLI version of Storage Navigator Modular 2 but has not changed significantly from Storage Navigator Modular.

HDS Confidential: For distribution only to authorized parties.

Management Tool

- Focus on usability and maintainability
- Wizard-based setup and configuration
- Simple licensing
- Host attachment
- Adding LUs
- Replication setup and management
- Optional feature management
- Multi-array management
- Common GUI for Simple Modular Storage 100 and Adaptable Modular Storage 2000
- Matching CLI and API



26

System and Data Security

- Account Authentication
  - System management access based on one of three roles (account administrator, storage administrator, or auditor)
  - Unique login and password for each user
- System Audit Logging
  - Records all system changes and all power ups and downs
  - Can send changes to Syslog server
  - Enables IT managers to meet compliance requirements including auditable trails
- Data Retention Utility
  - Makes selected data non-erasable and non-rewriteable
  - Read access can be limited
- Common Criteria Evaluation Assurance Level 2 Certification in process

27

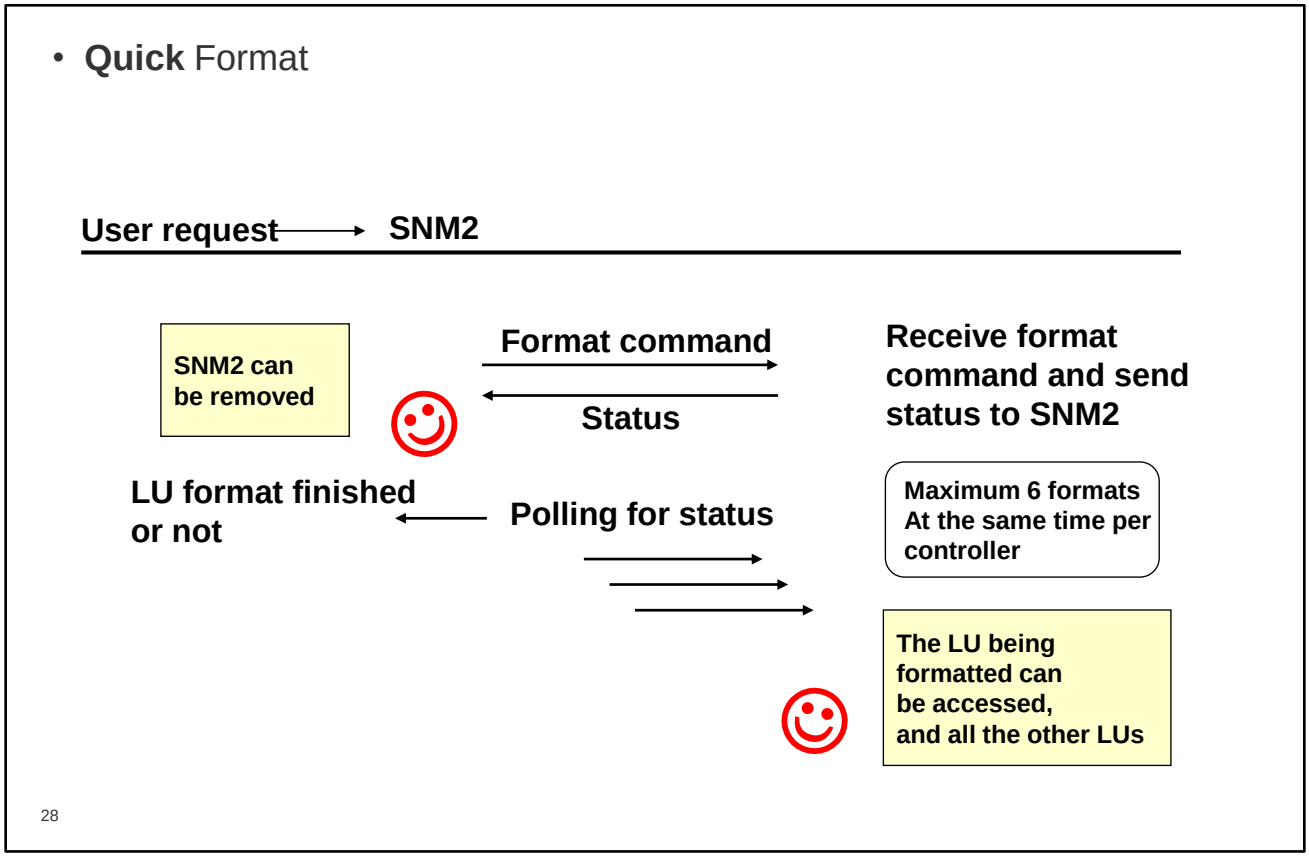
The Evaluation Assurance Level (EAL1 through EAL7) of an IT product or system is a numerical grade assigned following the completion of a Common Criteria security evaluation, an international standard in effect since 1999. The increasing assurance levels reflect added assurance requirements that must be met to achieve Common Criteria certification. The intent of the higher levels is to provide higher confidence that the system's principle security features are reliably implemented. The EAL level does not measure the security of the system itself, it simply states at what level the system was tested to see if it meets all the requirements of its Protection Profile.

EAL2: Structurally Tested

EAL2 requires the cooperation of the developer in terms of the delivery of design information and test results, but should not demand more effort on the part of the developer than is consistent with good commercial practice. As such, it should not require a substantially increased investment of cost or time. EAL2 is therefore applicable in those circumstances where developers or users require a low to moderate level of independently assured security in the absence of ready availability of the complete development record. Such a situation may arise when securing legacy systems.

HDS Confidential: For distribution only to authorized parties.

LUN Formatting

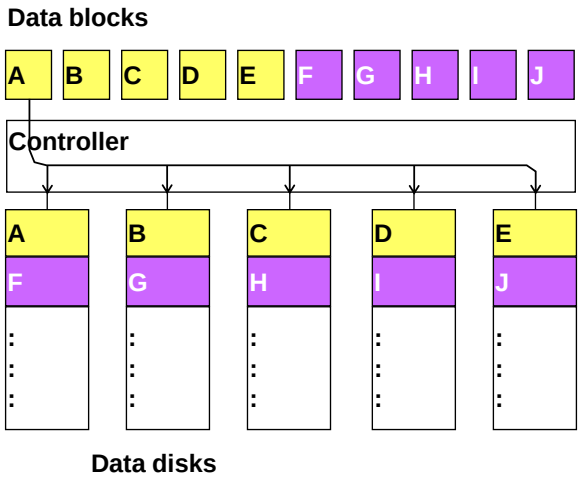


QuickFormat is the only formatting method available on Adaptable Modular Storage 2000 family. It allows issuing format commands for up to 512 LUNs. The queue is maintained storage system allowing the engineer's laptop to be disconnected. An Adaptable Modular Storage 2000 family, however, will only format a maximum of six LUNs concurrently per controller. The format is running in the background and the LUNs can almost immediately be accessed without the need of waiting for the full format to finish. (**Caution** – Do not POWER down the array before you check that the format is truly complete, otherwise the format will not finish).

In the diagram, **SNM2** stands for Storage Navigator Modular 2.

Supported RAID Levels

- RAID 0



Outline:
RAID 0 stripes the data across disk drives for higher throughput.

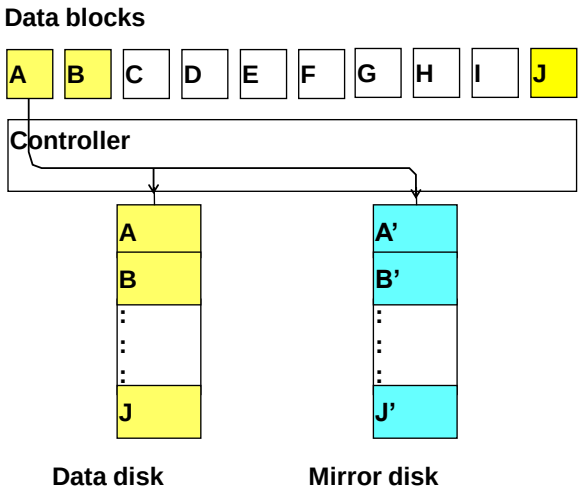
Pro:
Offers most available disk space to the user.

Con:
All data is lost in case of a disk failure.

29

HDS Confidential: For distribution only to authorized parties.

- RAID 1



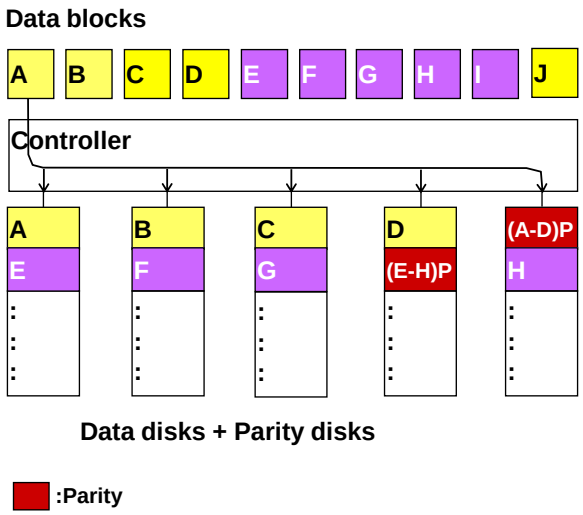
Outline:
RAID 1 mirrors the data.

Pro:
If a disk drive fails the data is not lost and the performance is not affected.

Con:
More expensive.

30

- RAID 5



Outline:
RAID 5 consists of three or more disk drives; one drive in round robin mode contains the parity.

Pro:
Striping offers higher reading throughput.

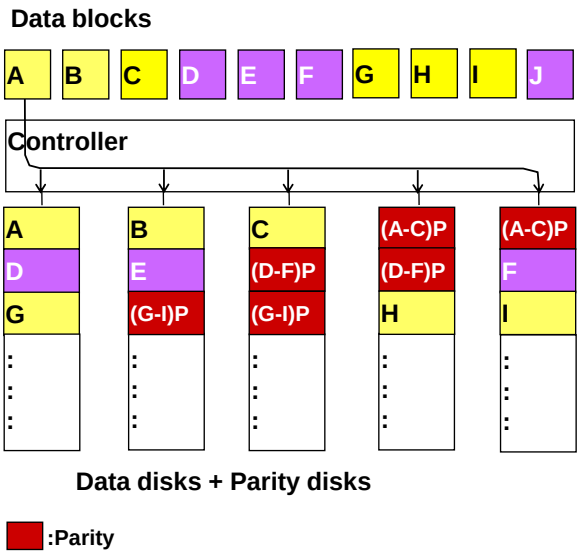
Con:
Lower performance on random writes and in case a drive fails.

31

At least three disks are required to implement RAID 5. RAID 5 will not sustain a double-disk failure and is more likely to occur with SATA drives.

HDS Confidential: For distribution only to authorized parties.

- RAID 6



Outline:
RAID 6 consists of four or more disk drives; two independent drives in round robin mode contain the parity.

Pro:
Allows recovery from a double disk failure.

Con:
Lower performance than RAID 5.

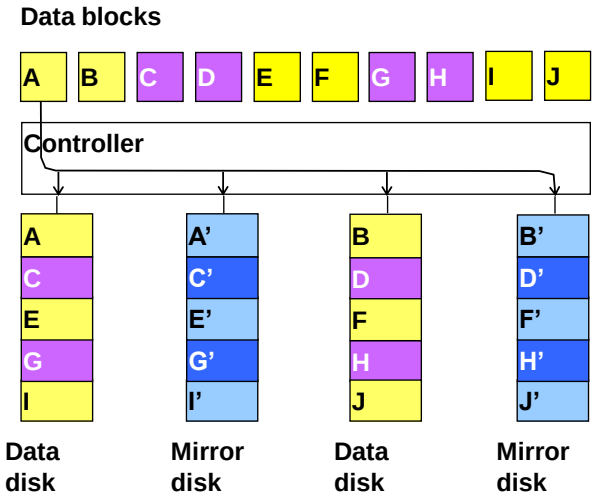
32

At least four disks are required to implement RAID 6. This configuration is very similar to RAID 5, with an additional parity block, allowing block level striping with two parity blocks. The advantages and disadvantages are the same as the RAID 5, except the additional parity disk protects the system against double-disk failure. This feature was implemented to ensure the reliability of the SATA drives.

Key value: Two parity drives allow a customer to lose up to two hard disk drives (HDDs) in a RAID group without losing data. RAID groups configured for RAID 6 are less likely to lose data in the event of a failure. RAID 6 performs nearly as well as RAID 5 for similar usable capacity. RAID 6 also gives the customer options as to when to rebuild the RAID group. When an HDD is damaged, the RAID group must be rebuilt immediately (since a second failure may result in lost data). During a rebuild, applications using the volumes on the damaged RAID group can expect severely diminished performance. A customer using RAID 6 may elect to wait to rebuild until a more opportune time (night or weekend) when applications will not require stringent performance.

HDD roaming allows the spare to become a part of the RAID group; no copy back is required saving rebuild time.

- RAID 1+0



Outline:
RAID 1+0 (four or more disk drives) is similar to RAID 1 but now the data is striped.

Pro:
Striping offers higher random write access performance compared to RAID 1, 5 or 6.

Con:
More expensive.

33

HDS Confidential: For distribution only to authorized parties.

Relative RAID Performance

- Relative Performance

  The table shows the performance levels for different RAID configurations relative to RAID 1+0 (2D+2P), which is assumed to be 100% for reads and writes.

  The illustration is not meant to be a recommendation or imply that one RAID Group is better than another RAID Group, but to illustrate how performance of a RAID Group should be matched to the type of performance that an application needs.

RAID Level	Random Read, Sequential Read	Sequential Write	Random Write
RAID 1+0 (2D+2P)	100%	100%	100%
RAID 5 (3D+1P)	100%	150%	50%
RAID 5 (7D+1P)	200%	350%	100%
RAID 6 (6D+2P)	200%	300%	66.7%
Note	Proportional to the number of disks	Proportional to the number of data disks	See notes

34

Note: The values listed in the table were not collected using an AMS 2000, but using an earlier AMS 500. The table is meant to illustrate that different applications may require different performance characteristics from its RAID Group.

This table shows the performance levels for different RAID configurations relative to RAID1+0 (2D+2P), which is assumed to be 100% for reads and writes.

The second column shows that the random read and sequential write performance is proportional to the number of disks, because the disks can be accessed simultaneously.

With sequential writes, there are no reads involved, as with random writes, therefore the performance is proportional to the number of data disks.

Random writes. The reason for the performance difference between RAID6 (6D+2P) and RAID5 (7D+1P) is that RAID6 (6D+2P) must process 1.5 times (see below) more disk I/Os than RAID5 (7D+1P), therefore the random write performance in RAID6 (6D+2P) is 33% lower than with RAID5 (7D+1P).

The number of disk I/Os in RAID5 random writes: Four (old data/old parity reads, new data/new parity writes).

The number of disk I/Os in RAID6 random writes: Six (old data/old parity (P)/old parity (Q) reads, new data/new parity (P)/new parity (Q) writes).

Supported RAID Levels

- Overview

RAID Level	Supported Range for Parity Groups Models 2100, 2300, and 2500
0	2D to 16D (RAID 0 only supported on SAS HDDs)
1	1D + 1P
5	2D + 1P to 15D + 1P
6	2D + 2P to 28D + 2P
1+0	2D + 2P to 8D + 8P

35

The above table shows the minimum and maximum size of what is called a Parity Group which is the building block of RAID Group. There is no restriction as to how many Parity Groups can be in a RAID Group. All Parity Groups within a RAID Group can be the same size.

HDS Confidential: For distribution only to authorized parties.

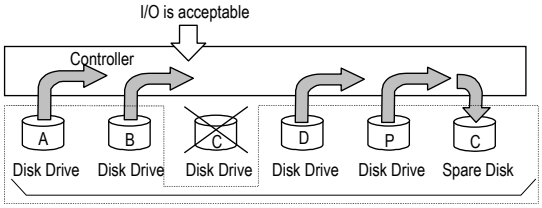
Sparing Out

- Two methods exist that support the *sparing-out* of RAID Group data:
  - **Correction copy**
    - Occurs when a drive in a RAID Group fails and a compatible spare drive exists.
    - Data is reconstructed on the spare drive.
  - **Dynamic sparing**
    - Occurs if the **online verify** process (built-in diagnostic) determines that the number of errors has exceeded the specified threshold of a disk in a RAID Group.
    - Data is moved to the spare disk (a much faster process than reconstruction).

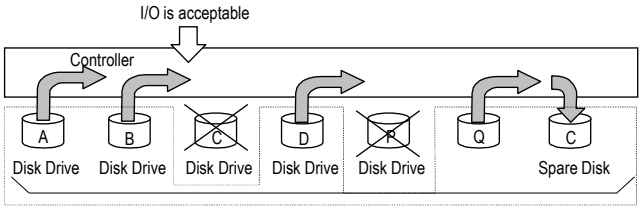
36

An AMS 2000 system supports up to 30 Spare Drives. Any disk can be configured as a spare drive. Although there are certain rules such as size and SAS and SATA mix.

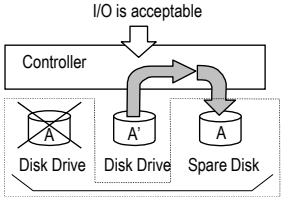
Correction Copy



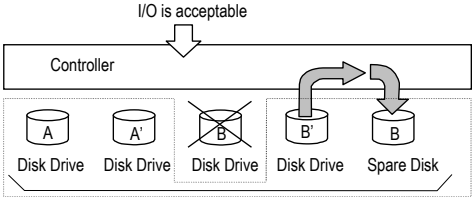
RAID 5 (in the case of 4D+1P)



RAID 6 (in the case of 4D+2P)



RAID 1



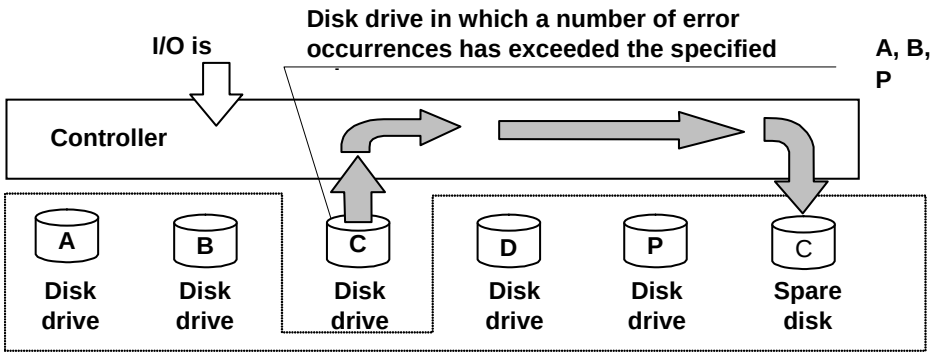
RAID 1+0 (in the case of 2D+2D)

A, B, C,… : data
P : parity data
A',B', C',… : mirror data

37

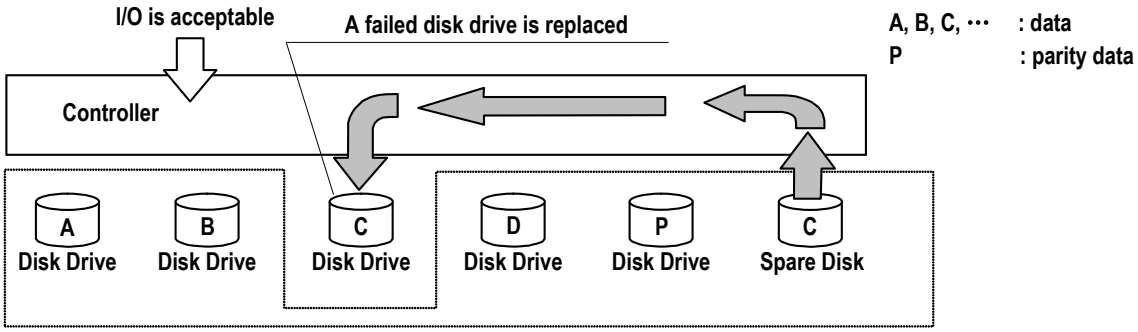
HDS Confidential: For distribution only to authorized parties.

Dynamic Sparing



38

Copy Back



Note: The Copy Back operation is controlled by the **Spare Drive Operation Mode** parameter. This parameter has two settings:

Variable (the default) = **No copy back**

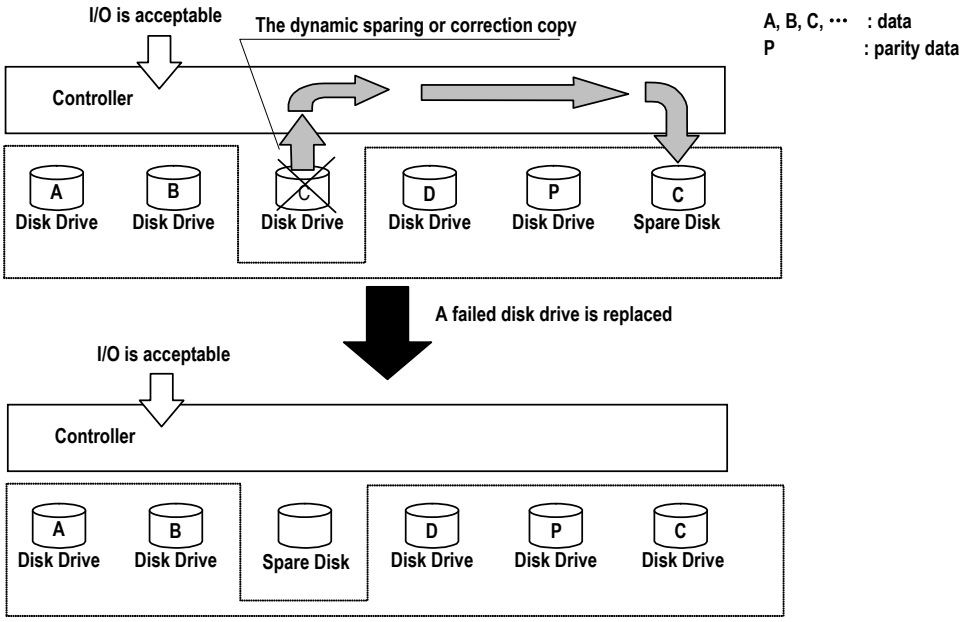
Fixed = **Copy Back**

39

In case a copy is required, the system simply copies the data from the spare disk that was being used to the new replacement disk C in this example.

HDS Confidential: For distribution only to authorized parties.

No Copy Back



Note: The Copy Back operation is controlled by the **Spare Drive Operation Mode** parameter. This parameter has two settings:

Variable (the default) = **No copy back**

Fixed = **Copy Back**

40

If the HDD Roaming rules allow this process, after a dynamic sparing process or correction copy, the spare disk will become a permanent member of the RAID Group and a copy back will not occur.

Sparing

- Rules for Sparing

Priority item	Description	Degree of priority
Drive type	It is judged whether the drive that requires restoration is SAS or S-ATA and only a drive of the same type is used.	High
Model name	A Spare Disk Drive with the same model name as the failed Disk Drive is used preferentially.	
Drive capacity	A spare disk with the same capacity as a RAID group to be restored is used preferentially. When a drive with the same capacity does not exist, a drive with the most approximate and larger capacity is selected preferentially.	
Rotational number	A Spare Disk Drive with the same capacity as the failed Disk Drive is used preferentially.	
The loop of the drive interface	A Spare Disk in the same loop as the failed Disk Drive is used preferentially.	Low

Note: In case the spare is not the same specification (speed, size) as the failed drive, a Copy Back will be performed after failed drive replacement.

41

HDS Confidential: For distribution only to authorized parties.

Hi-Track Monitor

- Installed on a Windows/Solaris/Linux workstation which is on the same network as the AMS storage system.
- Monitors the operation of the HDS storage at all times
  - Collects hardware status and error data
  - Transmits this data through a modem to the Hitachi Data Systems Support Center
    - The Support Center analyzes the data and implements corrective action as needed
- Hi-Track Monitor FTP Transport is available and greatly reduces time transferring dumps to the Support Center.

42

Hi-Track Monitor is a software utility program that is optionally installed on a server provided by the customer.

The Hi-Track Monitor software monitors the operation of the AMS storage systems at all times, collects hardware status and error data, and transmits this data through a modem to the Hitachi Data Systems Support Center. The Support Center analyzes the data and implements corrective action as needed. In the unlikely event of a component failure, Hi-Track Monitor service calls the Hitachi Data Systems Support Center immediately to report the failure, without requiring any action on the part of the user. Hi-Track tool enables most problems to be identified and fixed prior to actual failure and the advanced redundancy features enable the system to remain operational even if one or more components fail.

Hi-Track Monitor enables error analysis, case creation, and error/information data browsing functions. When Hi-Track Monitor is installed and the storage system is configured to allow it, Hitachi support staff can remotely connect to the storage system.

Note: Hi-Track Monitor does not have access to any user data stored on the storage system. The Hi-Track Monitor requires a dedicated RJ-11 analog phone line.

- Performs Hi-Track functions of error monitoring, configuration data acquisition, and remote SVP mode.
- Transport to the Hi-Track Center can be through either HTTPS or FTP (SSL or standard) via the public internet or via dialup modem.
- Can send email alerts to customer (user definable destinations).

43

HDS Confidential: For distribution only to authorized parties.

4. Active-Active I/O Architecture

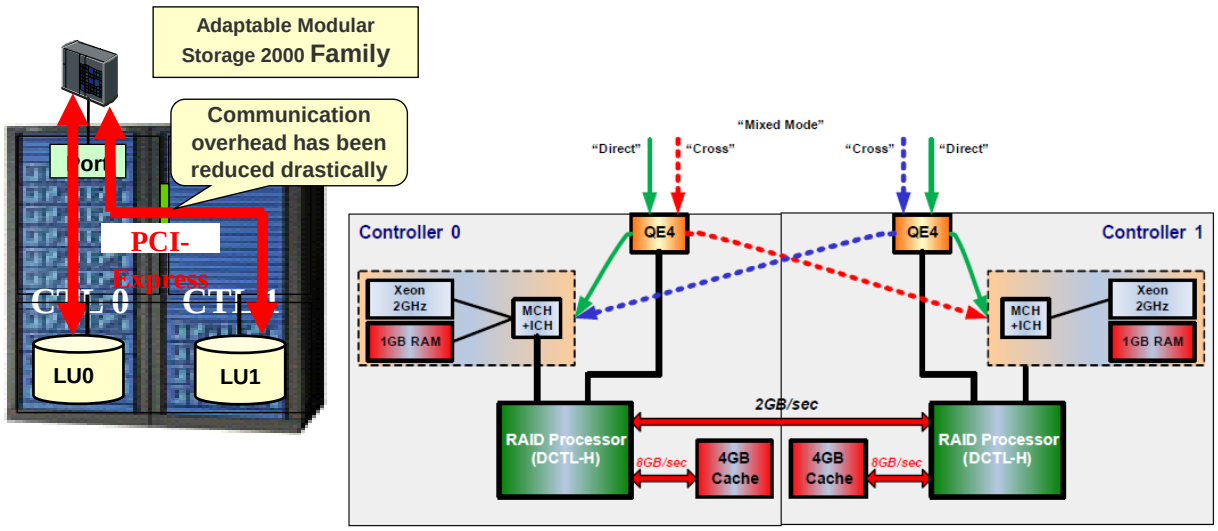
Module Objectives

- Upon completion of this module, the learner should be able to:
  - Explain active-active controller architecture
  - Identify the benefits in the Adaptable Modular Storage 2000 family in terms of hardware performance

2

Overview

- Adaptable Modular Storage 2000 family



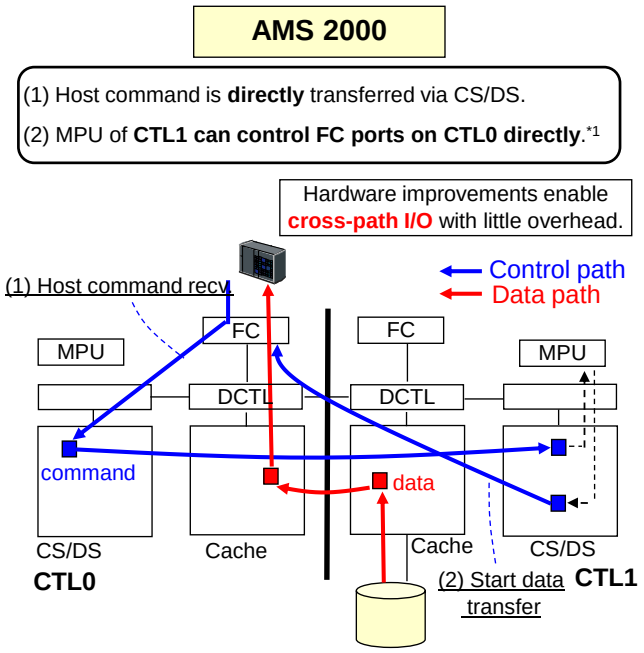
In a traditional controller architecture — referred to here as an active/passive architecture — one controller is actively available to process I/O requests while the other is idle in standby mode, ready to take over I/O activity if the active primary controller fails or is taken offline. The shortcoming of this mode has mainly to do with the fact that all the user data residing in LUNs must be assigned manually to a primary controller; this process is tedious and time consuming, and involves human effort, which in turn involves the possibility of human error. In the event of hot spots or utilization imbalance issues (such as a Microsoft® Exchange change of I/O pattern or VMware Virtual Servers moving by effect of VMotion), the administrator must manually reassign the workload to the secondary controller. If this does not happen, the application will experience worst performance, which will have a direct impact on the business. Sometimes this process involves a system being rebooted, which is a disruptive operation.

Hitachi AMS 2000 family products offer symmetric active/active architecture. Instead of having an active and a passive controller, both controllers are active, able to work simultaneously, and able to dynamically and automatically assign the access paths from the back of the controller to the LUN. All LUNs are accessible regardless of the physical port or the server from which the access is requested. Utilization rates of each controller are monitored and can be automatically brought into balance whenever necessary. This dramatically simplifies management and offers a number of administrative benefits.

HDS Confidential: For distribution only to authorized parties.

Internal Transaction

- Enables the MPU of one controller to access the cache, disk drives, and other devices of the other controller directly.
- Cross-path I/O is greatly improved.



4

In this diagram:

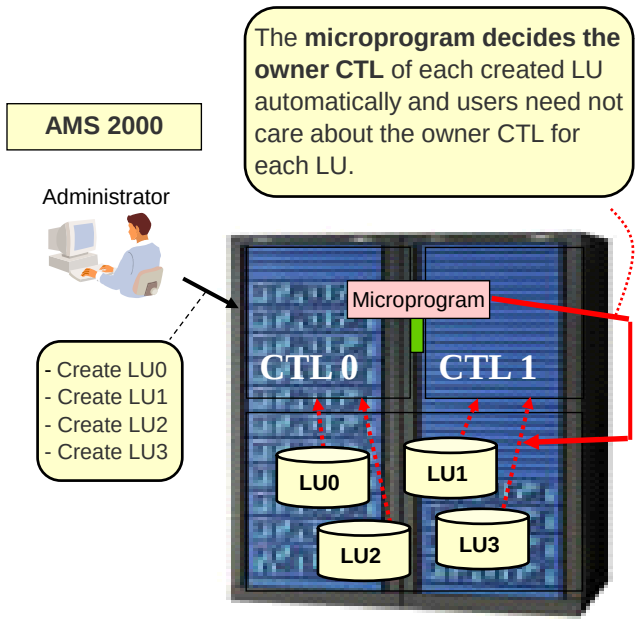
AMS/WMS stands for previous Hitachi modular storage systems.

CS/DS stands for local memory of the controller.

MPU – Microprocessor.

LU Ownership

- Owner controller of LUs or operations is not an issue.
- The firmware assigns ownership.



5

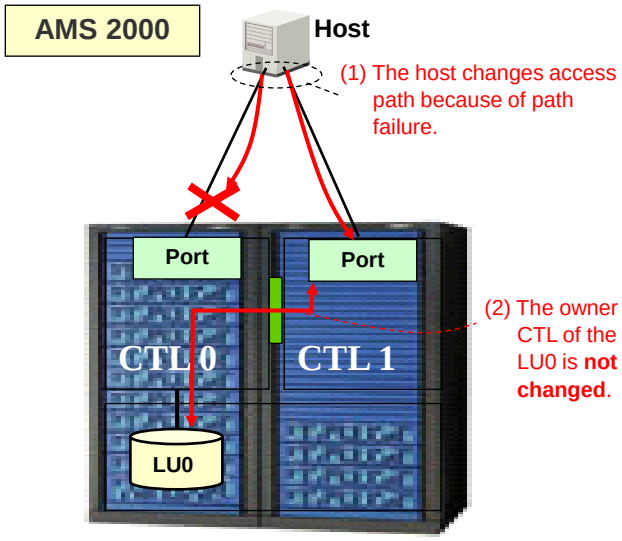
The user need not consider which controller should be the owner when he creates each LU or for all operations of the array.

Therefore, the non-owner controller of the target LU may receive I/O commands from hosts. But it is not a problem because such commands are processed by high performance cross-path.

The manual setting mode (like previous modular systems) is also available in the Storage Navigator Modular 2 GUI.

HDS Confidential: For distribution only to authorized parties.

- LU ownership is **not changed** because of path failure

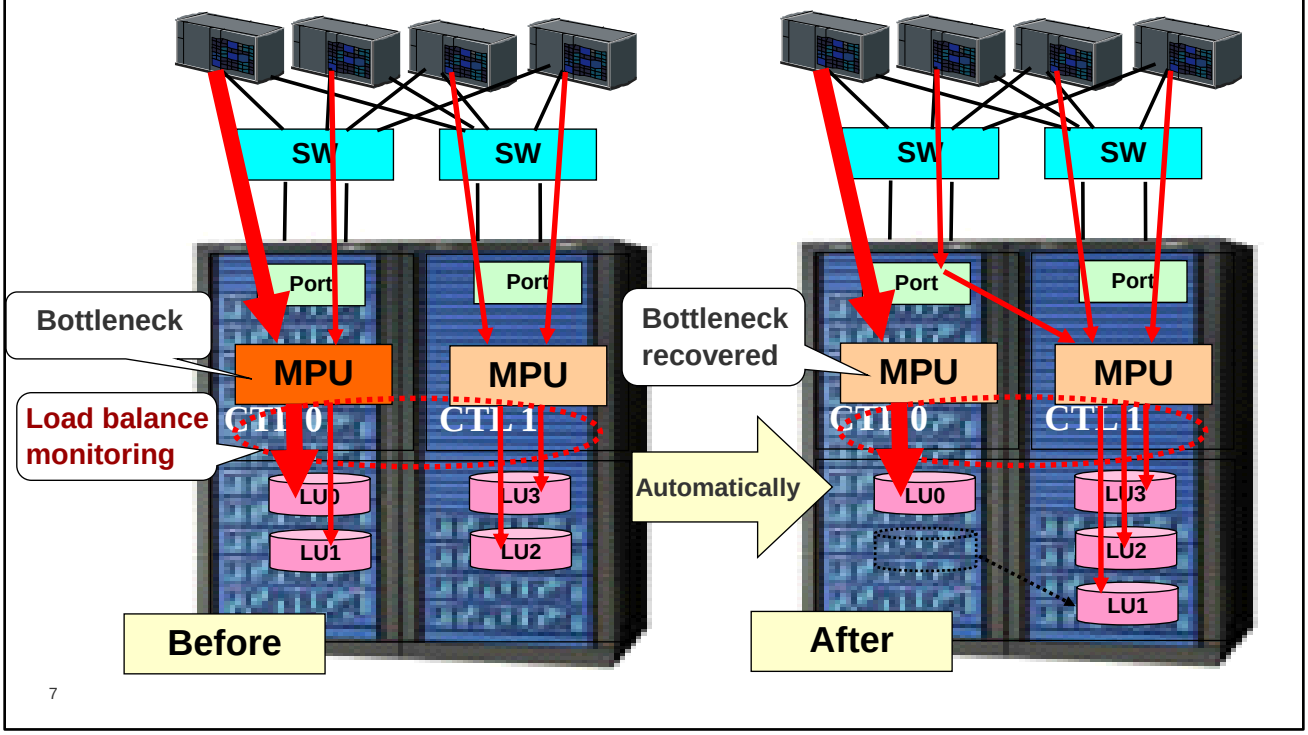


6

Hosts can send commands to storage via any path of any controller for the purpose of path load balancing. This is possible because cross-path I/O is high performance and ownership of each LU is stable.

Controller Load Balancing

- The owner controller of each LU may be changed automatically for the purpose of load balancing of processors on two controllers.

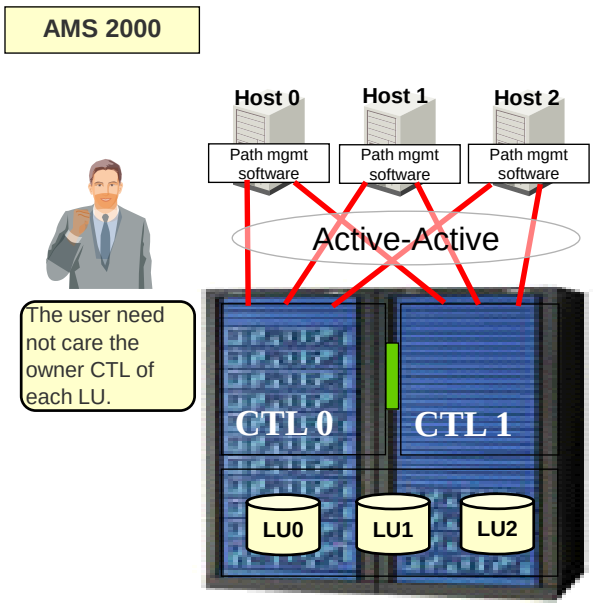


The load balancing function can be enabled and disabled. It must be disabled when using the Cache Partition Manager so as not to change the partition setting for each LU automatically.

HDS Confidential: For distribution only to authorized parties.

Performance Design

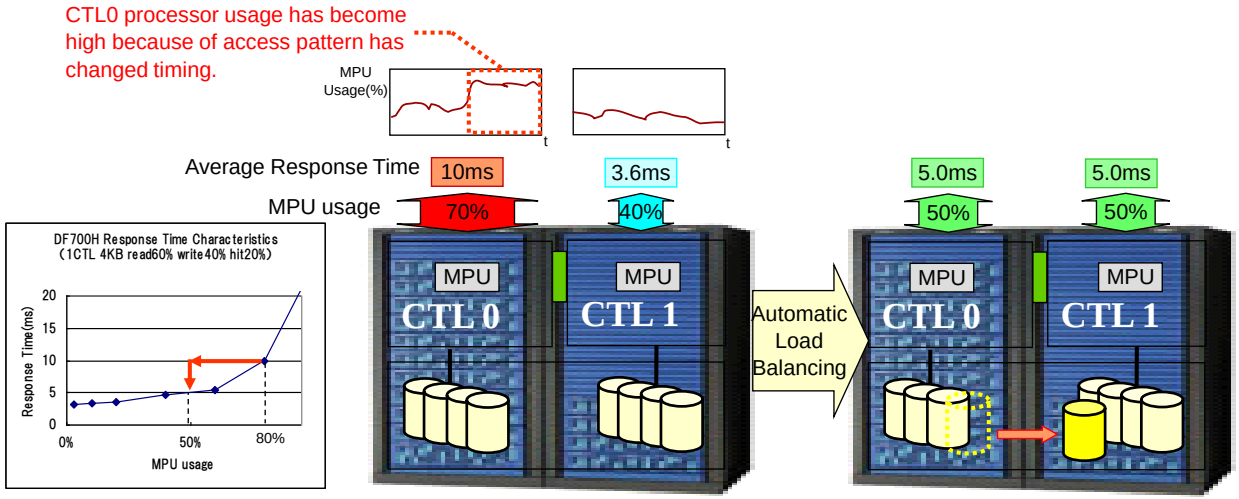
- Load balancing automatic with active-active mode



8

The user does not need to consider the load balance of controllers and ports when doing the performance design. The user must just set the path management software of all hosts as the **Active-Active** mode, and then the load of controllers is automatically balanced.

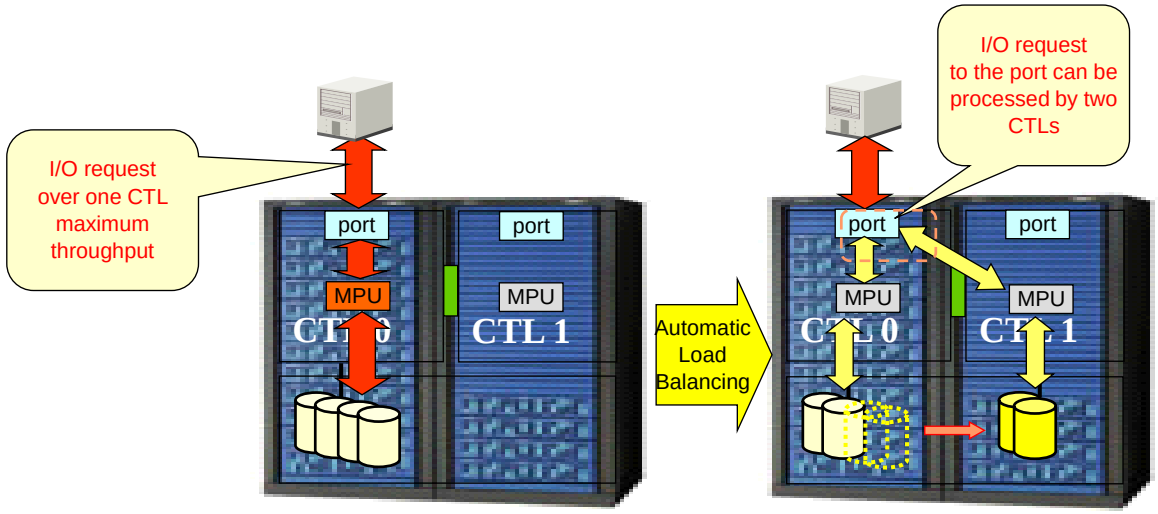
- Automatic Optimization for Performance
  - When the rate of access for the processor of one controller becomes high, the average of response time for the controller becomes long.
  - If the processor usage of the two controllers is balanced by automatic load balancing, the response time remains good.



9

HDS Confidential: For distribution only to authorized parties.

- Performance Aggregation
  - I/O processed on two controllers
  - Port performance can equal maximum for *both* controllers

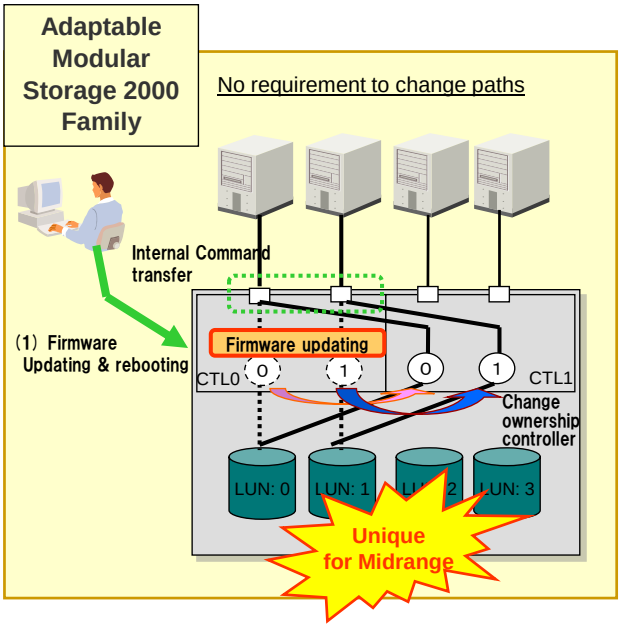


10

I/O requests to the port can be processed on two controllers by using a "cross-path" mechanism. So the port performance can exceed the maximum performance of a single controller and it can be expanded to the maximum performance for both controllers.

Microcode Updates

- Benefits
  - Non-disruptive firmware updates are easily and quickly accomplished.
  - Firmware can be updated without interrupting I/O.



11

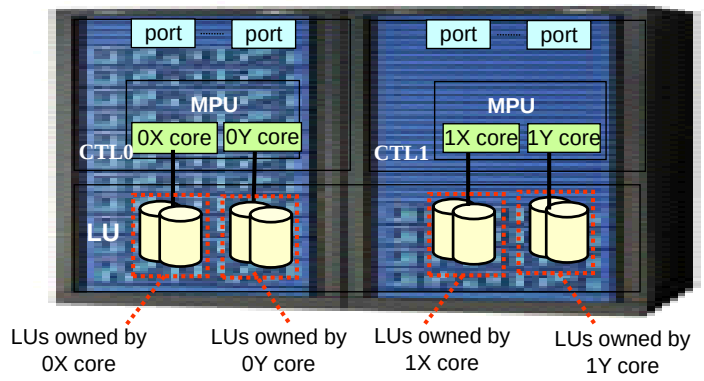
For firmware updates:

- No need to use host path management software.
- No need to change path from firmware-updating CTL to other CTLs.

HDS Confidential: For distribution only to authorized parties.

Dual Core Support

- Adaptable Modular Storage 2500 contains dual-core MPU on each controller.
- Designed with little shared resources.
- Each LU is owned by either one of two MPU cores, and each MPU core can access any LUs owned by it without any exclusion control.



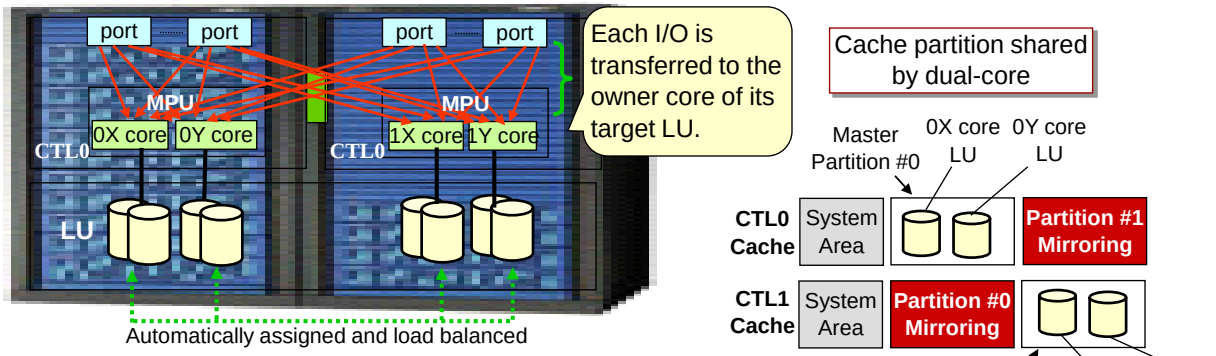
12

Options 3 and 4 for models 2100 or 2300, and 1,2,3,4 for model 2500 – Adaptable Modular Storage 2500 supports the dual core MPUs.

Automatically assigned to each LU based on round-robin order:

1. CTL1core1
2. CTL0core1
3. CTL1core0
4. CTL0core0 (1. 2.: only dual core)

- Users are unaware the owner core of each LU.
- Any MPU core can process I/Os received by any port on the storage system with same mechanism as active-active.



- Difference between dual core and dual controller:
  - Cache area is not divided into two areas for two cores externally. Though each core has independent cache directory, any free segments in the cache area of the controller are shared by two cores.
  - Each core is not represented as an independent component, that is, there is no status such as “one core partial stoppage.”

13

HDS Confidential: For distribution only to authorized parties.

Active-Active Summary

Items	Active-Active (AMS 2000)
Support way	Default function
Setting the LU ownership	Automatically assigned to each LU based on **round-robin order**: 1. CTL1 core1 2. CTL0 core1 3. CTL1 core0 4. CTL0 core0 (1. 2. : dual core only)
Exceptional commands for cross-path I/O	None
TrueCopy or TrueCopy Extended	- **Any** P-VOL and S-VOL can be a pair, so there is no need to change owner controller when a pair is created.

14

Items	Active-Active (AMS 2000)
ShadowImage or Modular Volume Migration	The owner controller of the S-VOL is automatically changed to the same controller as P-VOL.
Copy-on-Write or TrueCopy Extended	Though P-VOL and POOL must be in the same owner controller internally, the owner of the P-VOL is automatically changed to the same controller as the POOL.

HDS Confidential: For distribution only to authorized parties.

5. Software Feature Overview

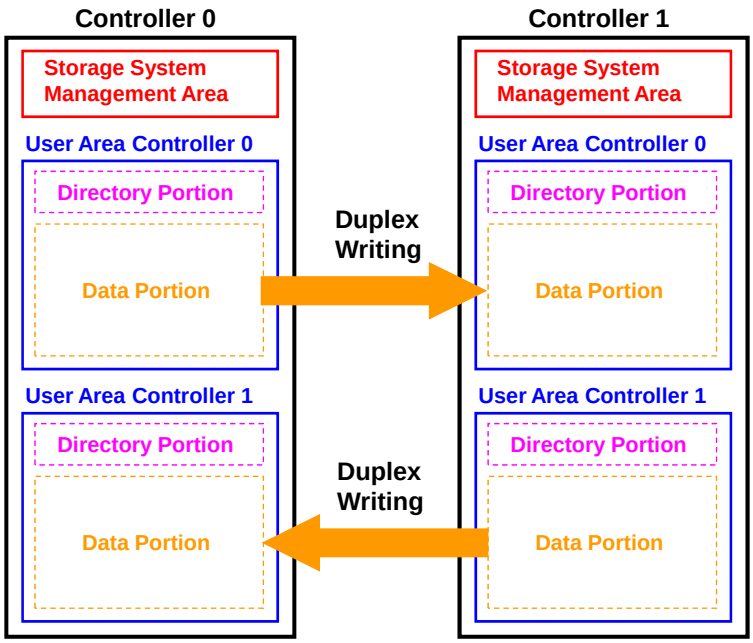
Module Objectives

- Upon completion of this module, the learner should be able to:
  - List the software features that are included with the Adaptable Modular Storage (AMS) 2000 family
  - Describe which program products can only be accessed with Storage Navigator Modular 2 Advanced Setting applet
  - Describe, enable, and configure Cache Partition Manager feature
  - Load a specific internal LUN into cache as a cache-resident LUN using Cache Residency Manager feature
  - State the purpose of the Performance Monitor feature and identify the types of performance data that can be collected
  - Describe the LU Grow & Shrink feature
  - Describe the Online RAID Group Expansion feature
  - Describe the AMS2000 Hitachi Dynamic Provisioning feature

2

Cache Memory Configuration

- Cache Memory Configuration

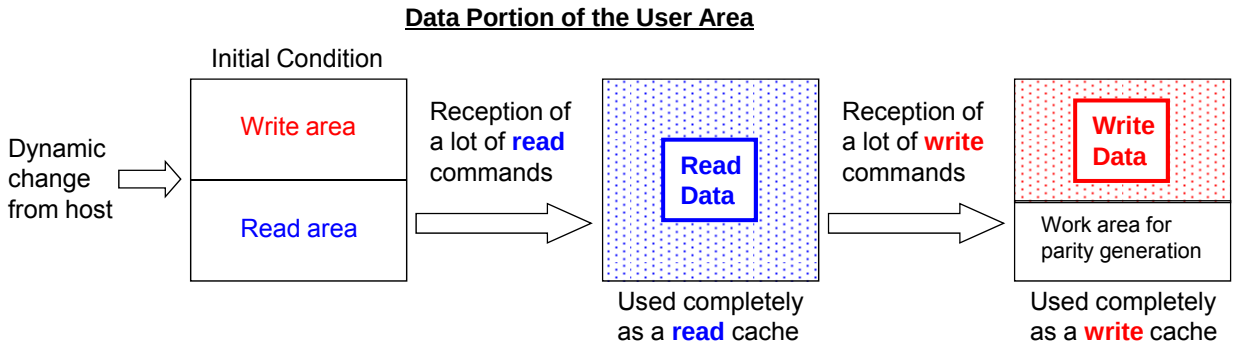


3

HDS Confidential: For distribution only to authorized parties.

Basic Operation of Cache Control

- Dynamic Optimizing is performed on the cache per the following:
  - **Read/Write cache area assignment** – The read and write cache areas are not fixed but are dynamically assigned according to the type of I/O from the host.
  - **Destaging algorithm** – This is selected automatically according to the write pattern occurring from the host.
  - **Staging algorithm** – This is selected automatically through a study of the read commands from the host.



4

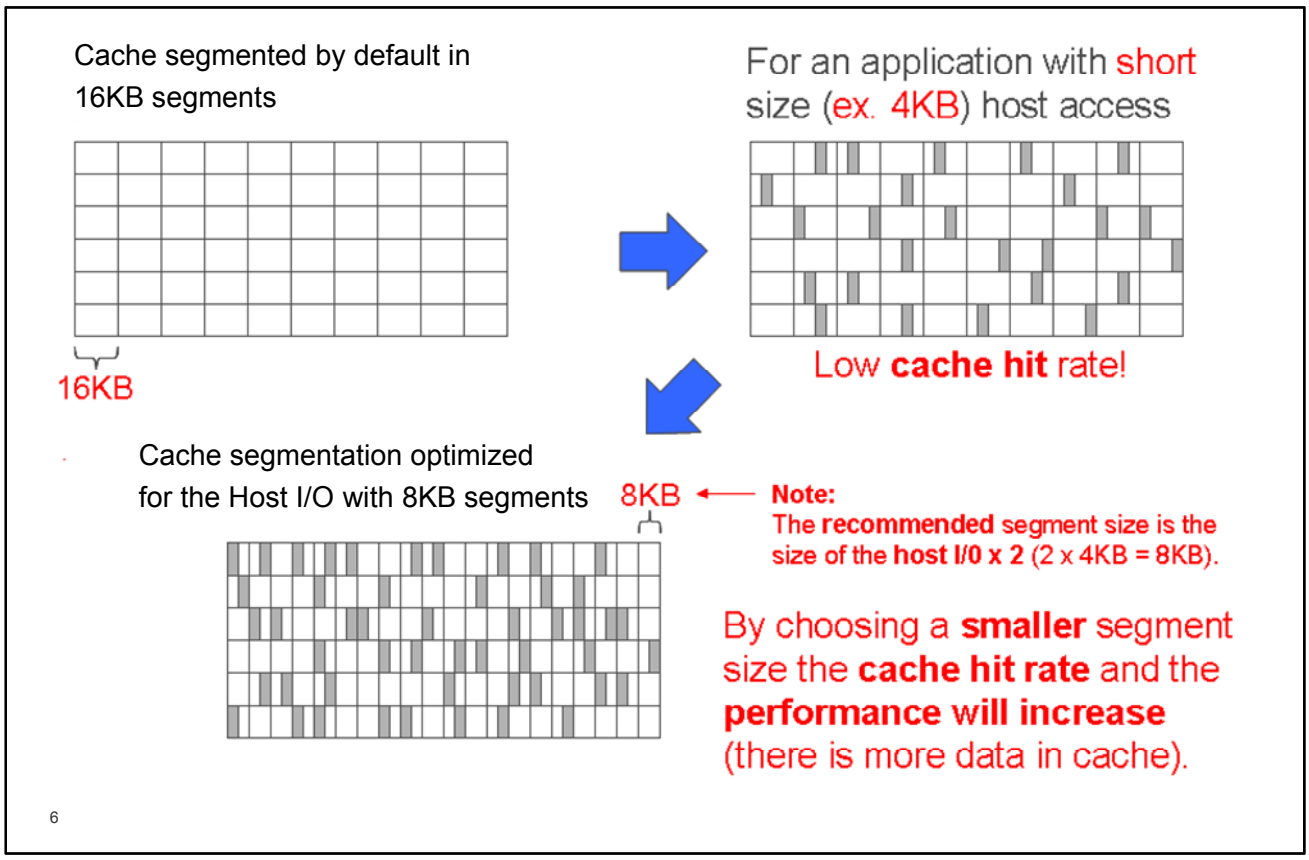
Cache Partition Manager Feature

- Cache Partition Manager (CPM) allows for **performance tuning** by allowing the segregation of workloads within the system.
- It includes the following:
  - **Selectable segment size** – Customize the cache segment size for a user application
  - **Partitioning of cache memory** – Separate workloads by dividing cache into individually managed, multiple partitions
    - A partition can then be customized to best match the I/O characteristics of its assigned LUs
  - **Selectable stripe size** – To increase performance by customizing the disk access size

5

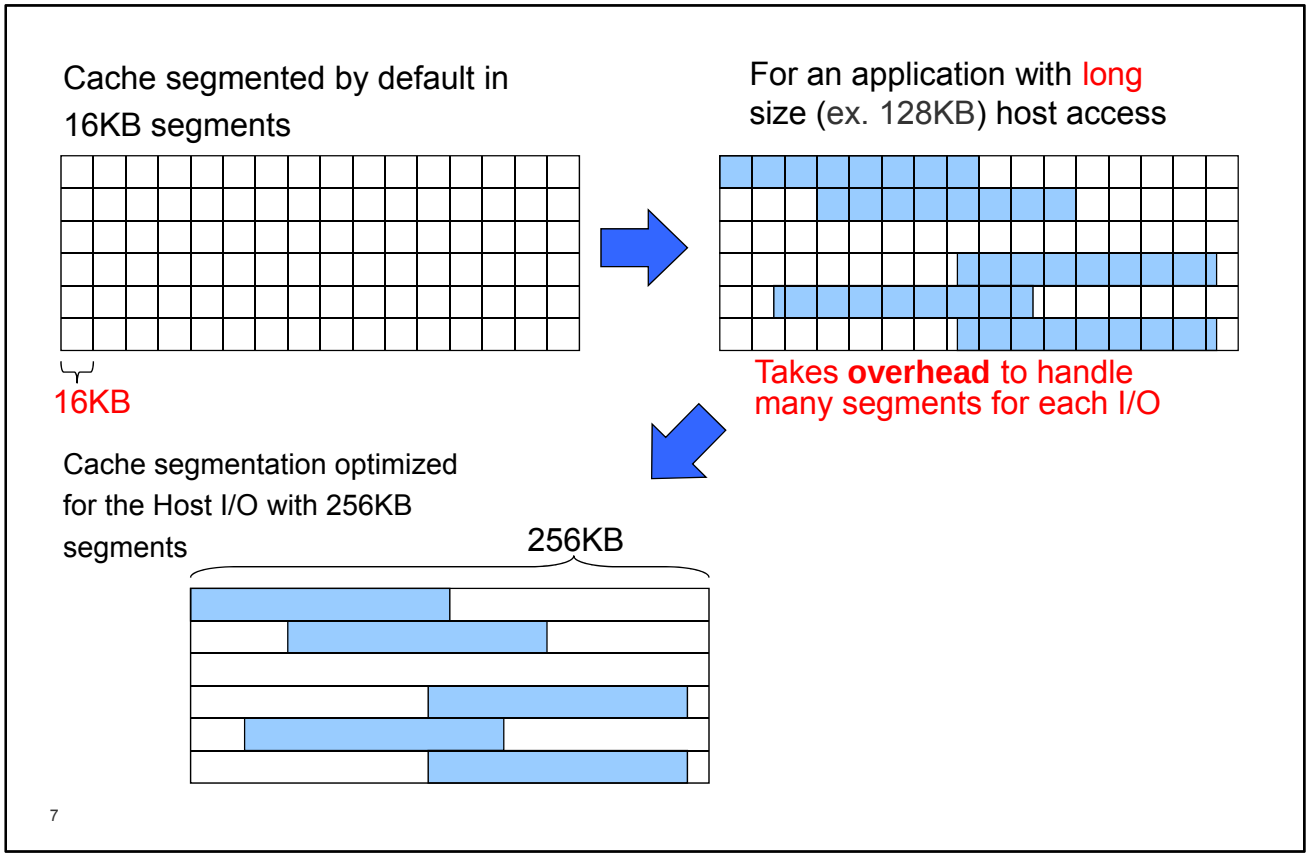
HDS Confidential: For distribution only to authorized parties.

Advantage of Selectable Segment Size – Small I/O



Recommended segment size is the host I/O size times two. Setting the segment size in this example to 4KB, indicates that up to two segments will be used for cache processing overhead.

Advantage of Selectable Segment Size – Large I/O

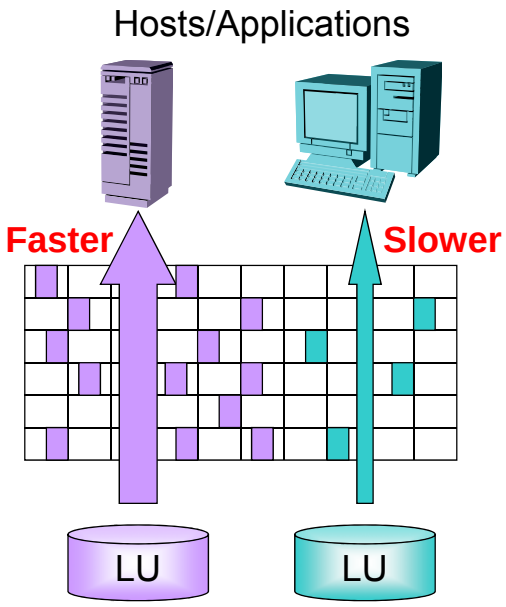


HDS Confidential: For distribution only to authorized parties.

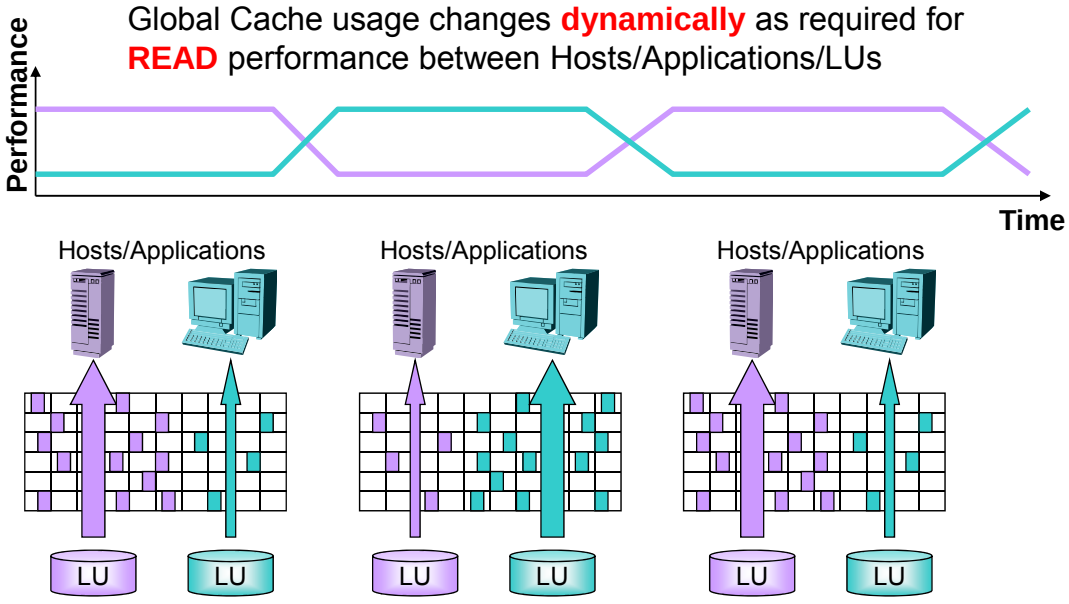
Advantage of Global Cache

- For **READ** access:
  - Cache memory is used on demand by multiple applications.
  - Therefore:
    - **Faster hosts** (applications) do not have a negative effect on **slower hosts** (applications).
    - A **faster LU** does not have a negative effect on a **slower LU.**

Model 2000 system without Cache Partition Manager or different cache partitions



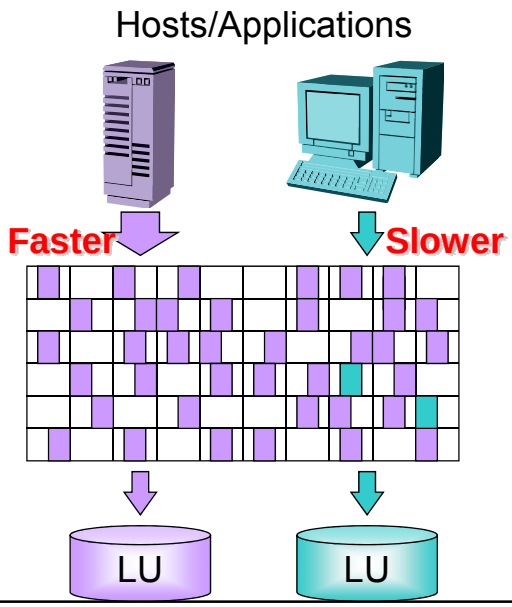
- Storage system without Cache Partition Manager or different cache partitions



9

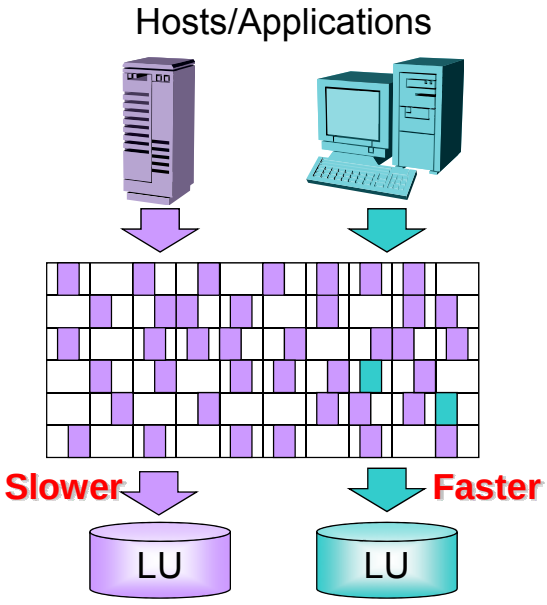
Disadvantage of Global Cache

- Storage system without Cache Partition Manager or different cache partitions



- High-intensity **WRITE** access from faster/slower hosts (applications):
  - Faster hosts (applications) use **much more** cache than slower hosts (applications), and fill almost all the cache memory.
  - Faster hosts (applications) have a **negative** effect on slower hosts (applications).

- Storage system without Cache Partition Manager or different cache partitions



- High-intensity **WRITE** access to faster/slower **LUs**:
  - Hosts (applications) uses same amount of cache memory, but data for **slower LU** remains **longer** in cache memory.
  - Slower LUs have **negative** effects on faster LUs.
  - Faster/slower LUs can be caused by:
    - SAS/SATA HDD
    - RAID level
    - The Parity Group

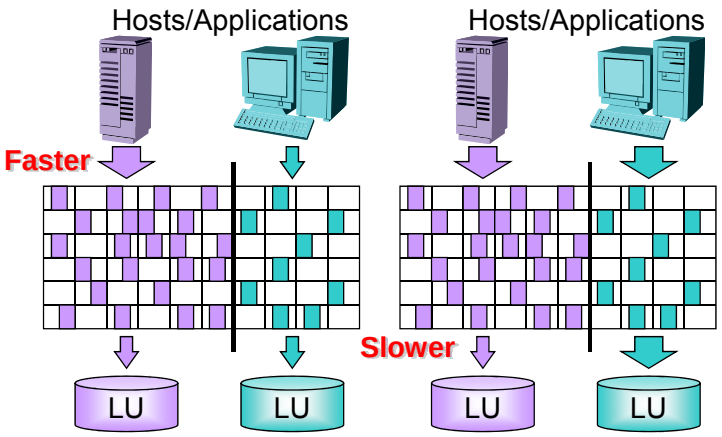
11

HDS Confidential: For distribution only to authorized parties.

Advantage of Partitioned Cache

- Selectable from partitioned cache and global cache

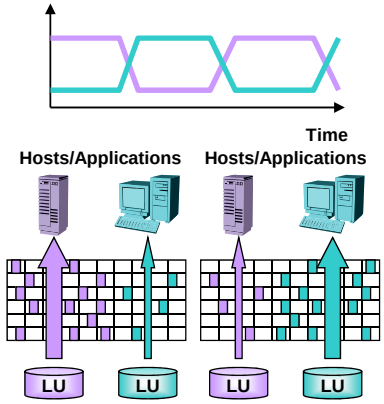
Partitioned Cache



Negative effects between faster/slower hosts, applications, and LUs will decrease.

Global Cache
is selectable for temporary changes in performance

or

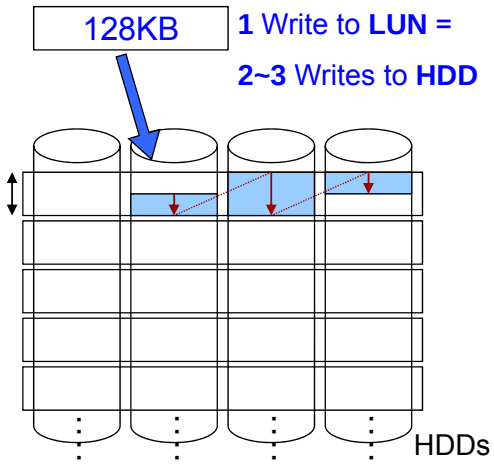


12

Configuring the cache for partitions is a static adjustment that will not dynamically change afterwards.

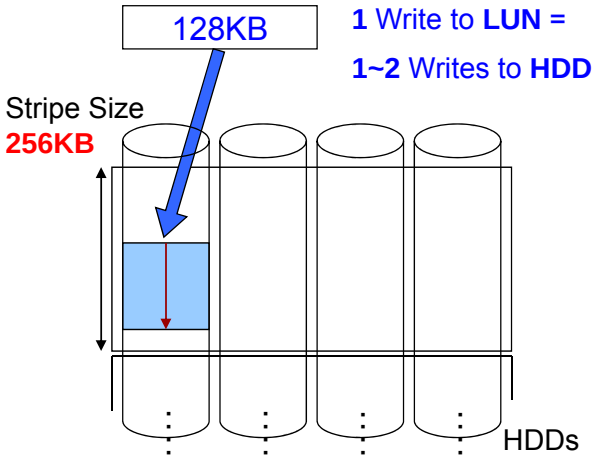
Advantage of Selectable Stripe Size

HDD usage with a LU Stripe Size 64KB



- High throughput with concurrent I/Os to HDDs
- Good for applications with transaction I/Os (Database systems)

HDD usage with the I/O optimized LU Stripe Size 256KB



- Lower overhead because of less HDD I/Os
- Good for applications with sustained I/Os

13

By selecting the most appropriate stripe size, the number of HDD I/Os can be brought back to the minimum which will improve the performance.

HDS Confidential: For distribution only to authorized parties.

Partitioning Cache

- Cache can be divided into partitions that can be exclusively used by assigned LUNs.
  - Maximum number of partitions:
    - Model 2100: 16
    - Model 2300: 16
    - Model 2500: 32
  - Partition 0/1 are the master partitions (fixed at 16KB only).
  - Partition 2 to n have selectable-size segments of: 4, 8,16, 64, 256, and 512K.
  - Partition sizes are flexible (each partition has a certain minimum).

14

Although proper use of the Cache Partition Manager can contribute to improving an application's performance, an incorrect configuration can easily achieve the opposite effect.

One partition can be used by one or more LUNs.

Cache Partition Manager

- Operational Considerations
  - Cache Partition Manager settings can be changed from the **Advanced Settings** panel in Storage Navigator Modular 2.
    - Create New Partition
    - Delete Partition
    - Change Size
  - When making changes to cache partitions, you must restart the storage system.
  - A LUN can be assigned to a cache partition from Storage Navigator Modular 2.

15

HDS Confidential: For distribution only to authorized parties.

Cache Residency Manager Overview

- User-specified LUN resident in the Cache Memory
  - 100% hit ratio for the resident LUN
  - Improve system throughput
    - Apply to a LUN that contains data being accessed frequently

Required conditions for Cache Residency Manager Feature

Item	Specification
Controller Configuration	Dual Controller
RAID level	RAID1+0, 5, 6
Max number of LUNs	1 per Controller

16

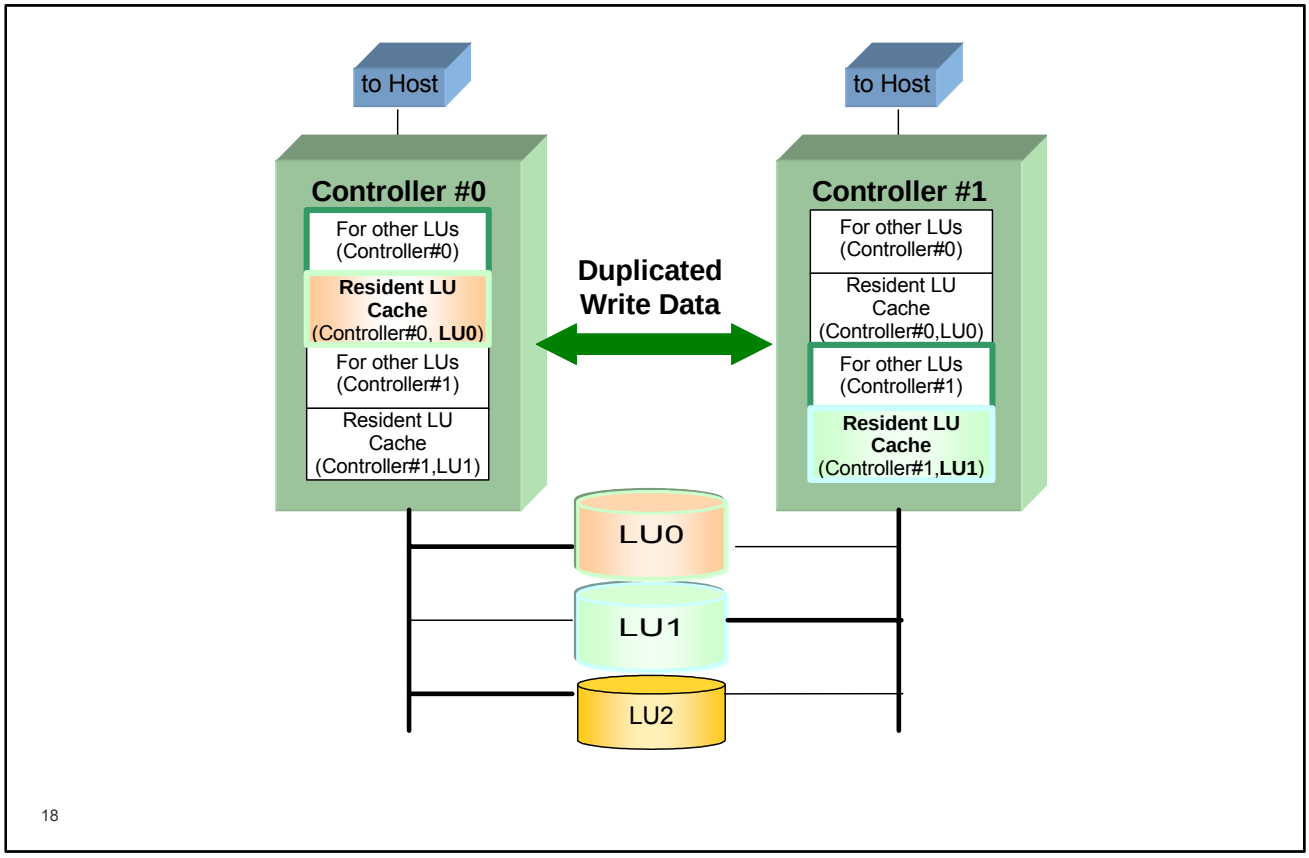
Installing Cache Residency Manager Feature

- Cache Residency Manager feature must be installed or uninstalled using a software license key.
- The storage system must be rebooted in order for the Cache Residency Manager changes to take effect, including installing, uninstalling, enabling, or disabling.

17

HDS Confidential: For distribution only to authorized parties.

Functionality



Conditions that Terminate Cache Residency Manager

No	Conditions	Remarks
1	When the setting of Cache Residency Manager is cleared	All of these conditions are caused by the operator.
2	When Cache Residency Manager is disabled or uninstalled	
3	When the LUN used for Cache Residency Manager is deleted or the RAID group in which the LUN is created is deleted	

19

HDS Confidential: For distribution only to authorized parties.

Maximum LUN Size for Cache Residency Manager

- The maximum size of a LUN that is used for the Cache Residency Manager depends on the capacity of the installed cache memory.
- Cache Residency Manager feature uses part of the cache memory.

20

Performance Monitor Feature Overview

- The Performance Monitor feature enables the user to collect and analyze performance information from the following functional areas:
  - Port information
  - RAID Group and Logical Unit information
  - Cache information
  - Processor information
  - Drive information
  - Drive operating information
  - Back end information
- Performance Monitor presents the information in **chart** or **table format** (.txt file)
  - Chart gives a real-time two dimensional view of the system in graph format.
  - Table format gives snapshots over a specified period of time.
    - Set an interval from 1 minute to 23 hours 59 minutes
    - Set number of repetitions from 1 to 20,000
    - Snapshots can be concatenated into a single, compressed into a ZIP file
- Requires a license

21

HDS Confidential: For distribution only to authorized parties.

Monitored Statistics

Functional Area	Measurement Item
Port Information **Logical Unit Information**	Command Count
	Command Hit Count
	Hit Rate
	I/O Rate
	Transfer Rate
	Initiator Command Count
	Initiator Command Transfer Size
	Initiator Command Time
	Initiator Command I/O Rate
	Initiator Command Transfer Rate
RAID Group Information	Command Count
	Command Hit Count
	Hit Rate
	I/O Rate
	Transfer Rate

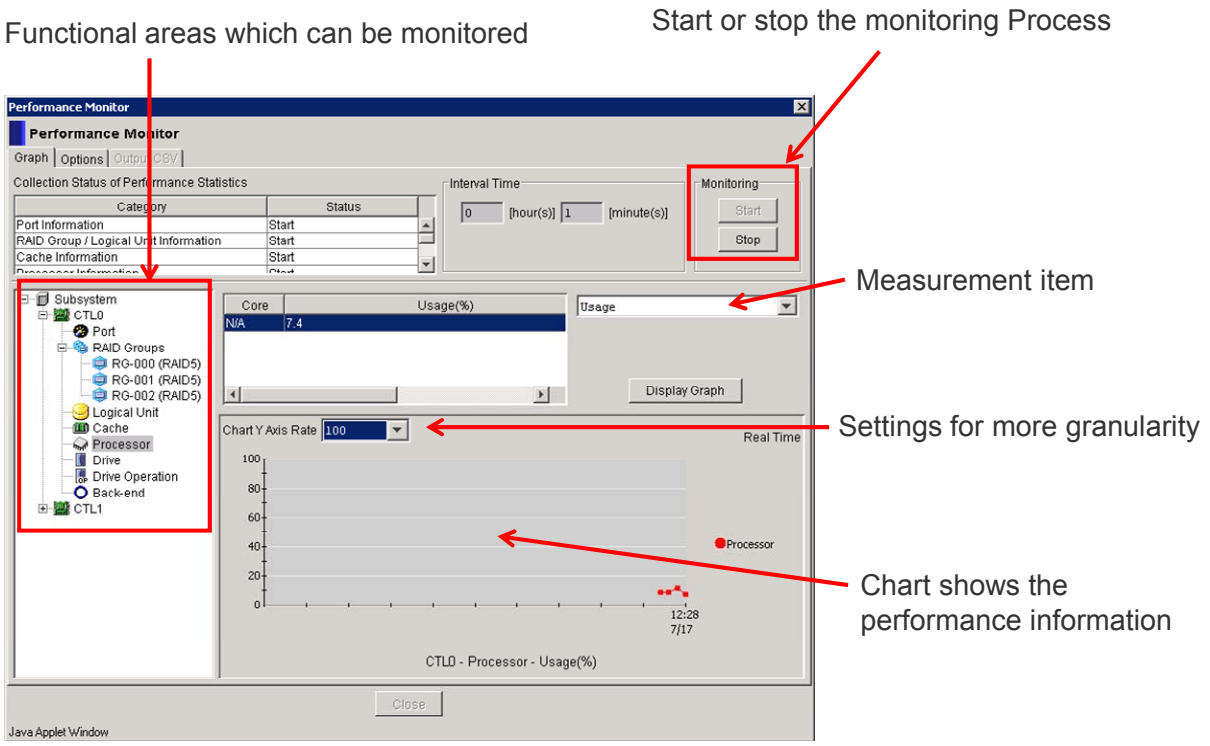
22

Functional Area	Measurement Item
Cache Information	Cache Write Pending Rate
	Cache Usage Rate
Processor Information	Processor Operating Rate
Drive Information **Back end Information**	Command Count
	I/O Rate
	Transfer Rate
	Online Verify Command Count
Drive Operation Information	Drive Operating Rate
	Drive Loading Information

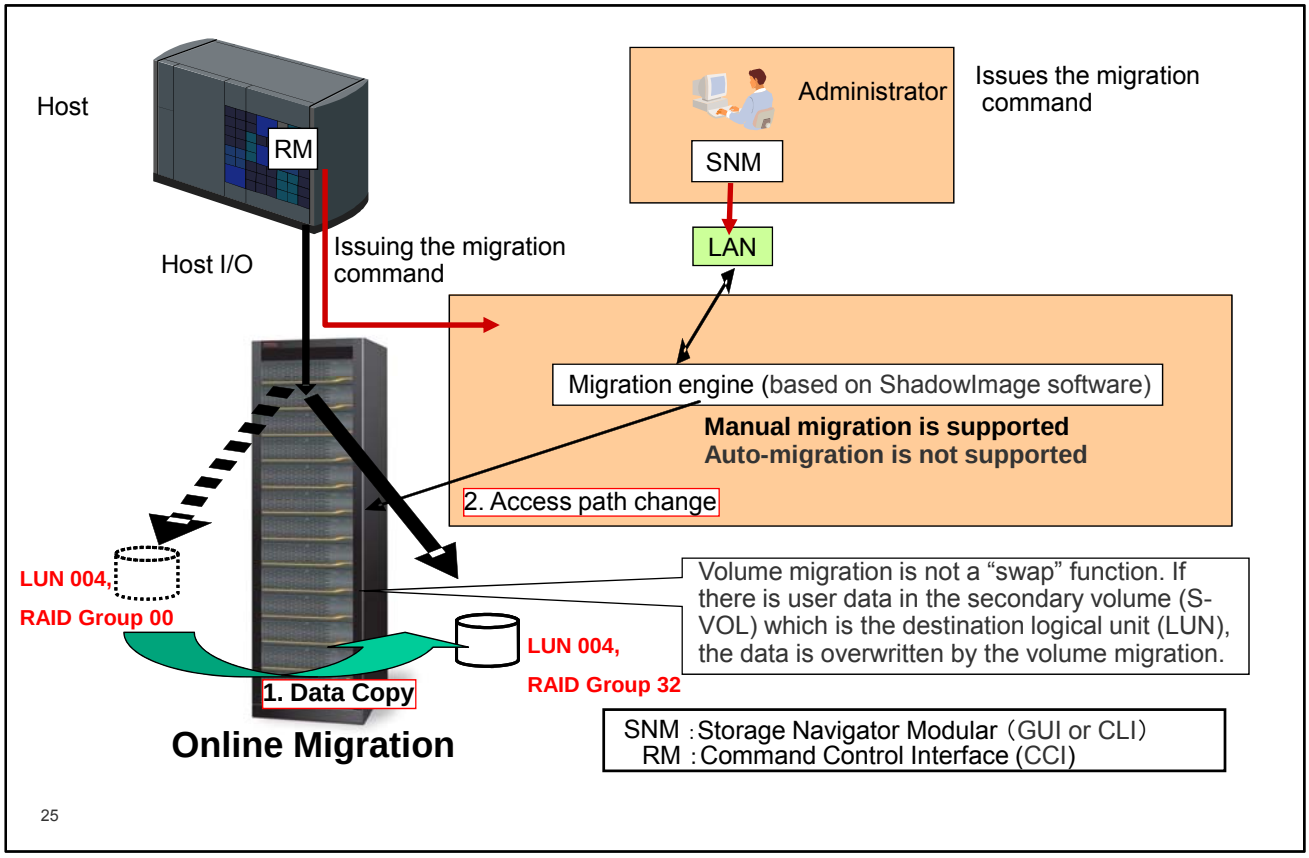
23

HDS Confidential: For distribution only to authorized parties.

Performance Monitor Window

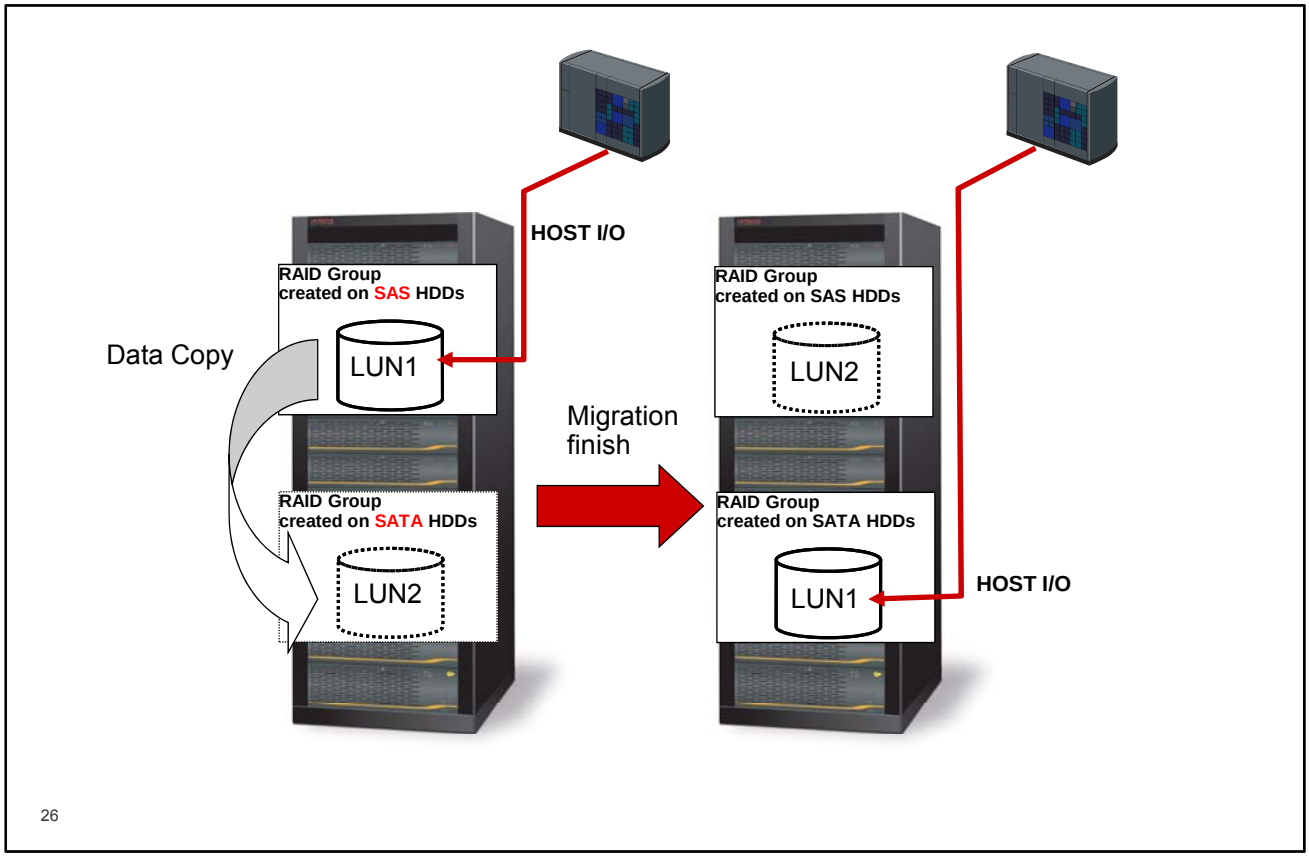


Overview of Modular Volume Migration

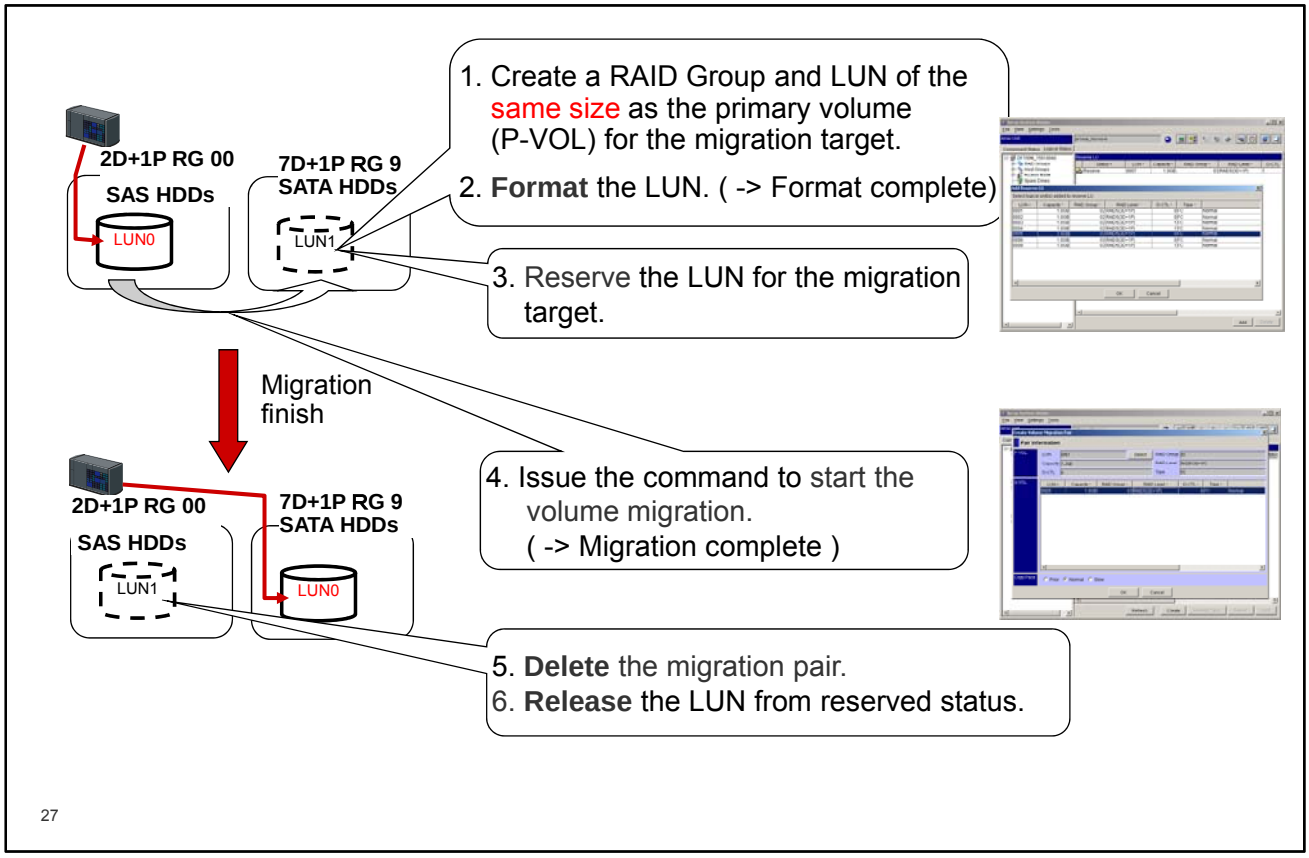


HDS Confidential: For distribution only to authorized parties.

Migration from SAS Drives to SATA Drives



Migrating Volumes for Performance

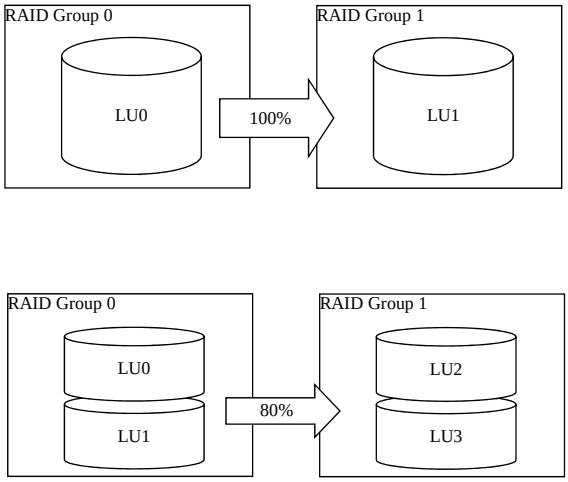


HDS Confidential: For distribution only to authorized parties.

Performance Considerations

- The recommended Copy Pace is **Normal**, but when the host I/O load is heavy, performance may deteriorate considerably. In that case, select **Slow** to maintain performance levels.
  - For example, the LUN is migrated from a SAS disk drive to a SAS disk drive:
    1. Random Read I/O declines from 65% to 45% when the Copy Pace is **Normal**.
    2. Random read I/O declines from 80% to 60% when the Copy Pace is **Slow**.
- Not only does the RAID structure of the P-VOL affect the performance of the host I/O, but it also affects the performance of the S-VOL.
  - For example, when a LUN is migrated from a SAS disk drive to a SATA disk drive, the random write I/O declines from 60% to 40% when the Copy Pace is **Normal**.

- Do not execute a volume migration to multiple LUNs in the same RAID Group at the same time.
  - For example, when the LUN is migrated from SAS disk drives to SAS disk drives, copy performance declines to 80%.



29

Change LU Capacity — LU Expansion

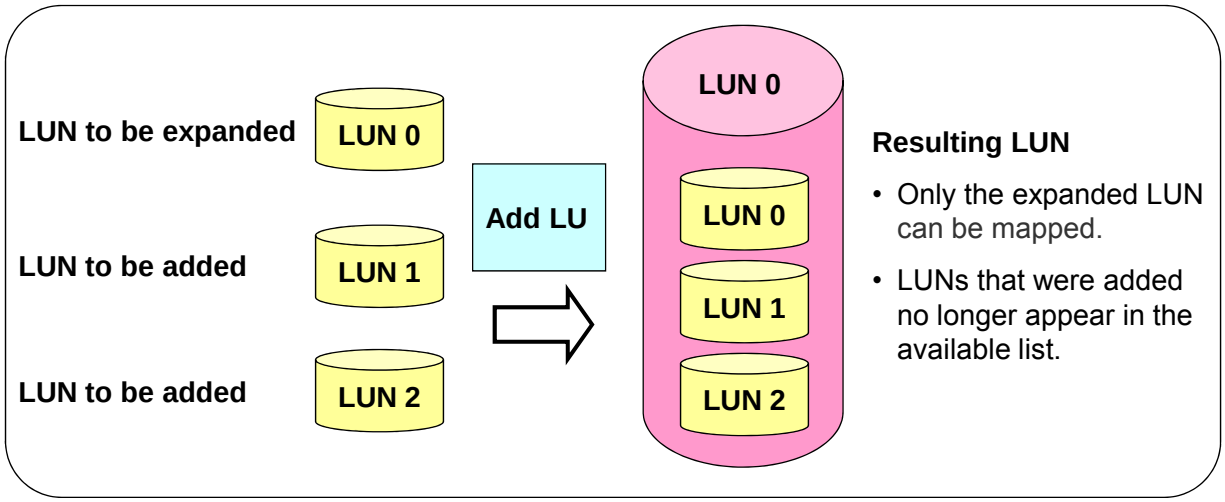
- Relative to LU expansion by unifying LUs, **Change LU Capacity** is comprised of the following functions:
  - **Add LUs** – Add additional LUNs to an existing unified LUN
  - **Separate Last LU** – Separate the last LUN from the unified LUN
  - **Separate All LUs** – Separate all the LUNs from the unified LUN
- Change LU Capacity is part of the LUN Manager function.
- Change LU Capacity operations can be performed using Storage Navigator Modular 2 GUI or CLI.
- This function used to be called Logical Unit Size Expansion (LUSE).

30

HDS Confidential: For distribution only to authorized parties.

Overview of Adding LUs together

- Two or more LUNs are unified (concatenated) into one larger LU
- LUNs can come from different RAID Groups
- 128 LUNs maximum
- RAID 5/RAID 6/RAID 1/RAID 1+0 (RAID 0 is not supported)
- LUN concatenation up to 60TB



31

The LUSE function refers to the LUNs in the process as:

The top LUN is the **Main LUN.**

The other LUNs are called **Sub-LUNs.**

General Conditions and Restrictions

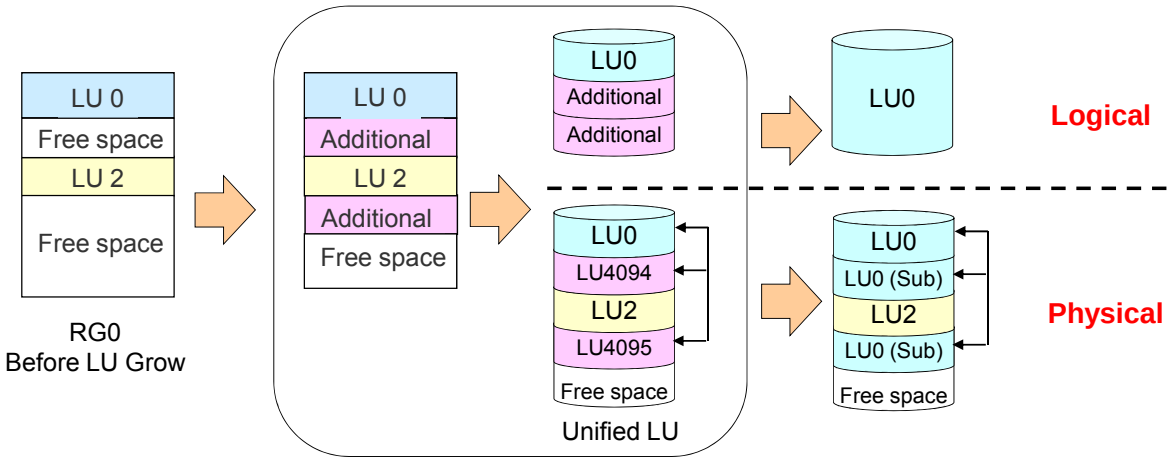
- LUNs must be in normal status.
- You can only use LUNs in formatted state.
- You cannot expand a LU if drive restoration is in progress on the affected LU.
- You cannot expand the following LU types:
  - ShadowImage replication and Copy-on-write SnapShot pairs
  - TrueCopy remote replication or TrueCopy Extended Distance pairs
  - LUs or reserved LUs for Modular Volume Migration
  - LUs in which Cache Residency Manager is set
  - LUs that are set for a command device
  - Differential Management Logical Units (DMLUs)
  - LUs that are registered in the data pool
  - LUs that are in the RAID Group during a RAID Group expansion

32

HDS Confidential: For distribution only to authorized parties.

Change LU Capacity — Grow LU Capacity from Free Space

- Grow (add) a specific capacity to an existing LU from disk free space.
- Free space has to be in the same RAID Group.

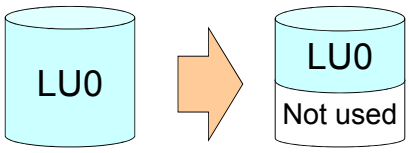


Grow **LU 0** by specifying a growth capacity and SNM2 automatically selects the space.

33

Change LU Capacity: LU Shrink

- This function reduces the capacity from an LU.



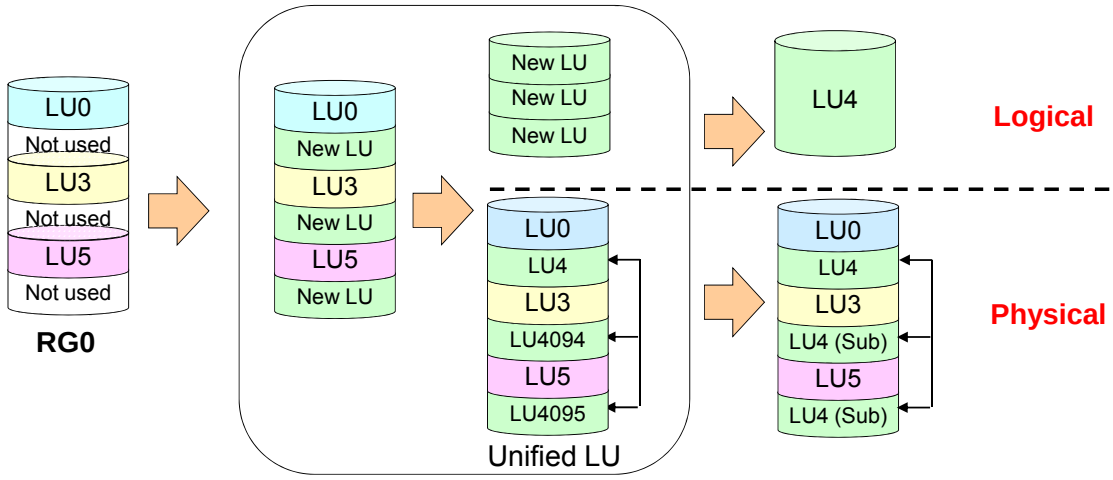
Note:
The Host OS must support volume shrinking if you use LU Shrink, and you must execute the host OS side volume shrink first, then execute the storage array side LU Shrink.

34

HDS Confidential: For distribution only to authorized parties.

Manual LU Creation Out of Free Space

- Create a new LU by manually selecting free space segments that exist in the same RAID Group.



Create a new **LU 4** from existing free space segments.

35

General Specifications

No	Functions	Specifications
1	Supporting Model	2100, 2300, 2500
2	LU Grow	Available
3	LU Shrink	Available
4	Easy LU creating	Available
5	Supported RAID Level	RAID 5/RAID 6/RAID 1/RAID 1+0 (RAID 0 is not supported)
6	Allocation of LUN for Add LU	Allocating of the last LUN which is not used automatically
7	Max. Number of Internal Unified LU	128 LUN
8	RG that the LU Grow can be executed	LU expansion will be done by only using unused space in the same RG.
9	Maximum capacity for shrinking the LU	Less than the original LU capacity (not deleting the LU)

36

HDS Confidential: For distribution only to authorized parties.

Online RAID Group expansion

- The online RAID Group expansion feature allows you to expand an existing RAID Group in a storage system.
- To expand an RAID Group you can use existing unused hard disk drives or a newly installed one.
- All RAID levels are supported for expansion, but you can not change the RAID level, for example from RAID 5 to RAID 6 or from RAID 5 to RAID 1+0.
- You cannot shrink a RAID Group.
- You can remove the expansion on a Waiting expansion RAID Group.
- You can remove the expansion on an Expanding RAID Group[1].

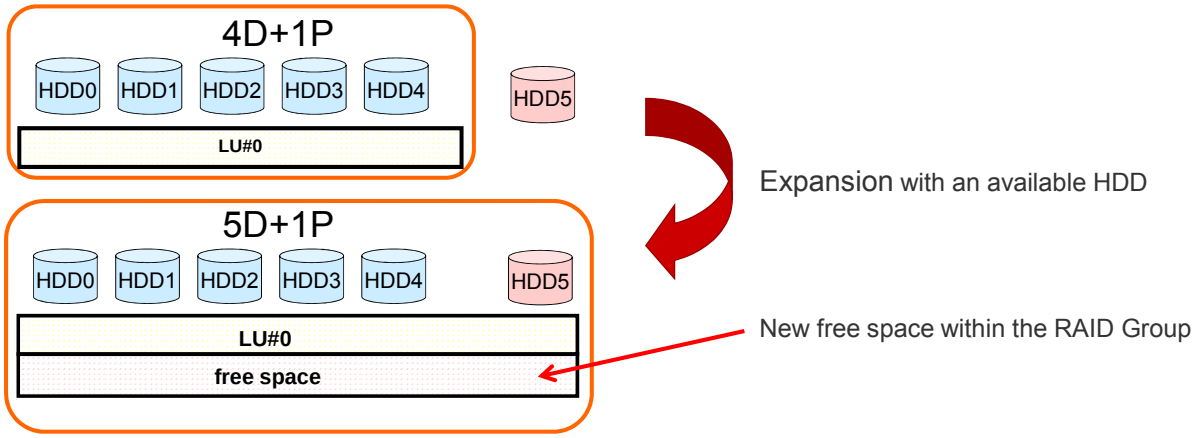
Note 1:

This is a disruptive procedure, requires Storage Navigator Modular 2 in Maintenance Mode and assistance from Hitachi Data Systems Technical Resource Center.

37

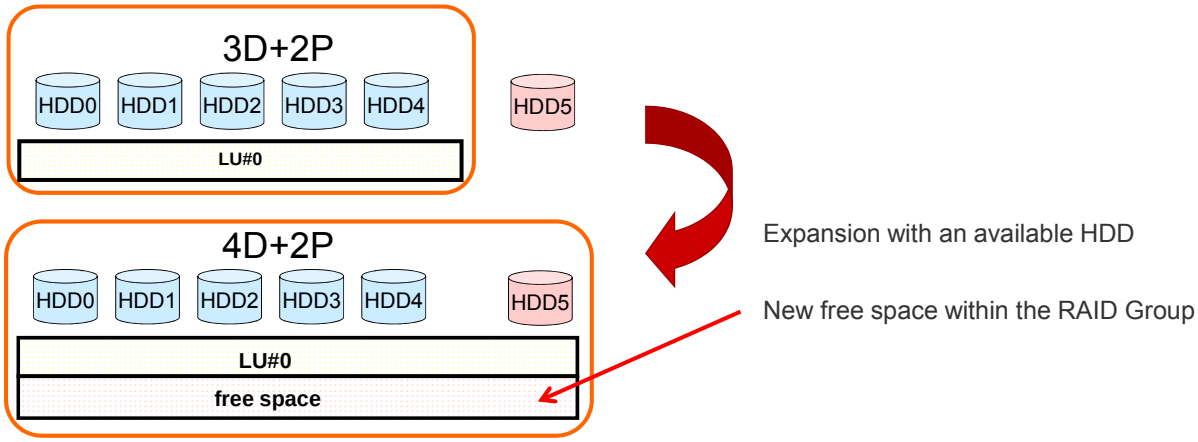
Example

- Example of an RAID 5 (4+1) RAID Group expansion



- After the expansion, you can either expand LU #0 by using the LU Grow feature or you can create a new LU.

- Example of an RAID 6 (3+2) RAID Group expansion

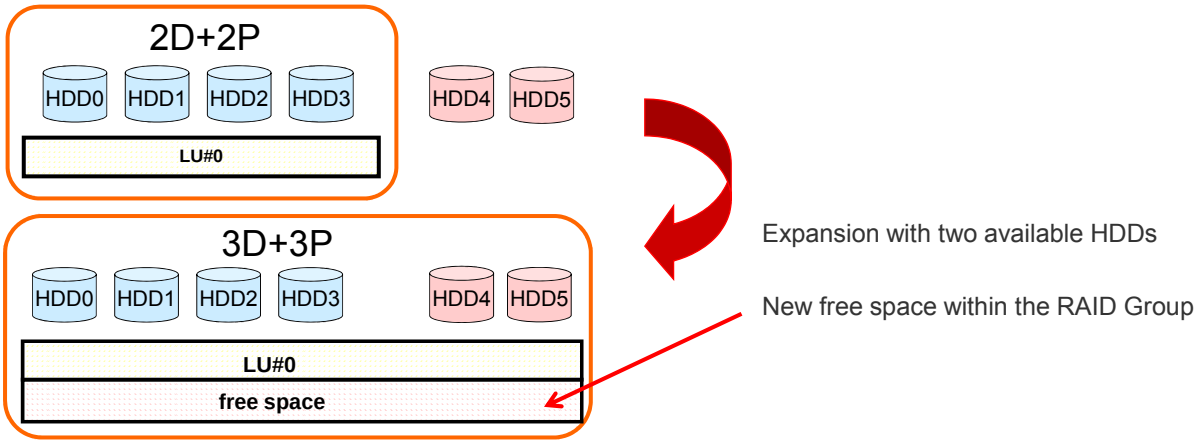


- After the expansion, you can either expand LU #0 by using the LU Grow feature or you can create a new LU.

39

HDS Confidential: For distribution only to authorized parties.

- Example of an RAID 1+0 (2+2) RAID Group expansion



- After the expansion, you can either expand LU #0 by using the LU Grow feature or you can create a new LU.

Note: When expanding RAID 1 or RAID 1+0 RAID Groups, you must expand these RAID Groups with a multiple of two HDDs.

40

Priority Mode

- RAID Group expansion affects front end performance.
- The user must select one of following two priority modes, which can be changed at any time:
  - Host Access
    - If the Host I/O **exceeds** a specific threshold, the RAID Group expansion transaction is aborted and the Host I/O transaction is prioritized.
    - If the Host I/O rate is **low**, the RAID Group expansion transaction is executed normally.
  - RAID Group Expansion
    - The host I/O gets a lower priority than the RAID Group expansion transaction.

41

HDS Confidential: For distribution only to authorized parties.

Remove Expansion

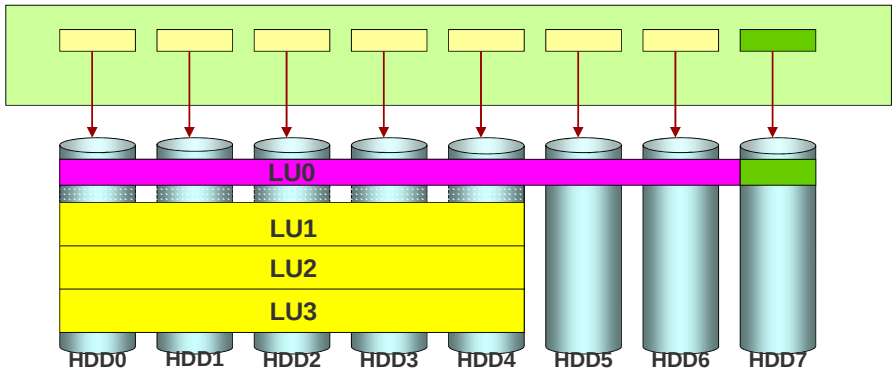
- The Remove (cancel) Expansion feature can only be used for RAID Groups, where the status is **Waiting** expansion.
- If the status is **Expanding**, you must use the Forcible suspending function.
  - This function can only be executed when Storage Navigator Modular 2 is in Maintenance Mode.
  - Internally, the LUs of the expanding RAID Group are in half finished status, so the operator needs to select one of the following two options.
    - Change the status of the LUs in the expanded area to **unformat**.

      OR
    - Change the status of the LUs in the unexpanded area to **unformat**.

Note:

The Forcible Suspending Function causes data loss on the affected LUs. Ask permission from the Technical Resource Center.

- Change the status of the LUs in the expanded area to **unformat**
- Change the status of the LUs in the unexpanded area to **unformat**



43

HDS Confidential: For distribution only to authorized parties.

Hitachi Dynamic Provisioning on AMS2000

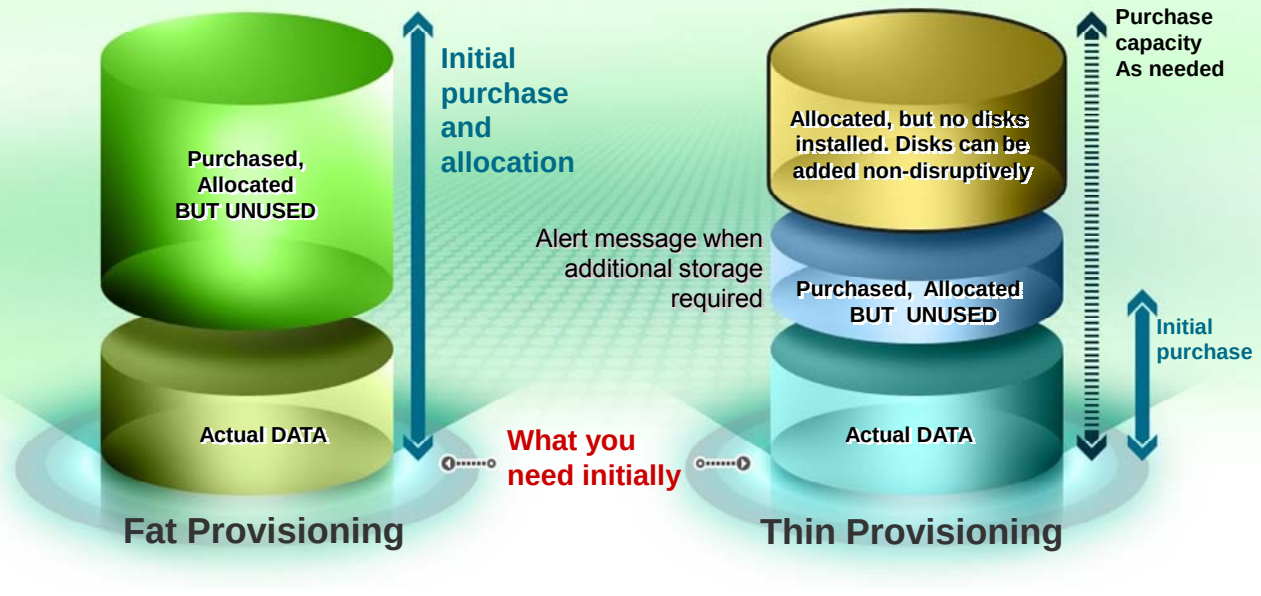
- What is Dynamic Provisioning Software?
  - It uses thin provisioning technology that allows customers to allocate ***virtual*** storage capacity based on their anticipated future capacity needs, using ***virtual volumes*** instead of physical disks.
  - Overall storage utilization rates may improve because customers can potentially provide more virtual capacity to the application while fully utilizing fewer physical disks than would be formerly required.

44

HDS Confidential: For distribution only to authorized parties.

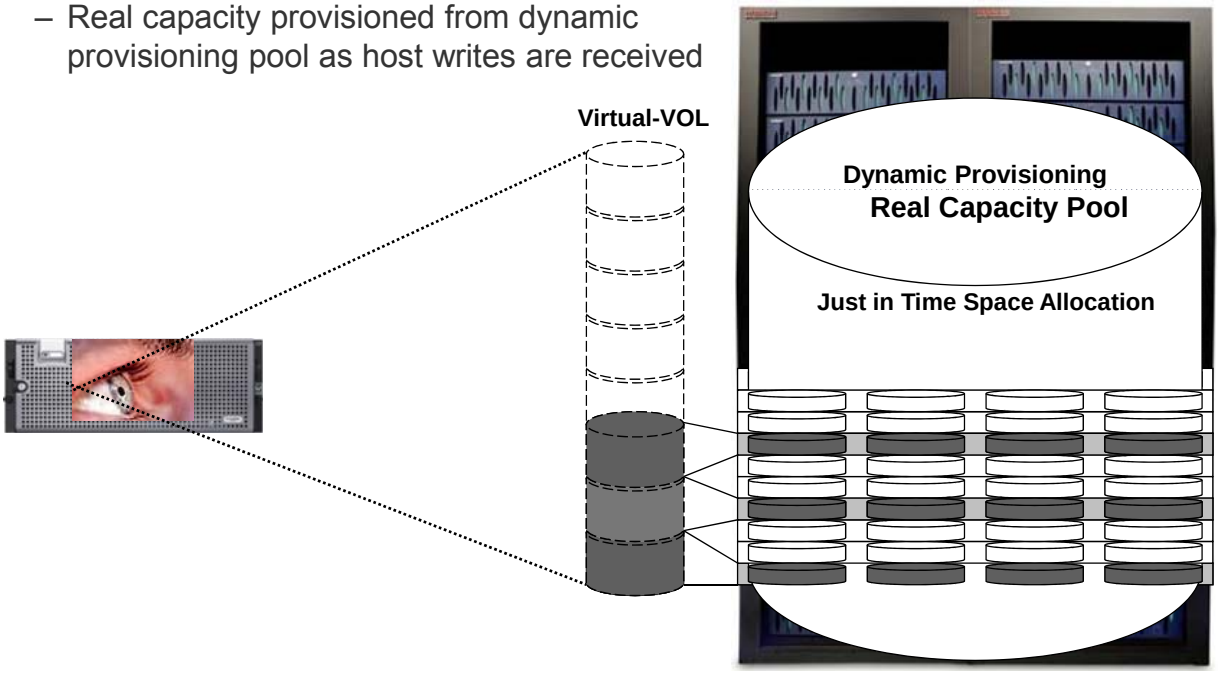
What is Dynamic Provisioning?

- To avoid future service interruptions, today it is common to over-allocate storage by 75% or more.
- With dynamic provisioning, disk capacity can be added as needed, when needed.



45

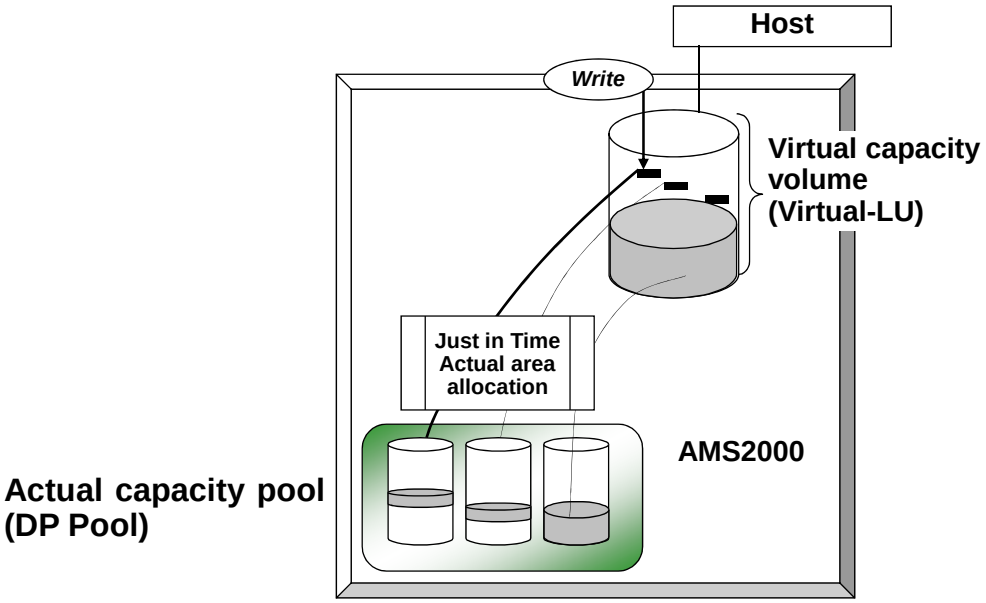
- Provision Virtual Capacity to Hosts/Application
  - Virtual maximum capacity specified and provisioned
  - Real capacity provisioned from dynamic provisioning pool as host writes are received



46

HDS Confidential: For distribution only to authorized parties.

Hitachi Dynamic Provisioning provides virtual capacity volume to the Host. Also, it assigns the actual capacity from the Pool to the Host Write area within "Just in Time."



47

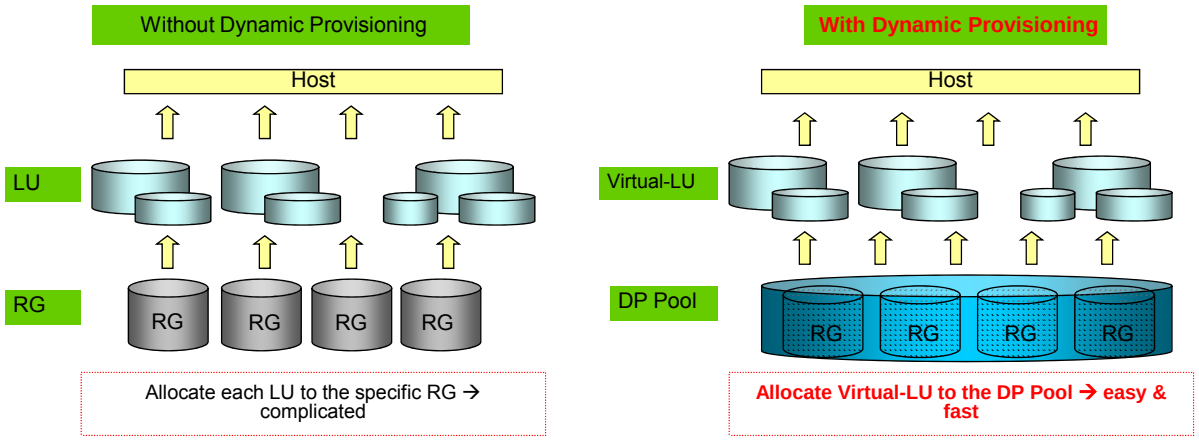
Benefit1: Easier and faster Volume Configuration

Benefit2: Reducing Volume initial cost and increasing Storage utilization

Benefit3: Simplifying Performance optimization

Benefit 1: Easier and Faster Volume Configuration

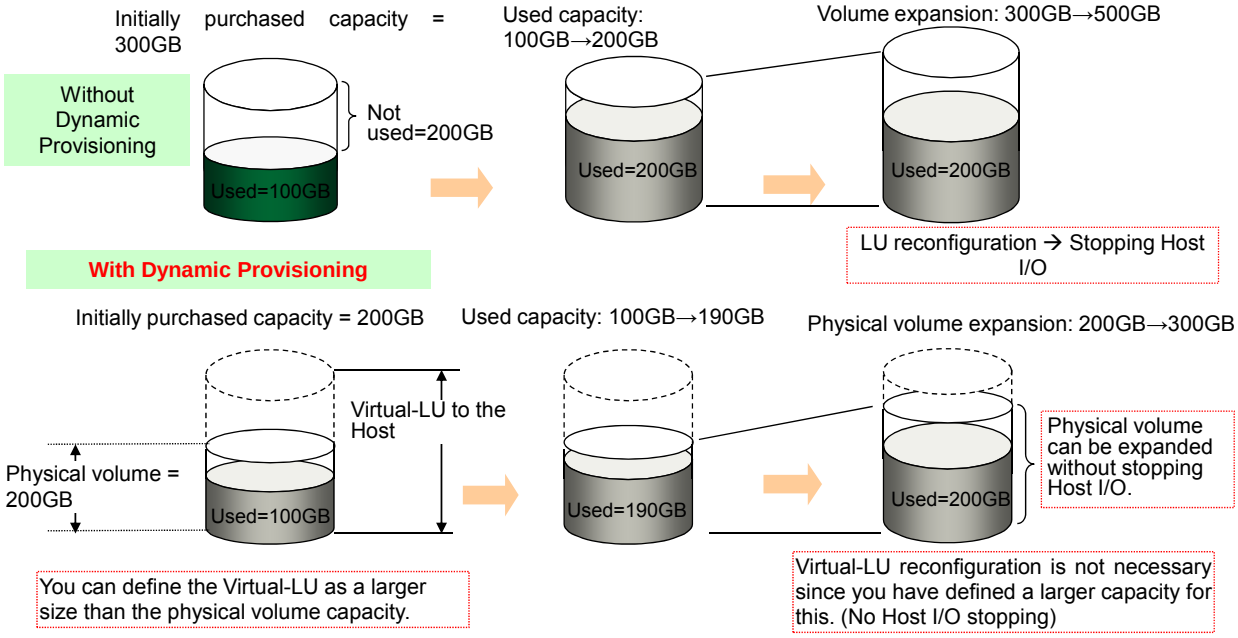
#	Without Dynamic Provisioning	With Dynamic Provisioning
1	When you create an LU, you must take care of each RG configuration and capacity, then allocate the LU to a certain RG. This is a little complicated.	When you create an LU, you need not take care of their RG configuration and capacity, but just allocate the LU to the DP Pool.
2	According to total physical volume, you must estimate and decide the LU size which actually will be required in the future. This is difficult.	You need not take care of the actual physical volume when you create a Virtual-LU which is not related to the physical volume capacity.
3	When you create the LU, then the LU formatting starts automatically in the background. During this formatting, there are some performance degradation for the LU.	Once you create the DP Pool, the formatting for each Virtual-LU is not necessary. (When you create the DP Pool, their whole area is formatted automatically on the background.)



48

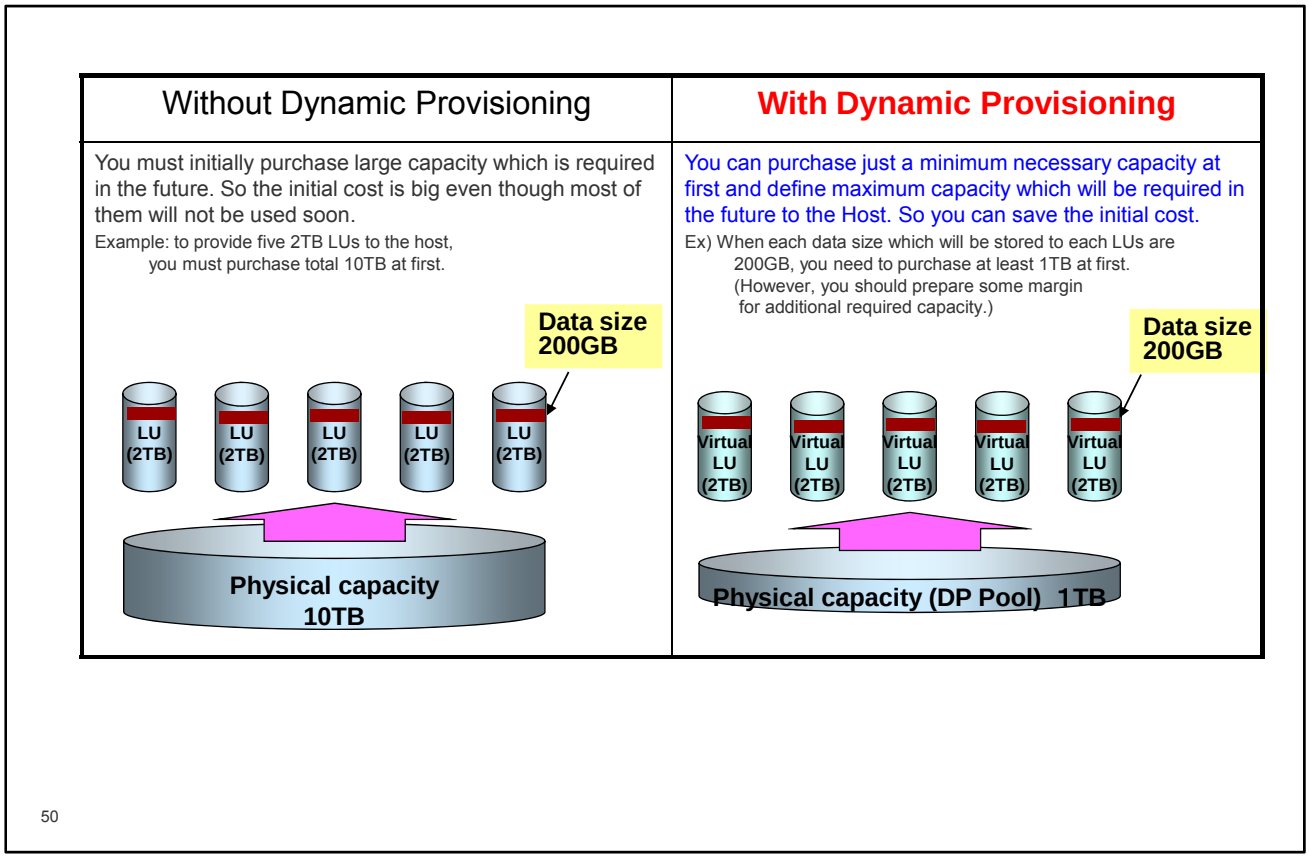
HDS Confidential: For distribution only to authorized parties.

#	Without Dynamic Provisioning	With Dynamic Provisioning
4	When you expand the LU capacity, generally you must stop the Host I/O, then reconfigure the LU.	Because you can create the large capacity of Virtual-LU which is not related to the actual volume capacity, you need not reconfigure the LU size.



49

Benefit 2: Reducing Volume Initial Cost and Increasing Storage Utilization

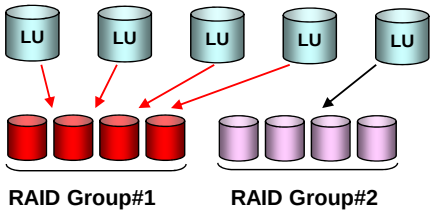


Note: To add physical capacity at a later time, you should analyze the actual daily data consumption for each DP Pool in past year. Thereby, you can plan the schedule to add physical capacity to the DP Pool. Also, when the actual data consumption is increasing rapidly, you can get notification by alerts sent from the storage system.

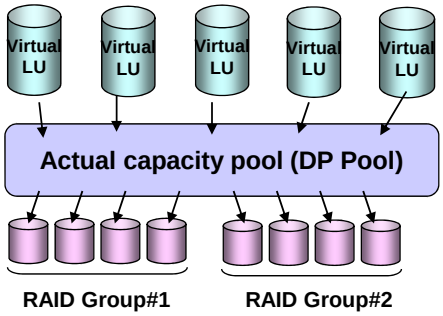
HDS Confidential: For distribution only to authorized parties.

Benefit 3: Simplifying Performance Optimization

Without Dynamic Provisioning	**With Dynamic Provisioning**
When Host I/O gathers to the specific RAID Group, it may be the performance bottleneck. To avoid this, you must configure manually RG and LU appropriately.	You can allocate plural Virtual LUs to specific DP Pool which consists of plural RGs. So their Host I/O load can be evenly diversified to plural RGs and specific RG does not become a performance bottleneck.



RAID Group#1 becomes performance bottleneck.



Host I/O load is balanced across RG#1 and RG#2.

51

Summary of Hitachi Dynamic Provisioning Benefits

- Summary of Hitachi Dynamic Provisioning benefits
  - Allows customers to allocate ***virtual*** storage capacity based on their anticipated future needs, but with fewer physical disks actually installed
  - Provides an easier and faster provisioning turnaround to applications
  - Additional physical disks can be purchased later and installed transparently without application service interruptions
  - Monitors storage resources and proactively alerts the storage administrator if additional physical storage is required
  - Simplifies performance optimization by transparently spreading hosts' individual I/O patterns across many physical disks
  - Increases storage utilization while reducing storage initial cost (total cost of ownership)

52

HDS Confidential: For distribution only to authorized parties.

6. Setup and Configuration of Storage Navigator Modular 2

Module Objectives

- Upon completion of this module, the learner should be able to:
  - Explain the purpose and benefits of Storage Navigator Modular 2 (SNM2) program
  - Register an Adaptable Modular Storage 2000 family system in SNM2
  - Use Add Array wizard
  - Use Initial Setup wizard
  - Create RAID Groups and use LU Wizard to create and format logical units (LUs)
  - Create Host Groups, enable Host Group Security and register the World Wide Name (WWN) of attached host bus adapters
  - Map internal LUNs to Host Group LUNs

Storage Navigator Modular 2 is a component of Hitachi Resource Manager™ utility package.

Features

- Storage Navigator Modular 2 allows you to:
  - Centralize storage management
  - Configure the storage system using:
    - Add Array wizard
    - Initial Setup wizard
    - Create and Map Volume wizard
  - Allocate user volumes to the host server
  - Replicate logical units using standard Hitachi replication software
    - Replication wizard
  - Obtain system configuration and status information
  - Automatic error alert monitoring
    - Error alert email system
    - Error alert window for failed components, listing all customer replaceable unit (CRU) details

3

The Adaptable Modular Storage houses the *actual* disks for storage. SNM2 manages those disks to provide end users access to the disk space to manage their storage environment. Through the GUI, users can centralize their storage management, replicate data, manage performance, and easily diagnose and fix errors.

HDS Confidential: For distribution only to authorized parties.

- Simple
  - Easy to use graphical user interface
  - Installs in minutes
  - Wizard-guided setup, configuration and data replication
  - Intuitive error detection and messaging
  - Non-disruptive firmware updates
  - Online Help

4

Supported Operations

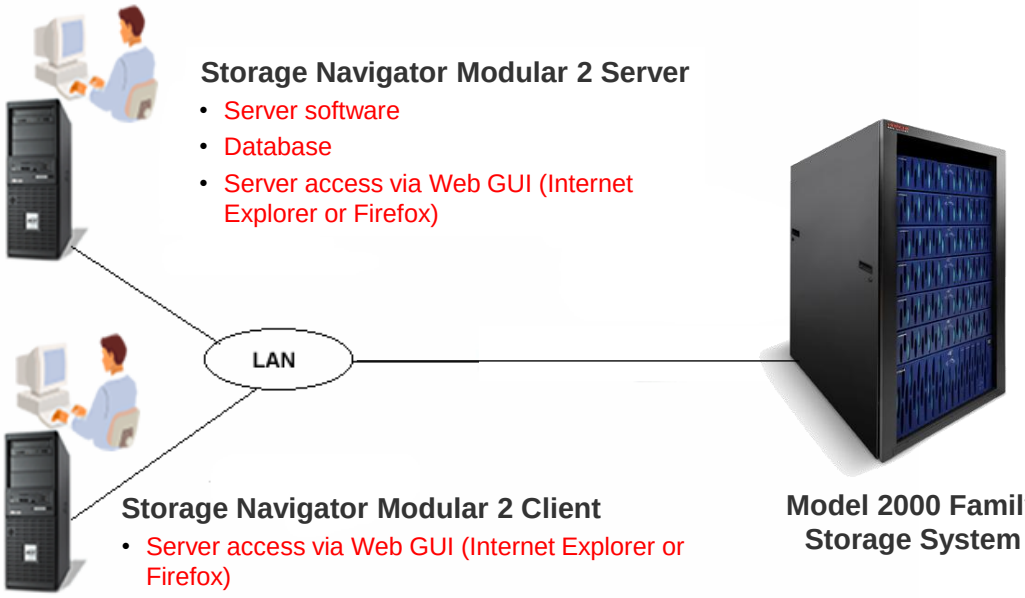
- Configuration of storage system
  - Create and Map Volume Wizard
- License key enablement of additional supported software
- Installation and updating of firmware
- Data replication
  - Replication Wizard
  - Snapshots (Copy-on-Write Snapshot software)
  - Full-volume cloning (ShadowImage Replication software)
  - Disaster Recovery
    - TrueCopy Synchronous Software
    - TrueCopy Extended Distance (Asynchronous)
- Alert settings
- Storage system health and performance monitoring

5

HDS Confidential: For distribution only to authorized parties.

Architecture

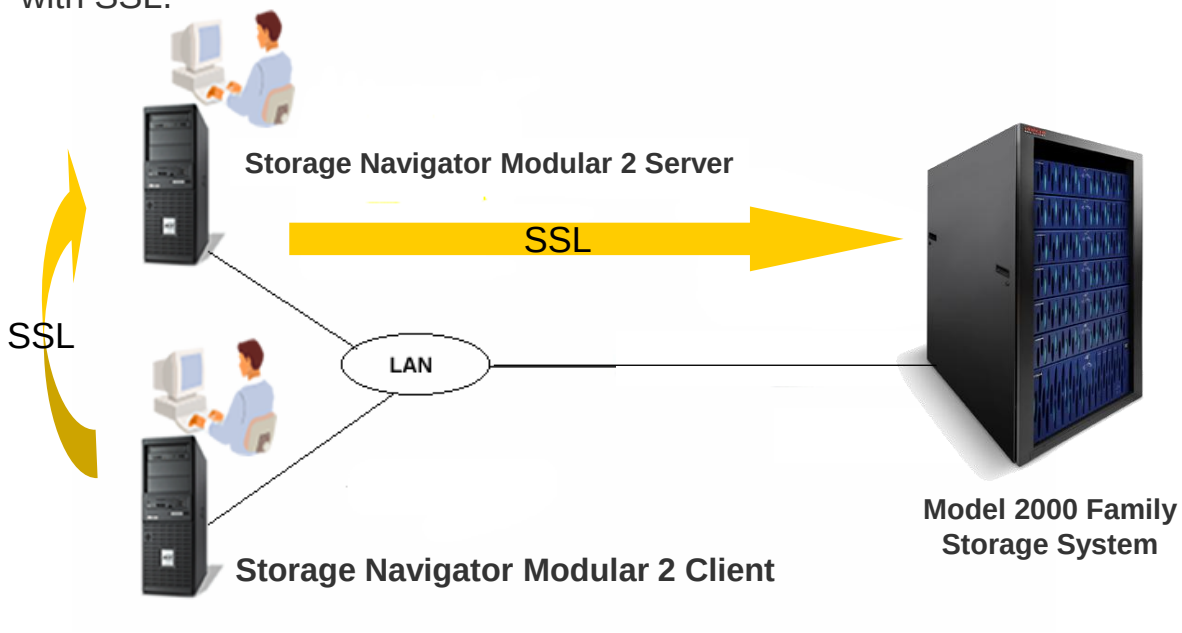
- Web GUI
- Client-Server design



6

Storage Navigator Modular 2 runs from your primary management server or client PC. It is designed on common web-based client-server technology using a standard IP network. In other words, you can attach your model 2100 or 2300 and Storage Navigator Modular 2 primary management server to your existing LAN environment. Storage Navigator Modular 2 communicates with the storage system through a web browser. If client PCs are attached to the network, they can connect to the Storage Navigator Modular 2 primary management server and remotely configure the storage system.

- Communication between the SNM2 server and the storage system, and between the SNM2 server and client can be either unsecured or secured with SSL.



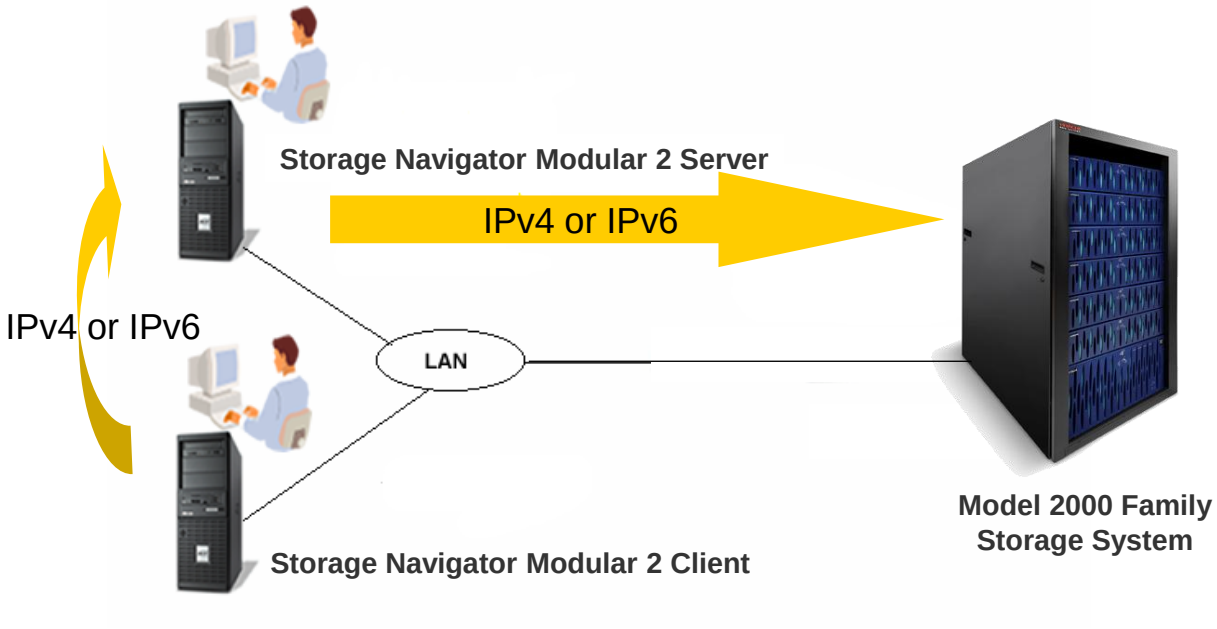
7

To get more details regarding the setup and configuration please refer to the appropriate user documentation.

The procedures may differ depending on the storage system firmware level and the SNM2 version.

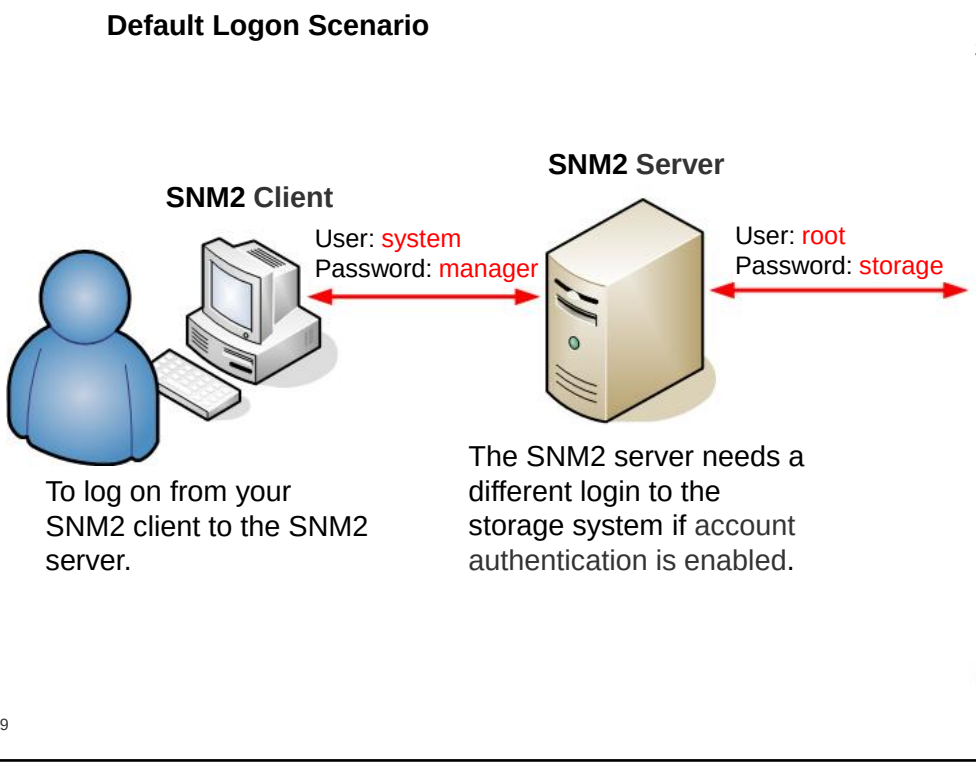
HDS Confidential: For distribution only to authorized parties.

- IPv4 or IPv6 is supported



8

Default Authentication



HDS Confidential: For distribution only to authorized parties.

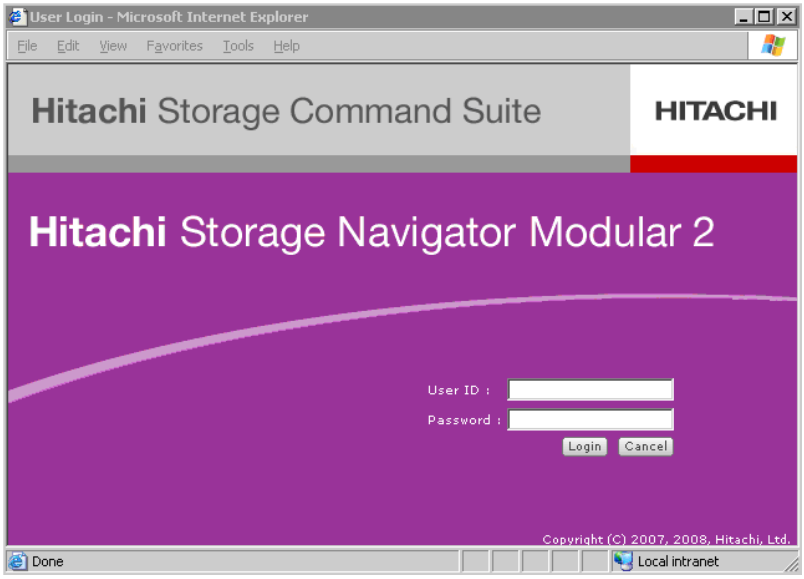
Configure

- Default Controller IP settings
  - Controller 0 & 1 - **User** LAN port – **192.168.0.16 & 192.168.0.17** – factory default
    - **LAN1** port
  - Controller 0 & 1 - **Maintenance** LAN port - **10.0.0.16 & 10.0.0.17** – factory default
    - **LAN0** port

10

Start From Web Brower

1. Open a Web browser.
2. Access the Storage Navigator Modular 2 software from the browser.
   Set URL = http:// <*IP address of host*>:**23015**/StorageNavigatorModular**/**
3. Log in.
   - User Name: **system**
   - Password: **manager**

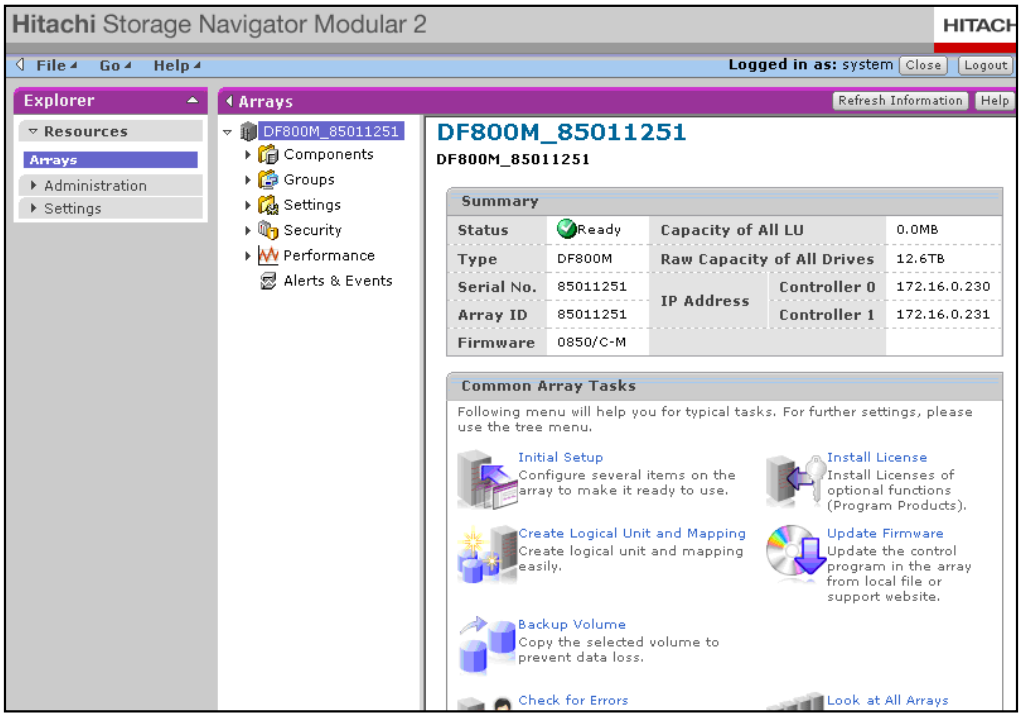


11

Since this is the first time you are running Storage Navigator Modular 2, the **Add Array** wizard appears, and prompts you to add your storage system.

HDS Confidential: For distribution only to authorized parties.

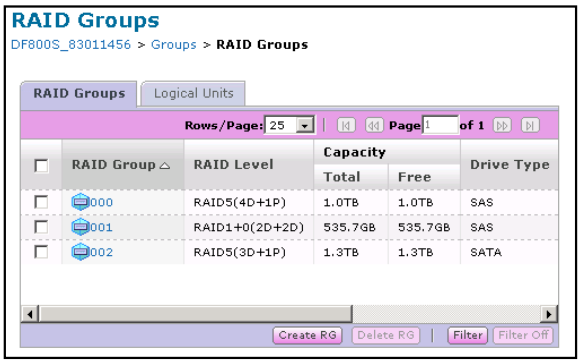
GUI Main Window



12

RAID Group Creation

- RAID Groups can be created in two ways:
  - **Automatic Selection** of disk drives
    - Takes drives in numerical order
    - Selects drives of same type, size
  - **Manual Selection** of disk drives (HDD roaming)
    - Drives can be selected at random
    - You **cannot** select drives of different **densities**
    - You **can** select drives with different **RPMs**

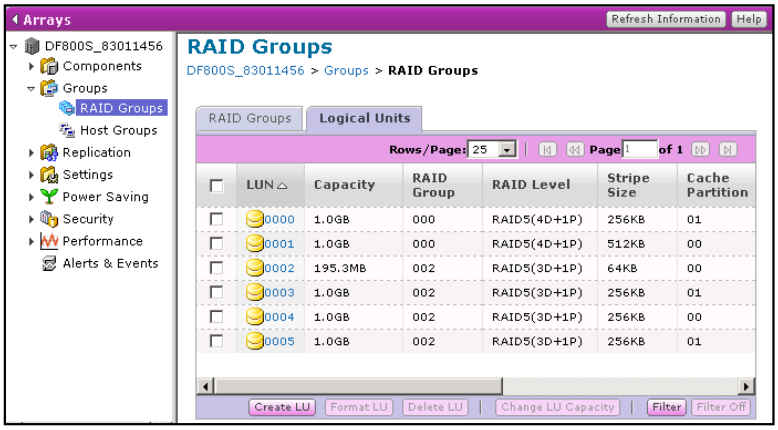


13

HDS Confidential: For distribution only to authorized parties.

Logical Units Creation

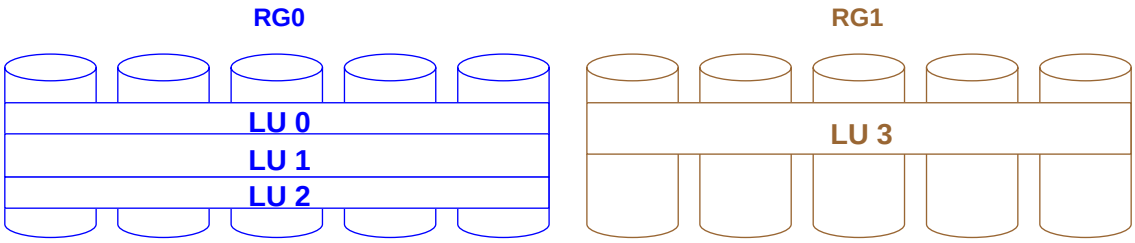
- **Logical Units** (LUs) can be created in two ways:
  - **Create LU** (one at a time)
    - **Automatically**: User specifies the size, SNM2 creates LU from available free space
    - **Manually**: The user selects one or more segments from a list of available free space segments
  - Create and Map Volume Wizard (create several LUs at once)



14

LU Configuration

- LUs are **slices** from the user data area of a RAID Group:
  - 3 LUs form RG0
  - 1 LU from RG1
- Maximum LUs
  - Model 2100 = **2,048**
  - Model 2300 = **4,096**
  - Model 2500 = **4,096**



15

HDS Confidential: For distribution only to authorized parties.

Host Connectivity Basics

- Basic steps for connecting the model 2000 systems to open systems hosts:
  1. Verify that the host has a compatible OS and HBAs:
     - Required patches and software installed
     - HBAs running with the correct version of firmware
     - You can refer to **HiFire** for this information.
  2. Connect ports to HBAs (direct connect - loop) or Switch (point-to-point).
  3. Set correct Topology and Transfer Rate for each HBA and AMS 2000 port.
  4. Verify Fibre Channel Link and Login.
  5. Enable Host Group Security on desired storage system ports.
  6. Create RAID Groups and LUs.
  7. Create Host Groups and associate the correct HBA WWNs to each Host Group.
  8. Map Internal storage system LUs to Host Group LUs.
  9. Verify the Host detection of LUs and prepare LUs for use.

16

Host Groups and LUN Mapping

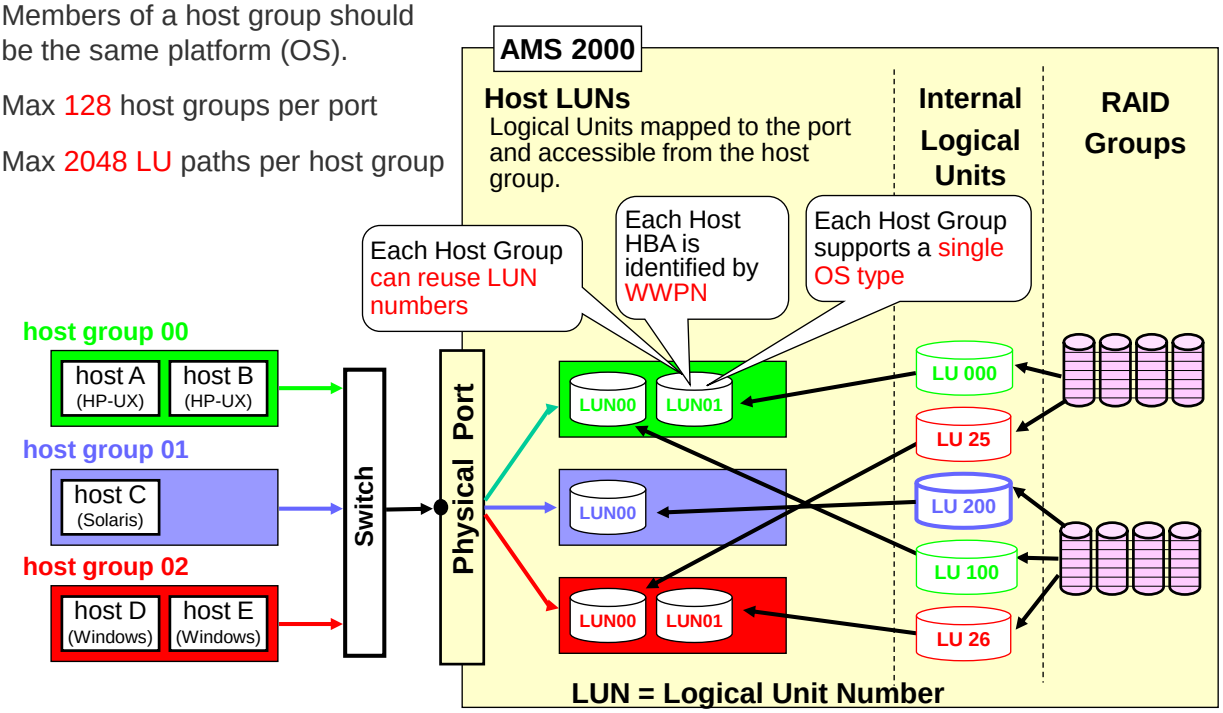
- Host Groups allow for the heterogeneous multi-host connections.
- **LUN Manager** is the Program Product that manages access paths between hosts and logical units for each port.
- Illegal access to LUNs from any host system can be prevented using the **LUN Security** feature.
- Host Group is a **virtual port** created on the physical port.
  - Host Group can be tuned to the connected host OS.
  - Each physical port has **one** default Host Group (Host Group **000**).
- **Internal** LUs are mapped to a Host Group as **Host LUNs.**
  - Two of more hosts that require a LUN 0 can be connected to the same port.

17

HDS Confidential: For distribution only to authorized parties.

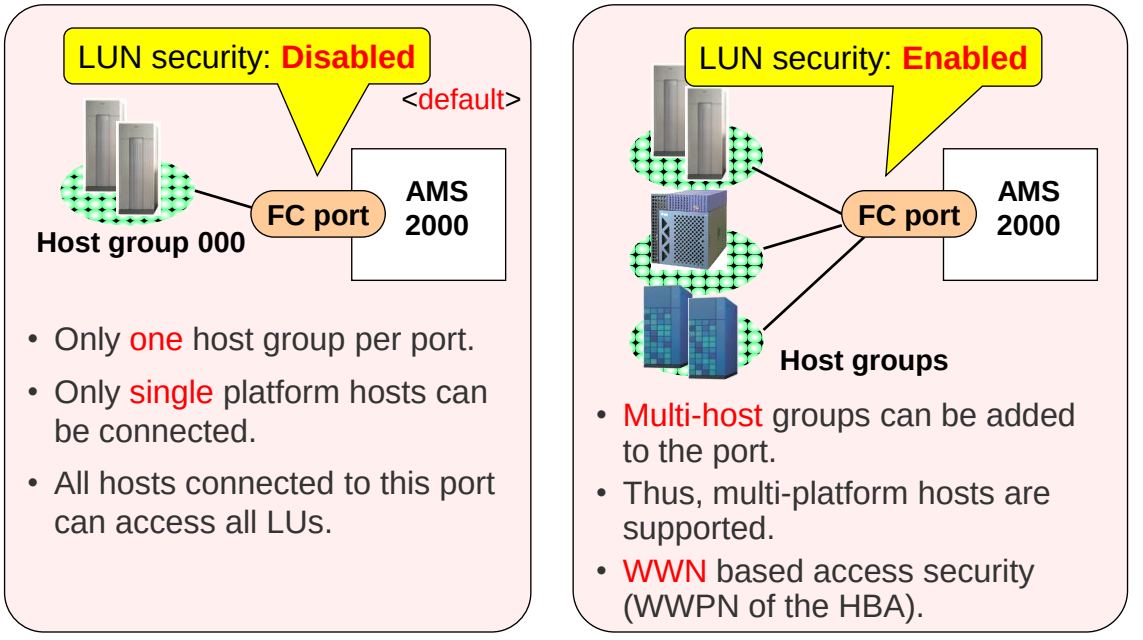
Host Groups

- Members of a host group should be the same platform (OS).
- Max 128 host groups per port
- Max 2048 LU paths per host group



18

Host Group Security (LUN Security)



To protect mission-critical data in your disk storage system from unauthorized access, you should implement LUN security. LUN security allows you to prevent unauthorized hosts from either seeing or accessing the data on the secured LU. If LUN security is applied to a particular port, a host has only access to the LUNs inside a Host Group it was configured for.

HDS Confidential: For distribution only to authorized parties.

Host Group - Options Settings

- Options Settings
  - **Platform**:
    - HP-UX, Solaris, AIX, Linux, Windows, and VMware
    - ***not specified*** is the default
  - **Middleware**:
    - VCS and TruCluster
    - ***not specified*** is the default

Note: Refer to the System Parameter section of the latest Maintenance Manual or the Adaptable Modular Storage 2000 Family Host Installation Guides for a complete list of all options, settings, and combinations.

20

VCS = Veritas Cluster Server.

TruCluster = HP high availability cluster server.

World Wide Names

- Adaptable Modular Storage World Wide Names (WWNs)
  - Each Hitachi storage system sets its own WWN for each port.
  - Each WWN is the result of the following:
    - Specific bits are determined by the storage family and model.
    - The system serial number is part of the WWN.
    - The last digit of the WWN identifies the port.

The first **three** digits identify the **model**, the last **five** digits the **Serial Number**:

830XXXXX = model 2100
850XXXXX = model 2300
870XXXXX = model 2500

For Example: 85011251$_{10}$ (a decimal value)

11251
+ 10000 (add this offset)
21251$_{10}$

5 0 0 6 0 E 8 0 1 0 **4** **5 3 0 3** **0**$_{16}$

Fixed

Model descriptor
2 = 2300
4 = 2100
0 = 2500

Serial Number
In Hexadecimal

Refers to the physical FC Host Connector Port
See next slide

22

HDS Confidential: For distribution only to authorized parties.

2100	2300	2500
CTL0 Port A = 0 CTL0 Port B = 1 CTL1 Port A = 2 CTL1 Port B = 3	CTL0 Port A = 0 CTL0 Port B = 1 CTL0 Port C = 2 CTL0 Port D = 3 CTL1 Port A = 4 CTL1 Port B = 5 CTL1 Port C = 6 CTL1 Port D = 7	CTL0 Port A = 0 CTL0 Port B = 1 CTL0 Port C = 2 CTL0 Port D = 3 CTL0 Port E = 4 CTL0 Port F = 5 CTL0 Port G = 6 CTL0 Port H = 7 CTL1 Port A = 8 CTL1 Port B = 9 CTL1 Port C = A CTL1 Port D = B CTL1 Port E = C CTL1 Port F = D CTL1 Port G = E CTL1 Port H = F

23

Host HBA Persistent Binding

- If your host systems are connected to a Fibre Channel switch, HDS strongly recommends that host HBA Persistent Binding be configured.
  - User associates WWPN of the connected storage array port to a specified Target ID (TID).
  - LUNs are identified and associated to the TID.
  - Persistent binding prevents the TID from changing when devices join or leave the fabric.
  - If the TID changes, the host will lose access to the LUNs on that path.

24

Persistent Binding is also known as LUN mapping or mapping.

Recorded Product Demonstration

Now… Recorded Product Demonstration

From your Web Browser ….
http://edemo.hds.com/edemo/vault1/AMS2000FamilyInstallation.swf

HDS Confidential: For distribution only to authorized parties.

7. Hitachi Dynamic Link Manager

Module Objectives

- Upon completion of this module, the learner should be able to:
  - Describe the benefits of Hitachi Dynamic Link Manager software
  - Identify the key components and functions

2

Overview

- Purpose of the Dynamic Link Manager software (HDLM)
  - A server based path management software that provides path failover and load balancing capabilities

  Provides:
  - Support for Fibre Channel connectivity
  - Automatic path discovery
  - Automatic path failover and failback
  - Command Line Interface (CLI) and GUI
  - Integration with Hitachi Storage Management software

3

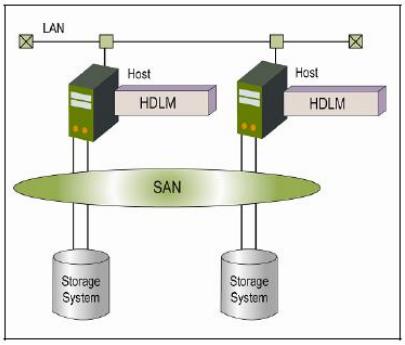
The Dynamic Link Manager is a family of Hitachi provided middleware *server based* software utilities. Dynamic Link Manager enhances the availability of RAID systems by providing automatic error recovery and path failover from server-to-RAID connection failures. It provides load balancing in addition to path failover by redirecting I/O activity to the least busy path using complex algorithms.

Just because a system is RAID-protected does not mean that it is protected against connection bus failures, which is why the Dynamic Link Manager is required for true nonstop operations. It allows System Administrators to take advantage of the multiple paths by adding redundant connections between data servers and RAID systems. The Dynamic Link Manager software, therefore, provides increased reliability and performance.

HDS Confidential: For distribution only to authorized parties.

Benefits

- Removes HBA as single point of failure
- Automatically detects failed path and reroutes I/O to alternate path
- Automatically discovers host bus adaptors (HBAs) and LUNs in SAN environment
- Has multiple load balancing options to balance I/Os across available paths
- Provides tools to control and display the path status
- Supports the most popular cluster technologies
- Supports HBA vendors and standard open drivers
- Provides GUI and CLI support
- Provides error logging capability



4

The Dynamic Link Manager provides the ability to reduce the server's HBA as the single point of failure in an open environment. One of the strengths of Dynamic Link Manager is its ability to configure itself automatically. It is designed to enhance the operating system by putting all alternate paths offline in Disk Administrator. It functions equally well in both iSCSI and Fibre Channel environments. It supports the most popular cluster technologies like HACMP, MSCS, MC Service Guard, Sun Cluster, and VERITAS Cluster Server. It has GUI/CLI support for configuration management, performance monitoring and management and authentication of user IDs using the Hitachi Storage Management facility.

Features

Multi-pathing

- Enables multiple paths from the host to devices, allowing access to the device even if a specific path is unavailable.
- Multiple paths can also be used to share I/O workloads and improve performance.

Path Failover

- Automatically redirects I/O operations to alternate paths if a failure occurs, allowing processing to continue without interruption.
- With the threat of I/O bottlenecks removed and your data paths protected, performance and reliability increase.

Failback –

- When a failed path becomes available, HDLM places the recovered path back online.
- This ensures that the maximum number of paths is always available for load balancing and failover.

5

Load Balancing

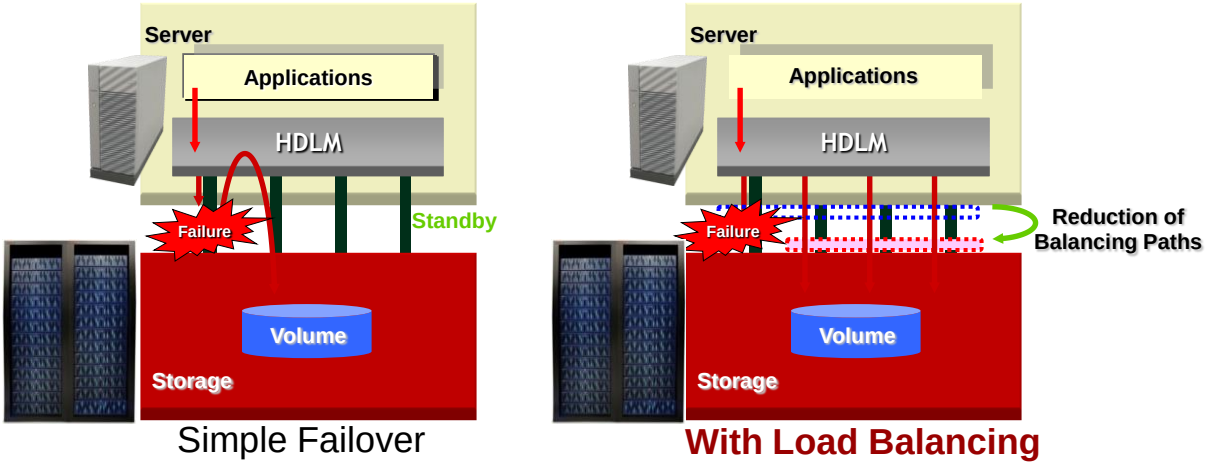
- Intelligently allocates I/O requests across all available paths to prevent a heavily loaded path from adversely affecting processing speed.
- Load balancing ensures continuous operation at optimum performance levels, along with improved system and application performance.

Path health checking

- Automatically checks the path status at regular user-specified intervals, eliminating the need to perform repeated manual path status checks.
- Proactive identification of issues.

HDS Confidential: For distribution only to authorized parties.

Path Failover and Failback

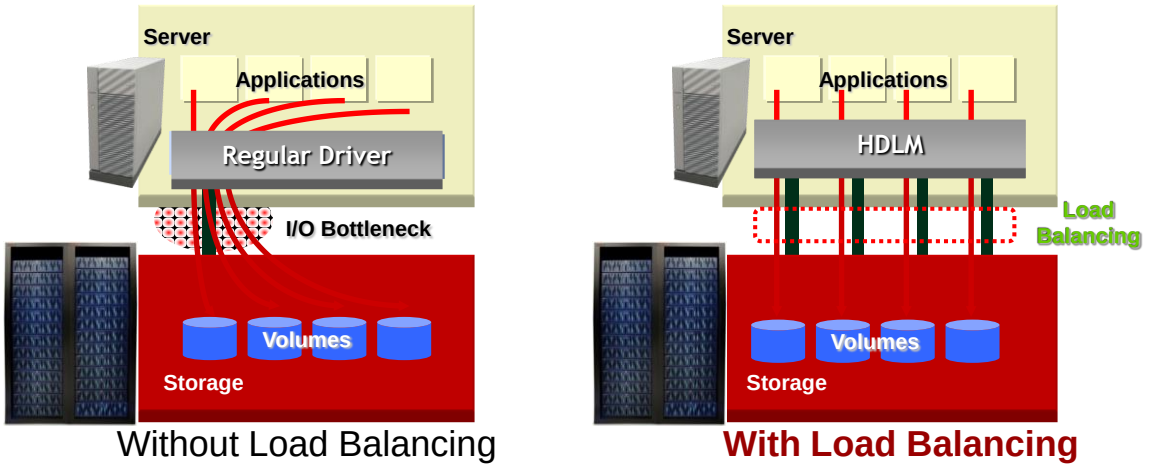


- Dynamic Link Manager software provides continuous storage access and high availability by distributing I/O over multiple paths
- Failover and fallback in either manual or automatic modes
- Automated path health checks
- Allows dynamic LUN addition and deletion without a server reboot *

* O/S and array dependent, check system requirements for details

7

Load Balancing



- Dynamic Link Manager software distributes storage access across multiple paths to improve I/O performance with load balancing.
  - ***On modular storage 2000 Family, Dynamic Link Manager software does allow load balancing through two controllers.***
- Bandwidth control at the HBA level — and in conjunction with Global Availability Manager at the LUN level.

8

When there is more than one path to a single device within a logical unit (LU), the Dynamic Link Manager can distribute the load across the paths by using the paths to issue I/O commands. Load balancing prevents a heavily loaded path from affecting performance of the entire system.

HDLM stands for Dynamic Link Manager software.

HDS Confidential: For distribution only to authorized parties.

Load Balancing Algorithms

Load Balancing Algorithm	Description
round-robin	Distributes all I/Os among multiple paths
extended round-robin	Distributes I/Os to paths depending on whether the I/O involves sequential or random access: • For sequential access, a certain number of I/Os are issued to one path in succession. The next path is chosen according to the round robin algorithm. • For random access, I/Os will be distributed to multiple paths according to the round-robin algorithm.

9

Load Balancing Algorithm	Description
least I/Os	I/O operations are issued to a path that has the least number of I/Os being processed (regardless of I/O block size).
extended least I/Os (Default)	I/O operations are issued to a path that has the least number of I/Os being processed (regardless of I/O block size): • For sequential access, a certain number of I/Os are issued to one path in succession. The next path is chosen according to the least I/Os algorithm. • For random access, I/Os are issued to multiple paths according to the least I/Os algorithm.

10

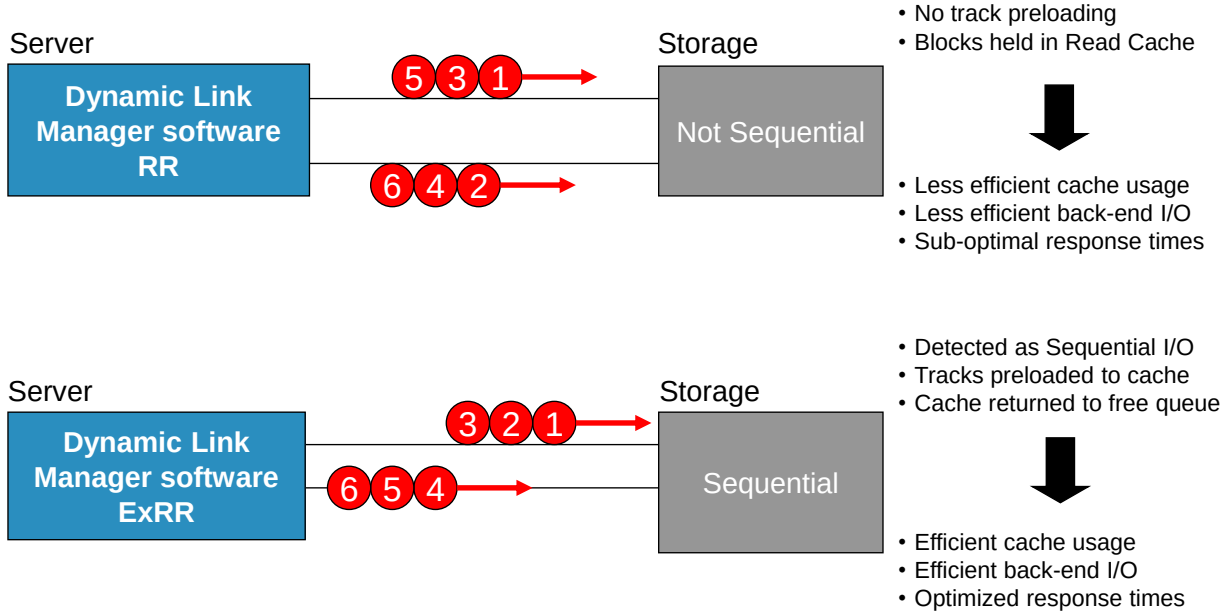
Load Balancing Algorithm	Description
least blocks	I/O operations are issued to a path with the least pending I/O block size (regardless of I/O count).
extended least blocks	I/O operations are issued to a path with the least pending I/O block size (regardless of I/O count): ▪ For sequential access, a certain number of I/Os are issued to one path in succession then the next path is chosen according to the least blocks algorithm. ▪ For random access, I/Os are issued to multiple paths according to the least blocks algorithm.

11

HDS Confidential: For distribution only to authorized parties.

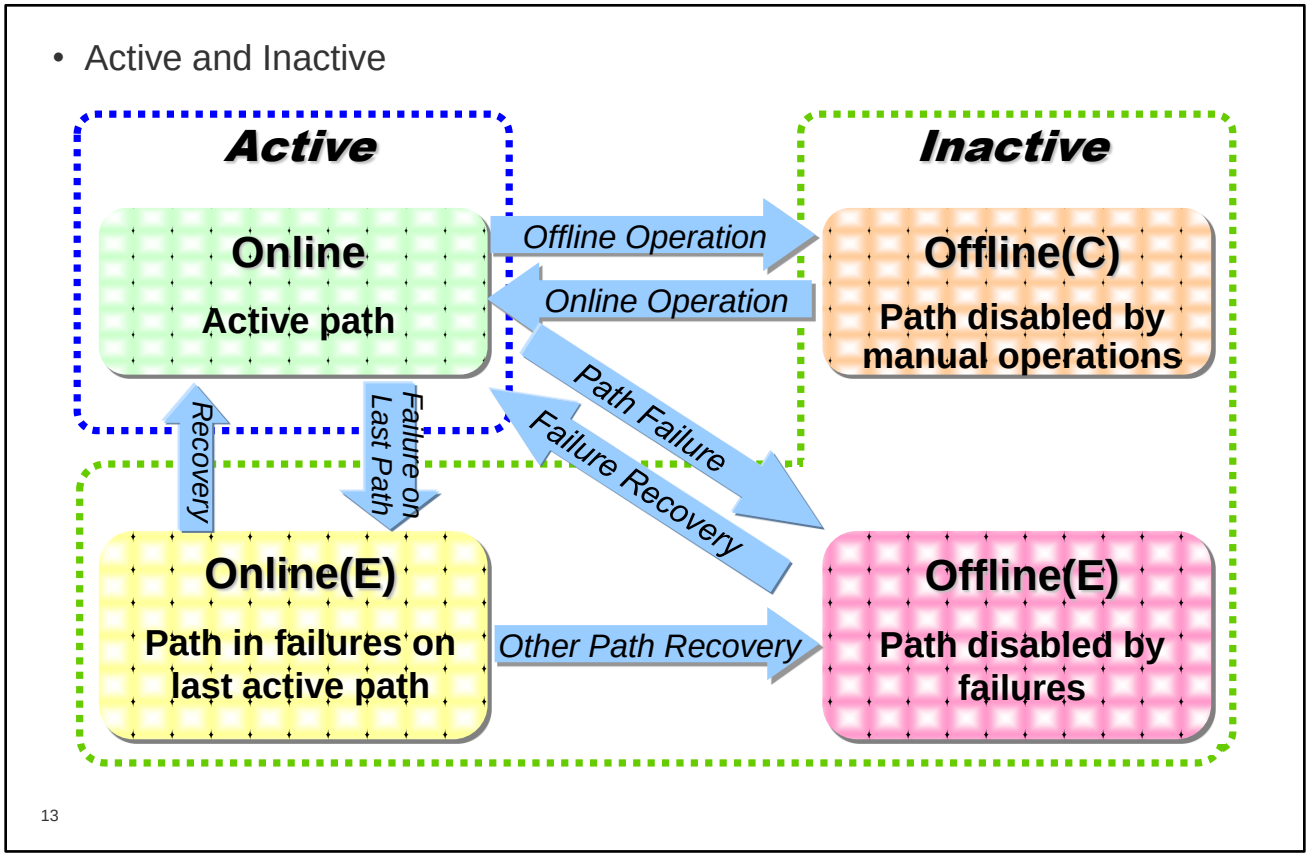
Round Robin – Extended Round Robin

- Impact of Extended Round Robin versus Round Robin on Sequential Detect



12

Path State Transitions



Active is the state where the path is healthy. **Inactive** is the state where all access is disabled. **Online** is the state where the Dynamic Link Manager software allows applications access to the path. **Offline** is the state where Dynamic Link Manager software does not allow applications to the path.

This illustration shows the path status transition.

Path has two status; Online and Offline status.

Online status path is what Dynamic Link Manager software uses for failover and load balancing. I/O can basically be issued to the path.

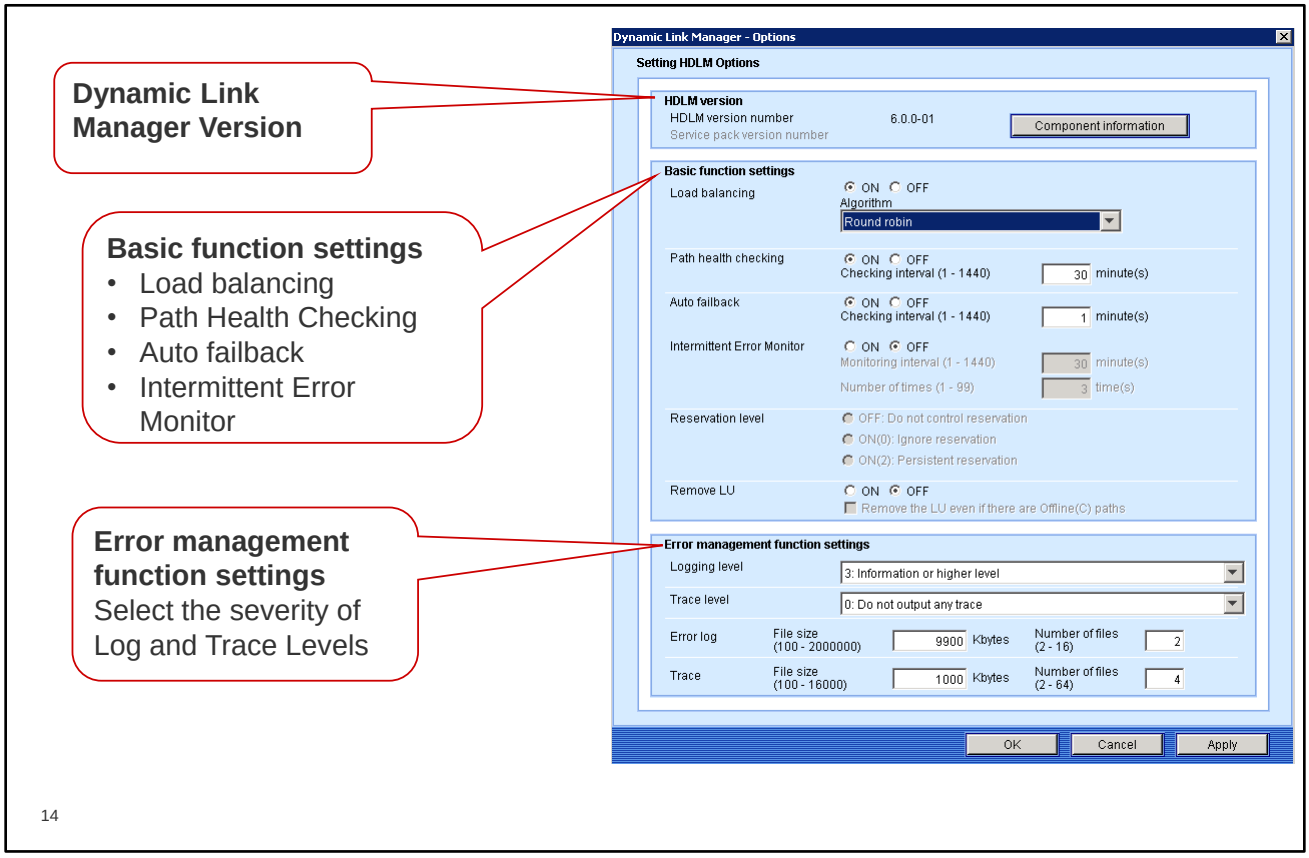
Offline status path is what Dynamic Link Manager software does not use for failover and load balancing. I/O cannot be issued to the path.

When an error occurred or the **offline** command is executed, Dynamic Link Manager software places the online path offline.

When the **online** command is executed, Dynamic Link Manager software places the offline path online.

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

Graphical User Interface



Optional Parameters

Load Balancing

Path Health Check: When enabled (default), Dynamic Link Manager monitors all online paths at specified interval and puts them into **Offline(E)** or **Online(E)** status if a failure is detected.

- There is a slight performance penalty due to extra probing I/O
- The default interval is 30 minutes.

Auto Failback: When enabled (not the default), Dynamic Link Manager monitors all Offline(E) and Online(E) paths at specified intervals and restores them to online status if they are found to be operational. The default interval is one minute.

Intermittent Error Monitor

- Auto Failback must be On
- Parameters are Monitoring Interval and Number of Times, for example:
  - Monitoring Interval = 30 minutes
  - Number of Times = 3

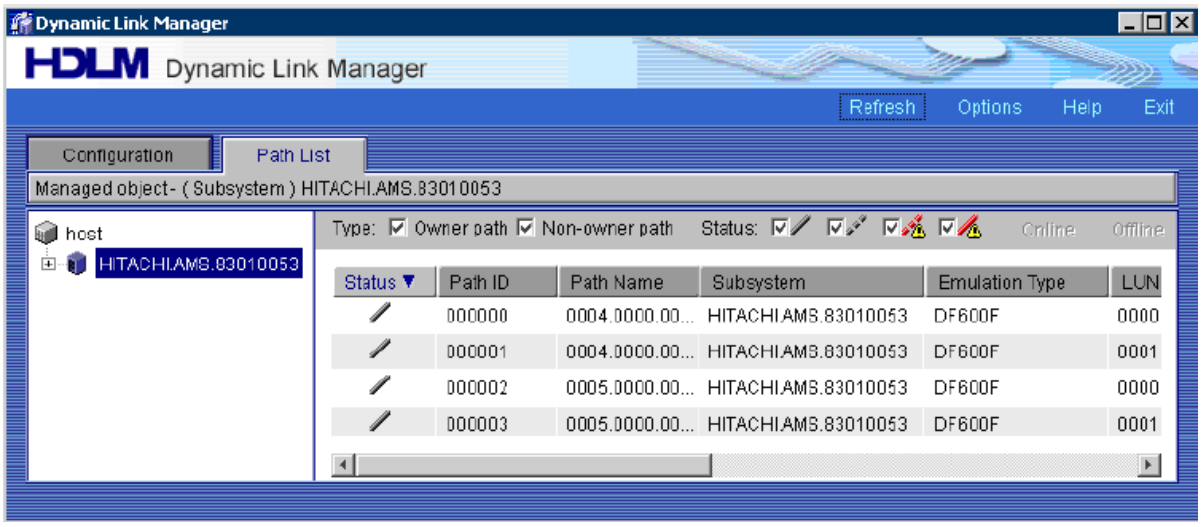
If an error occurs three times in a 30 minute period, the path is determined to have an intermittent error and is removed and subject to automatic failback. The path will display error status until problem is corrected.

Reservation Level: Used only by AIX and specific paths can be reserved for I/O

Remove LU

- Used with Microsoft Windows 2000 and Server 2003
- Removes the LUN when all paths to the LUN are taken offline

- View the path information and change the path status through the Path Management Window.



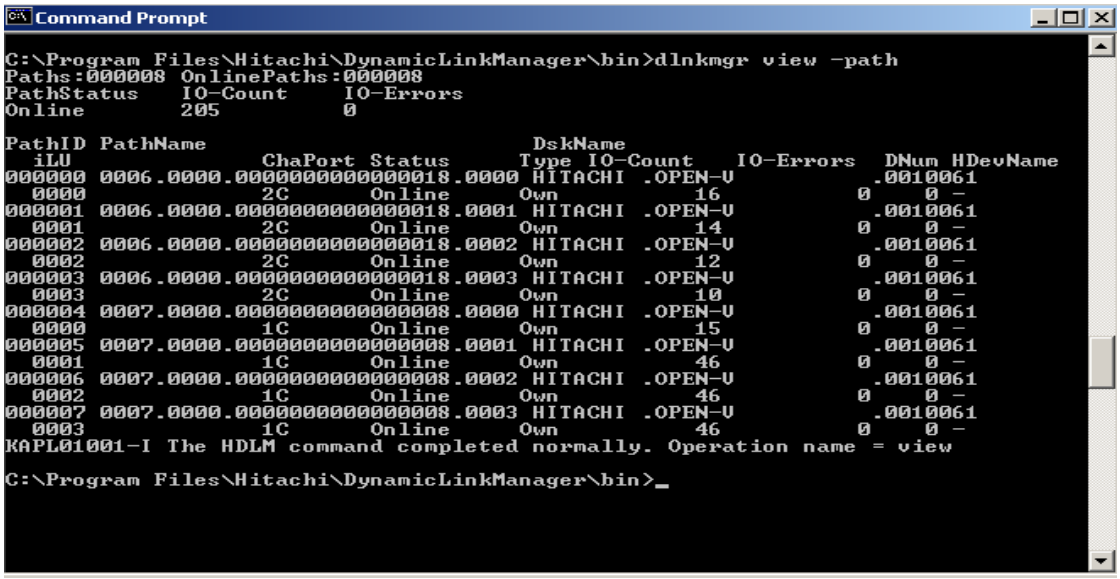
15

You can switch a path by manually placing a path online or offline. Manually switching a path temporarily is useful for maintenance of the system.

HDS Confidential: For distribution only to authorized parties.

Operations

- Path View Operation for the CLI
  - Allows you to see information about data paths
  - dlnkmgr view -path



16

View allows you to see two types of information:

- Information about your data paths -path
- Information about the Dynamic Link Manager system settings – sys

Both of these parameters can take several different values.

Recorded Product Demonstration

Now… Recorded Product Demonstration

From your Web Browser …. edemo.hds.com/edemo/vault1/HDLM.swf

Module Review

1. Which two load balancing algorithms are supported by Dynamic Link Manager software? (Choose two)
   a. Active/Standby
   b. Round Robin
   c. Active/Active
   d. Extended round robin
   e. Least busy
2. What is the maximum number of paths per LU that Dynamic Link Manager software supports?
   a. 8
   b. 16
   c. 32
   d. 64

18

HDS Confidential: For distribution only to authorized parties.

8. Hitachi Replication Software Overview

Module Objectives

- On completion of this module, the learner should be able to:
  - Describe the Business Continuity solutions and the available software

2

Hitachi Replication Products

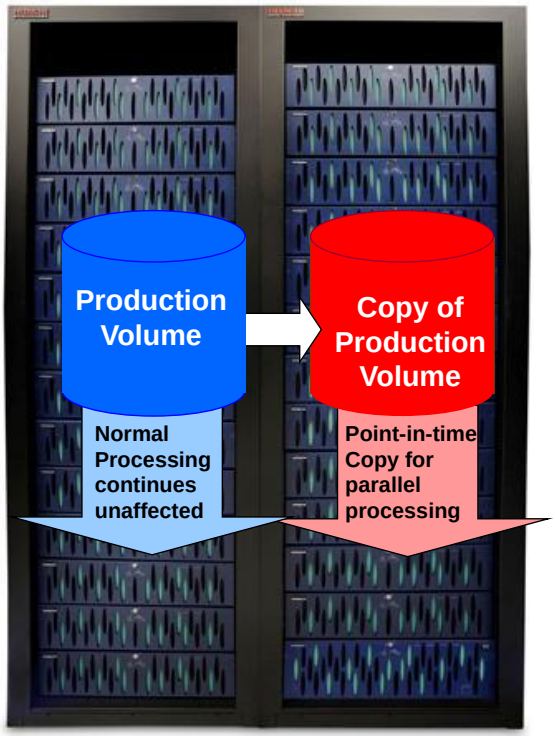
In-System Replication	Remote Replication
Hitachi **ShadowImage** Replication Software	Hitachi **TrueCopy** Remote Replication Software (TC)
Hitachi **Copy-On-Write Snapshot** Software	Hitachi **TrueCopy Extended Distance** Remote Replication Software (TC)

3

HDS Confidential: For distribution only to authorized parties.

ShadowImage Replication Software

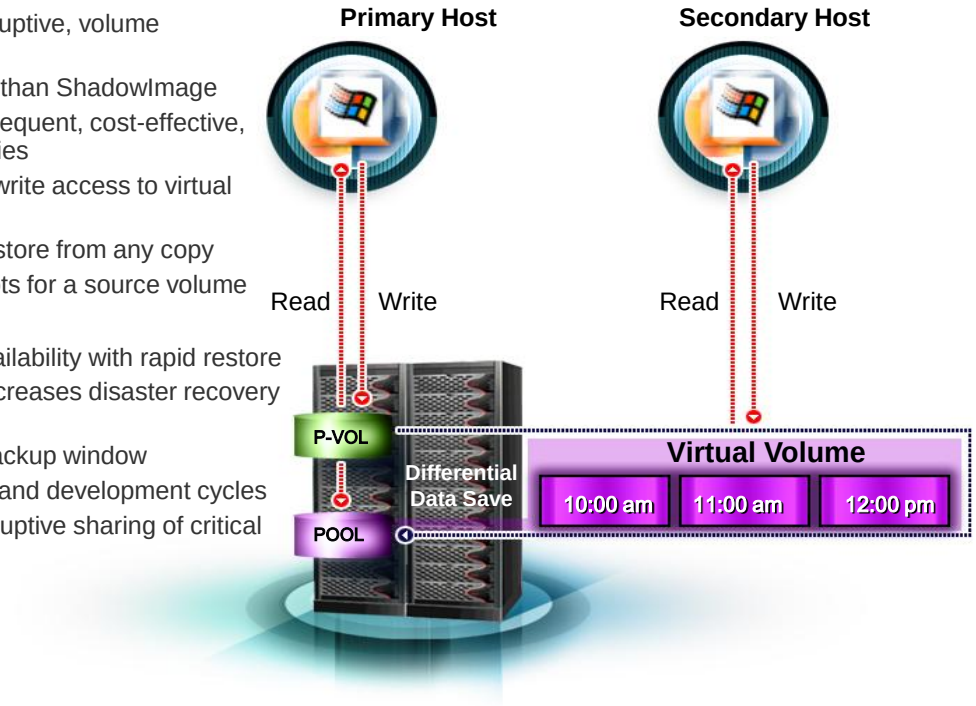
- Features
  - Full physical copy of a volume at a point in time
  - Immediately available for concurrent use by other applications
  - No host processing cycles required
  - No dependence on operating system, file system or database
  - All copies are additionally RAID protected
  - Up to eight copies for a source volume
- Benefits
  - Protects data availability
  - Simplifies and increases disaster recovery testing
  - Eliminates the backup window
  - Reduces testing and development cycles
  - Enables non-disruptive sharing of critical information



4

Copy-on-Write Snapshot Software

- Features
  - Provides nondisruptive, volume "snapshots"
  - Uses less space than ShadowImage
  - Allows multiple frequent, cost-effective, point-in-time copies
  - Immediate read/write access to virtual copy
  - Nearly instant restore from any copy
  - Upto 32 snapshots for a source volume
- Benefits
  - Protects data availability with rapid restore
  - Simplifies and increases disaster recovery testing
  - Eliminates the backup window
  - Reduces testing and development cycles
  - Enables non-disruptive sharing of critical information

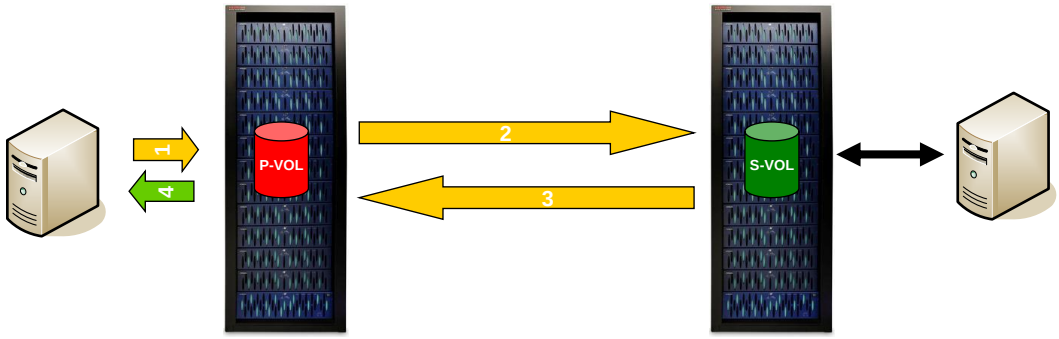


5

HDS Confidential: For distribution only to authorized parties.

TrueCopy Remote Replication Software

- Features
  - Synchronous solution
  - Consistency Group support
  - The remote copy is always a "mirror" image
  - Provides fast recovery with no data loss
  - Supported between AMS 2000 and AMS 1000
- Benefits
  - Improves customer service by reducing downtime of customer-facing applications
  - Increases the availability of revenue producing applications



6

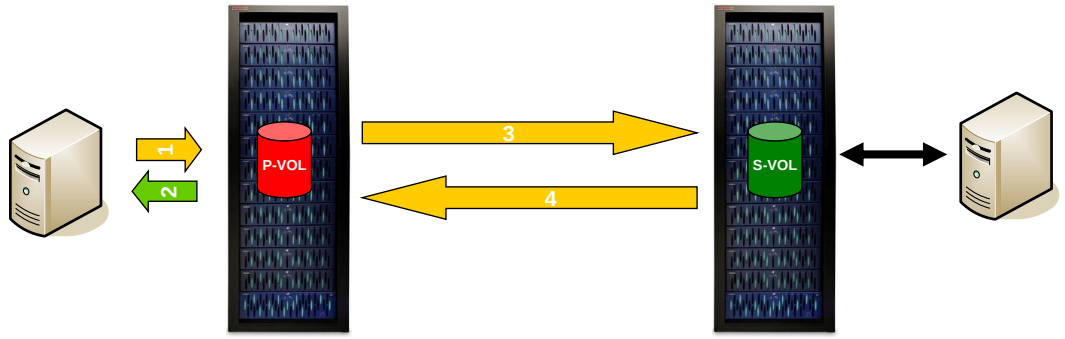
In the diagram and following slides,

P-VOL stands for primary volume.

S-VOL stands for secondary volume.

TrueCopy Extended Distance Replication Software

- Features
  - Asynchronous solution
  - Consistency Group support
  - Write order fidelity preservation
  - Bi-directional
  - Local or remote site back-up capability
- Benefits
  - Long distance replication
  - Improves competitiveness by distributing time-critical information anywhere and anytime



7

In the diagram and following slides,

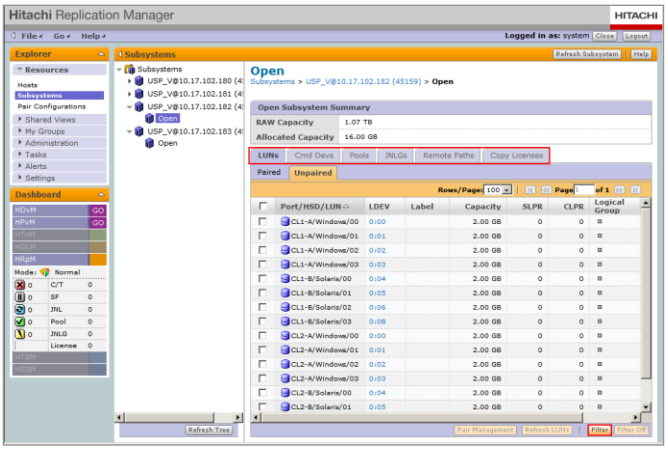
P-VOL stands for primary volume.

S-VOL stands for secondary volume.

HDS Confidential: For distribution only to authorized parties.

Hitachi Replication Manager Software

- Single interface for performing all replication operations including:
  - Managing replication pairs
    - ShadowImage
    - Copy-On-Write
    - TrueCopy Synchronous
    - TrueCopy Extended Distance
  - Configuring
    - Command Devices
    - Differential Management LUs
    - Copy-On-Write Pools
    - TrueCopy Ports
  - Creating Alerts
  - GUI representation of replication environment



8

Tools for Setting up Replication

Graphical User Interface

- Storage Navigator Modular 2
  - Storage Centric
- Hitachi Device Manager
  - Data Centre view of resources, limited or no monitoring options, primary focus is provisioning
- Hitachi Replication Manager
  - Geographically spread Data Centre and Site views, enhanced monitoring and alerting features, primary focus is replication

Command Line Interface

- Used to script replication process
- Command Control Interface software (CCI or RAID Manager)
  - HORCM configuration Files
  - Command Device
- Storage Navigator Modular 2 CLI

9

HDS Confidential: For distribution only to authorized parties.

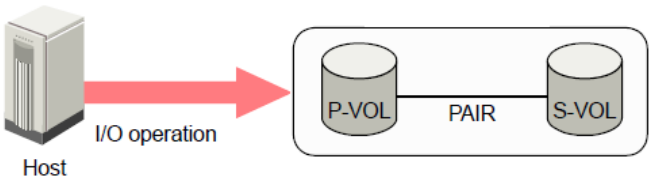
Requirements for all replication products

- Any volumes involved in replication operations (source and destination) should be:
  - Same size (in blocks)
  - Must be mapped to a port
    - Source can be online and in use.
    - Destination must not be in use/mounted.
- Intermix of RAID levels and drive type is supported.
- Licensing
  - Purchase the license with enough licensed capacity to accommodate the pairs you are going to create.
    - If you exceed the licensed capacity,
      - Existing operations will continue
      - You will not be able to perform new replication operations
  - Source, destination, and reserved volumes require licenses for volume capacity and management.

10

Replication Operations

- Basic operations when working with replication products
  - Paircreate
  - Pairsplit
  - Pairresync
  - Pairsplit –S
- Commands are consistent across products (in-system or remote replication) but implementation varies depending on the product
  - In-system – All operations with LDEVs pairings within the same storage system
  - Remote – All operations with LDEVs pairings across different storage systems
  - Use manual to identify product specific operations with above commands
- A volume of original data is called a Primary Volume (P-VOL), and a volume to which the data is copied is a Secondary Volume (S-VOL).

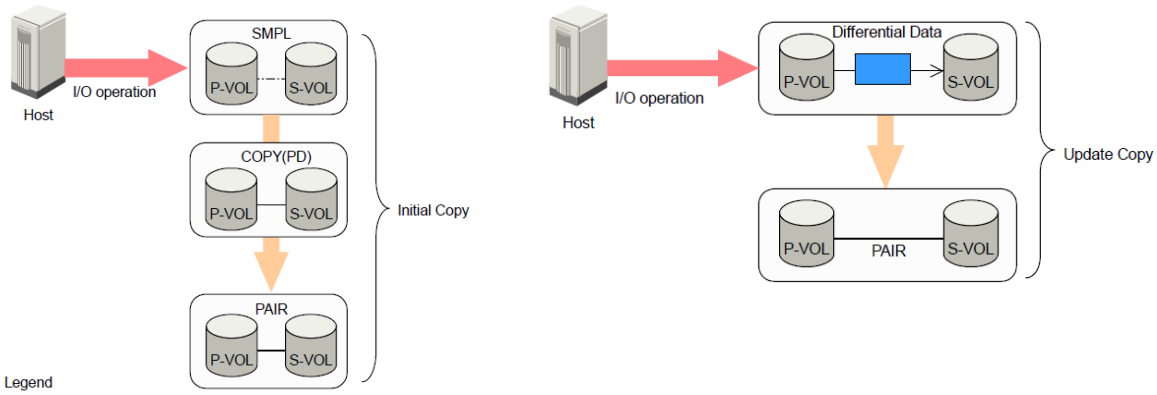


11

HDS Confidential: For distribution only to authorized parties.

Copy Operations

- Data copy operations
  - Initial Copy
    - Results in all data being copied from PVOL to SVOL
    - Copies everything including empty blocks.
  - Update Copy
    - Only differentials are copied



12

Replication Operations

- Paircreate
  1. Select a volume and issue **paircreate.**
  2. Initial copy (full copy) takes place. Block by block copy of P-VOL to S-VOL happens regardless of the amount of data in P-VOL.
  3. Volume status changes from SMPL to PAIR.

	P-VOL		S-VOL	
Time frame	**Vol Status**	**Host Access**	**Vol Status**	**Host Access**
Before	SMPL	R/W	SMPL	R/W OR R
During	COPY(PD)	R/W	COPY(PD)	NA
After	PAIR	R/W	PAIR	NA

13

The **paircreate** command generates a new volume pair from two unpaired volumes. The **paircreate** command can create either a paired logical volume or a group of paired volumes.

When issuing **paircreate**, you can select the pace for the initial copy operation:

- Slower
- Medium
- Faster

The slower pace minimizes the impact of software 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.

HDS Confidential: For distribution only to authorized parties.

- Pairsplit
  - Update copy takes place to flush all pending changes.
  - Volume status changes to PSUS.
  - S-VOL is the PiT copy and now available to applications for read/write.
  - Differential bitmaps track changes to P-VOL and S-VOL while pair is split.

	P-VOL		S-VOL	
Time frame	**Vol Status**	**Host Access**	**Vol Status**	**Host Access**
Before	PAIR	R/W	PAIR	NA
During	COPY(SP)	R/W	COPY(SP)	NA
After	PSUS	R/W	SSUS	R/W

14

The **pairsplit** command stops updates to the secondary volume of a pair and can either maintain (status = PSUS) or delete (status = SMPL) the pairing status of the volumes. The **pairsplit** command can be applied to a paired logical volume or a group of paired volumes. The **pairsplit** command allows read access or read/write access to the secondary volume, depending on the selected options.

- Pairresync
  - S-VOL is no longer available to host.
  - S-VOL and P-VOL differential bitmaps are merged.
  - Changed tracks are marked and written from P-VOL to S-VOL.
  - Volume status changes to pair.

	P-VOL		S-VOL	
Time frame	Vol Status	Host Access	Vol Status	Host Access
Before	PSUS	R/W	SSUS	R/W
During	COPY(RS)	R/W	COPY(RS)	NA
After	PAIR	R/W	PAIR	NA

15

Replication software allows you to perform **pairresync** operations on split and suspended pairs.

Pairresync for split pair. When a pairresync operation is performed on a split pair (status = PSUS), the system merges the S-VOL bitmap into the P-VOL bitmap and then copies all flagged blocks from the P-VOL to the S-VOL. This also greatly reduces the time needed to resynchronize the pair.

HDS Confidential: For distribution only to authorized parties.

- Pairsplit –S (delete)
  - Delete the pair and stop replication operations for the pair.
  - Immediate access to S-VOL.
    - No update copy (pending changes are lost/ignored)
  - Changes volume status back to simplex.

	P-VOL		S-VOL	
Time frame	**Vol Status**	**Host Access**	**Vol Status**	**Host Access**
Before	PAIR	R/W	PAIR	NA
After	SMPL	R/W	SMPL	R/W

16

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

When a pair is deleted, the pending update copy operations for the pair are discarded, and the status of the P-VOL and S-VOL is changed to SMPL.

HDS Confidential: For distribution only to authorized parties.

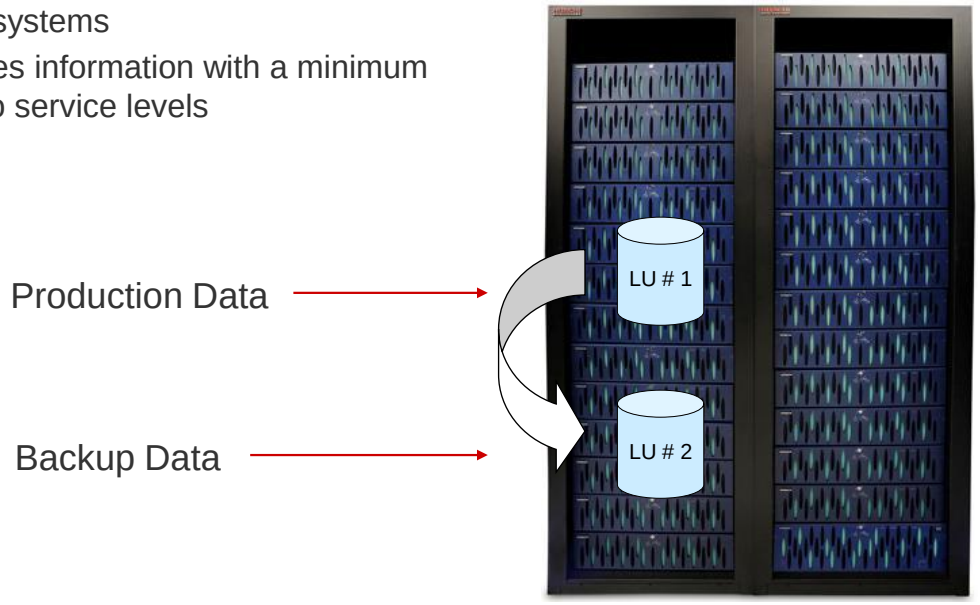
9. ShadowImage Replication and Copy-on-Write Snapshot Software

Module Objectives

- Upon completion of this module, the learner will be able to:
  - Describe the key features and operations of ShadowImage Replication software
  - List the rules associated with the operations
  - Describe the key competitive advantages
  - Identify commands and the operations of each command
  - Describe the purpose of Copy-on-Write Snapshot software
  - List key specifications
  - Compare the functionality of Copy-on-Write Snapshot to ShadowImage Replication software
  - Describe typical Copy-on-Write Snapshot operations

Overview

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



3

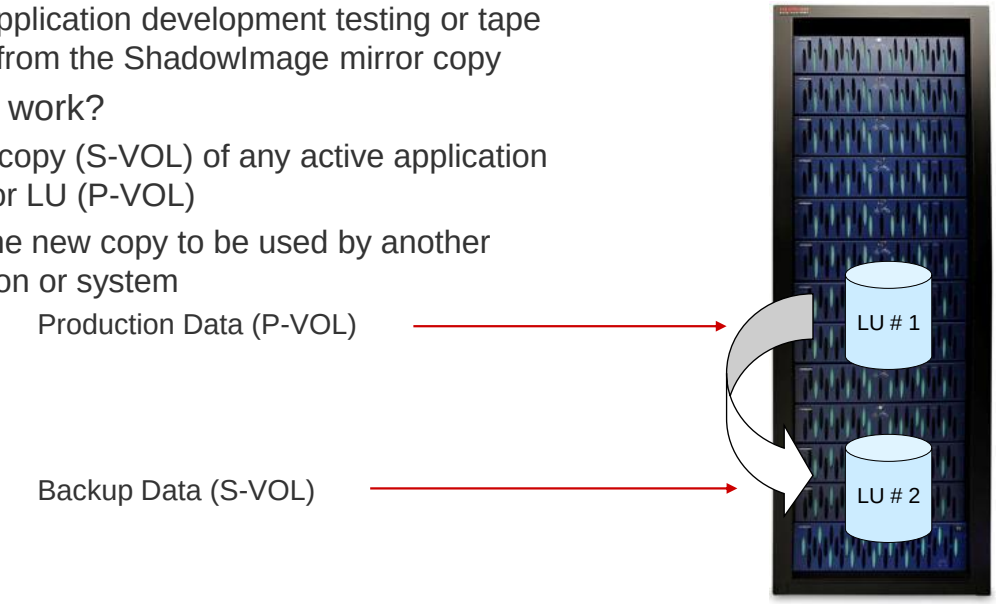
ShadowImage software is the in-system copy facility for Adaptable Modular Storage 2000 family. It enables server-free backups, which allows customers to exceed service level agreements (SLAs). It fulfills two primary functions:

- 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.

HDS Confidential: For distribution only to authorized parties.

- 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 up to eight secondary volumes (S-VOLs) 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

HDS Confidential: For distribution only to authorized parties.

Benefits

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

6

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 high performance active-active 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.

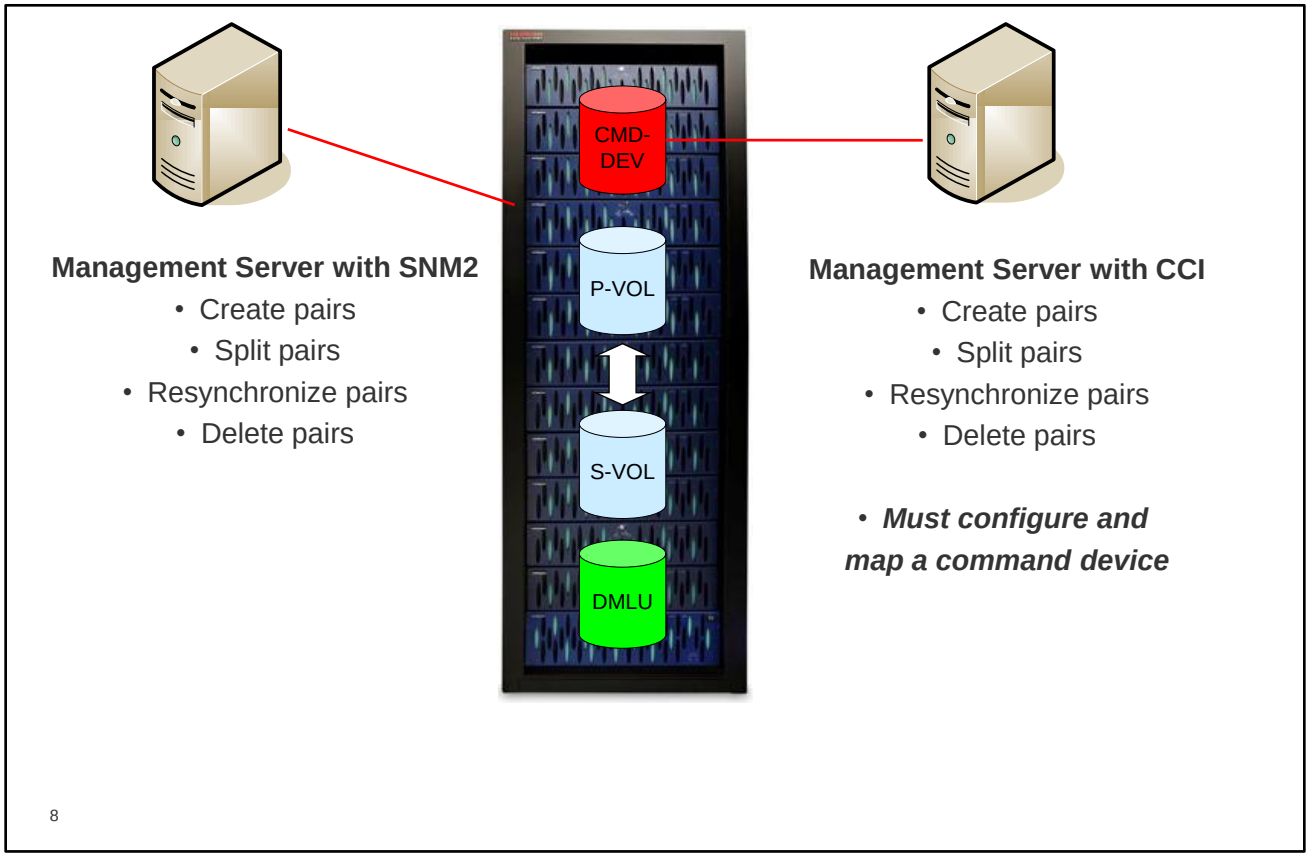
General Replication Terminology

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

7

Note: The DM-LU does not store the actual changed data. It just has the metadata for the differential data.

Management Resources



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

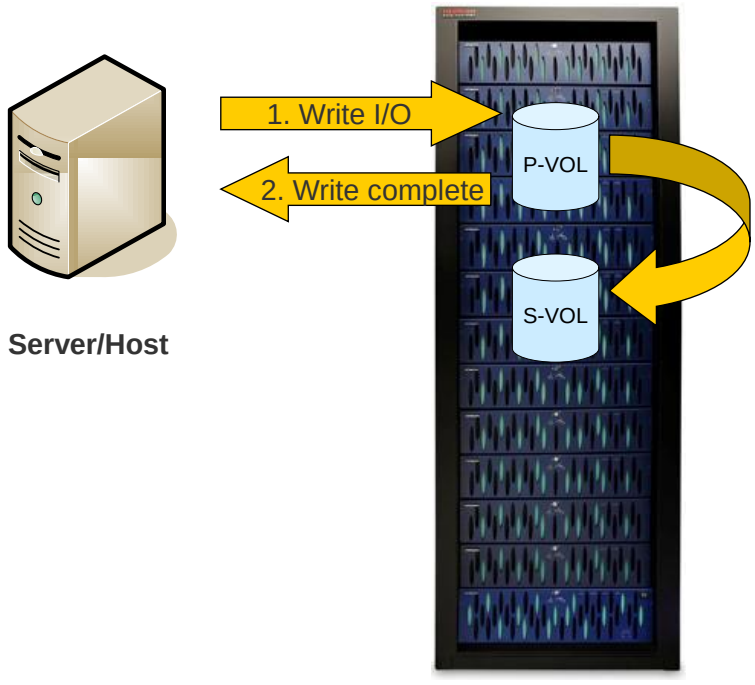
HDS Confidential: For distribution only to authorized parties.

Rules

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

9

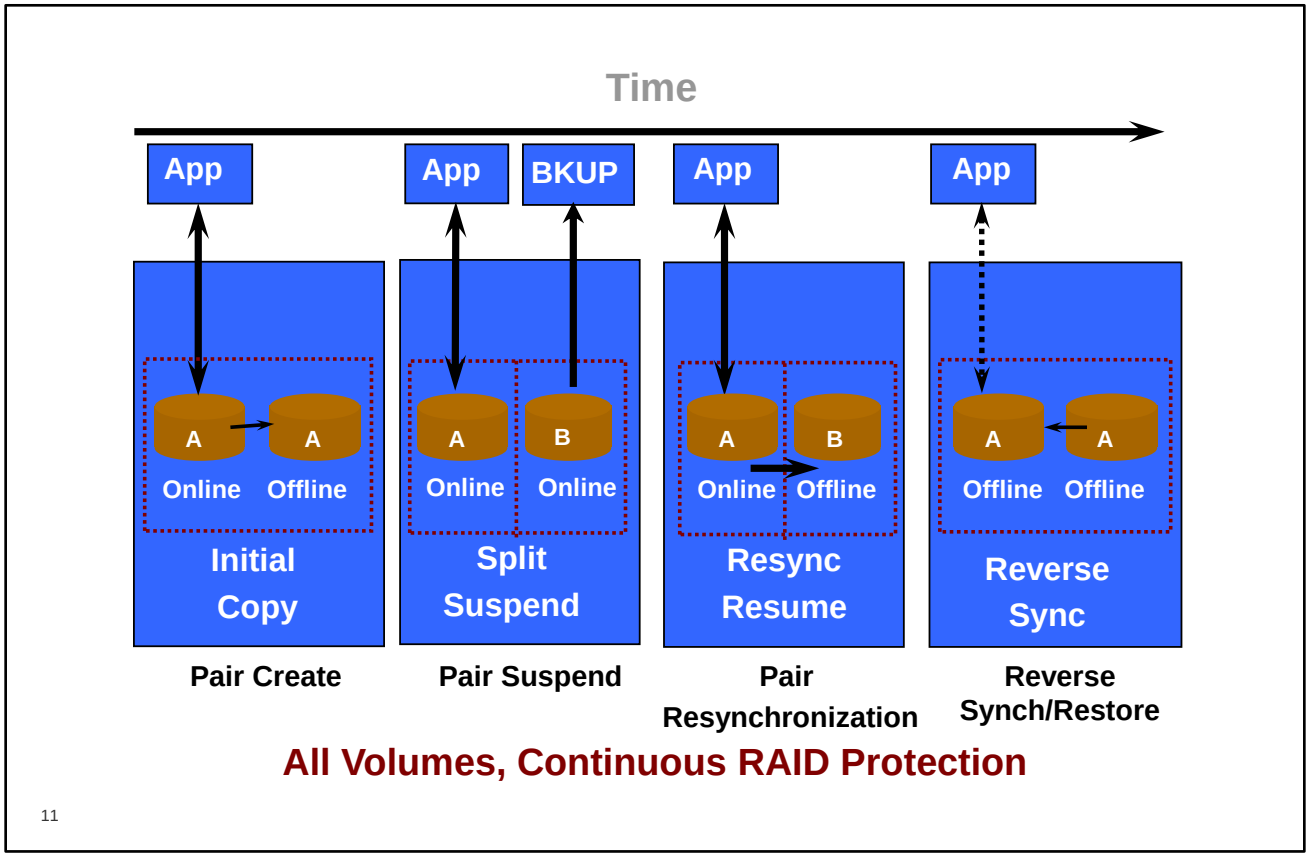
Internal ShadowImage Operation



10

HDS Confidential: For distribution only to authorized parties.

Operations



ShadowImage operations include:

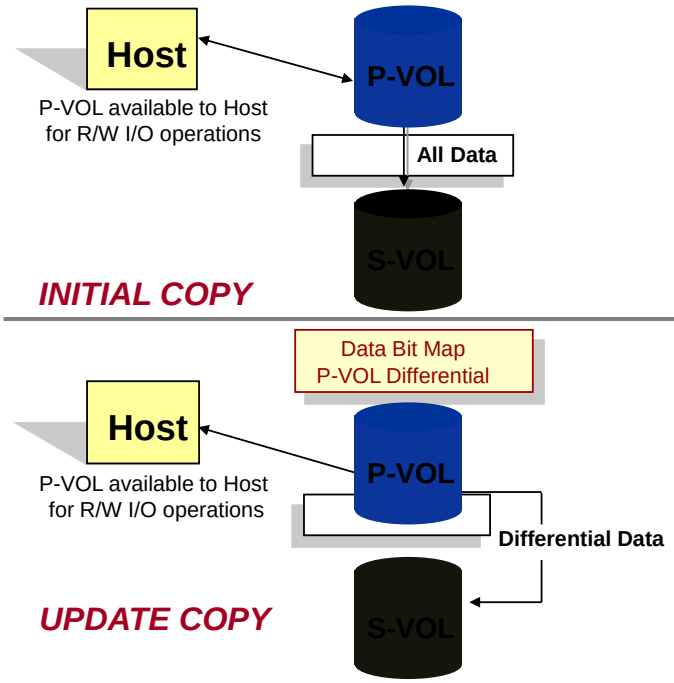
paircreate

pairsplit

pairresynchronize

paircreate Command

- Establishes a new ShadowImage pair



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:

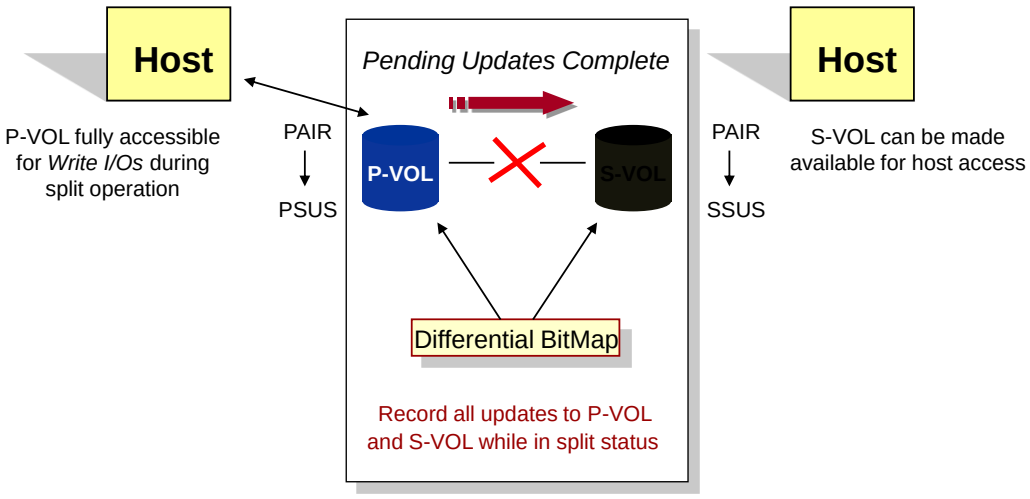
- Slower
- Medium
- Faster

HDS Confidential: For distribution only to authorized parties.

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



13

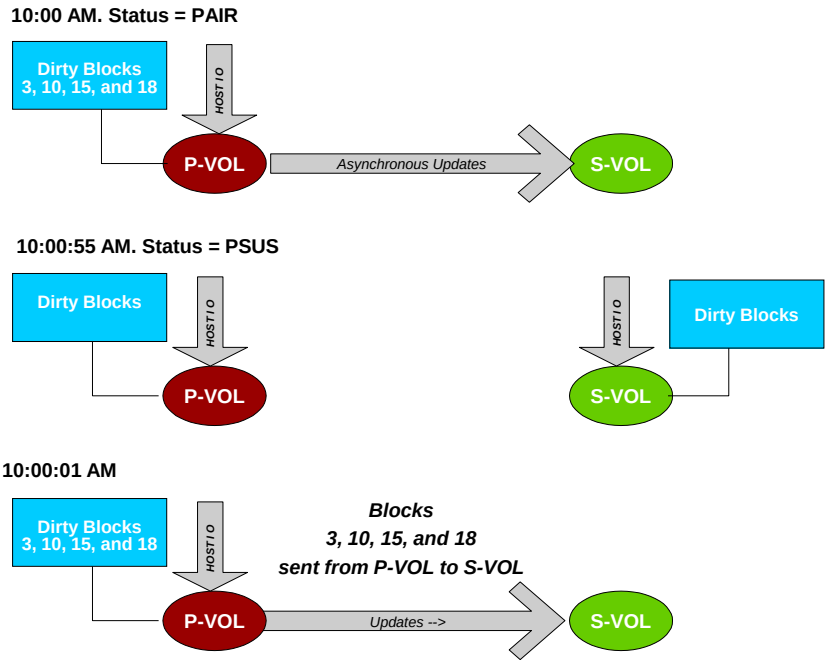
The ShadowImage **pairsplit** operation performs all pending S-VOL updates (those issued prior to the split command and recorded in the P-VOL bitmap) to make the S-VOL identical to the state of the P-VOL when the suspend command was issued and then provides full read/write access to the split S-VOL.

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 system establishes a bitmap 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.

HDS Confidential: For distribution only to authorized parties.

- pairsplit Illustration



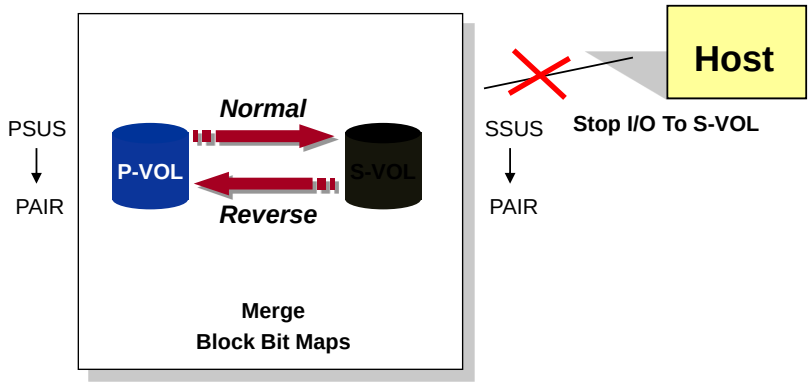
14

The P-VOL and S-VOL are in PAIR status as of 10:00 AM. Blocks 3, 10, 15, and 18 are marked as dirty because of host I/O.

The status of the P-VOL and S-VOL is changed to PSUS. During this state there are block bitmaps attached to both the P-VOL and the S-VOL. These bitmaps keep block of changes on both the P-VOL and the S-VOL.

Blocks 3, 10, 15, and 18 are sent across to the S-VOL from the P-VOL.

pairresync Command



Sync S-VOL with P-VOL

NORMAL

Copy direction = P-VOL to S-VOL

S-VOL bitmap merged into P-VOL bitmap — copies all "flagged" blocks 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 bitmap merged into S-VOL bitmap — copies "flagged" blocks from S-VOL to P-VOL

P-VOL is unavailable for R/W I/Os

15

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.

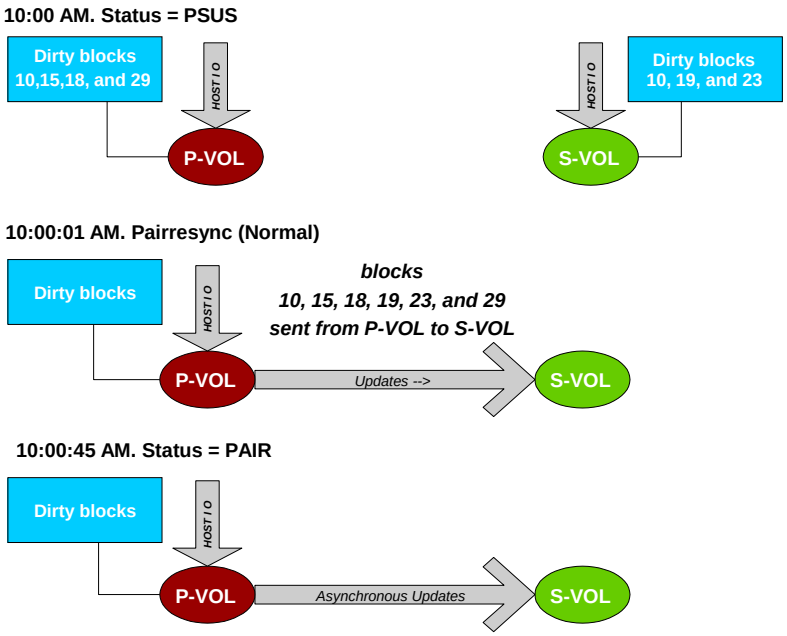
Reverse. The reverse pairresync (**pairresync –restore**) operation synchronizes the P-VOL with the S-VOL. The copy direction for a reverse pairresync operation is S-VOL to P-VOL. The pair status during a reverse resync operation is COPY(RS-R), and the S-VOL becomes inaccessible to all hosts for write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

HDS Confidential: For distribution only to authorized parties.

When a pairresync operation is performed on a suspended pair (status = PSUS), the storage system merges the S-VOL differential bitmap into the P-VOL differential bitmap 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 bitmap into the S-VOL differential bitmap and then copies all flagged blocks 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



16

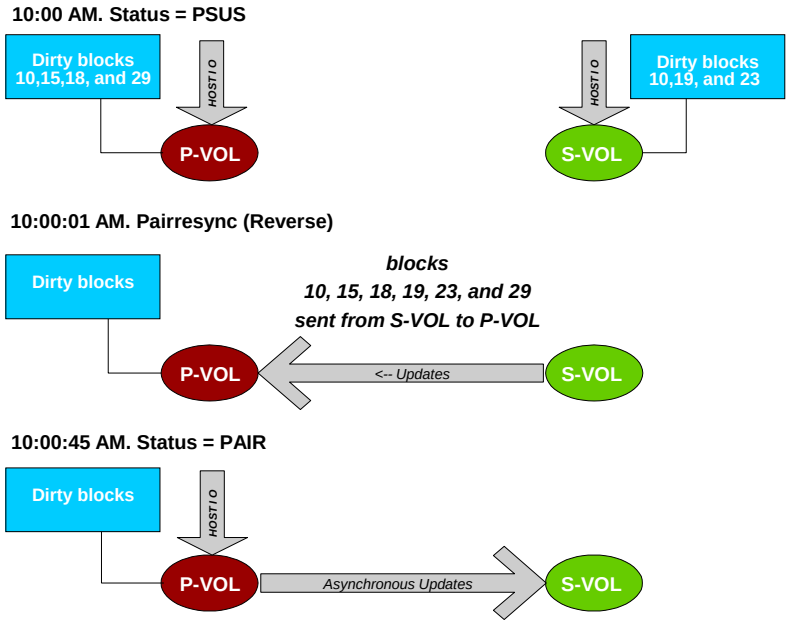
The status of the P-VOL and the S-VOL is PSUS as of 10:00 AM. blocks 10, 15, 18, and 19 are marked as dirty on the block bitmap for the P-VOL. blocks 10, 19 and 23 are marked as dirty on the block bitmap for the S-VOL.

At 10:00 AM, a **pairresync** (normal) command is issued. The block bitmaps for the P-VOL and S-VOL are merged. The resulting block bitmap has blocks 10, 15, 18, 19, 23, and 29 marked as dirty. These blocks are sent from the P-VOL to the S-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.

HDS Confidential: For distribution only to authorized parties.

- Reverse resync Illustration



17

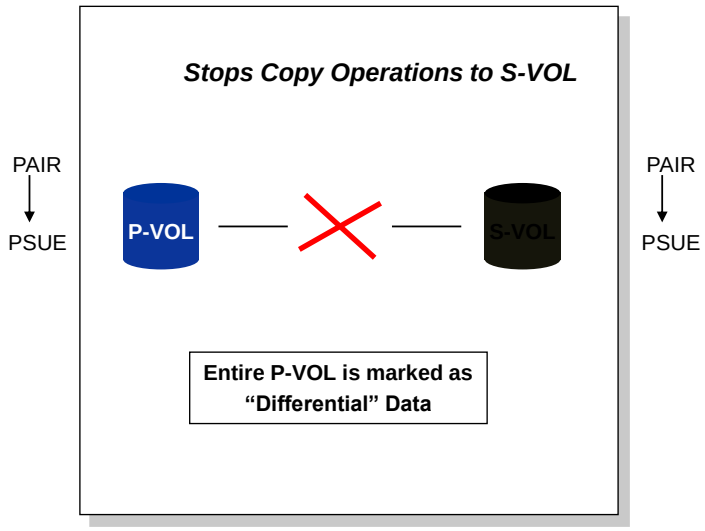
The status of the P-VOL and the S-VOL is PSUS as of 10:00 AM. blocks 10, 15, 18, and 19 are marked as dirty on the block bitmap for the P-VOL. blocks 10, 19, and 23 are marked as dirty on the block bitmap for the S-VOL.

At 10:00 AM a **pairresync** (normal) command is issued. The block bitmaps for the P-VOL and S-VOL are merged. The resulting block bitmap has blocks 10, 15, 18, 19, 23, and 29 marked as dirty. These blocks 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



18

When a ShadowImage pair is suspended due to error or manually using **pairsplit –E** option (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 as differential 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:

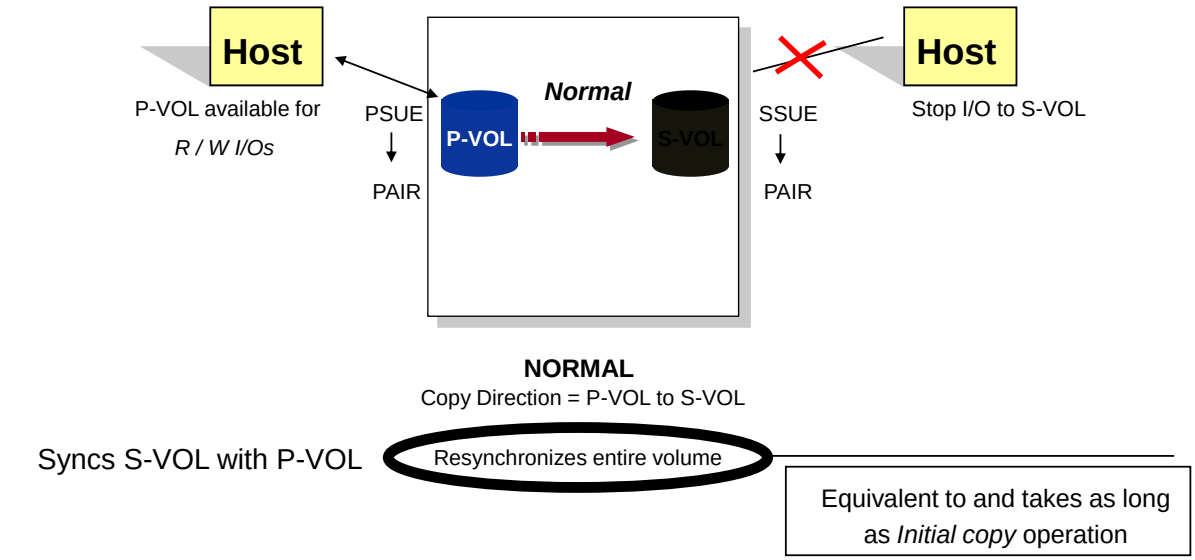
When the ShadowImage volume pair has been suspended or deleted.

When the storage system detects an error condition related to an initial copy operation.

HDS Confidential: For distribution only to authorized parties.

Resynchronize Suspended Pair

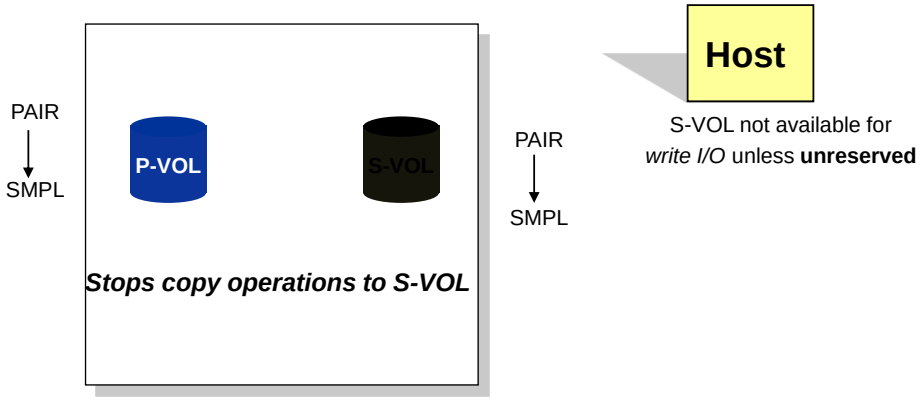
- pairresync on a pair suspended due to error (PSUE)



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)



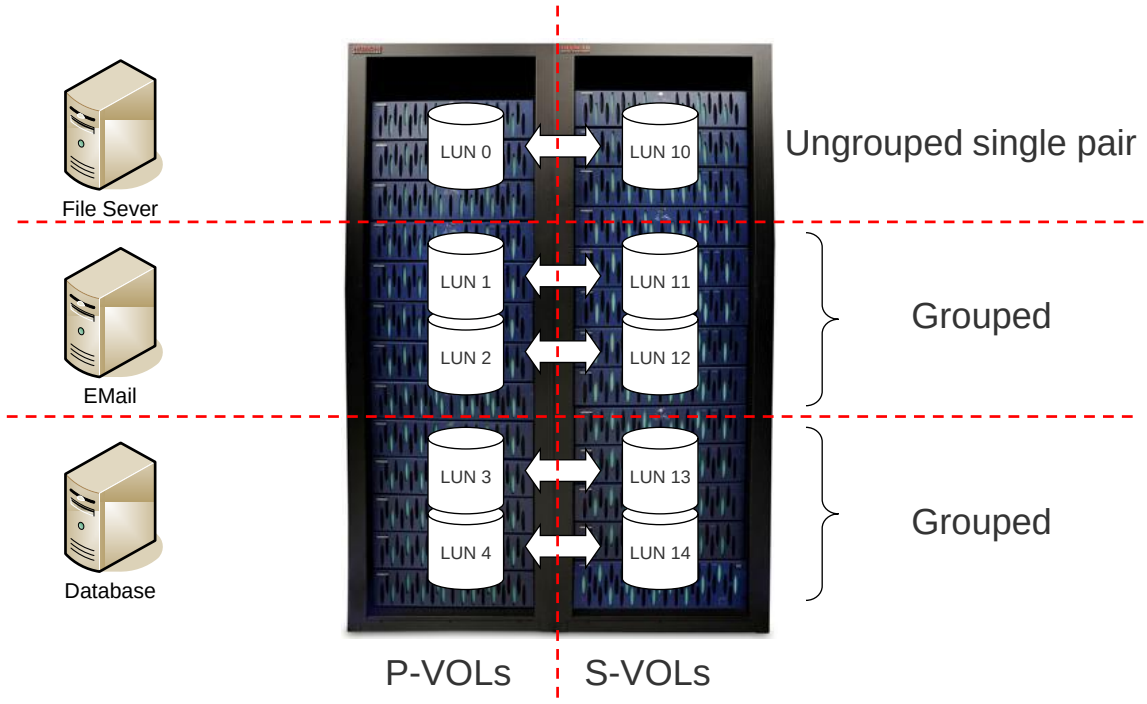
20

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.

HDS Confidential: For distribution only to authorized parties.

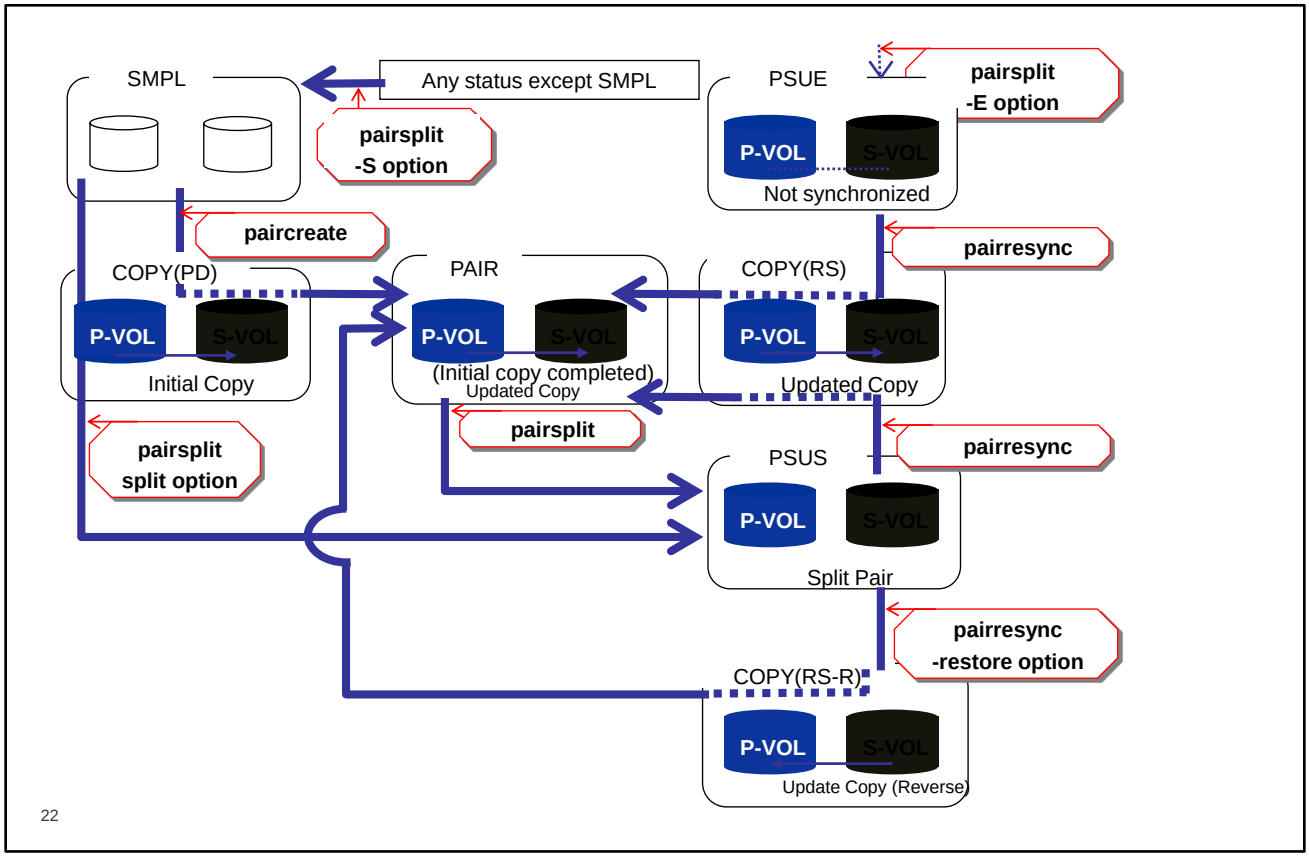
Volume Grouping

- Perform operations for pairs or a group of pairs



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.

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 storage system 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.

HDS Confidential: For distribution only to authorized parties.

Copy-on-Write Snapshot Software

Copy-on-Write Snapshot Software

Overview

- Copy-on-Write Snapshot software is a storage-based hardware solution for point-in-time backups, or snapshots.
- Operations can be performed from SNM2 GUI/CLI
- A volume by Copy-on-Write Snapshot consists of physical data stored in the primary volume and differential data stored in the Data Pool.
- 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.

24

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?

ShadowImage software for backup solutions raises the entire system price as the storage requirement grows based on number of copies required.

Snapshot function requires less capacity to maintain virtual copies.

- 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 Storage Navigator Modular 2 GUI, CLI, or CCI.

25

- SnapShot creates a virtual duplicate volume of another volume.
- This volume “pair” is created when you:
  - Select a volume that you want to replicate
  - Identify another volume that will contain the copy (virtual volume)
  - Associate the primary and secondary volumes
  - Create a snapshot of primary volume data in the virtual (secondary) volume.
    - Once a snapshot is made, it remains unchanged until a new snapshot instruction is issued. At that time, the new image replaces the previous image.

26

Components

- Primary Volume (P-VOL)
  - The storage volume in a volume pair. It is used as the source of a copy operation.
  - Can be cascaded with TrueCopy Remote Replication Software
- 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.
  - A data pool can be shared by multiple SnapShot pairs.
  - Up to 64 LUs can be assigned to a data pool.
  - Hitachi recommends a minimum of 20 GB for data pool size.
    - If the pool is not sized properly and becomes full during operations. The pairs will fail and all snapshots using the pool will be lost.

28

HDS Confidential: For distribution only to authorized parties.

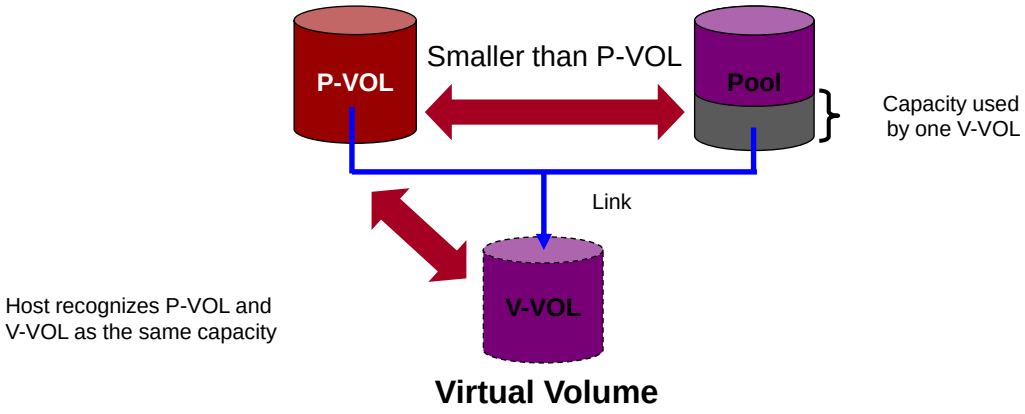
Comparison with ShadowImage

	Copy-on-Write Virtual Volume	ShadowImage Clone
Features	Provides a very quick copy because it is a virtual volume	Superior protection because it is a complete copy Immediate restore if P-VOL becomes corrupted
Time to Create	Instantaneous	From minutes to hours, depending on the size of the P-VOL
Disk Space Used	Size will vary depending upon rate of data change but will be far less than the P-VOL	Same amount as P-VOL
Data Recovery Time After P-VOL is corrupted	Two stage restore: 1. V-VOL verify, data not available to host 2. Copy back from a tape storage or else, data available	Instantaneous
Size of Physical Volume	P-VOL ≥ Pool for one V-VOL	P-VOL = S-VOL
Pair Configuration	P-VOL : V-VOL = 1 : 32	P-VOL : S-VOL = 1 : 8
Restore	P-VOL can be restored from any V-VOL	P-VOL can be restored from S-VOL

29

Volume Size

- 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.



30

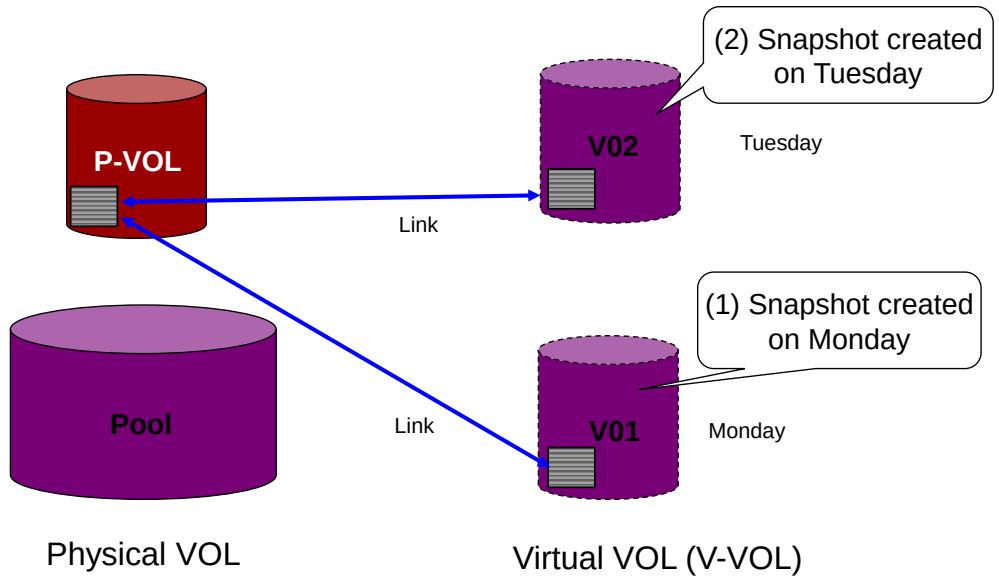
In Copy-on-Write Snapshot, the **V-VOL** does not physically exist but is a representation of a **set of pointers** that point to the locations where the data is physically located, partly in the Pool and partly in the P-VOL. Since only part of the data belonging to the V-VOL is located in Pool (and 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.

HDS Confidential: For distribution only to authorized parties.

Operation Scenarios

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

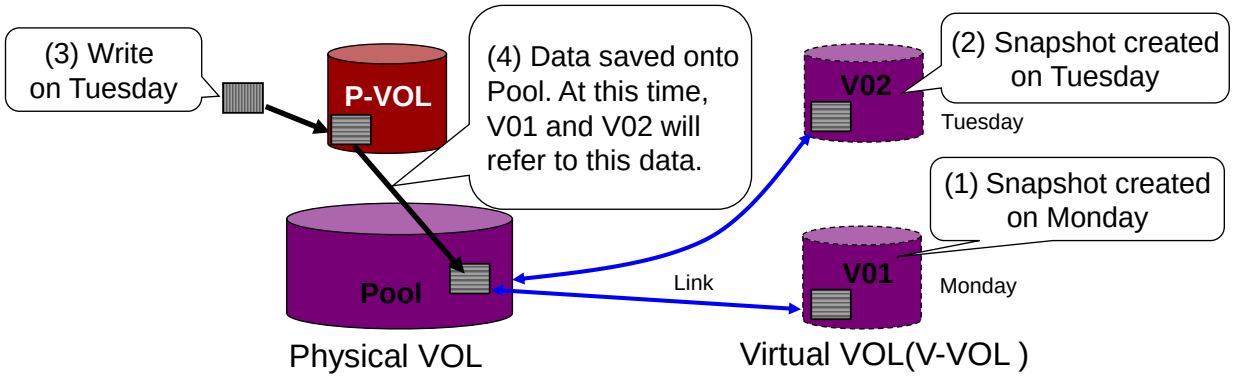


31

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.

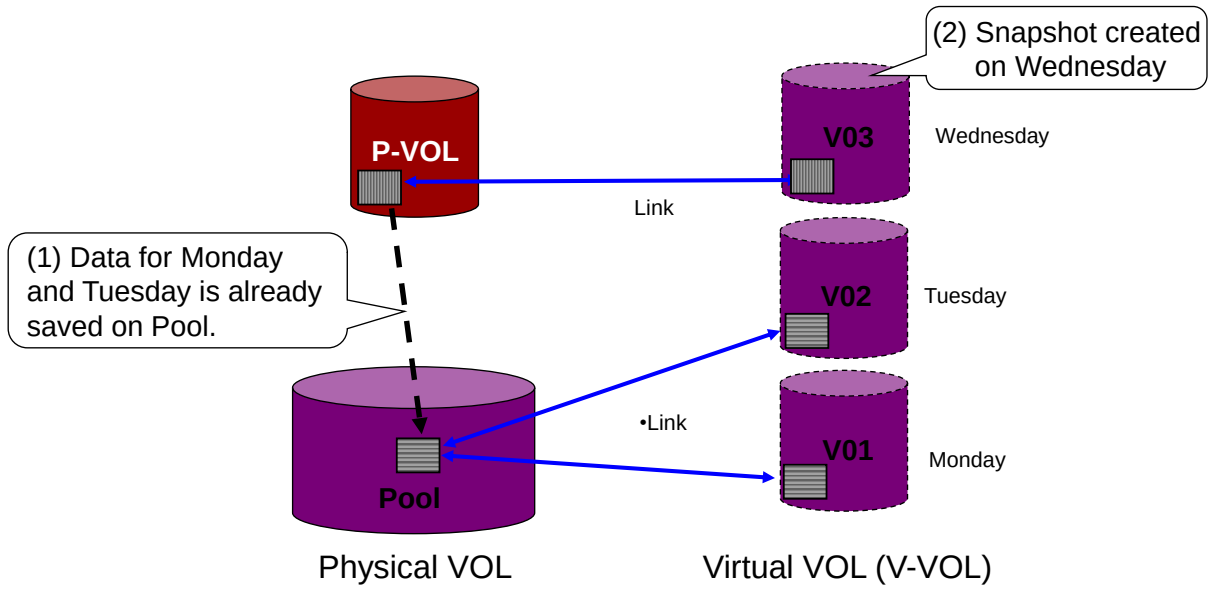
32

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.

HDS Confidential: For distribution only to authorized parties.

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



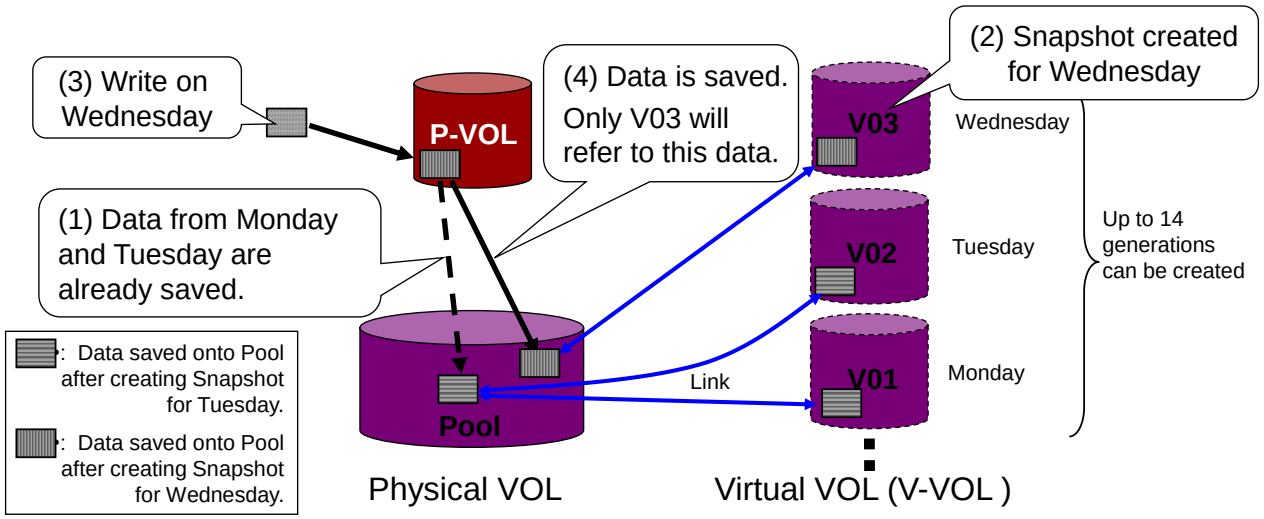
33

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.

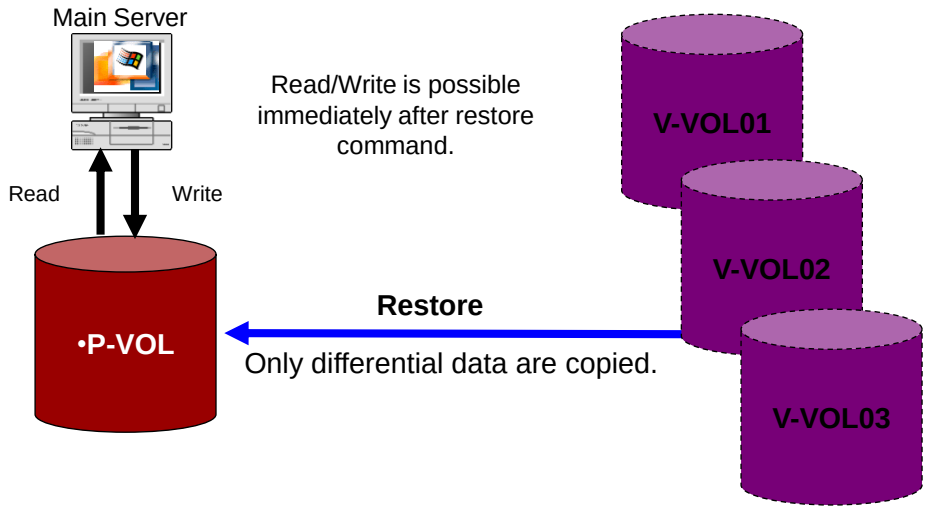


34

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.

HDS Confidential: For distribution only to authorized parties.

- 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



35

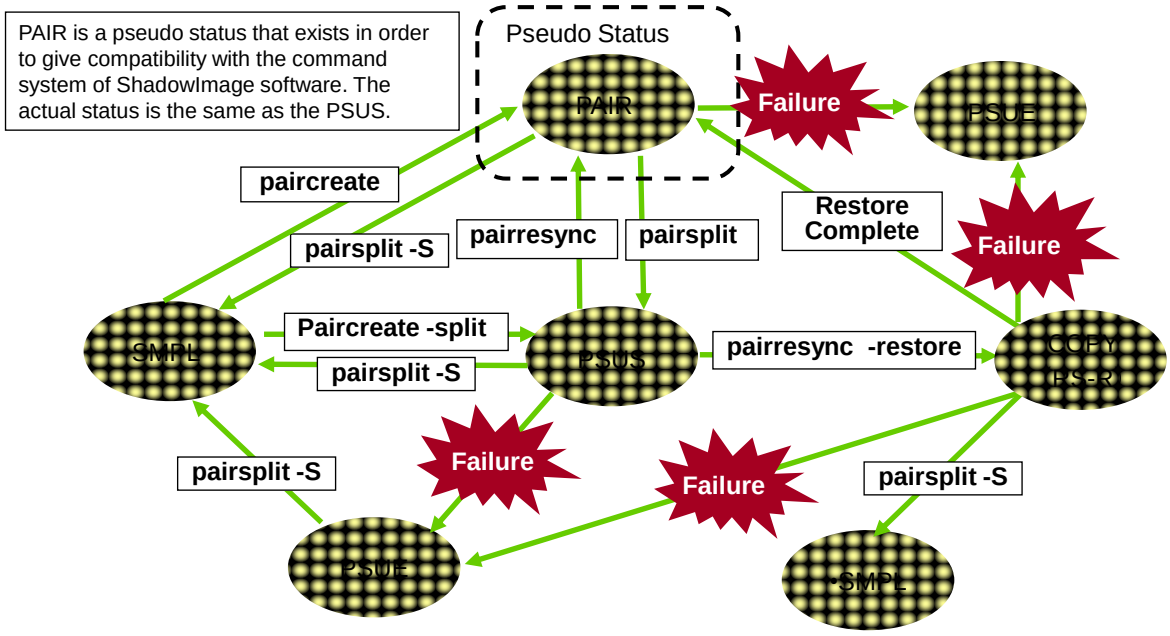
Restoration of a primary volume can be done instantly from any V-VOL because it does not involve immediate moving of data from Pool to P-VOL only pointers need to be modified.

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.

Status Transitions

- Status and commands for Copy-on-Write Snapshot



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.

HDS Confidential: For distribution only to authorized parties.

Recorded Product Demonstration

Now… Recorded Product Demonstration

From your Web Browser …
http://edemo.hds.com/edemo/vault1/InSystemReplication.swf

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.
4. What does Copy-on-Write Snapshot software do?
5. Why does Copy-on-Write Snapshot require less space than the source volumes?
6. Why is a Pool necessary for Copy-on-Write Snapshot?

38

10. Hitachi TrueCopy Remote Replication Software

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

Description

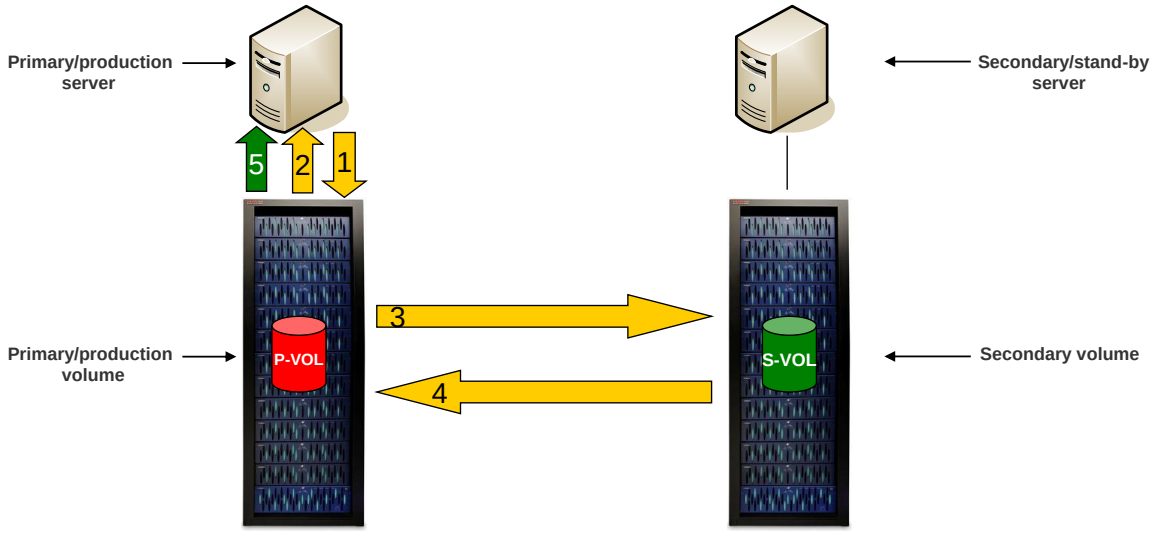
- TrueCopy Synchronous Software (at LU level)
  - Purpose:
    - Data migration
    - Data replication
    - Disaster recovery (DR)
  - Benefits:
    - Reliable, seamless management
    - High performance, cost effective solution

3

Data migration is typically a one time occurring event. Data replication occurs repeatedly at a set schedule.

HDS Confidential: For distribution only to authorized parties.

- 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

4

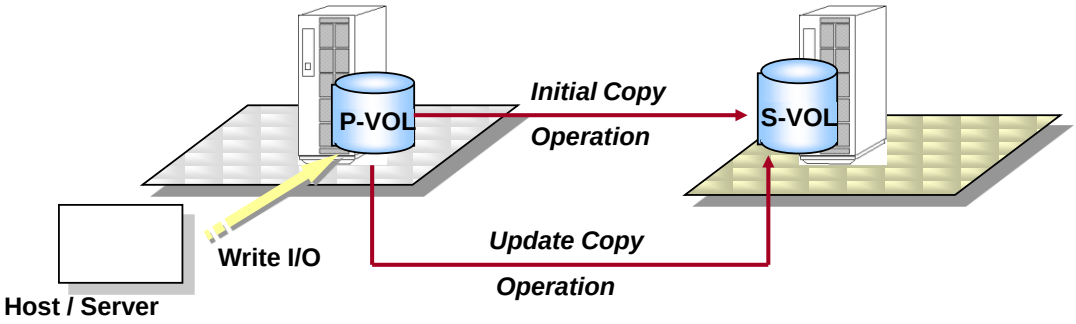
Advantages

- 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

5

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.

HDS Confidential: For distribution only to authorized parties.

Specifications

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

6

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
  - Preparation
    - 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.

7

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

HDS Confidential: For distribution only to authorized parties.

Preliminary Steps

- Check the licenses on the local and the remote storage system.
- You must set up and configure, on the local *and* the remote system:
  - DM-LU
  - Command Devices
  - Remote copy paths
    - One remote path must be set up per controller, two paths per array.
    - With two paths, an alternate path is available in the event of link failure during copy operations.

8

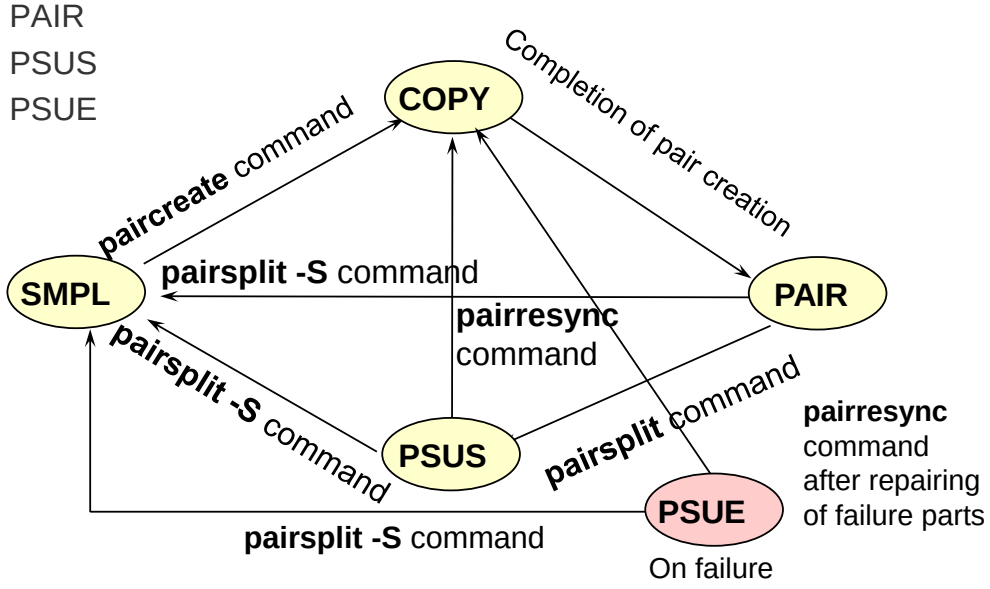
- Detailed Preliminary Steps
  - Find the following information from the customer:
    - Identify the master and the remote systems.
    - Use SNM2 to identify the system serial numbers.
  - Use SNM2
    - At Local Site – record the P-VOL port, LUN, host group and size.
    - At Remote Site – record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

9

HDS Confidential: For distribution only to authorized parties.

Volume States

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



10

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).

PSUE (Pair suspend-error) – This pair is not synchronized, because the MCU or RCU has suspended this pair due to an error condition.

TrueCopy Extended Distance

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

11

HDS Confidential: For distribution only to authorized parties.

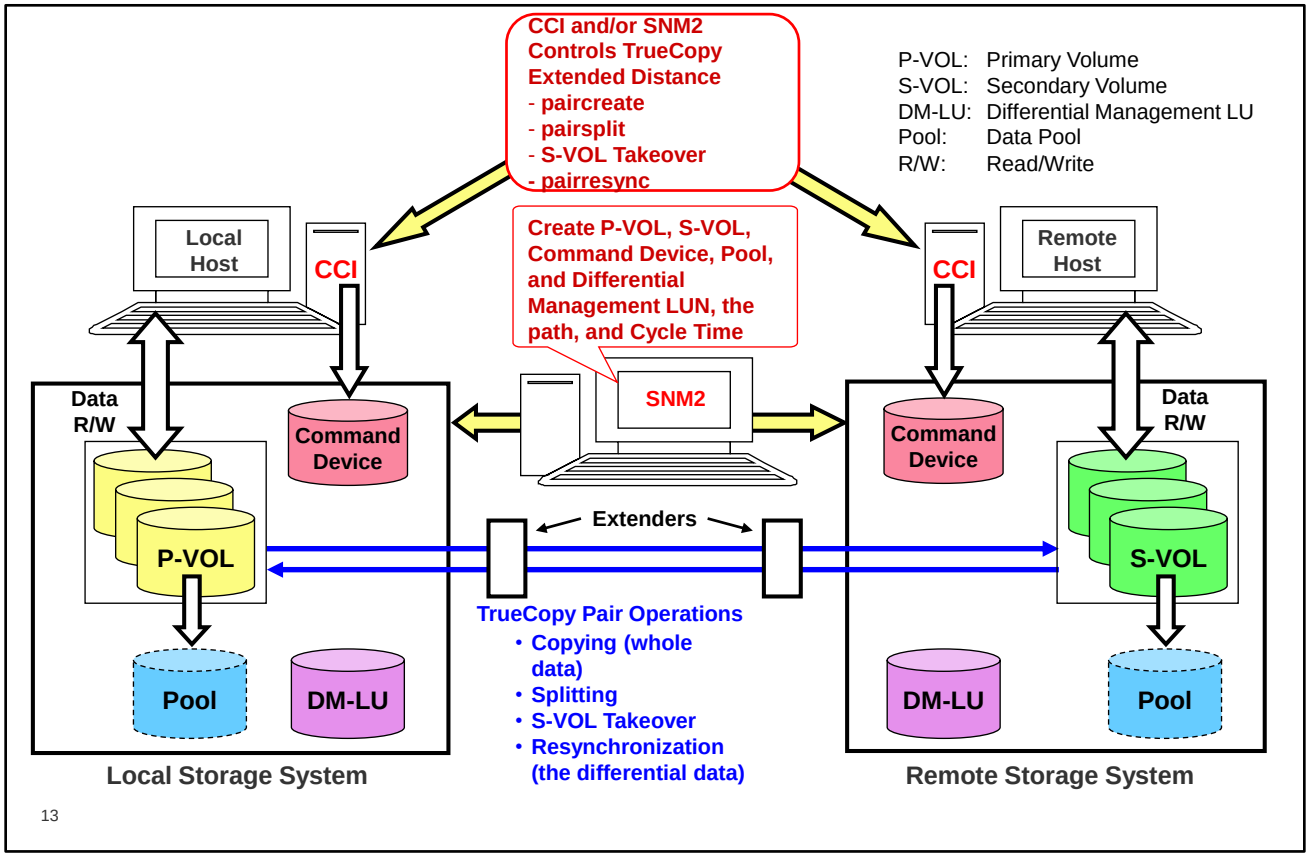
Business Needs

- 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 storage systems.
    - 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

12

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

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.

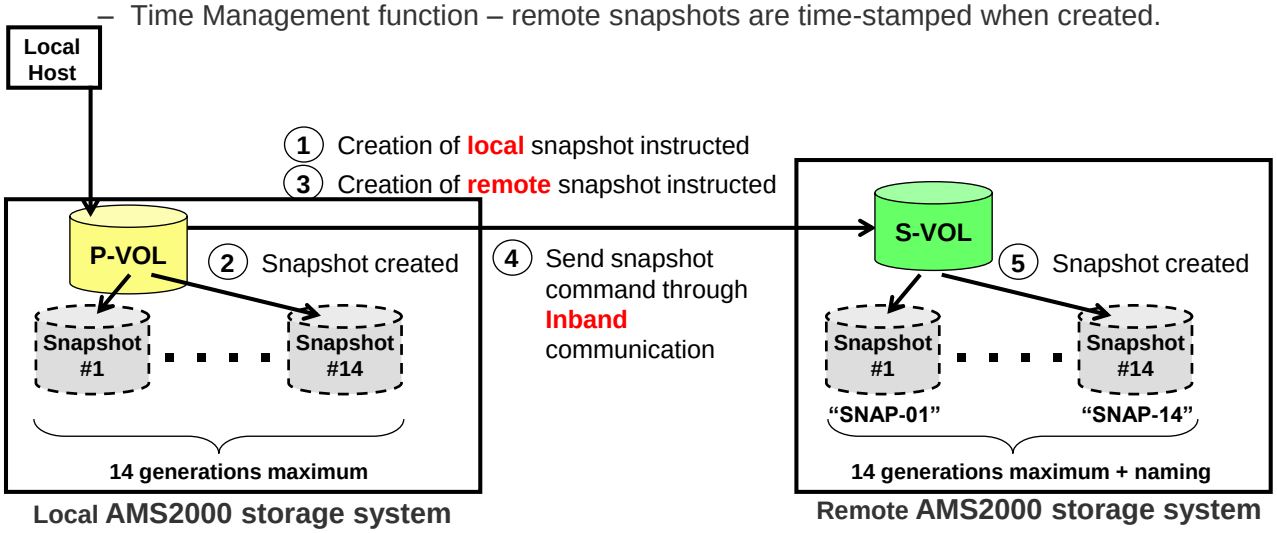
Pool in the remote storage system is used for saving consistent, internally determined data.

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

HDS Confidential: For distribution only to authorized parties.

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.



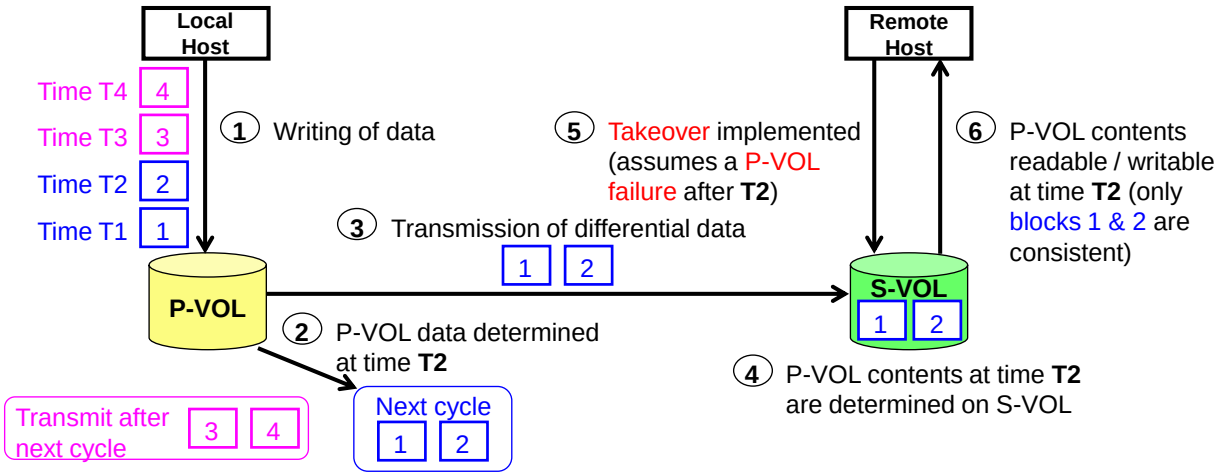
Time Management function

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** DF800 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.
- 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.

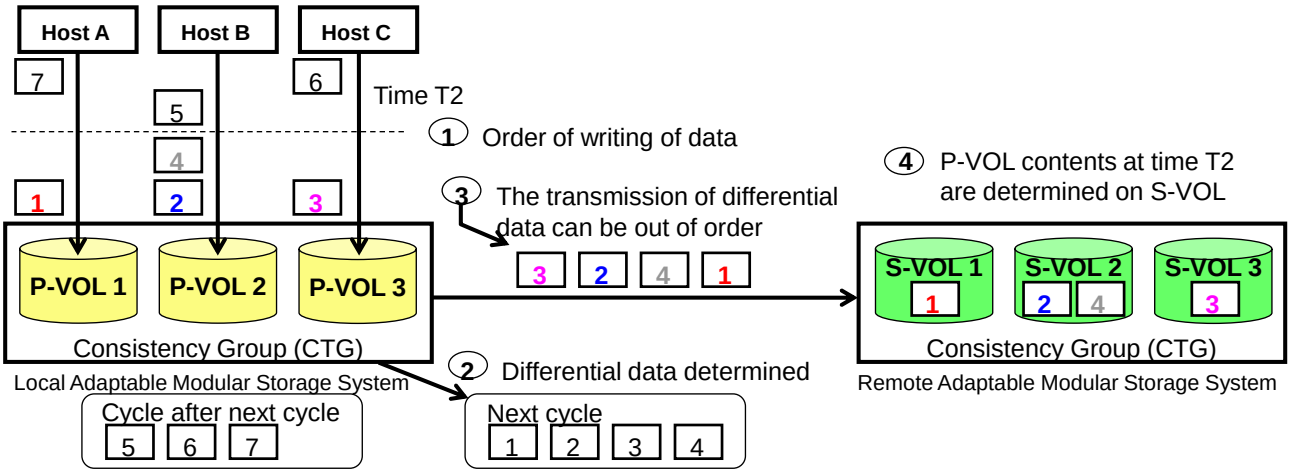


15

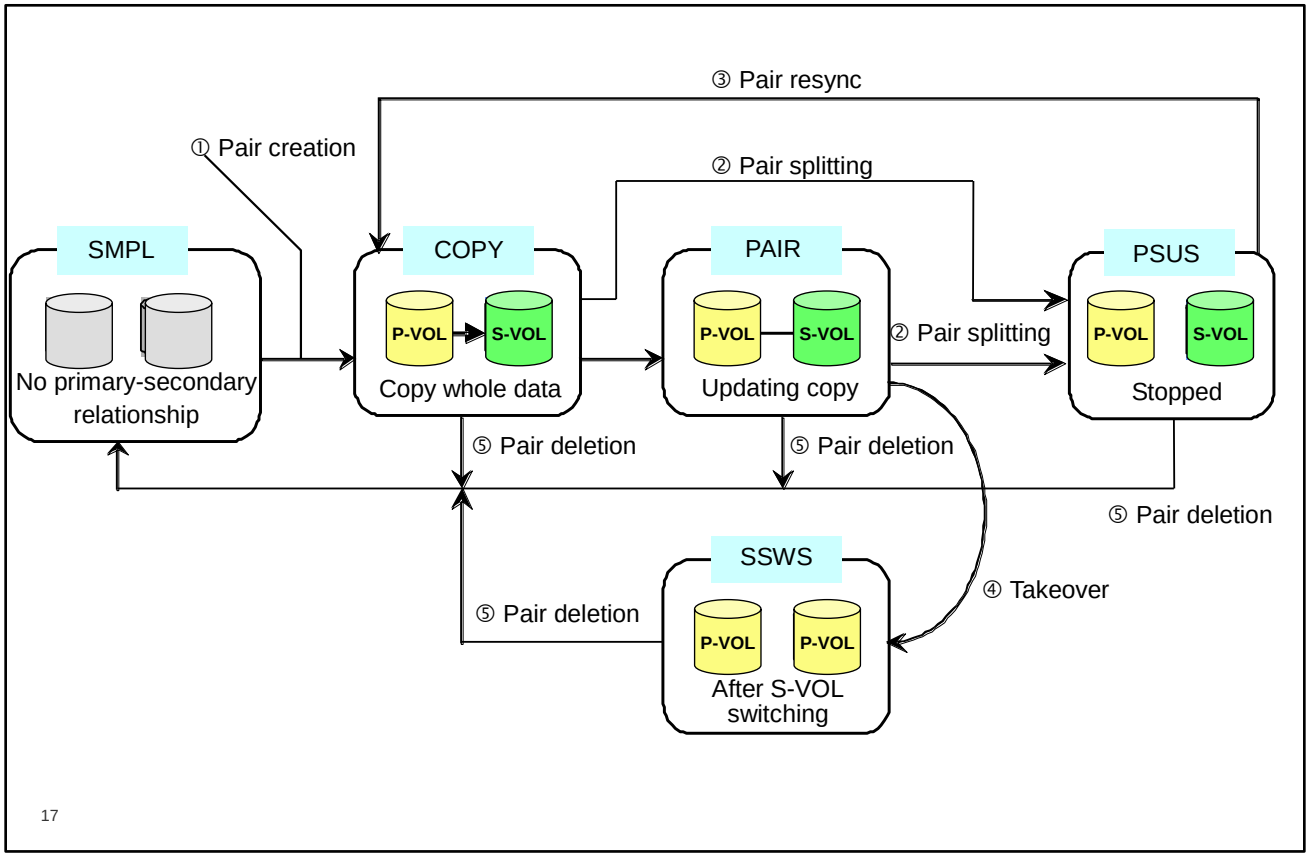
HDS Confidential: For distribution only to authorized parties.

Consistency Groups

- TrueCopy Extended Distance software supports Consistency Groups (CTG) which manage multiple logical units as a group and guarantee 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.



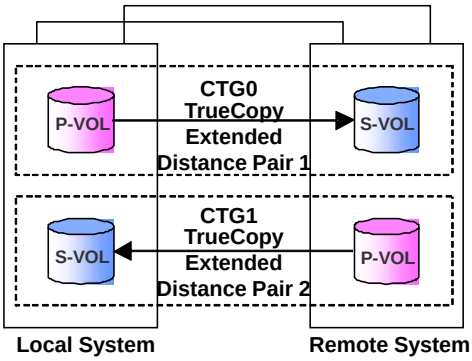
Pair Status Transition

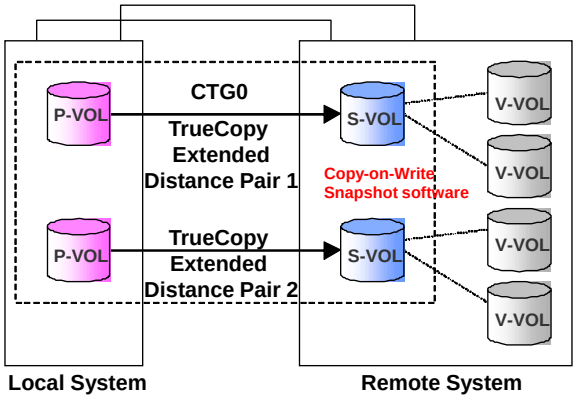


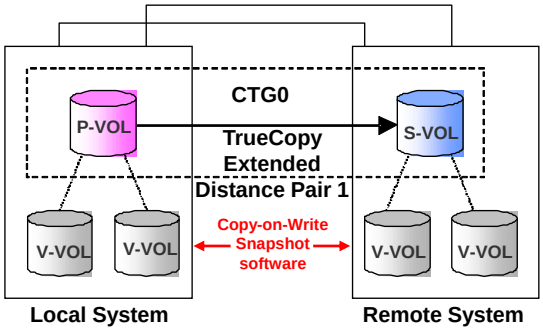
HDS Confidential: For distribution only to authorized parties.
Hitachi Data Systems

Supported Configurations

- Examples of Supported Configurations

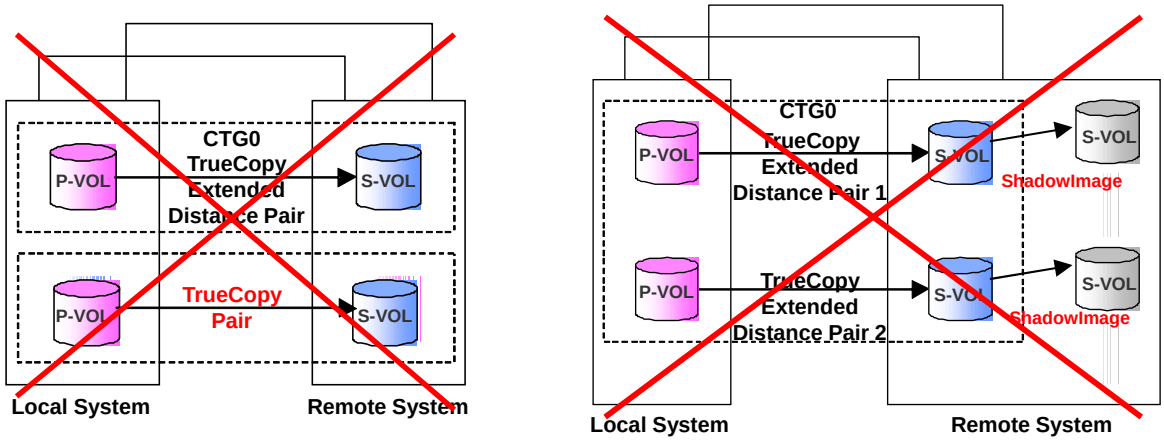






18

- Examples of Unsupported Configurations



19

HDS Confidential: For distribution only to authorized parties.

Recorded Product Demonstration

Now… Recorded Product Demonstration

From your Web Browser …. http://edemo.hds.com/edemo/vault1/TCED.swf

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?

21

HDS Confidential: For distribution only to authorized parties.

11. Command Control Interface Overview

Module Objectives

- Upon completion of this module, the learner should be able to:
  - Identify the Command Control Interface (CCI, or commonly called RAID Manager) components

2

Overview

- What is CCI/RAID Manager?
  - Provides a command line interface to all Hitachi replication products
    - ShadowImage Replication software
    - Copy-on-Write Snapshot software
    - TrueCopy Remote Replication software
    - TrueCopy Extended Distance
  - Is executed on a host
  - Is an in-band admin tool
- For what do you use RAID Manager?
  - To control or script ShadowImage software operations
  - To control or script Copy-on-Write Snapshot software operations
  - To control or script TrueCopy software operations
  - To control or script TrueCopy Extended Distance
  - To control or script Data Retention Utility tasks

3

RAID Manager is also used for configuration of a few other products in the enterprise area.

Note: For the Adaptable Modular Storage 2000 family:

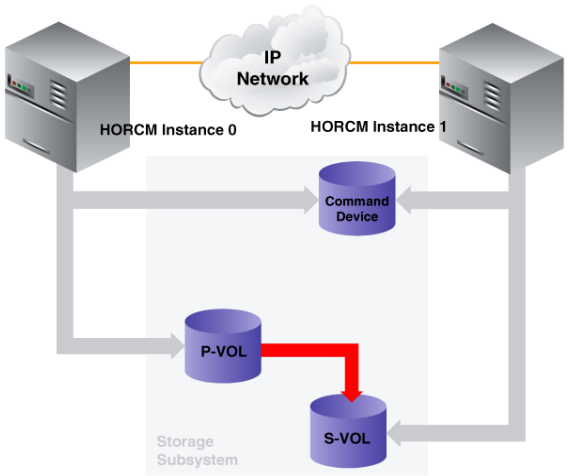
- Must use RAID Manager (CCI) when replicating to or from previous AMS models
- Can use RAID Manager (CCI) when replicating to or from AMS 2000 models
- Can use SNM2 GUI/CLI when replicating to or from AMS 2000 models

HDS Confidential: For distribution only to authorized parties.

- Components
  - Hitachi Open RAID Configuration Manager (HORCM) Command Device
    - Used to accept commands from HORCM
    - Used to report command results to HORCM
  - HORCM Instance
    - Service or Daemon
    - Used to communicate with storage system and with remote server
  - HORCM Configuration File
    - Defines communication paths
    - Defines LUNs (volumes) to be controlled
  - HORCM Commands
    - Monitor and control remote copy operations

These are four of the CCI components needed for the replication products. ShadowImage license must already be installed and the required S-VOLs created.

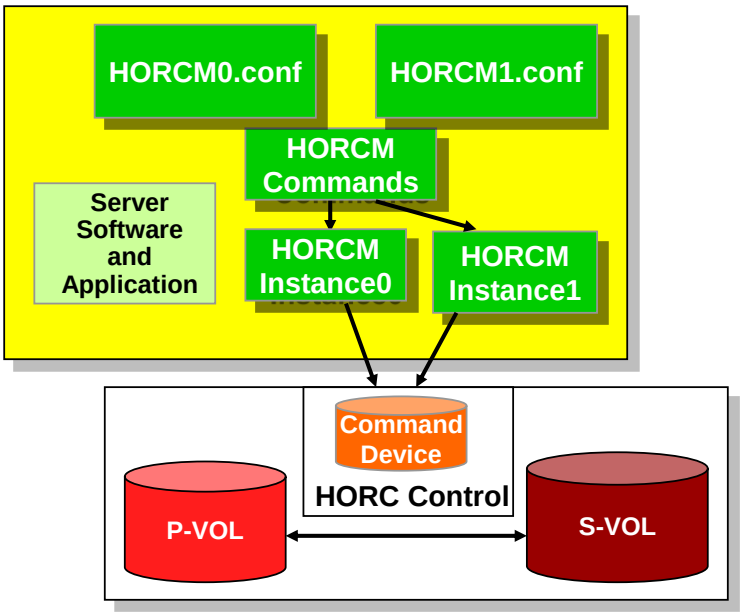
- RAID Manager environment establishes a “conversation.”
- Instances negotiate out of band (TCP/IP).
- In-band communication with storage system
  - Through SCSI Channel
  - Command Device



5

Example — ShadowImage

- ShadowImage software — 1 Server — 2 HORCM Instances



6

ShadowImage example

The relationship between these four components includes:

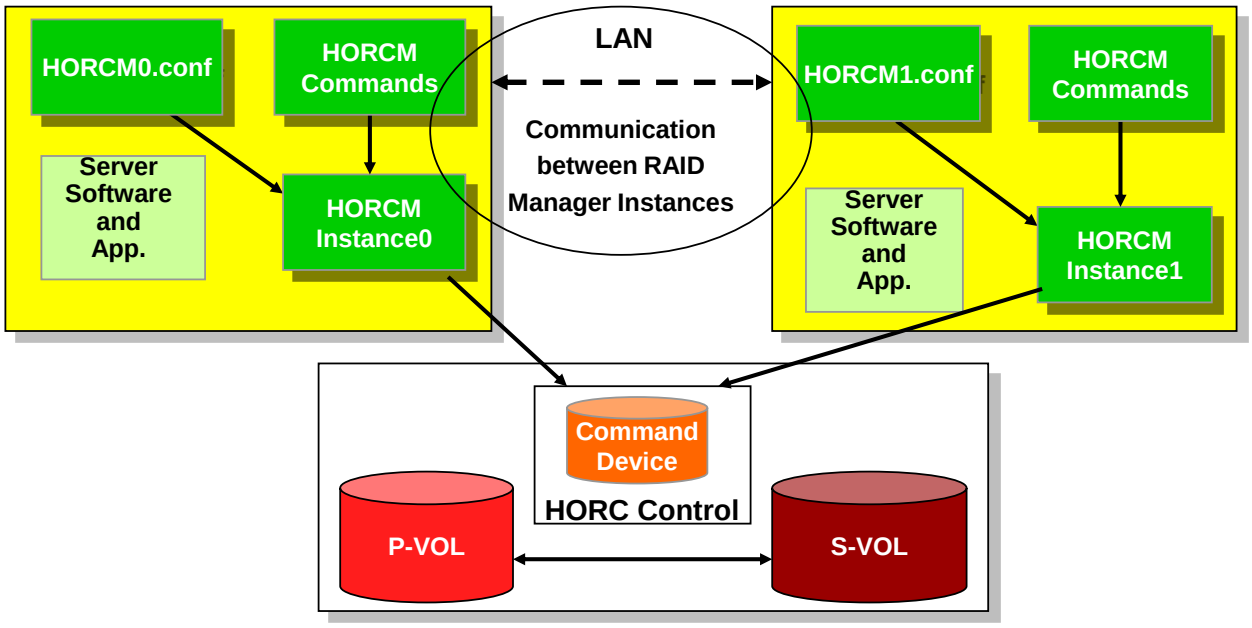
- There are always at least two instances, one controlling the P-VOL and the other one controlling the S-VOL.
- Each instance relies on a configuration file in order to communicate with the other instance, as well as, to communicate with the storage system.
- The configuration file defines the volumes that will be “paired up”.
- If you have two instances, you will have two corresponding configuration files.

An environment variable has to be set in order to identify which HORCM instance will be the 'command interpreter'.

When a command is issued, usually via a script, the "command interpreting" instance sends the command to the CMD Device. The subsystem will then actuate the command.

HDS Confidential: For distribution only to authorized parties.

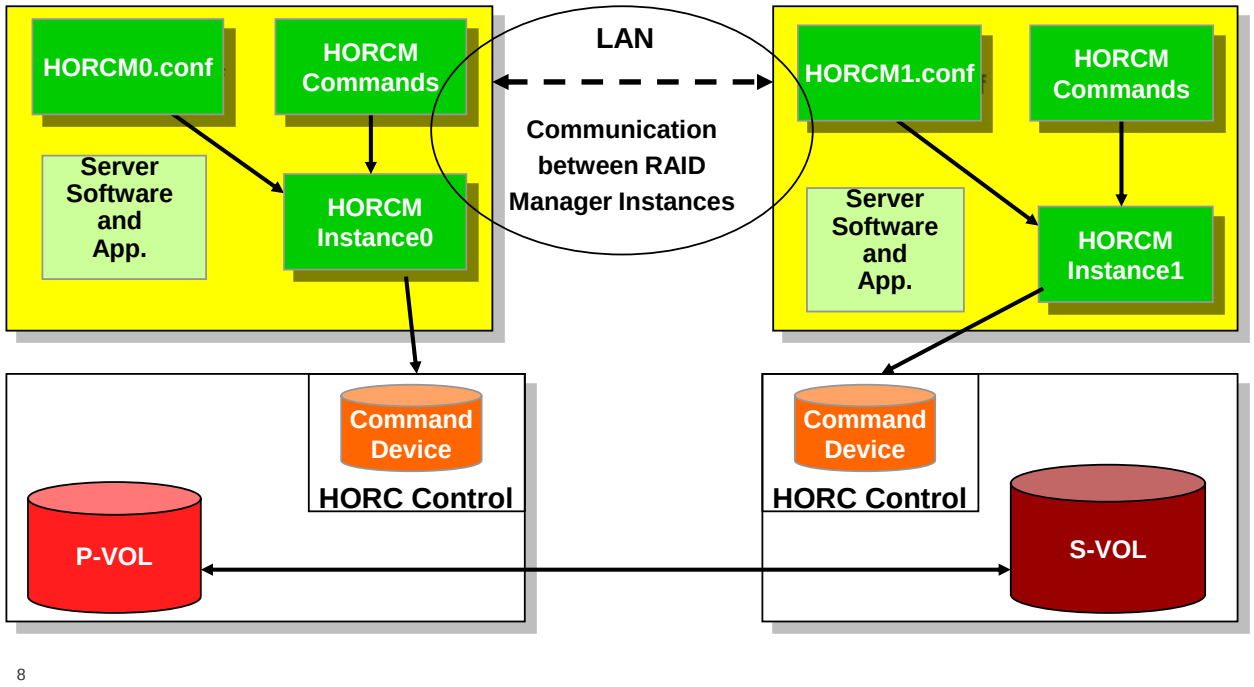
- ShadowImage software — 2 Servers — 2 HORCM Instances



7

Example — TrueCopy

- TrueCopy software — 2 Servers — 2 HORCM Instances



HDS Confidential: For distribution only to authorized parties.

HORCM Configuration Files

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

9

- 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:
  - 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.

10

- HORCM_DEV — Defines the device groups and creates device names.
  - Uses the port ID, target ID, and LUN to identify P-VOLs and S-VOLs
  - 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.

11

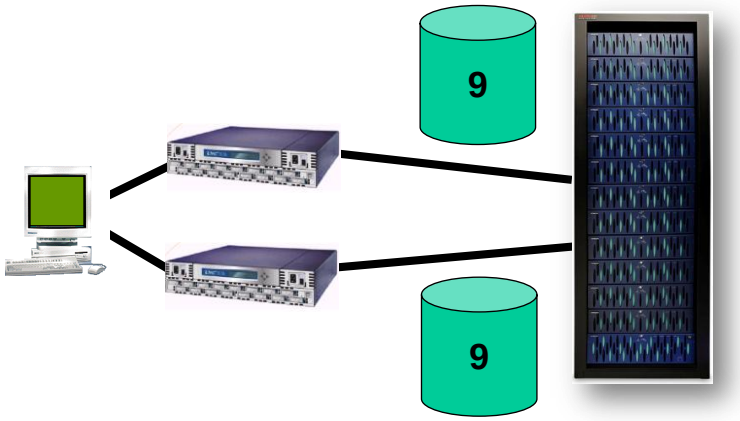
HDS Confidential: For distribution only to authorized parties.

HORCM_DEV

- Multipath configuration
  - LUN 9 is mapped to CL1-B and CL2-B

HORCM_DEV

#dev_group	*dev_name*	*port #*	*target ID*	*LUN#*	*MU#*
ora1	ora_tab1	CL1-B	0	9	0
Or					
ora1	ora_tab1	CL2-B	0	9	0



12

Managing One Volume with horcm0.conf

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-SERIALNO **# Windows 2000 and 2003**

HORCM_DEV

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

HORCM_INST

#dev_group	*ip_address*	*service*
oradb1	SVR1	horcm1

13

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_DE 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.

Managing the Other Volume with horcm1.conf

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**
\\.\CMD-SERIALNO	**# 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

14

Port CL1-A translates to:

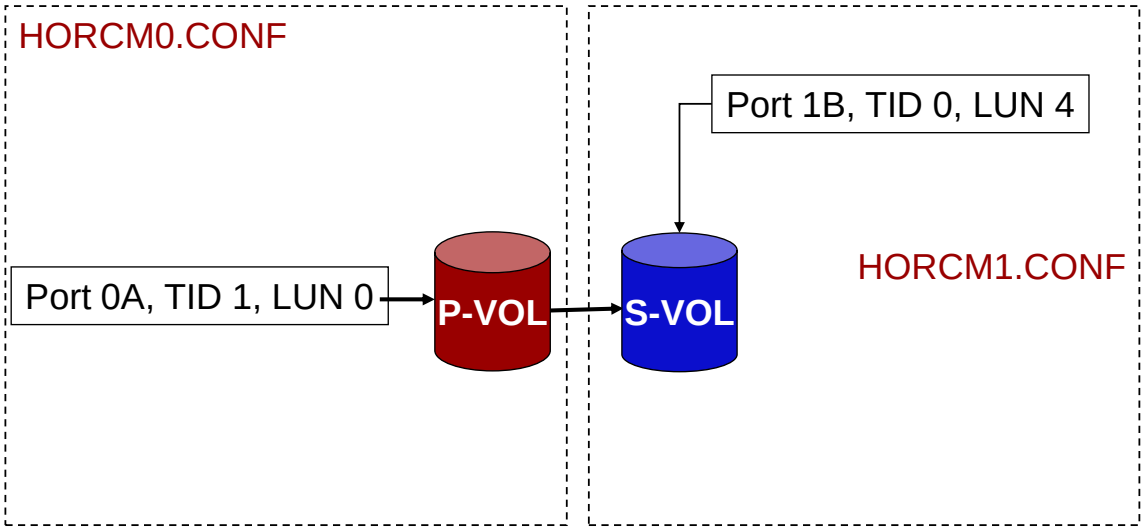
0A on the Adaptable Modular Storage 2000/Workgroup Modular Storage.

Port CL2-B translates to:

1B on the Adaptable Modular Storage 2000/Workgroup Modular Storage.

Configuration

- Proposed configuration



With the **paircreate** command it is determined which volume will become the P-VOL. The other one will become the S-VOL.

15

Commands

- pairdisplay - show pair status
- pairvolchk - checks the attributes and status of a pair volume
- paircreate - create a pair
- pairsplit - split a pair, temporarily or permanent
- pairresync - resynchronize a pair after split
- raidscan - find LUNs and show their status
- setenv - set environment variables
- sync - flush system buffers to disk

16

HDS Confidential: For distribution only to authorized parties.

Module Review

1. Which Hitachi replication or copy products can be managed by RAID Manager CCI?
2. Which four major components are necessary to manage Hitachi replication or copy products with RAID Manager CCI?

17

HDS Confidential: For distribution only to authorized parties.

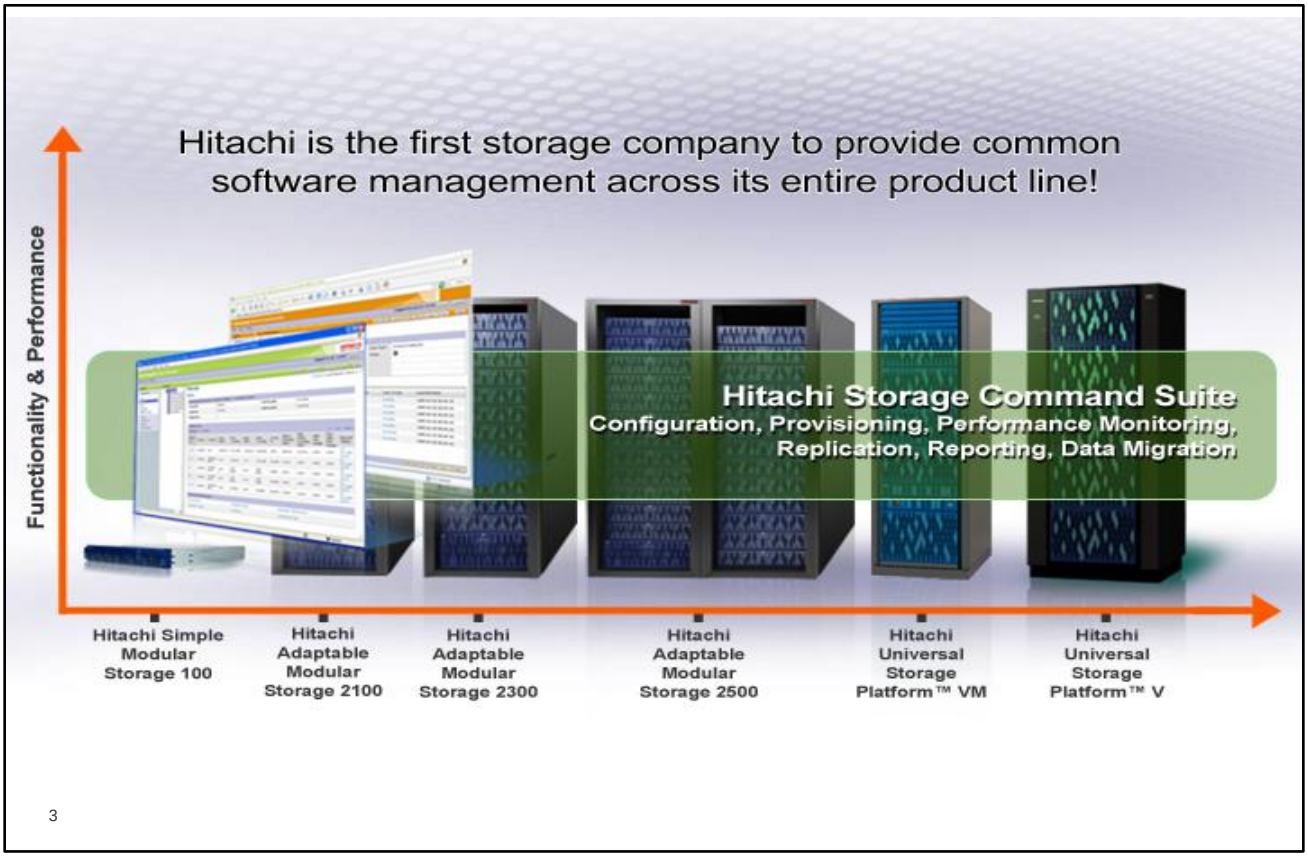
12. Hitachi Storage Command Suite

Module Objectives

- Upon completion of this module, the learner should be able to:
  - List the products in the Hitachi Storage Command Suite
  - Describe the key features

2

Common Software Management Framework



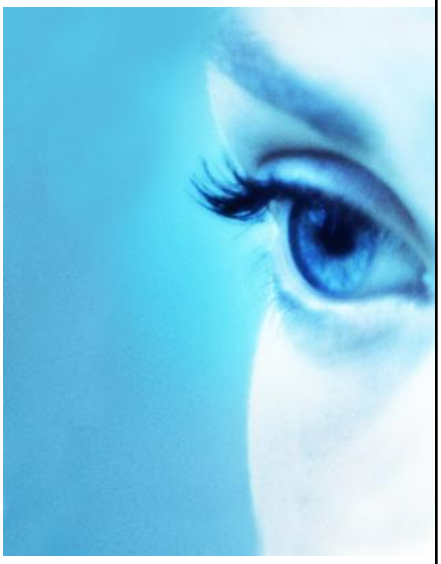
- Supports storage configuration and management across Hitachi's entire range of storage systems
- Enables storage management tasks including configuration, provisioning, performance monitoring and performance management, replication, reporting and data migration
- Presents a view of your complete environment
  - Storage Systems
  - Servers
  - Switches*
- Allows you to logically arrange resources
- Supports reporting

* Hitachi Tuning Manager only

HDS Confidential: For distribution only to authorized parties.

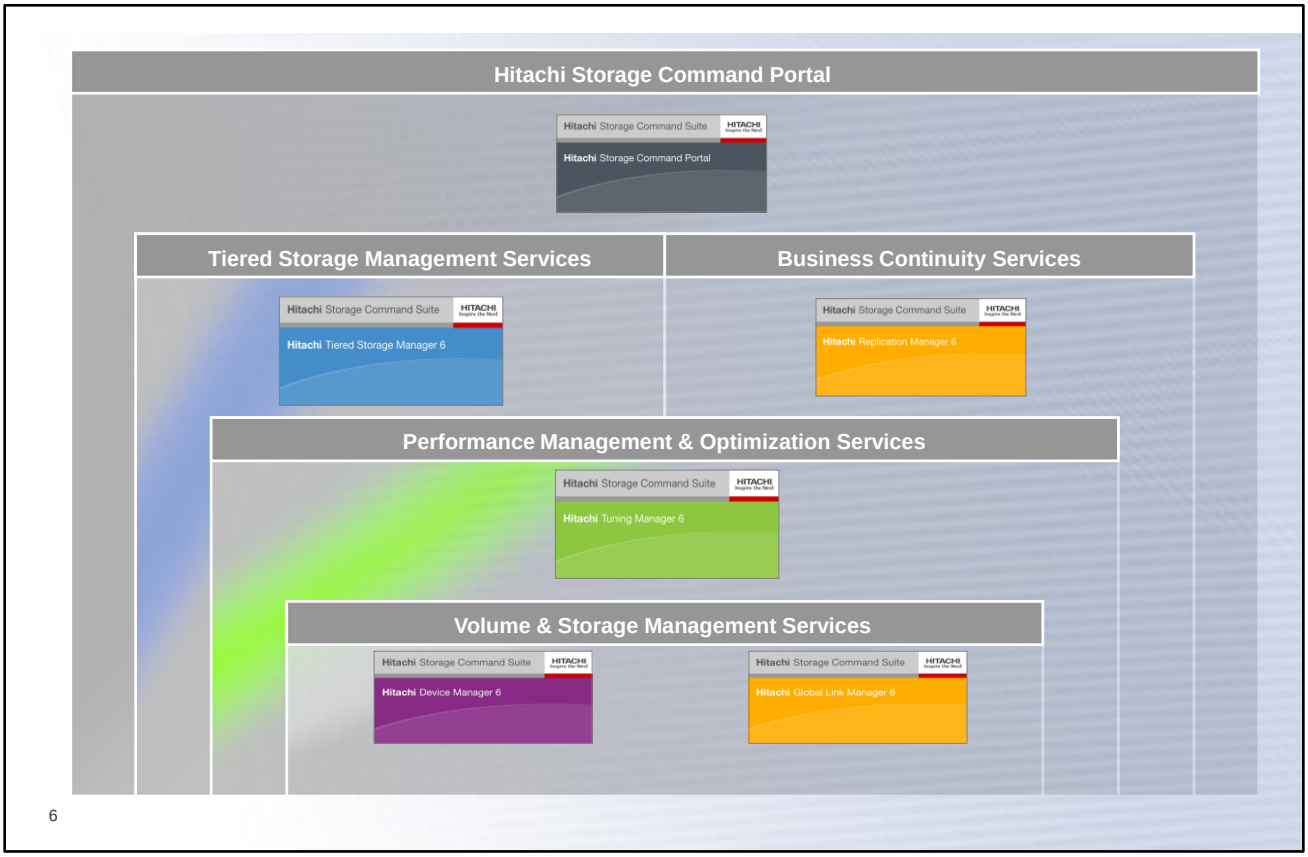
Vision

Hitachi Storage Management software enables **business applications** to most **efficiently** leverage all required storage resources to meet their unique **service level requirements**.



5

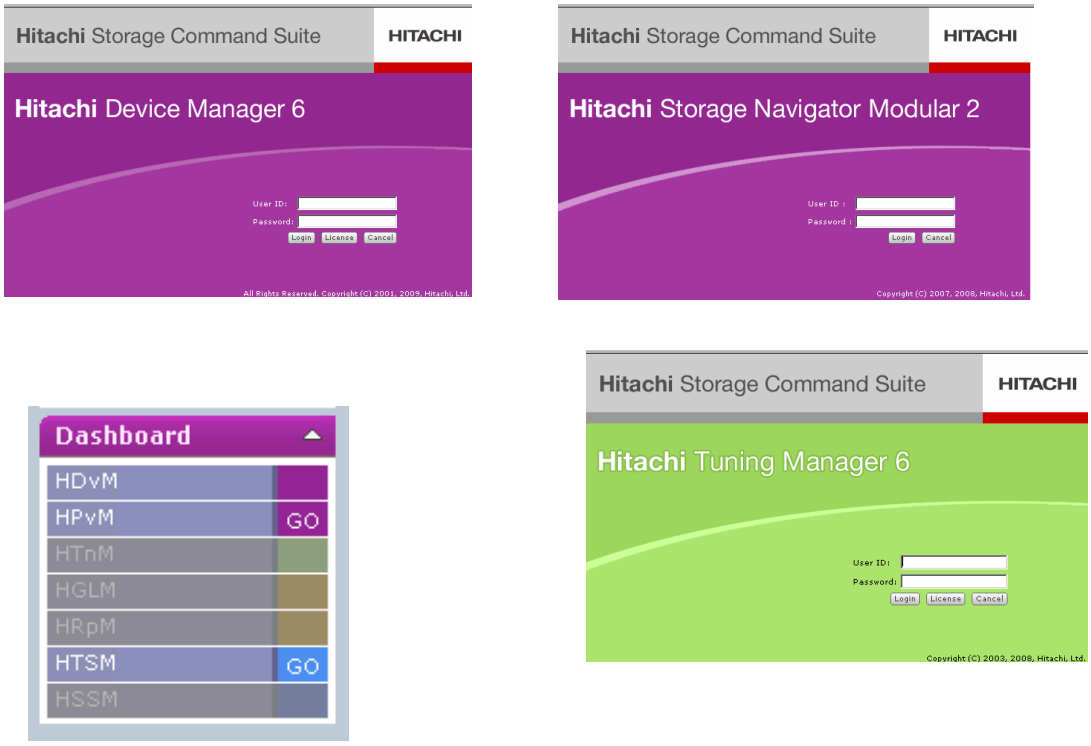
Components



- Provisioning and Path Management
  - **Hitachi Device Manager**
  - Hitachi Global Link Manager
- Replication Management
  - **Hitachi Device Manager**
  - Hitachi Replication Manager
- Performance and Capacity Management
  - **Hitachi Device Manager**
  - Hitachi Tuning Manager
  - Hitachi Storage Capacity Reporter
  - Hitachi Storage Command Portal
- Virtualization
  - **Hitachi Device Manager**
  - Hitachi Tiered Storage Manager

HDS Confidential: For distribution only to authorized parties.

Graphical User Interface

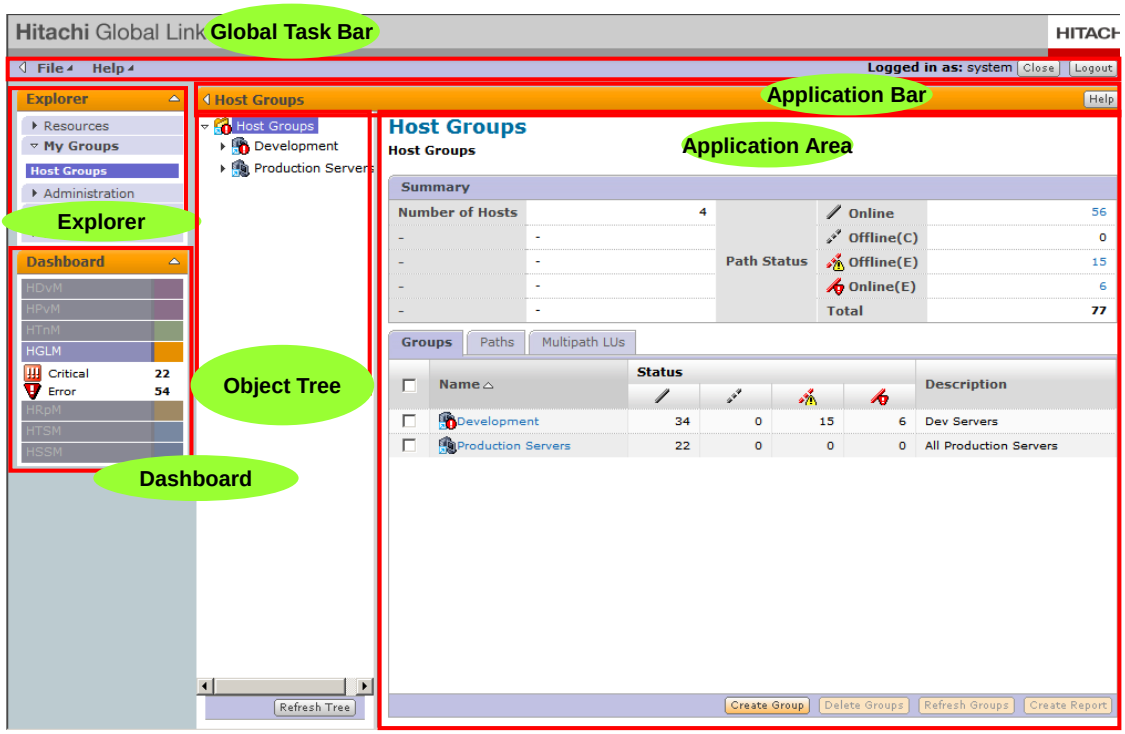


The products in the Storage Management Command Suite have a common look and feel to the user interfaces.

This slide shows examples of the login screens for Device Manager, Storage Navigator Modular 2 and Tuning Manager.

You can see the consistency in format, font and style. In some cases, a different color scheme is used. Tuning Manager uses a green color scheme as compared to the purple colors of Device Manager and Storage Navigator Modular 2.

Access to all the Storage Management Command Suite products is tightly integrated in the easily accessible “link and launch” Dashboard area.



9

Global Tasks Bar Area — Displays the menus and action buttons as well as information about the user.

Explorer — Displays task menus. Choosing a menu item displays the corresponding information in the navigation and application areas.

Dashboard — Displays a list of Hitachi Storage Command Suite products. A product displayed in the active mode can be started by using the **Go** menu.

Application Bar Area — Displays the menu name and action button selected from the **Explorer** menu. You can use the application bar to change the navigation area's view/hide setting.

Application Area — Displays information about the object chosen from the **Object Tree** or the **Explorer** menu.

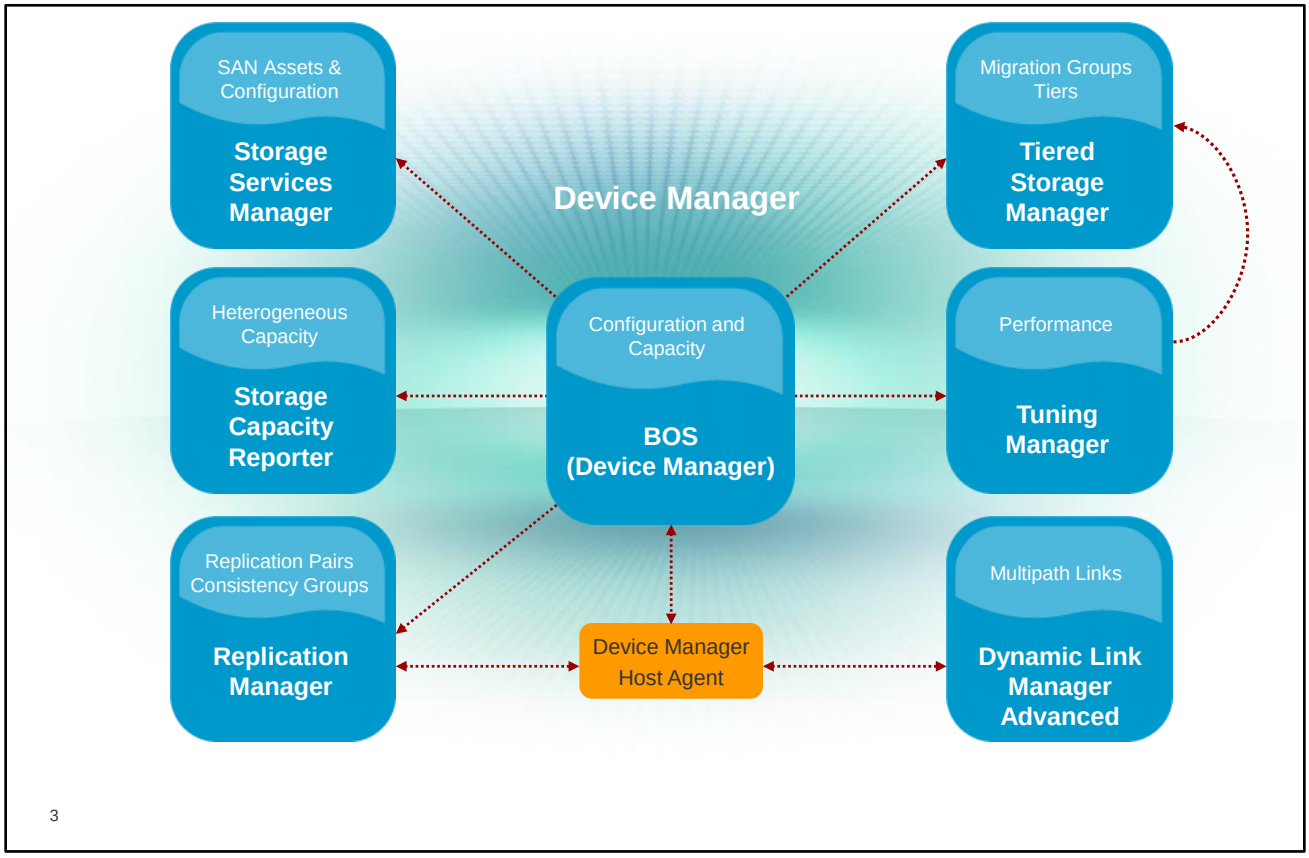
13. Hitachi Device Manager Software

Module Objectives

- Upon completion of this module, the learner should be able to:
  - Articulate key storage management challenges for Storage Administrators, and how the Hitachi Device Manager solution can help
  - Describe the hardware and software components of the Device Manager architecture
  - Describe the features and operations for configuring and managing Hitachi storage systems
  - Understand Device Manager storage provisioning
  - State the purpose and advantages of using the Command Line Interface (CLI)

2

Overview



Hitachi Storage Command Suite:

Hitachi Device Manager software is the primary Basic Operating System (BOS) interface,, has an easy-to-use and intuitive graphical user

Device Manager is the core software product for managing and maintaining storage configuration and capacity information. The Device Manager Host Agent provides all relevant server information. The other products all leverage the Device Manager database of configuration information.

Replication Manager software and Dynamic Link Manager/Global Link Manager software leverage the Device Manager host agent information– not requiring their own.

Tuning Manager software uses the Device Manager configuration information rather than polling all the storage systems, or requiring Tuning Manager agents to get host/correlation information.

Note: Tuning Manager host agents are still required to obtain performance information of the host server and application.

Tiered Storage Manager software gets Array Group performance information from Tuning Manager software.

- Proactive management of complex storage environments
- Single interface for all Hitachi storage (enterprise or modular)
- Simple and secure storage management
  - Storage classification based on usage
  - Distributed management–“MyStorage”
- Storage provisioning and capacity reporting
  - Host-oriented
  - Storage-oriented
  - Application Oriented
- Replication operations

4

Device Manager quickly discovers the key configuration attributes of storage systems and allows users to begin proactively managing complex and heterogeneous storage environments quickly and effectively using an easy-to-use browser-based GUI.

By using *Logical Group* configuration effectively, you can manage storage systems easily and effectively, corresponding to actual storage usages and configuration.

The concept of Logical Group is not only for classifying certain group of storages by corresponding purposes, but the appropriate logical group configuration may be useful for access control privilege management on Device Manager software and Storage Pool management on Device Manager software.

Consolidate and rationalize storage management points

- Use a single Device Manager server to manage all your Hitachi storage. Get access from anywhere using secure protocols.
- Improves IT productivity through consistent administrative operations across storage and operating systems.
- Centralized Administrator definitions. Define different authorities for each administrator. These administrators can only “see” their own storage and their Hosts.
- View all hosts and storage from a single pane. Drill down on hosts and storage as required.

- Display File System mount points and utilization.
- Manage all your storage from a common pool.

Single view of Open and Mainframe storage

- See all LUs/LUNs assigned and available for all Windows, UNIX/Linux and mainframes.
- Allocate from a common pool.

Integrate Hitachi Storage with Enterprise Management Systems

- Use SNMP MIB and traps to integrate with enterprise management systems such as CA Unicenter, HP Openview, and IBM Tivoli.
- Use SMI-S interface to integrate into IBM and other CIM-based management tools.

HDS Confidential: For distribution only to authorized parties.

Problems Solved

- Who benefits from Device Manager software?

  Anyone who:
  - Has a storage infrastructure which is growing in complexity, and requires a simpler, less time-intensive method of managing their storage environment
  - Need to configure Hitachi Data Systems business continuity products
    - ShadowImage Replication software
    - Copy-On-Write Snapshot software
    - TrueCopy Remote Replication software
    - TrueCopy Extended Distance software
  - Would like to prepare for deploying other modules of the Hitachi Storage Command Suite:
    - Tuning Manager software
    - Protection Manager software
    - Replication Manager software
    - Tiered Storage Manager software
    - Storage Capacity Reporter software
    - Storage Command Portal software

5

Unified Interface

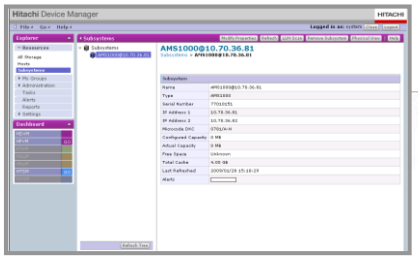
Single console, single product for managing all tiers of Hitachi storage

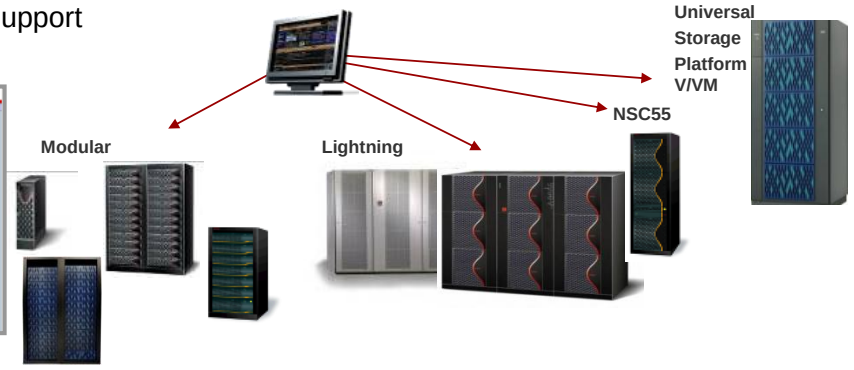
Capabilities

- Full server awareness to discover, configure, provision, monitor, report
- One common interface, browser and CLI
- One or multiple storage systems
- Centrally manage Host Group and LUN security
- Manage Hitachi replication software
- Open standards interface support

Benefits

- Improved productivity of IT resources
  - Single point management
- Improved alignment with business functions using logical grouping
- Risk mitigation
  - Proactive alerts on storage systems to prevent outages
  - Reduced manual error prone storage processes





6

It manages all HDS arrays, enterprise and modular, with the same interface. This is unique. It can also manage multiple arrays in a network environment. It is path-aware and can uniformly provision all types of servers, but only manages HDS arrays.

Targeted for users managing multiple storage arrays in open or shared environments, Device Manager software quickly discovers the key configuration attributes of storage systems and allows users to begin proactively managing complex and heterogeneous storage environments quickly and effectively using an easy-to-use browser-based GUI.

HDS Confidential: For distribution only to authorized parties.

Storage Management

- Multiple views for managing storage—Logical Groups

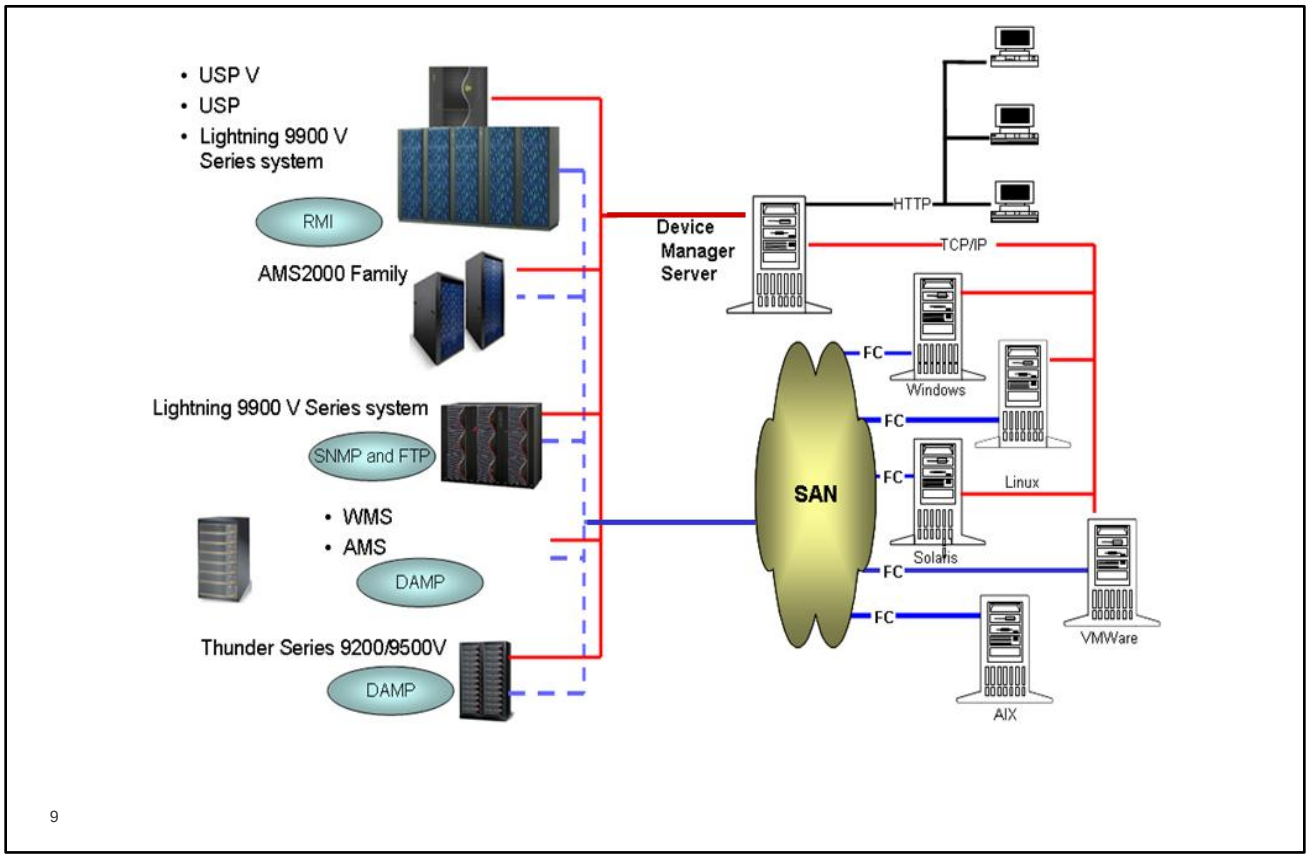
- Logical Groups
  - Logical groups allow the storage team to create nested hierarchical structures to keep track of storage utilization at LUN level.
  - The structures are logical and are local to Device Manager.
  - The hierarchy is user defined. For example, base on:
    - Application
    - Location
    - Storage
  - Reports can be generated based on logical groups.

- Multiple views for managing storage—Resource Groups

- Resource Groups
  - Resource Groups allow you to implement distributed management.
  - A user can be restricted to the subset of the resources managed by Device Manager.
  - Resource Group consists of:
    - Logical Groups
    - Hosts
    - LDEVs (from one or more storage systems)
  - Resource Group is assigned to a aser. User will be able:
    - To see ONLY the allocated resources on the GUI
    - View and perform Device Manager operations against these resources

8

Architecture



The Device Manager (HDvM) server becomes the core of the managed storage environment.

Storage is connected to the SAN and the storage consumers (hosts) using the Fibre Channel (FC) network.

Storage management tools, including HDvM, communicate with the storage systems and storage consumers (hosts) using TCP/IP connectivity. This is also called the Management LAN.

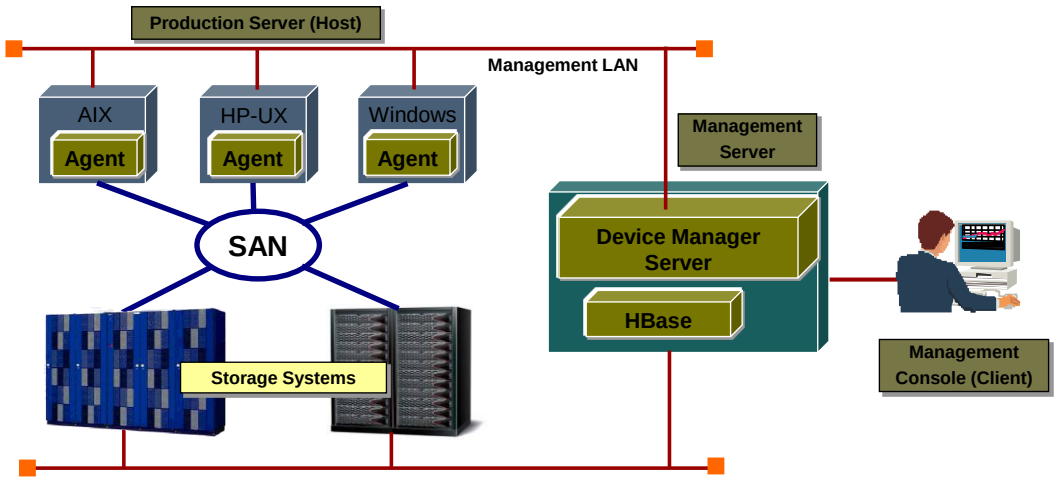
This slide shows the managed components of the storage environment including the storage systems and the hosts or storage consumers. The diagram shows that the storage and the hosts are connected through the FC network which includes the SAN. Storage management occurs using the TCP/IP network connectivity. The storage systems and the hosts all must have TCP/IP connectivity to the Device Manager Server over the Management LAN.

Any PC which has network access to the Device Manager Server over TCP/IP with HTTP and that runs a supported browser can access and use Device Manager Web Client. This is the thin client component.

HDS Confidential: For distribution only to authorized parties.

Components

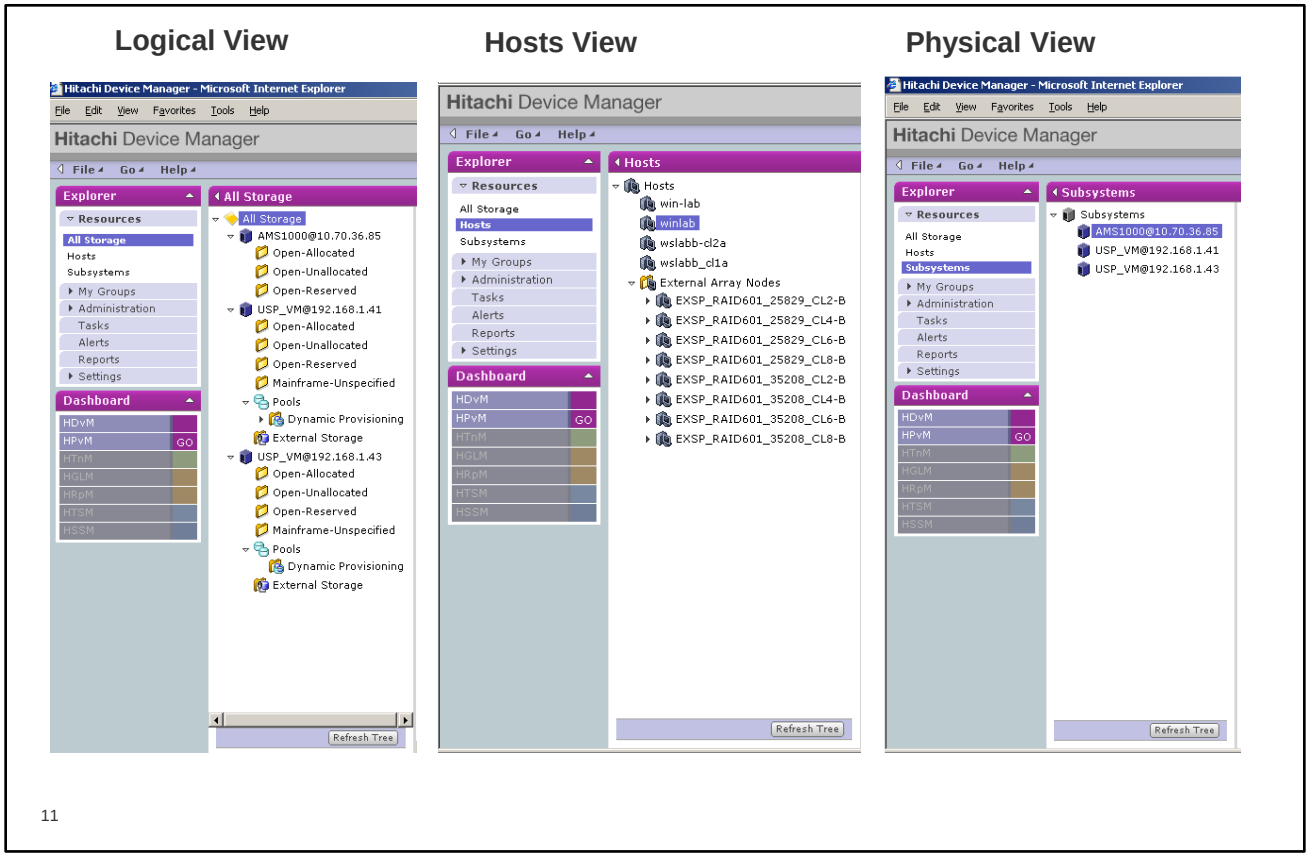
- Device Manager software consists of the following components:
  - Server and its subcomponent (on the management server)
  - Host agent (on the customer production server)
  - Management console (on the Web browser)



10

The Device Manager system includes the Device Manager server, the storage systems connected to the server, the (optional) Device Manager Agents, and Device Manager clients. The Device Manager Web Client provides a web-distributed client for real-time interaction with the storage systems.

Views of the Storage Environment



Device Manager is organized to provide three different views of the managed storage environment. Under the **Resources** heading are three categories: **All Storage**, **Hosts** and **Subsystems**. When a Device Manager user signs on, the **All Storage** view within the **Resources** section of the Explorer area is the default view.

The All Storage view is also referred to as the Logical View. The All Storage view presents information about storage based on whether or not it is in use (Allocated) and if it is being used, how it is being used. For example, in the Logical View example shown on this slide, you can see categories such as Open - Allocated, Mainframe — Unspecified, and Pools. These groupings are termed “logical” because they are not attributes of the physical storage such as RAID structure, or drive capacity. Device Manager will automatically move storage between these logical categories as storage is used (allocated or mapped) and freed from use. Try to remember this description of the All Storage view as a logical view. Later, you will see the use of Logical Groups which is different from this Logical View of All Storage.

In the Hosts View, you can see that the Navigation Area shows entries for individual hosts and also a group heading **External Array Nodes**. The External Array Node entries will be discussed in detail later.

The third high-level view provided by Device Manager is the Physical View. In this view, Device Manager provides one entry for each registered storage system. You will learn more about the Physical View and the important functions accessed from the Physical View.

Architecture Components

- Management Server
  - Device Manager server software is installed here
  - Consists of:
    - Device Manager software
    - Web Server
    - HiRDB Database
    - CLI
    - Distributable files for Agent and CLI

- Device Manager Agent
  - Optional Component
  - If Installed on a host, automatically discovers and provides the following information to Device Manager server –
    - Host HBA(WWPN) details
    - Hostname/IP Address
    - Filesystem mountpoint/drive letter information
    - Filesystem capacity utilization information
  - Is a required component if replication operations must be performed from Device Manager
  - Registers the host with Device Manager
    - If the agent is not installed, the Host must be registered manually or through an **Execute Host Scan** operation from the GUI.

13

HDS Confidential: For distribution only to authorized parties.

- Management Console (Device Manager Client)
  - Provides GUI-based operation console on remote PC
  - Uses thin client
  - Consists of following two types:
    - HTML GUI – Browser-based application
    - Java GUI – Java application deployed by Java Web Start technology



Management Console (Client)

14

The **Java GUI** is a stand-alone Java-based application that provides windows and dialog boxes for the storage system and resource group management features.

The **HTML GUI** offers simple, easy-to-use navigation.

Explorer Navigation Drawers – simple to use

Consistent navigation tree – for all Hitachi Storage Command Suite products

Dashboard – provides Launch capabilities to Hitachi Storage Command Suite products installed on the Device Manager server.

"Bread Crumb" trail – never gets lost, while browsing resources and metrics

Open and Close Frames – allowing good visibility

Operations

- Categories of Device Manager Operations
  - Provisioning storage
    - To host
      - From Host View
      - From All/My Storage View
  - Replication operations
  - Linking-and-launching Integration
  - Generating reports
  - Managing users and access restrictions
  - Using and customizing system security

15

HDS Confidential: For distribution only to authorized parties.

Report Operations

- Device Manager built-in reporting function generates reports in HTML format and comma-separated value (CSV) format.
- Reports include:
  - Physical configuration of storage system
  - Storage utilization by host
  - Storage utilization by Logical Group
  - Users and resource groups
  - Detailed array reports

16

Reporting on the Physical Configuration of the Storage System

The Physical Configuration of Storage System report displays and provides detailed information on the physical configuration of each storage system added to the Device Manager server. The HTML report shows a graphic representation of the storage system's configuration and components. The CSV report provides the detailed storage system information as comma-separated values.

Reporting on the Storage Utilization by Host

The Storage Utilization by Host report displays storage utilization statistics and detailed storage information for each host added to the Device Manager server. The storage utilization statistics for each host include capacity (GB) currently in use by host operating systems and percentage of capacity in use by hosts. The storage information for each host includes allocated storage displayed by storage group and LUN.

Reporting on the Storage Utilization by Logical Group

The Storage Utilization by Logical Group report displays the following information:

For each storage system added to the Device Manager server: storage system name, model, serial number, IP address, configured capacity, allocated capacity, and percentage of allocated capacity in use by the host operating systems.

For each logical group created to the Device Manager server: group name and level, capacity, and percentage of capacity in use by the host operating system.

Reporting on the Users and Resource Groups

The Users and Resource Groups report (see Figure 12.5) displays the following information for Device Manager users:

Resource groups managed by the Device Manager server. (The Peer group is reserved for use by the Device Manager Agent only.)

Number of users assigned individual resource groups

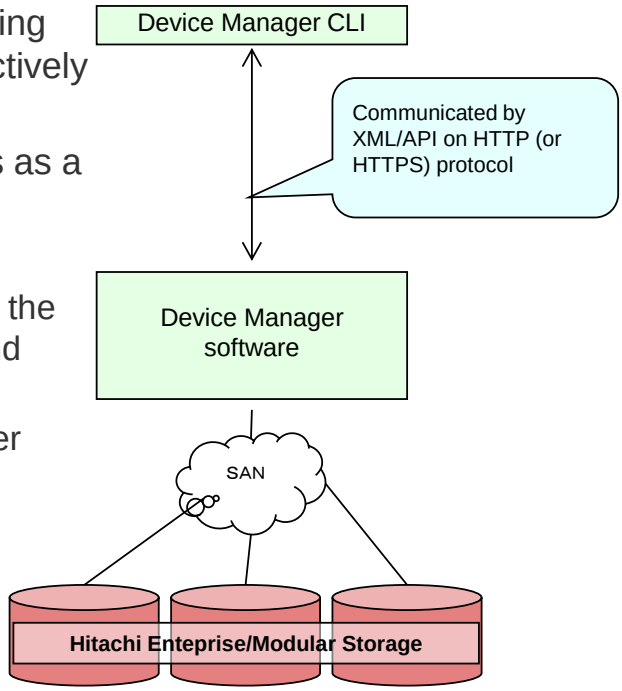
Properties of each user (user ID, full name, and description)

Reporting on the LDEV Details

You can create, download, and view a detailed report on the LDEVs in each storage system that has been added to the Device Manager server.

Command Line Interface

- The CLI enables you to perform Device Manager software operations by issuing commands from the CLI client interactively or using a script.
- The CLI communicates with and runs as a client of the Device Manager server.
- Device Manager in-band CLI is:
  - A new CLI which communicates to the storage system through a command device
  - Provided as part of Device Manager software bundle



17

Storage Navigator Modular 2 versus Device Manager Summary

Capability	Device Manager	Storage Navigator Modular 2
Span		
Modular Storage	■	■
Host / Server	■	
Provisioning		
Multiple Storage Systems	■	
RAID Storage Allocation	■	■
HLUN Labels and Logical Groups	■	
Administrative Security Subdomains	■	
Reporting		
Basic Reports	■	
Interfaces		
Physical View	■	■
HSCS GUI, shared HSCS services	■	
Command Line Interface (CLI)	■	■
SNMP	■	■

18

HDS Confidential: For distribution only to authorized parties.

Replication Operations

- Device Manager software provides the ability to create and manage volume pairs of the following replication products:
  - In-System replication
    - Hitachi ShadowImage In-System Replication software
    - Hitachi Copy-on-Write Snapshot software
  - Remote Replication
    - Hitachi TrueCopy Remote Replication software
    - Hitachi TrueCopy Extended Distance Software

19

Using the Device Manager software graphical user interface (GUI) you can perform the following tasks:

- HORCM configuration file creation
- Pair creation
- Modify pairs

- Requirements for pair volume management:
  - Device Manager Agent installed on hosts.
  - Command Control Interface (CCI) software installed on hosts.
  - Command device mapped to hosts.
  - Logical units (LUNs) used as replication software pairs must be mapped to a host and have LUN security enabled.

20

- Scenario 1: Device Manager software creates the HORCM configuration files only
  - Device Manager software creates the HORCM configuration files.
  - User scripts are used to take over the HORCM file and manage the pair volume operations.
  - After the user scripts have taken over the HORCM files, Device Manager software is used to monitor the pair volume status.
- Scenario 2: Device Manager software creates the HORCM configuration files and manages pair volumes
  - Device Manager software manages the pair volume throughout its life cycle.
  - Device Manager software is responsible for defining, creating, synchronizing, and deleting the pair volume.

21

HDS Confidential: For distribution only to authorized parties.

Pair Volume Maintenance

- Device Manager software provides the following functions for existing pair volumes:
  - Monitor the current status of the pair volume
  - Modify the pair volume status

22

Note: The **modify pair** function must be used carefully. It is not recommended that you mix operations with user written CCI scripts.

Recorded Product Demonstration

Now…..Recorded Product Demonstration

From your Web Browser …. edemo.hds.com/edemo/vault1/dm5USPV.swf

Module Review

1. Which of the three following reports does Device Manager software provide? (Choose three)
   a. Host
   b. Physical
   c. Switch
   d. SAN
   e. Logical Group
2. On which two operating systems can Hitachi Device Manager Server be installed? (Choose two)
   a. Sun Solaris
   b. HPUX
   c. AIX
   d. Microsoft Windows 2003
   e. Novell

24

HDS Confidential: For distribution only to authorized parties.

14. Hitachi Global Link Manager Software

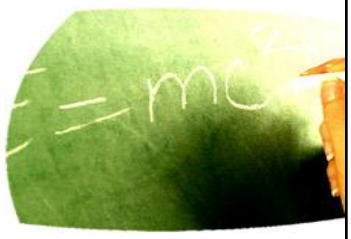
Module Objectives

- Upon completion of this module, the learner should be able to:
  - Identify the key components and functions of Global Link Manager software
  - Describe the architecture

2

Multipathing – Problems and Solutions

- Customer's Challenges – The Problem
  - Hitachi Dynamic Link Manager software offers fault tolerance and load balancing.
  - However, operationally, management can be time consuming and difficult in complex multipathing environments.
- Customer's Challenges – The Solution
  - Global Link Manager software eliminates many of the operational challenges associated with managing multipath I/O environments.



3

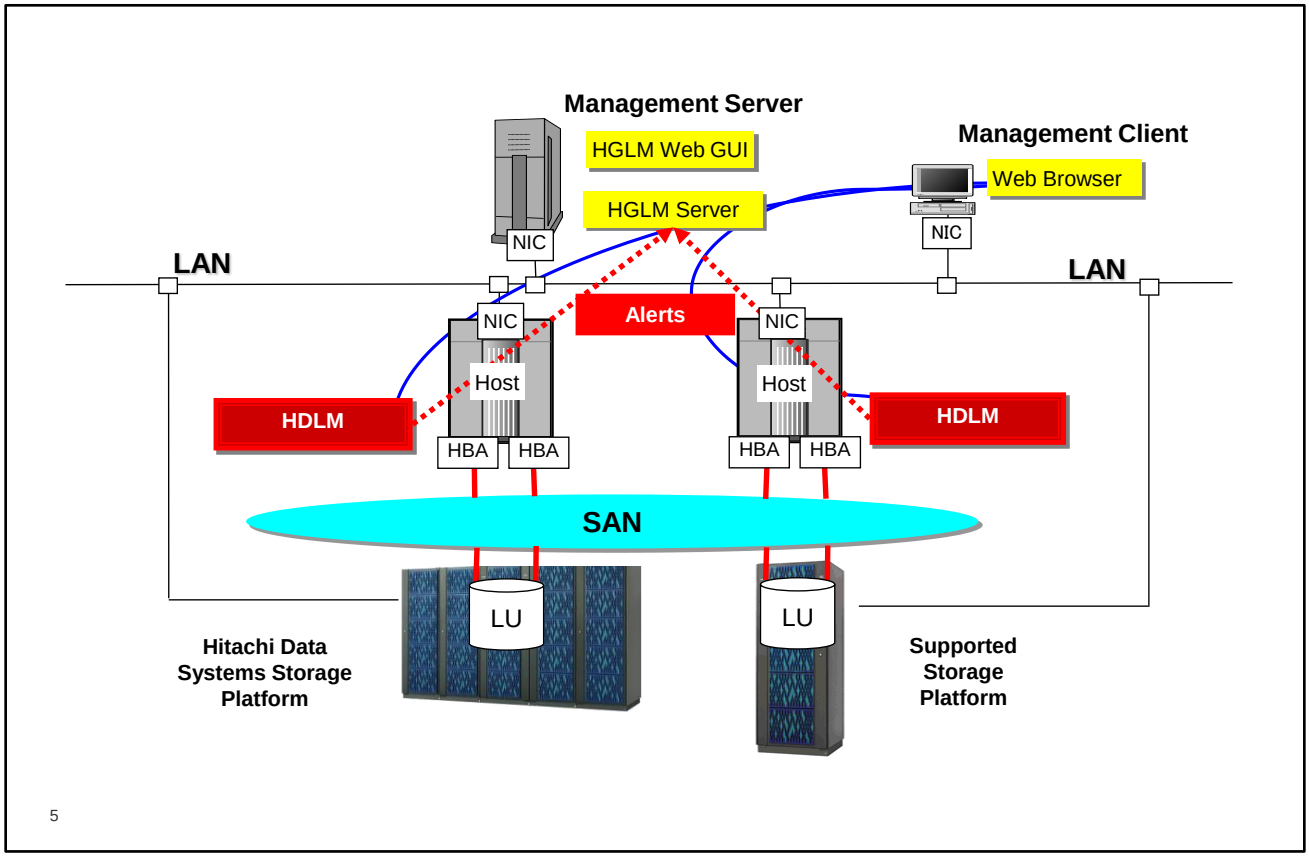
HDS Confidential: For distribution only to authorized parties.

Problems and Solutions

Dynamic Link Manager Software - Pain Points	Global Link Manager Software - Solution
• No unified GUI	• Manages many servers multi-path connections from a single console
• Views and configures only one instance at a time	• Simplifies storage maintenance activities
• Each instance must be monitored independently for alert notification of path status changes	• Keeps customers informed on the status of all their multipathed links

4

Configuration Overview



HGLM = Global Link Manager software

HDLM = Dynamic Link Manager software

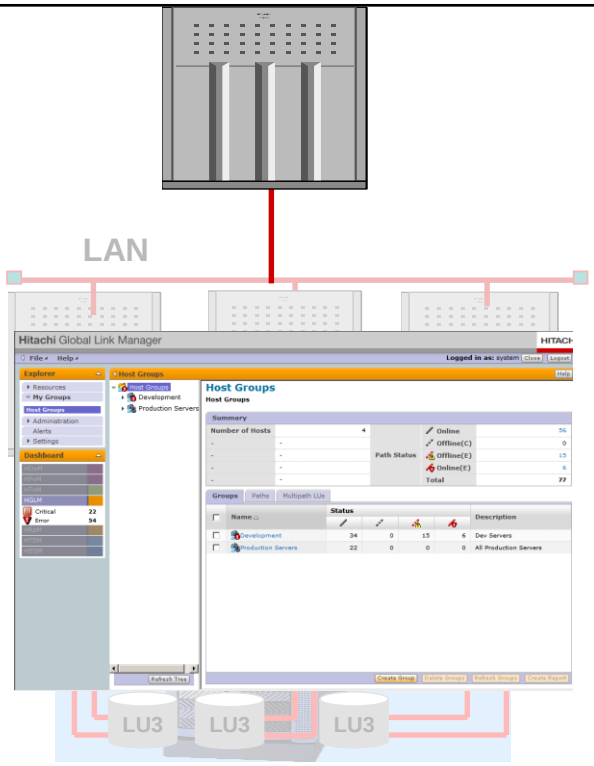
LU = logical unit

HBA = host bus adapter

HDS Confidential: For distribution only to authorized parties.

Features

- Single unified GUI
- Manages multiple HDLM instances
- Simplified Path maintenance activity
- Scheduling algorithms by HDEV (LU)
- Alert notifications
- Configure Dynamic Link Manager software instances
- Provides secure resource grouping



6

Event notification — Alerts generated by Dynamic Link Manager software are displayed by Global Link Availability Manager software near real-time.

Path Management — Global view of paths and Hidden devices (HDevs) for all Dynamic Link Manager software instances. Management capabilities are based on user role definitions.

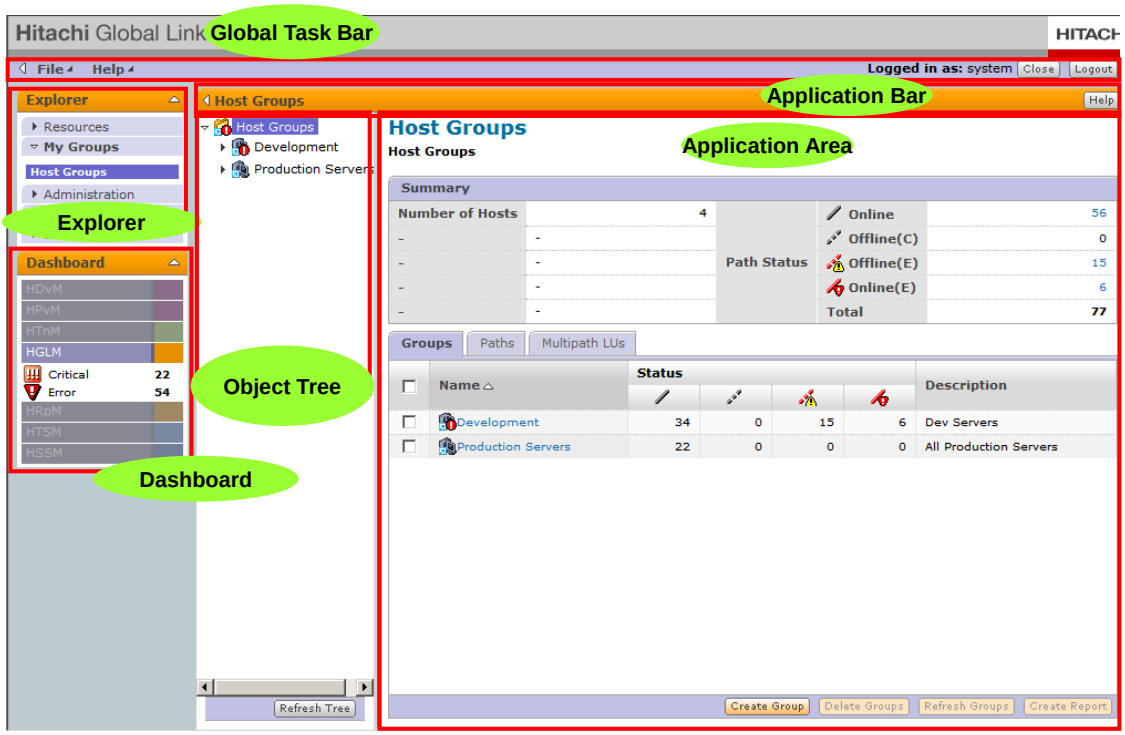
Host Management — Centrally manages configuration of all Dynamic Link Manager software instances.

Host group management — A customized grouping of hosts created by an individual user.

Resource group management — Administrator controls user's access to a specific group of hosts (subset of Dynamic Link Manager software instances).

Access control — User role definitions control operational and host resource access.

Graphical User Interface



Global Tasks Bar Area — Displays the function menus and action buttons used in Global Link Manager software as well as information about the user who is logged in.

Explorer — Displays Global Link Manager software task menus. Choosing a menu item displays the corresponding information in the navigation and application areas.

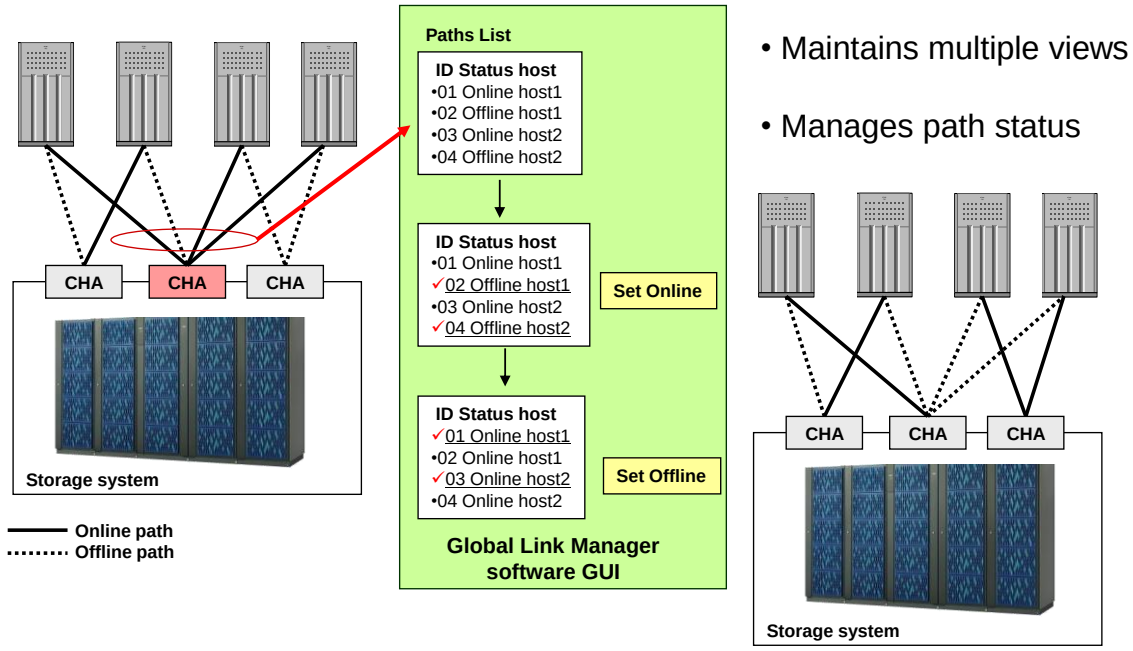
Dashboard — Displays a list of Hitachi Storage Command Suite products. A product displayed in the active mode can be started by using the **Go** menu. In **HGLM**, the number of alerts that are reported from the host being managed is displayed.

Application Bar Area — Displays the menu name and action button selected from the **Explorer** menu. You can use the application bar to change the navigation area's view/hide setting.

Application Area — Displays information about the object chosen from the **Object Tree** or the **Explorer** menu.

Software Path Management

- Global Link Manager software manages paths across multiple hosts



8

Manage the entire Dynamic Link Manager software multi-pathing environment from a single console:

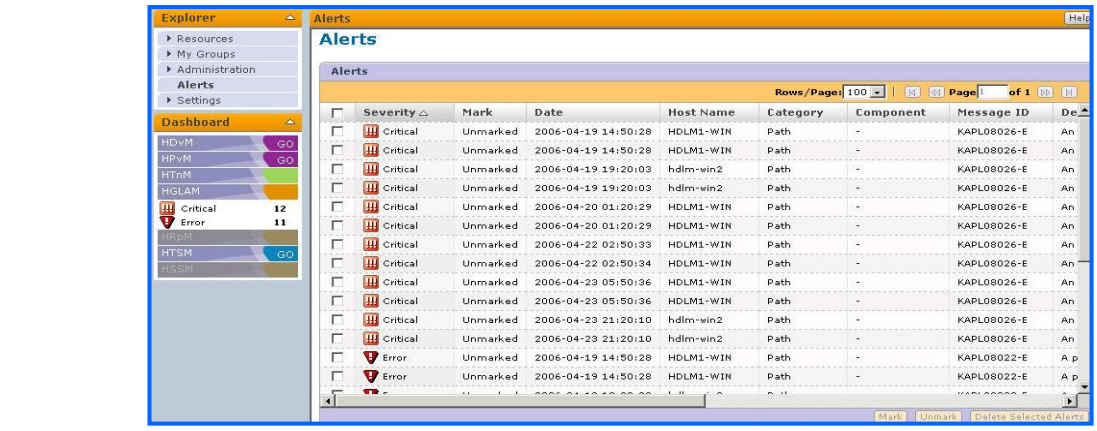
For each Dynamic Link Manager software instance list path information for all paths or for each host, HBA port, storage system, and storage port

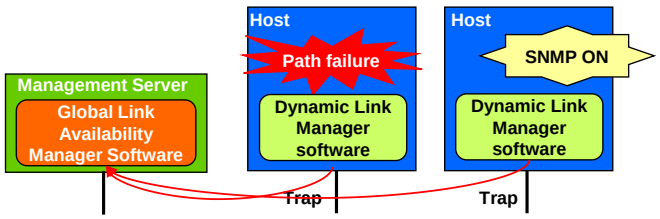
Aggregated path views corresponding to path status (online or offline) check the health of the entire multi-pathing environment

Adjust the online / offline path status for single or multiple hosts

Adjust load balancing for individual LUNs

SNMP-based Alert Management





Events sending alert

	KAPL08022-E	Path failure
	KAPL08026-E	No path available to a LU
	KAPL08027-E	Intermittent path error
	KAPL04042-I	Changed Dynamic Link Manager software SNMP trap configuration

9

SNMP setting can be configured only through Global Link Manager software.

The alert information detected in Dynamic Link Manager software is notified as SNMP trap to the server.

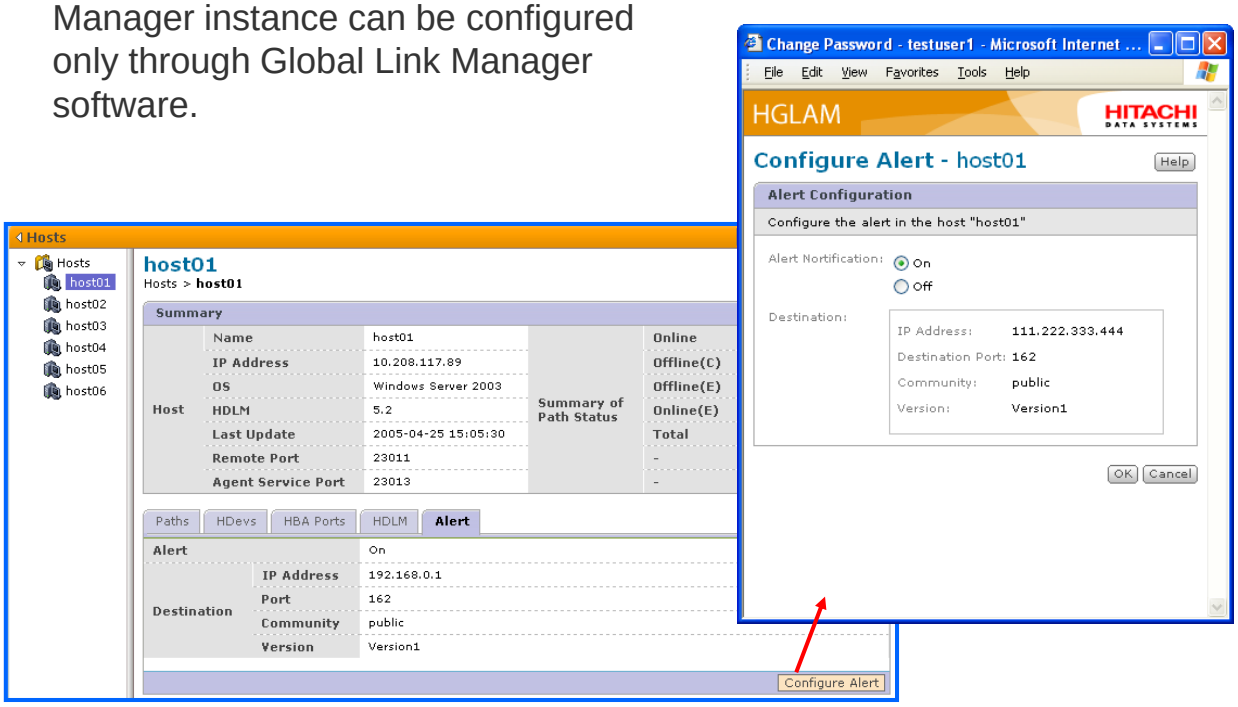
The alerts received by the server are stored in the database (DB).

Notification of the alerts can be identified from the **Dashboard** menu. The user can then view the details by displaying the alerts list.

Minimum setting is one minute.

Alert Management

- SNMP setting on each Dynamic Link Manager instance can be configured only through Global Link Manager software.



10

SNMP setting on each Dynamic Link Manager software instances can be configured during registration process or host refresh process using values specified in server.properties, if **auto_set** option in server.properties is enabled (default is enabled).

SNMP configurations (IP-address, port, community) for each host refer to server.properties on Global Link Manager software server.

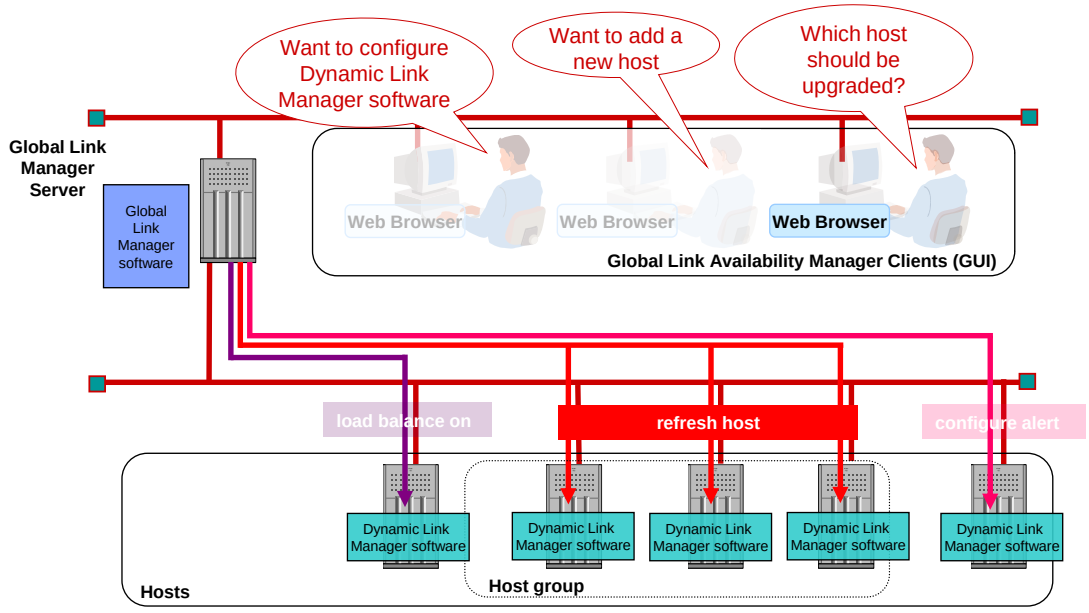
Note:

server.properties is located at **<HGLAM-install-dir>\HiCommand\HGLAM\conf**

Global Link Manager software restart (Single Sign On Service rebooting) is required after editing entries

Host Management

- Host management provides the ability to manage all aspects of the Dynamic Link Manager software host.
  - Host Information View



11

Host Information View displays host information such as the path and Dynamic Link Manager software version.

Dynamic Link Manager software configuration is used for load balancing, auto failback, path health check, and others.

HDS Confidential: For distribution only to authorized parties.

Recorded Product Demonstration

Now…..Recorded Product Demonstration - HGLM

From your Web Browser …. edemo.hds.com/edemo/vault1/hglam.swf

Module Review

1. Which two statements refer to Global Link Manager software? (Choose two)
   a. Views and configures only one instance at a time
   b. Manages many servers multipath connections from a single console
   c. Keeps customers informed on the status of all their multiple path links
   d. Each instance must be monitored independently for alert notification of path status changes

13

HDS Confidential: For distribution only to authorized parties.

15. Hitachi Tuning Manager Software

Module Objectives

- Upon completion of this module, the learner should be able to:
  - Describe how Tuning Manager software enables optimization of storage performance
  - Describe the architecture and the major components
  - Discuss how an I/O operation works
  - State the types of data monitored and reported by Tuning Manager software

2

Overview

- The Top Five Needs of Customers
  - Determine quickly *if* performance trouble is within the storage system or not.
  - Determine the *cause* of performance trouble within the storage system and logically organize storage resources.
  - Forecast *when* a storage system will saturate in terms of performance.
  - Determine *quickly* if close attention to storage performance is necessary.
  - Extract performance data *easily* (for use and analysis by user, user tools or ISV).

- Tuning Manager is a real-time software monitor that you can use to check the current state of all the host, file system, database, SAN, and storage resources that are being used by applications.

- It lets you contrast data about the monitored resources with a historical view of the normal, baseline performance of those resources that was previously stored in the database.

- The ability to query a historical database for performance and capacity trend analysis for each component of a SAN enables you to easily correlate the current changes in performance with recent changes to
  - Physical configuration
  - Software
  - Workload
  - Other environmental changes

4

HDS Confidential: For distribution only to authorized parties.

Product Positioning

Name	Description
Hitachi Tuning Manager	• Advanced reporting, analysis, and troubleshooting application for Hitachi storage systems and services. • Application-to-spindle visibility and correlation in near-time and historical. • Full storage path awareness and deep knowledge of Hitachi storage systems.
Performance Monitor	• Performance Monitor is a monitoring product that provides detailed point-in-time reporting within Hitachi storage.

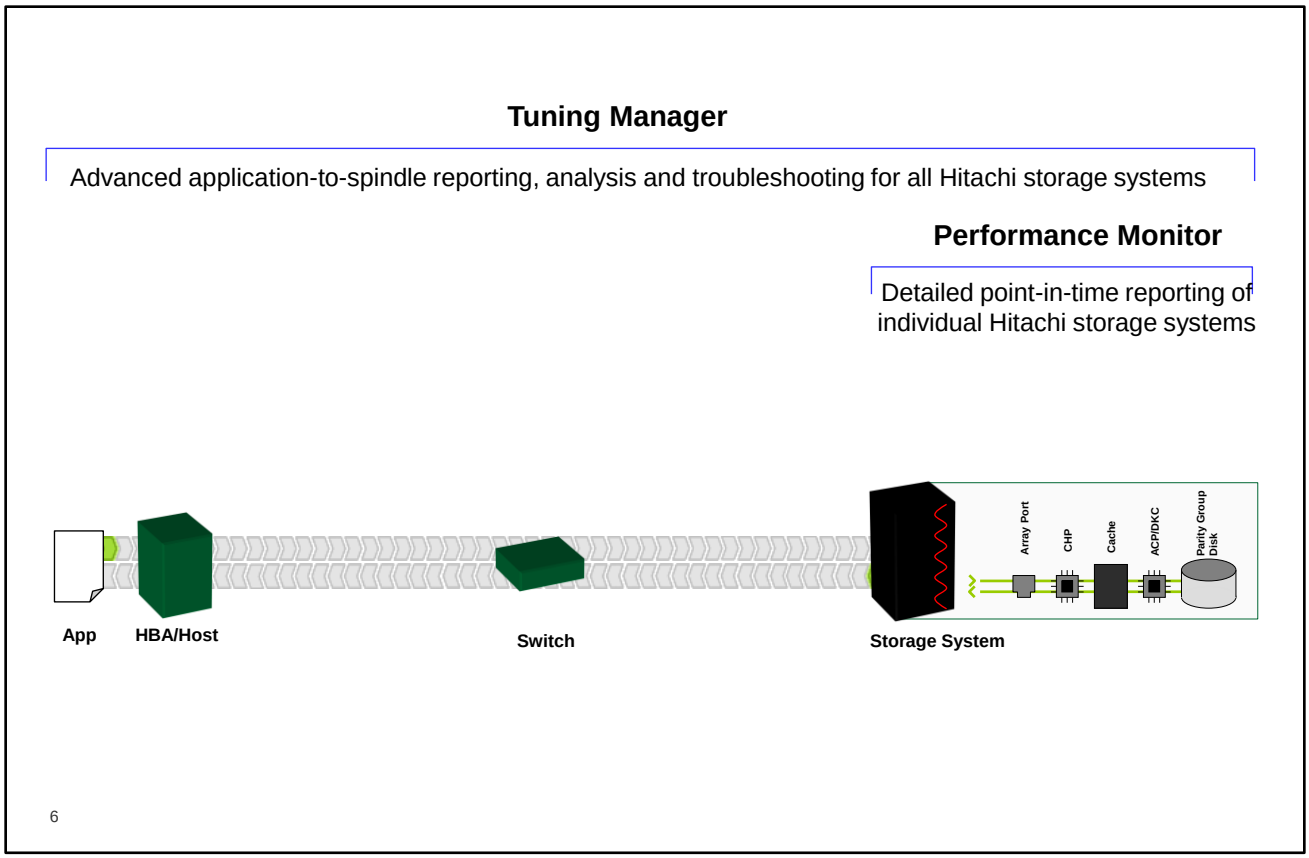
5

Tuning Manager is a premier performance and capacity analysis software tool.

- Its strength is its ability to view performance from the application, through the network and within the storage system.
- It is the most robust performance analysis tool

Performance Monitor

- It provides basic reporting and monitoring within a storage system, but only within the storage system. It has no knowledge of applications.
- It cannot correlate information outside the storage system.



This is a visualization of how these products work, and what they cover.

Performance Monitor provides in-depth, point-in-time information about performance within a Hitachi storage system. It does not provide any information about the network, the host, or the application. Nor does it provide any correlation to that information.

Tuning Manager software provides end-to-end visibility for storage performance. Though limited to Hitachi storage systems, it provides the most thorough view of the system, tracking an I/O from an application to the disk. This ability to correlate this information, and link from step-to-step in the I/O path provides the most efficient solution to identifying performance bottlenecks.

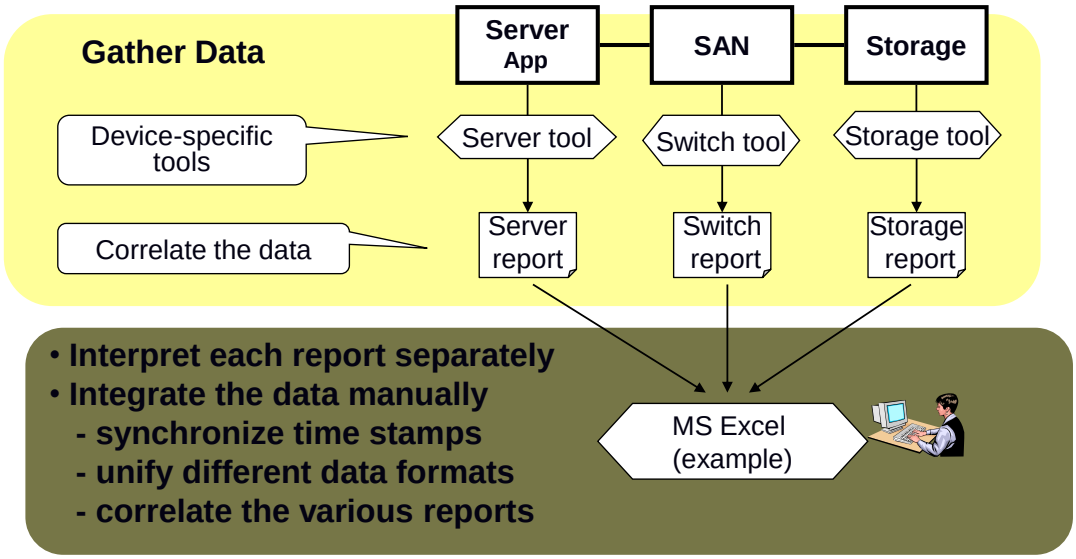
I/O response time, both host side and array side:

It provides the ability to monitor the round trip response time for troubleshooting and proactive service level error condition alerting results in improved application performance. On Hitachi Universal Storage Platform family products this ability includes and extends to round trip response to/from external storage.

HDS Confidential: For distribution only to authorized parties.

The Challenge

- Gather data from servers, databases, switches, and storage systems with device-specific tools, then consolidate, analyze, and correlate data that is presented in different formats



7

The Performance and Capacity Management Challenge of a Networked Storage Environment

Troubleshooting requires a view of the path from the application to the storage system. Without a tool that consolidates and normalizes all of the data, the system administrator has difficulty distinguishing between possible sources. When a performance problem occurs or the "DB application response time exceeds acceptable levels", they must quickly determine if the problem is in the application server.

Server / App Analysis — is the problem caused by trouble on the server? (DB, file system, HBA)

Fabric Analysis — is there a SAN switch problem? (Port, ISL, for example)

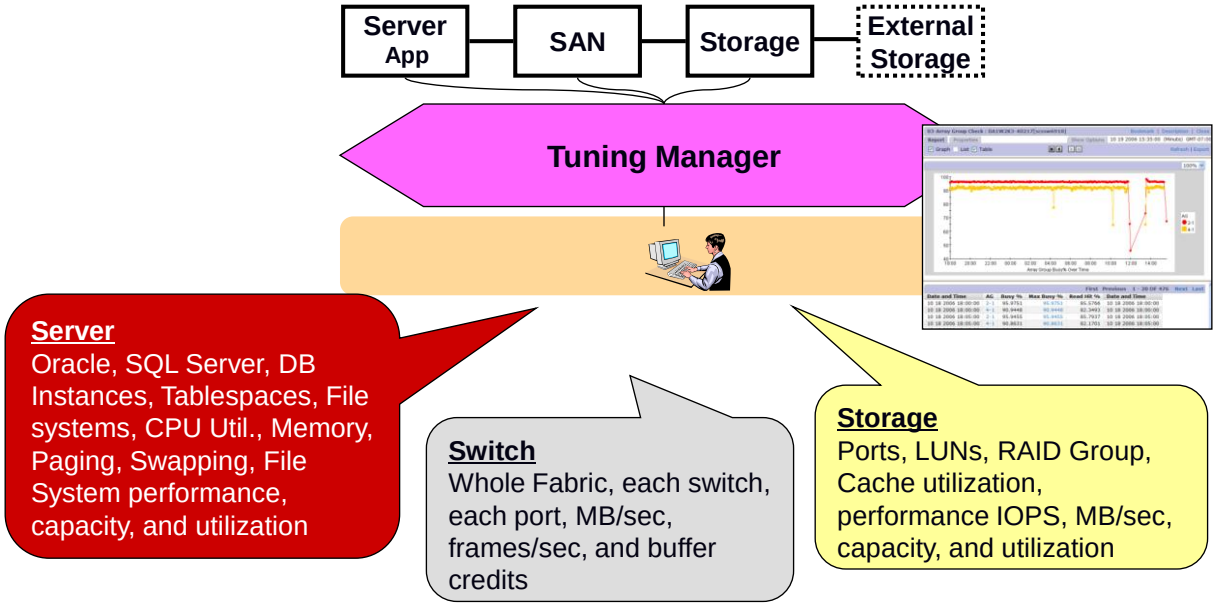
Storage Analysis — is the storage system a bottleneck?

All of the data from the components of the Storage network must be gathered by different device-specific tools and interpreted, correlated and integrated manually, including the timestamps, in order to find the root cause of a problem.

Some customers achieve this by exporting lots of data to spreadsheets and then manually sorting and manipulating the data.

Scope

- Consolidates and analyzes performance and capacity data while hiding platform-dependent differences



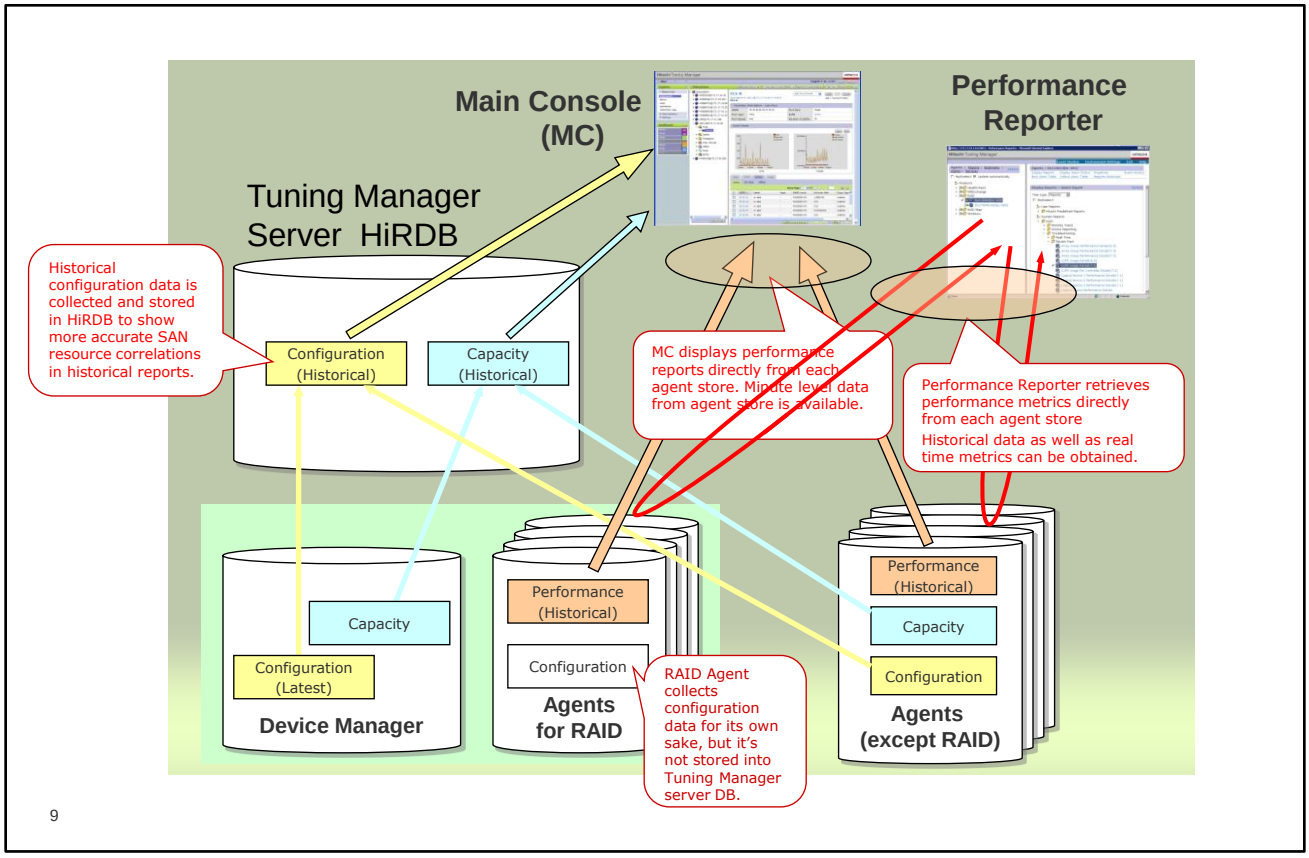
The Tuning Manager software is a collection of programs that provide information allowing for central management of a whole network including SAN. In specific, those program components collect performance and capacity data from various parts of the system, such as host OS, file system, switch and storage systems. Therefore Tuning Manager software is of high value simplifying network management and reducing maintenance costs.

Tuning Manager consolidates the data from the entire storage path. It hides much of the platform dependent differences in performance *and* capacity data from O.S. > DataBase > File System > switch port > storage port > LUN> RAID Group for historical, current, and forecast data.

- Eliminates the user tasks of gathering and integrating metrics
- Provides a single performance view for end to end resources
- Uses automated metrics gathering and various reporting

HDS Confidential: For distribution only to authorized parties.

High-Level Architecture



Storage configuration information is extracted from Device Manager software. This is copied to the Tuning Manager server repository Hitachi Relational Database (HiRDB) and stored as historical information. The storage view of capacity information is also collected from Device Manager and stored in HiRDB.

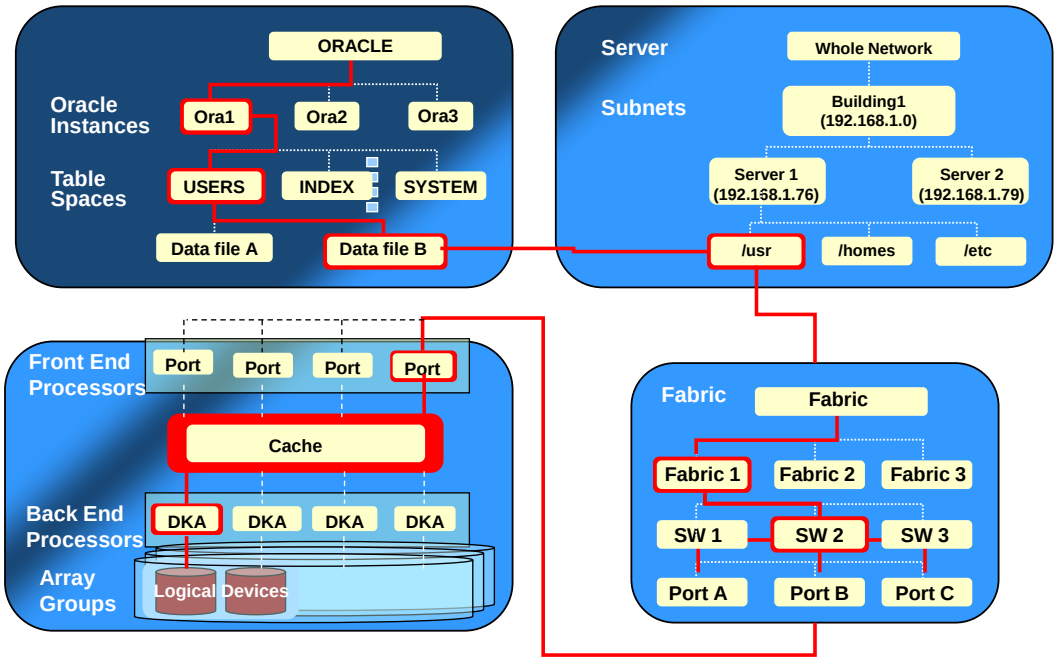
Server and switch configuration information is collected from their respective agents and historical data is stored in the server database; server-based capacity data is collected by the host agents.

Performance information, collected by the agents, is available through Main Console reports, directly from the agent Store databases without going through HiRDB.

The Tuning Manager Performance Reporter retrieves performance metric data from the agents on demand executing predefined system reports or user-defined reports. Depending on the type of report, data will either be taken from the Agent Store database (historical reports) or directly from the resources through the Agent Collector (Real time reports).

Components

- Server, application, and fabric down to storage view



10

Programs in the Tuning Manager software collect performance and capacity data from the OS installed on the servers and databases (such as Oracle) connected to a SAN. Metrics are collected from fabrics and storage systems. Displaying this data facilitates proper management and configuration of each resource in a SAN environment.

HDS Confidential: For distribution only to authorized parties.

- Device Manager
  - Device Manager collects and manages the configuration information of storage devices.
  - Device Manager software is a prerequisite for the Tuning Manager server.
- Tuning Manager server
  - The Tuning Manager server collects information from Device Manager and agents, organizes the information into formats that suit the levels of monitored resources, and then displays the operating status of the system on the screen.
- Agents
  - Agents manage as monitored resources Hitachi storage systems, SAN switches, file systems on hosts, operating systems, and applications such as Oracle according to their features.
  - Agents also collect performance information (such as the I/O count per second) and capacity information (such as logical disk capacity) from these resources as performance data.

11

Device Manager Software

Device Manager provides consistent, easy to use, and easy to configure interfaces for managing storage products. In addition to a command line interface (CLI) for scripting, Device Manager provides a Web-based GUI for managing storage products. Device Manager also provides maintenance commands for backing up and recovering the database that stores configuration information.

Tuning Manager - Distributed Metrics-Repository Architecture

Tuning Manager programs have the capability to manage the large-scale infrastructures of storage environments: This capability is provided by a distributed metrics-repository architecture consisting of the Tuning Manager server database for trend-analysis and forecasting, and distributed Store databases deployed to each Agent for troubleshooting and deeper analysis. The Tuning Manager server database stores configuration and capacity information about resources, while the Store databases of Agents store mainly performance information.

A Tuning Manager server has two viewers:

The Main Console, which allows you to view data retrieved from Device Manager and Agents from a perspective based on the storage environment

Performance Reporter, which allows you to view real-time data and detailed data retrieved from each Store database

Agents

Agents run in the background and collect and record performance data. A separate Agent must exist for each monitored resource. For example, Agent for RAID is required to monitor storage subsystems, and Agent for Oracle is required to monitor Oracle. The Agents can continually gather hundreds of performance metrics and store them in the Store databases for instant recall.

HDS Confidential: For distribution only to authorized parties.

- Main Console
  - The Main Console stores the configuration and capacity information that the agents and Device Manager collect from the monitored resources in the database.
  - The Main Console displays links to Performance Reporter.
- Performance Reporter
  - Performance Reporter displays performance data collected directly from the Store database of each agent.

12

Main Console

The Main Console stores the configuration and capacity information that the Agents and Device Manager collect from the monitored resources in the database. According to the specified time frame and interval, the Main Console displays reports in which the data accumulated in the database is mapped to the performance data managed by the Agents. The Tuning Manager server database is managed by the relational database system HiRDB.

The Main Console displays links to Performance Reporter.

Performance Reporter

Performance Reporter displays performance data collected directly from the Store database of each Agent and also provides a simple menu-driven method to develop your own custom reports. In this way, Performance Reporter enables you to display Agent-instance level reports and customized reports with a simple mouse click. Performance Reporter also enables you to display reports in which the current status of monitored targets is shown in real time. Performance Reporter does not connect to HiRDB.

- **Main Console** allow you to perform the following operations:
  - Obtain an overview of the system, perform comparative data analysis, and understand trends.
  - Display a summary of resource configuration and capacity data that is stored in the Tuning Manager server database and store databases of agents.
  - Display a summary of the resources related to the resources displayed in a report.
  - Simplify analysis tasks, such as comparing data at different points in time, by allowing each user to define and retain time ranges for displaying data.
  - Detailed capacity and performance reports show
    - Historical data
    - Future forecasts
    - Relationship between resources

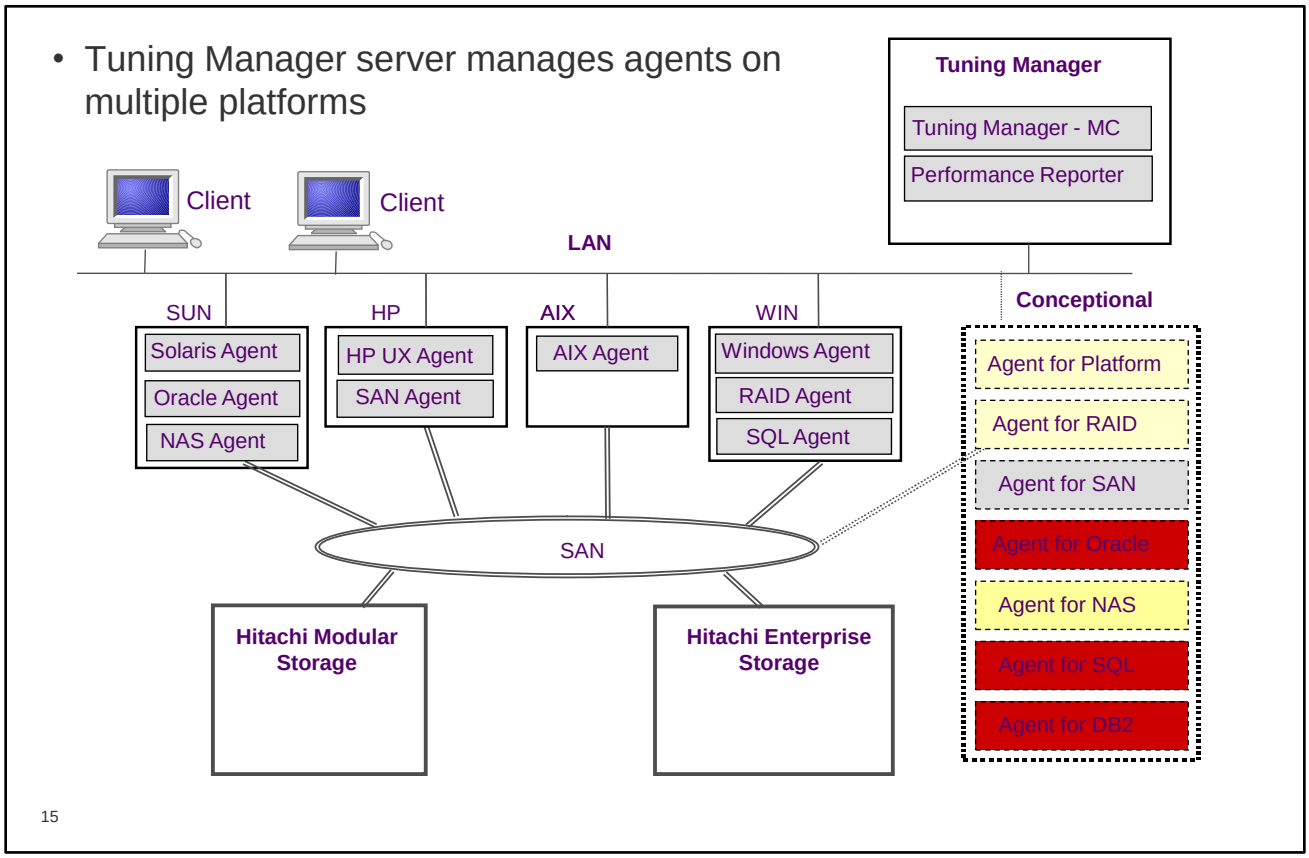
13

- **Performance Reporter** allows you to complete the following operations:
  - Analyze trends and performance in detail
  - Display near-real time reports (by the minute)
  - Specify the display formats of graphs and change time ranges
  - For historical reports, display data stored in store databases of agents
  - Display all data collected by agents
  - Display data from single instance records or multi-instance records in a single report
  - Output CSV or HTML files
  - Display different historical reports in one report

14

HDS Confidential: For distribution only to authorized parties.

Agents for Tuning Manager



Tuning Manager software consists of agents and a server. The agents collect performance and capacity data for each monitored resource, and the server manages the agents. This diagram shows an example system configuration.

Agents can run multiple instances to collect metrics from multiple application instances, fabrics, and storage systems.

The instances of the Agent for RAID collect metrics from enterprise storage systems using inbound Fibre Channel connection communicating to the CMD device in the array. Modular storage is accessed via LAN to collect metric data.

Tuning Manager server can concurrently serve as business server on SUN Solaris and Microsoft Windows in small environments. To be able to manage as many resources as possible with good performance, carefully consider the Tuning Manager system requirements.

Hitachi modular storage includes Hitachi Adaptable Module Storage and Hitachi Thunder series storage.

Agent	Installed On
Platform	Server being monitored
RAID	Enterprise Storage – Server with Fibre Channel connectivity to monitored storage system (uses command device) Modular Storage – Server with IP connectivity to monitored storage system
SAN	Server with IP connectivity to monitored proxy/seed switch and also the SMI Agent for the switch
Database	Server on which database is installed

All Agent instances communicate with Tuning Manager server over IP

16

- Considerations for installing Tuning Manager agents
  - Are all types of agents always required?
  - Are agents required for every application on the production server?
  - Is a server box required for each agent?

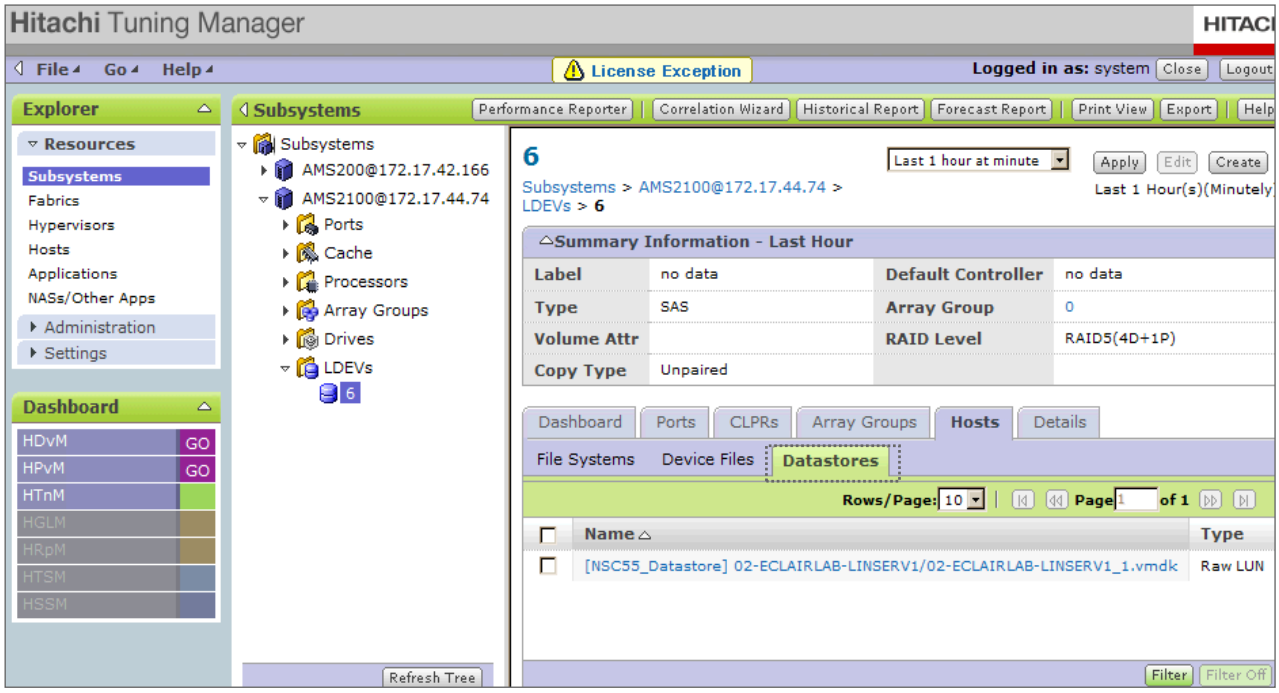
The answer is No…

  - The installation and set up requirements for agents depends on what and how a customer wants to manage their business.

17

HDS Confidential: For distribution only to authorized parties.

Graphical User Interface



The **Datastores** tab is available as Host option in case of ESX Hosts related resources are selected.

Now…..Recorded Product Demonstration

From your Web Browser …. edemo.hds.com/edemo/vault1/tuningmanager_16_part1.swf

Module Review

1. Which two statements most accurately reflect the functions of Tuning Manager software? (Choose two)
   a. Detailed point-in-time reporting of only individual Hitachi storage systems
   b. Advanced application-to-spindle reporting, analysis, and troubleshooting for all Hitachi storage systems
   c. Tracks an I/O from an application to the disk
   d. Advanced application-to-spindle reporting, analysis, and troubleshooting for only Hitachi enterprise storage systems
2. What is a key objective of Tuning Manager software?
   a. Forecasting future I/O usage
   b. Providing configuration and status information
   c. Optimizing LUNS for host attachment
   d. Automatically generating reports

20

HDS Confidential: For distribution only to authorized parties.

16. Hitachi Replication Manager Software

Module Objectives

- Upon completion of this module the learner should be able to:
  - Identify the positioning of Hitachi Replication Manager software
  - Describe the Replication Manager modules and other dependent software components
  - Describe the configuration for open systems
  - Identify the server requirements
  - List prerequisite software for setting up Replication Manager server
  - Configure the server

2

Overview

- Hitachi Replication Manager is the software tool that configures, monitors and manages Hitachi replication products in both open and mainframe environments for enterprise and modular storage systems from a “single pane of glass.”

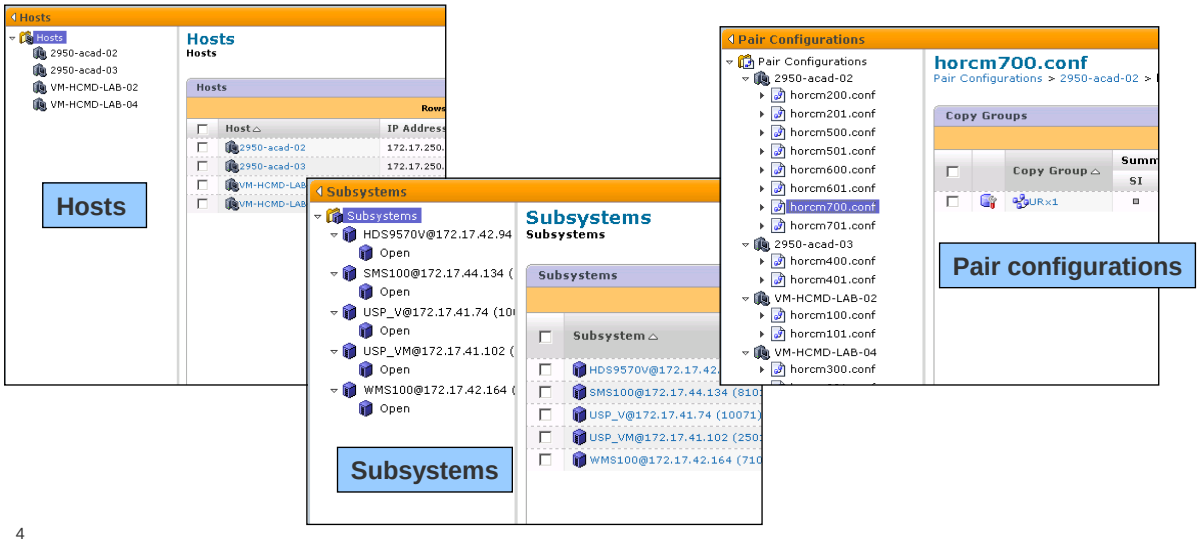
- Key Features
  - Centralized management of replication
  - Integrated Microsoft Exchange backup management
  - Visual representation of replication structures
  - Immediate notification of error and potential issues based on user-defined thresholds
  - Simple wizards for pair creations and changing pair statuses
  - Reports consistency deltas between source devices and their targets

3

HDS Confidential: For distribution only to authorized parties.

Features – Centralized Views

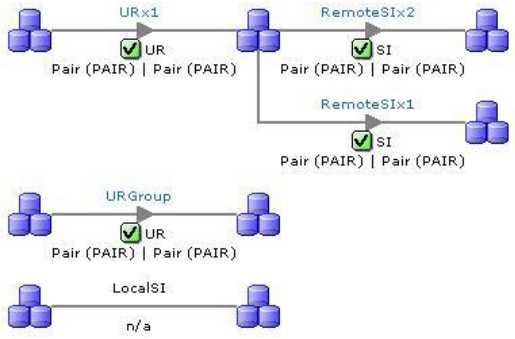
- Three types of views allows users to understand replication environment depending on the perspective
  - **Hosts** View – Perspective of hosts using the pairs
  - **Subsystems** View – Perspective of storage systems containing the pairs
  - **Pair Configurations** View – Perspective of hosts managing the pairs



4

Features – Single Interface

- "Single Pane of Glass"
  - Integrated console for multi-site replication pairs
  - Consolidated monitoring for Copy Pair status and remote copy metrics
- Visual Representation of Replication Structure
  - Copy groups, sites, all volume pairs



5

Visual Representation of Replication Structure

Copy Groups: A group of Copy Pairs created for management purposes, as required by a particular task or job. By specifying a Copy Group, you can perform operations such as changing the Pair Status of multiple Copy Pairs at once.

Using My Copy Groups feature, a user can register a Copy Group into My Copy Groups that are most important to monitor and see how Copy Groups are related and check Copy Pair statuses in a single window. My Copy Groups is also the default screen after you login to the Hitachi Replication Manager software interface.

Sites: With Hitachi Replication Manager software, you can define logical sites in the GUI just as you would define actual physical sites (actual data centers). It allows you to manage resources more efficiently if you setup separate sites because it is easier to locate a required resource among many resources displayed in the GUI.

HDS Confidential: For distribution only to authorized parties.

Features – Pair Management Functions

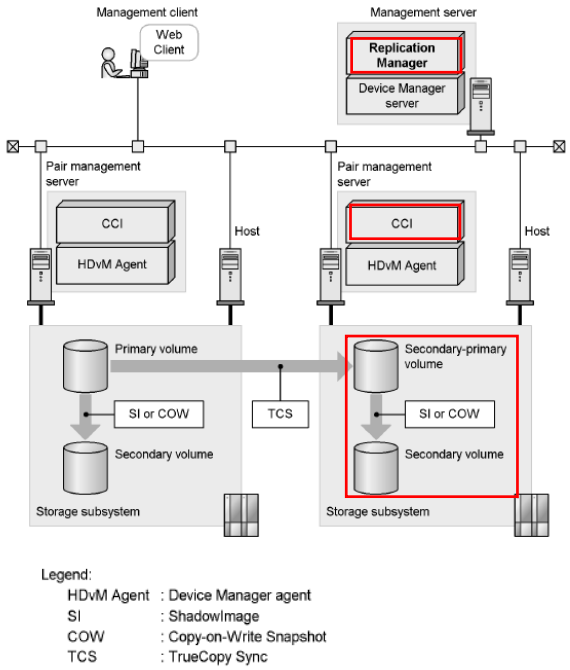
- Pair Volume Lifecycle Management
  - Simplified replication configuration from setup to deletion
    - **Setup > Definition > Creation (Initial copy) > Operation > Monitoring > Alerting > Deletion**
- Storage System Configuration Functions
  - Setup functionality required for Copy Pair management
    - Setting command devices, DMLU, Journal Groups and pools
    - Setting up Remote Paths for remote replication
- Copy Pair Creation/Deletion
  - Pair Configuration Wizard
    - Intuitive pair definition screen with topological view
  - Task Scheduler
    - Scheduler functionality that allows users to execute the copy operations off-peak time

6

System Configuration

System Components include:

- HRpM/HDvM Management Server
  - Provides information in response to requests from clients
  - HDvM is installed (pre-requisite)
- Pair Management Server
  - Collects information, including copy pair status and performance information for pair copy groups
  - Used by HRpM for pair operations
  - CCI is installed here and HORCM files are stored (pre-requisite)
  - HDvM Agent is installed (pre-requisite)
- Storage System
  - Pairs and Copy Groups
  - Replication Program Products are licensed here (pre-requisite)



7

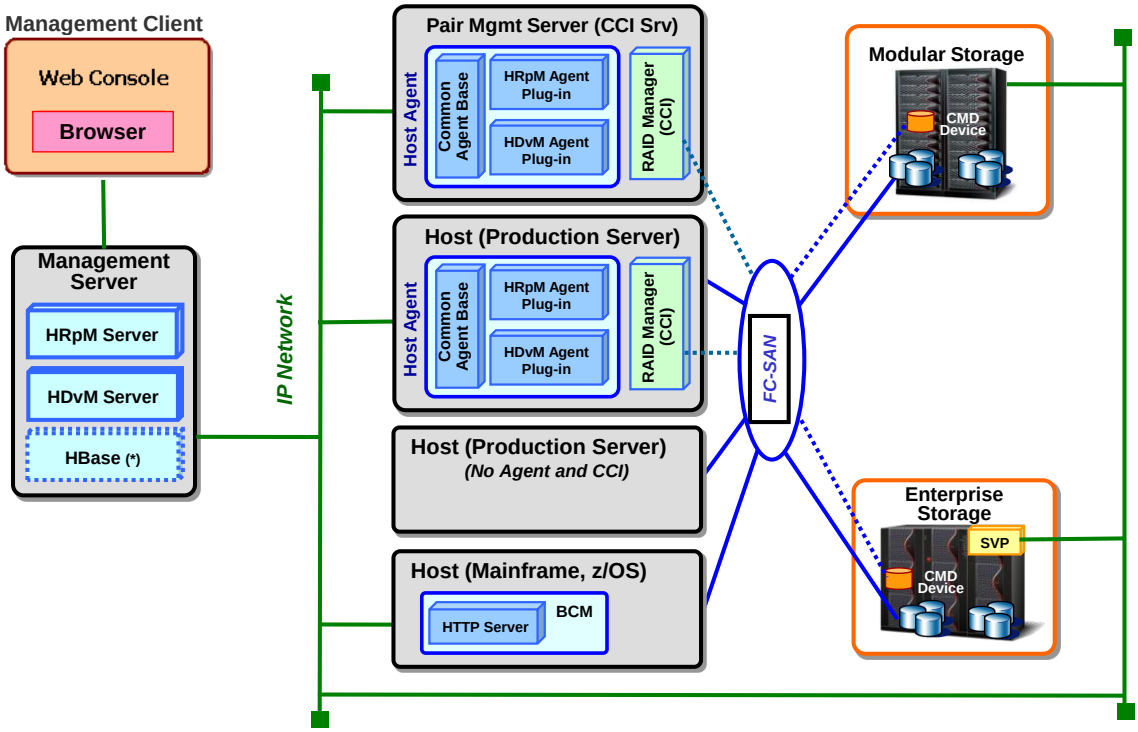
HRpM stands for Replication Manager software

HDvM stands for Hitachi Device Manager software

HDS Confidential: For distribution only to authorized parties.

Architecture in an Open System

- Standard Configuration of a Site



Overview of Prerequisite Software

- Before installing Replication Manager software, install and configure the following:
  - Device Manager software 6.x on the management server
  - Device Manager license key
  - CCI and Device Manager agents 6.x on the pair management servers for open systems

9

HDS Confidential: For distribution only to authorized parties.

Managing Users and Permissions

- Replication Manager software implements access control in two ways:
  - Uses management roles and user permissions that restrict the operations that users can perform
  - Uses *resource groups* that restrict the range of resources that specific users can access

10

- All users can set up personal profiles and Replication Manager licenses regardless of their permissions.
- The built-in user ID, **SYSTEM,** lets you manage all users in Storage Command Suite. You cannot change or delete this user, **SYSTEM,** or its permissions.



11

The user ID, **Peer**, is internally used by an agent.

HDS Confidential: For distribution only to authorized parties.

Concept of Resource Groups

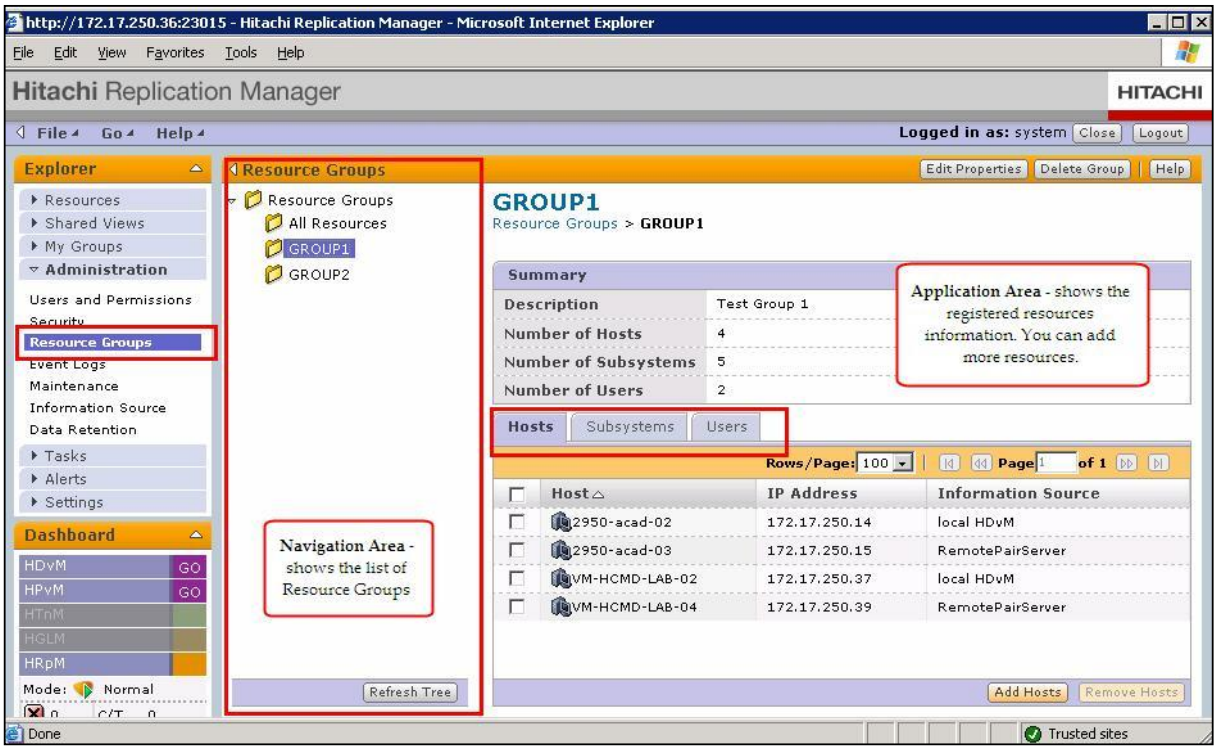
- Resource groups provide access control functionality.
  - A *resource group* is a collection of hosts and storage systems that are grouped by purpose and associated with a user for controlled access by the user.
  - Large environments require security management for managing resources such as:
    - Who can access this storage system?
    - An administrator is assigned to hosts and storage systems that are grouped by a site or department.
  - Users can only see the allocated resources on the GUI.
  - A user can be associated with multiple resource groups to increase the range of operations.
  - Default resource group – All Resources
    - User **SYSTEM** is added by default

- Types of Resource Groups
  - All Resources
    - System-defined, containing all the resources in the system
  - User-Defined
    - Users with administrative privileges can define a resource group and add resources, such as hosts and storage systems.

- Any new user created in the system should be explicitly associated with a resource group.
  - A given user can be associated with multiple resource groups.
  - Any user can be added to the All Resources group if they do not belong to another resource group.

13

Resource Group View



14

HDS Confidential: For distribution only to authorized parties.

17. Hitachi Tiered Storage Manager Software

Module Objectives

- Upon completion of this course, the learner should be able to:
  - Describe customer storage management needs
  - Describe the customer business value of Hitachi Tiered Storage Manager software
  - Describe functionality and architecture

2

Customer Storage Management Needs

- Organize storage resources into logically organized tiers based on business needs
- Simplify the management of heterogeneous storage devices
- Provide storage resources as logical pooled resources that hide the underlying complexity
- Group storage devices based on performance, availability, cost, and security
- Proactively manage resources per service level agreements (SLA) and operating level agreements (OLA)
- Dynamically move data following its natural lifecycle process

3

HDS Confidential: For distribution only to authorized parties.

Overview

- Hitachi Tiered Storage Manager (HTSM) software enables you to **non-disruptively migrate data** based on the principal of Service Oriented Storage Solutions (SOSS).
  - SOSS provides safe and efficient data storage by transferring data to the most suitable storage resource in response to changing storage requirements for applications.
- HTSM software repositions (migrates) data on storage system volumes to other volumes simply by specifying a migration target.
  - Migrations can be between volumes
    - Internal to internal
    - Internal to external or vice-versa
    - External to external

4

Note: Migrations between Normal LDEVs and DP-VOLs is also supported.

Benefits of Tiered Storage Manager

- Easy to use service-level based volume migration
- Non-interruptive, completely transparent volume movement
- Safely moves data when application requirements increase or decrease
- Can schedule automatic migrations through CLI
- Removes complexity in managing storage tiers
- Simplifies steps to migrate data prior to storage reassignment or removal, such as migrating hardware coming off lease
- Eliminates restrictive data migration windows and requirements
- Helps respond to data retention standards and compliance issues
- Can utilize LDEVs that are larger than the source LDEV in a migration as a Migration Target.[1]

5

[1] This function carves the same size LDEV out of a larger target LDEV. The Migration is still same size LDEV to LDEV. Space freed from carving out the target from a larger LDEV is returned to Array Group free space. This function increases the Target Candidates for Migrations and is helpful when the user does not have same size LDEVs available for Targets.

HDS Confidential: For distribution only to authorized parties.

Requirements

- Tiered Storage Manager software relies on other Hitachi products for its operations:
  - Hitachi Device Manager software
    - Communicates with the storage controller system
    - Distributes storage system information to HTSM software upon request
    - Provides an application programming interface (API) to HTSM software
  - Hitachi Volume Migration V2 (license installed on storage system)
    - Contains the migration engine embedded in the microcode of the Hitachi Universal Storage Platform (USP) family of systems
    - License automatically included with HTSM software license

6

Software Options

- Tiered Storage Manager software features can be expanded using optional Hitachi products for its operations.
  - Linkage with Hitachi Tuning Manager software enables additional search fields to select volumes (for tiers)
    - Array Group Busy Rate
    - Array Group Max Busy Rate
  - Hitachi Volume Shredder software is required to enable data shredding
    - The volume shredding engine is embedded in the USP family microcode
      - Can be licensed to enable volume shredding at the end of Migration Tasks or standalone shredding tasks in HTSM
  - Hitachi Data Retention Utility can be used for providing data protection for volumes involved in migrations using Tiered Storage Manager.

7

HDS Confidential: For distribution only to authorized parties.

Components

- Tiered Storage Manager server
  - Component that manages all operations of the product
  - Through Device Manager software sends instructions to the domain controller of the storage system
- Management Client
  - Component used to communicate with HTSM software server via GUI or the command line interface (CLI)
- Domain Control Storage system
  - USP family storage system which performs the migrations

8

Volumes Used in Tiered Storage Manager

- Migration source volume and migration target volumes conditions:
  - Emulation types of source and target volumes must be the same.
    Note: The emulation type is different for OPEN-V and OPEN-0V.
  - Number of blocks for source and target volumes must be the same.
  - Sizes of the source and target volumes must be the same.
  - If volumes are normal volumes and have the emulation type OPEN-V, migration is possible even if the target volume is larger than the source volume.
    - Target volume is deleted
    - Volume that has same capacity as source volume is re-created
  - If the emulation type is not OPEN-V, both source and target volumes must be either CVS or normal volumes.

9

When creating a migration plan, even if you specify a target volume with a greater capacity than that of the source volume, you can perform migration. If the capacity of a target volume is greater than that of the source volume, HTSM software re-creates the target volume when migration is performed, making the volume the same size as the source volume. This functionality is now standard. If you want to perform operations using only the older functionality, specify 0 for the property server.migrationPlan.candidateCapacityGroupDisplayMaxCount in the property file server.properties.

HDS Confidential: For distribution only to authorized parties.

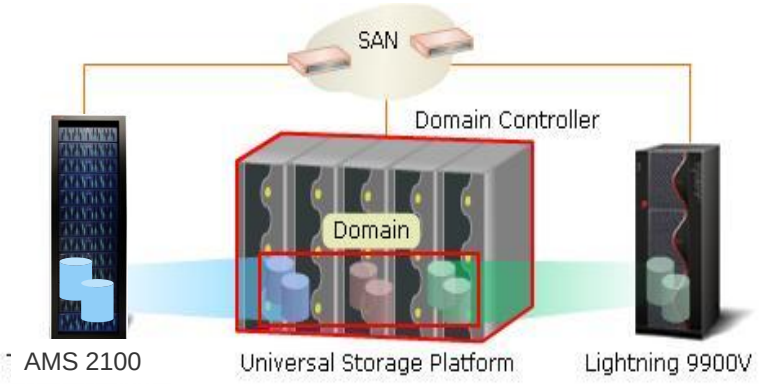
Concepts

- Tiered Storage Manager software introduces several concepts:
  - Storage Domain
  - Migration Group
  - Storage Tier
  - Migration Task

10

Storage Domain

- A Storage Domain includes all storage available through any USP family system.
  - Includes internal- and external-attached storage
  - One storage controller = one Storage Domain
  - HTSM migration can only be performed between volumes within a Storage Domain



11

HTSM operates and manages storage systems by Storage Domain. All volumes in the Storage Domain are under the control of a domain controller. For HTSM, you can use any USP family system as the domain controller.

Storage Domains and domain controllers are associated on a one-to-one basis. Therefore, for a single HTSM server, each Storage Domain should have a unique name.

HDS Confidential: For distribution only to authorized parties.

Migration Group

- A Migration Group is a group of volumes that can be migrated (or prevented from migrating) together.
  - The volumes in a Migration Group serve as migration source volumes.
  - Typically, the volumes share both a common application or host affinity as well as demands upon storage services.
    - Example: A Migration Group could be an application's volumes used for logging.
  - Once an actual list of volumes is created, supplemental volumes can be added as required.
    - Example: Volumes added to a given application.

HTSM enables the grouping of a specified set of volumes that need to be managed as a collection. Such a grouped set of volumes is called a **Migration Group**. A Migration Group is created within a Storage Domain. Several Migration Groups can be created within a single Storage Domain.

For creating a Migration Group, volumes could be categorized according to such criteria as:

- Usage by application
- Usage by different organizational units
- Other volume filter conditions

While performing migration, certain rules have to be considered:

- Migrations can occur immediately or can be scheduled for a future date and time.
- In order to schedule a Migration Task to occur in the future, you can use a scheduler such as AT or Cron.

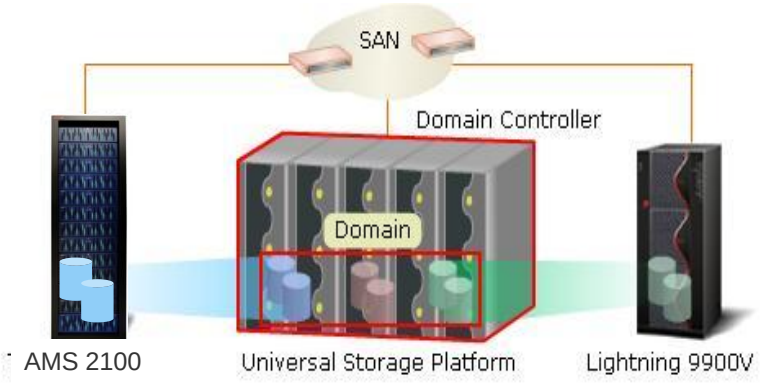
Specific migration can have three possible states:

- Standby (Migration can be cancelled if it is on a standby status).
- Active
- Success – failed

If the number of migration requests exceeds the maximum number of concurrent copies, the requests are queued in the domain controller.

Storage Tier

- A Storage Tier provides a collection of migration target volumes based on a set of search criteria.
  - Search filter is saved in the Storage Tier definition, not the list of volumes.
  - Search filter is applied when the Storage Tier is used while creating a Migration Task.
- The Storage Tier definition reflects the set of requirements a Migration Group wants to have upon storage assets.
  - RAID-5, Fibre Channel, Array Group Usage (if HTM is integrated) etc.



13

A group of volumes that serve as migration targets is called a Storage Tier. When you migrate data, you select a target volume from a specified Storage Tier.

Storage Tiers help you to view storage resources in a more organized manner. For example, creating Storage Tiers based on performance can help you differentiate available high performance resources from overused storage resources. This helps you decide the placement of different sets of data, and plan what data needs to be migrated where.

You can create several Storage Tiers in a domain. Also, one volume can qualify as a member of more than one Storage Tier, based on the parameters used for classification.

Using HTSM, it is imperative to know about the relationship between a Migration Group and Storage Tiers. When migrating data, the data stored on the volumes in the Migration Group is copied onto the volumes in the Storage Tiers. This means that, after the migration has completed, the source volumes in the Migration Group are relocated within the Storage Tier. In addition, the target volumes get excluded from the management target of the Storage Tier after the migration.

However, if any of the volumes in the source Migration Group is already a part of the target Storage Tier, then, that volume will not be migrated

HDS Confidential: For distribution only to authorized parties.

Migration Task

- Migration is the act of relocating data from one storage system volume to another.
- Migrations can be performed immediately or held for later execution
- After a successful migration, the old source volume has “stale” data
  - A Migration Task can be enabled to shred the stale data on the old physical location for the source volume

14

Migration Task refers to movement of data from the source volume to the target volume for each volume in a Migration Group. Migration Task is created by linking a migration source volume in the Migration Group with a migration target volume in the Storage Tier and specifying options

Note: LDEV sizes are not all created equally. Different Hitachi systems employ different methods when calculating capacities. As a result, LDEVs are slightly different from one array to another by a small but significant rounding size. As a consequence, great challenge is met when exact LDEV size matching is required to perform operations such as LDEV migrations.

Using CVS, any LDEV from any external subsystem can be migrated to the USP. The external LDEV must be attached as OPEN-V.

Storage Management Command Suite (SMCS) V6.0 provides the ability to utilize LDEVs that are larger than the source LDEV in a migration as a Migration Target. This is not migration to a larger LDEV. This function carves the same size LDEV out of a larger target LDEV. The migration is still same size LDEV to LDEV. Space freed from carving out the target from a larger LDEV is returned to Array Group free space. This function increases the Target Candidates for Migrations and is helpful when the user does not have same size LDEVs available for Targets.

Operations

- Steps
  - Create Storage Domain
    - Specify Domain Controller
  - Create Migration Groups
  - Create Storage Tiers
  - Create and Execute Migration Tasks
  - Check Task Status for completion

15

HDS Confidential: For distribution only to authorized parties.

Recorded Demonstration

Now…..Recorded Product Demonstration

From your Web Browser …. edemo.hds.com/edemo/vault1/HTSSolution.swf

Module Review

1. In Tiered Storage Manager software, which one of the following statements is true of a Storage Domain?
   a. Requires a Hitachi enterprise or Hitachi modular storage system
   b. The source and the target of a migration must be in the different domains
   c. Highest organizational level in Tiered Storage Manager software
2. Which three functions must be completed to perform volume migration in Tiered Storage Manager software? (Choose three)
   a. Reserving target volumes
   b. Creating a Storage Domain
   c. Creating a migration group
   d. Creating at least one cache logical partition (CLPR)
   e. Creating a storage tier

17

HDS Confidential: For distribution only to authorized parties.

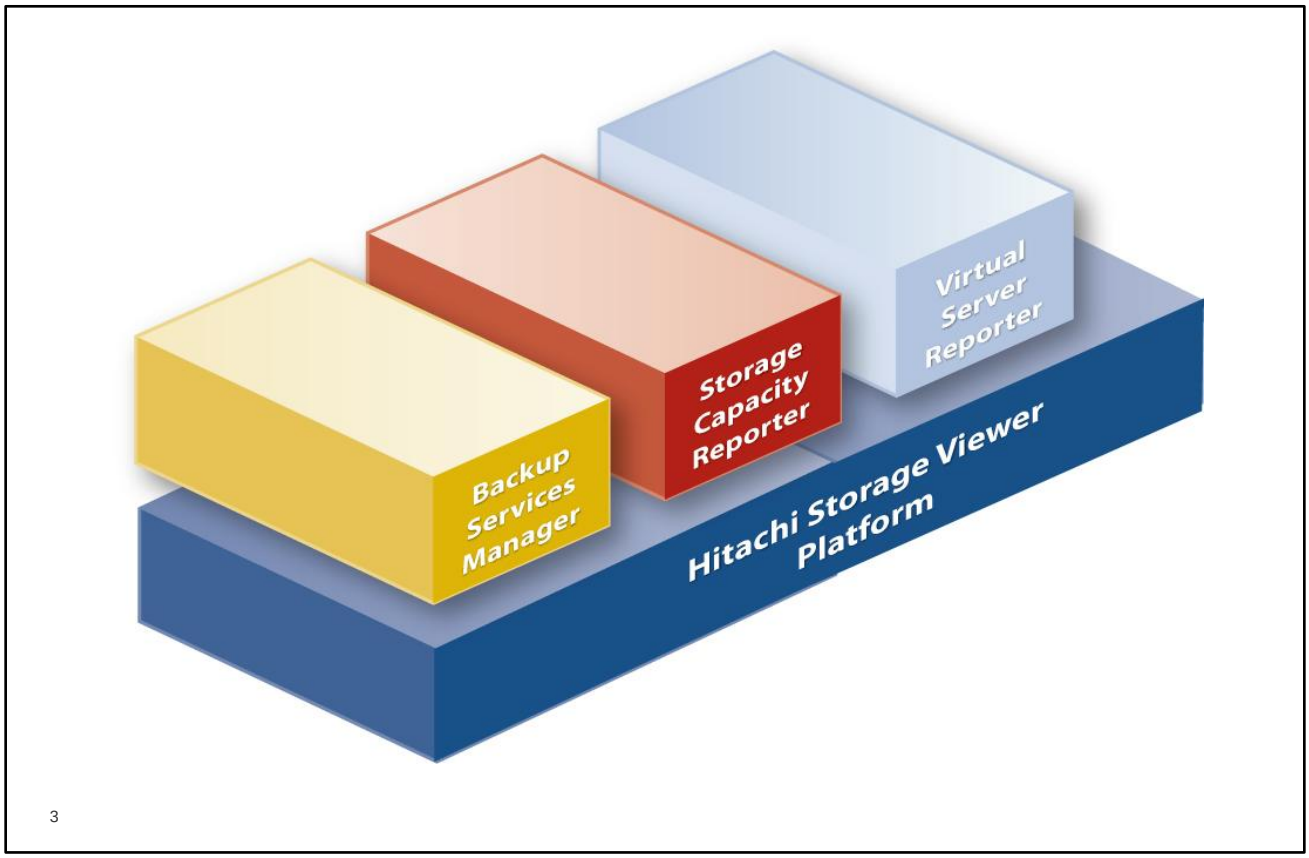
18. Hitachi Storage Capacity Reporter Software

Course Objectives

- Upon completion of this module, the learner should be able to:
  - Describe the key value messages of the Hitachi Storage Capacity Reporter software and how it supports capacity management and planning challenges
  - Describe the Storage Capacity Reporter architecture

2

Hitachi Storage Viewer Architecture Platform



HDS Confidential: For distribution only to authorized parties.

Features

- Agent-less Storage Capacity Reporting
  - Hitachi Storage Capacity Reporter software provides end-to-end storage capacity reporting from applications and hosts to heterogeneous storage arrays without the use of host based agents.
- Hitachi Storage Capacity Reporter
  - Storage capacity utilization reporting for storage arrays, hosts, and leading enterprise applications, including Oracle, Microsoft Exchange, Microsoft SQL Server
    - Current and historical views
    - Predictive analysis for future projected growth
  - Heterogeneous storage array support for Hitachi, EMC, Sun StorageTek, HP XP
  - Integrated with Hitachi Backup Services Manager software
  - Web 2.0 architecture
    - Ease of use
    - GUI speed
    - Multiple reporting options

4

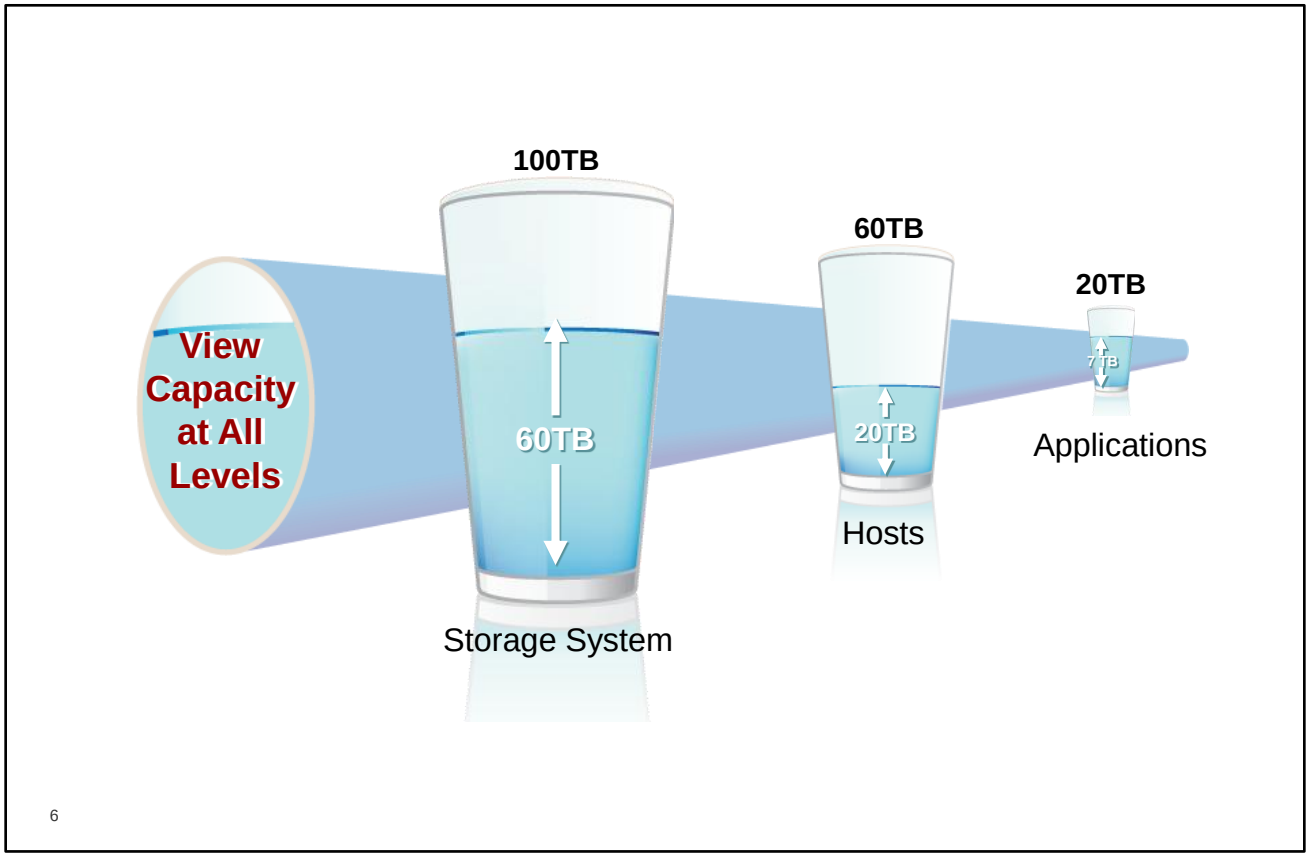
Key Benefits

- Increase Management Efficiency
  - Shows the relationship between each host and LUN in your environment to convey how the storage resources are being used.
  - Host utilization shows the amount of storage capacity allocated and used by each host.
  - Array utilization shows the amount of allocated storage from each storage system.
  - Application utilization identifies the amount of storage used versus allocated by application.

5

HDS Confidential: For distribution only to authorized parties.

Complete End-To-End Storage Capacity View



Hitachi Storage Capacity Reporter provides a host and application view to storage capacity use when host agent deployment is restricted that complements the storage array capacity viewpoint provided by the other products in the Hitachi Storage Command Suite such as Hitachi Device Manager and Hitachi Tuning Manager software.

Supported Storage Systems

- Hitachi Universal Storage Platform V
- Hitachi Universal Storage Platform VM
- Hitachi Universal Storage Platform
- Hitachi Network Storage Controller
- Hitachi Adaptable Modular Storage
- Hitachi Workgroup Modular Storage
- Hitachi Thunder and Lightning Storage Systems

- CLARiiON
- Symmetrix
- DMX

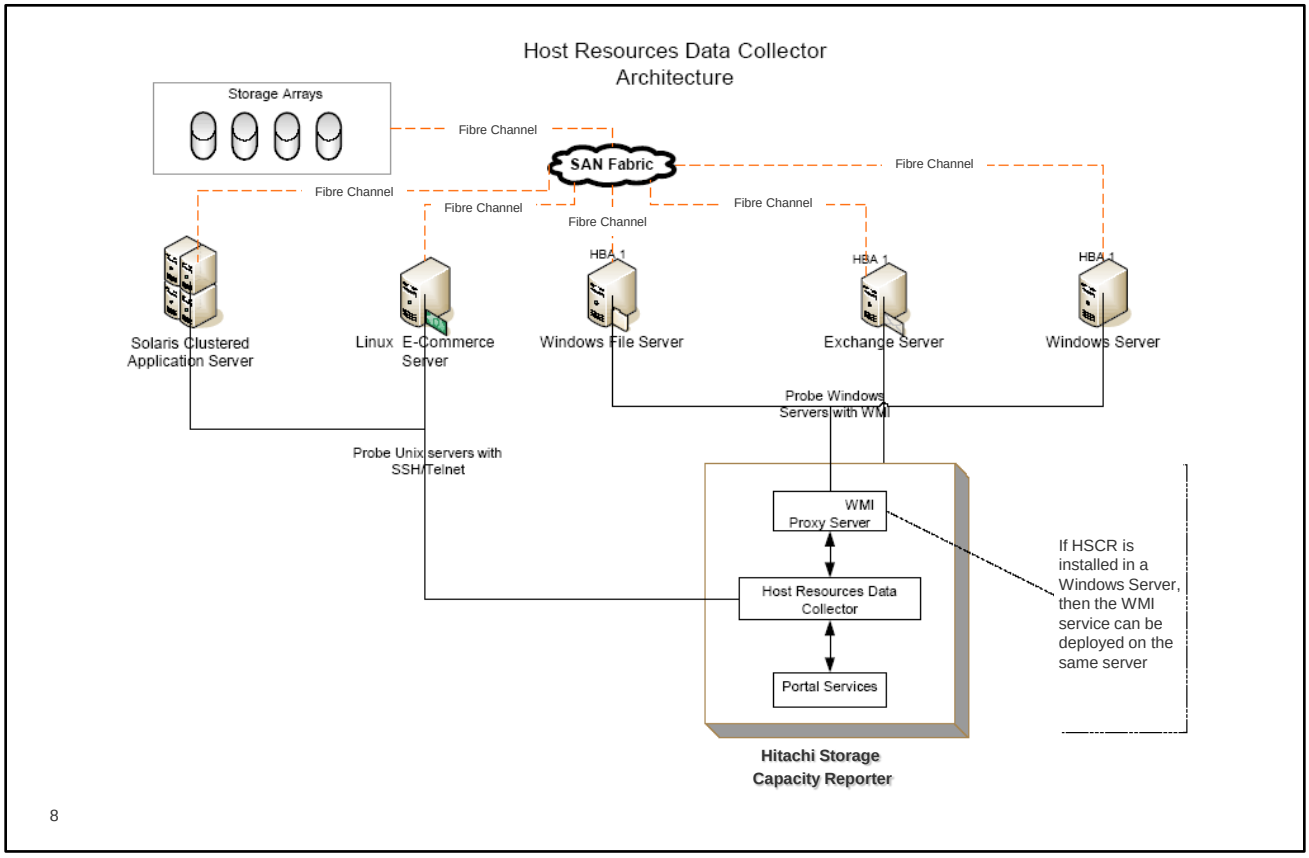
- XP Series

- Sun StorageTek 9900 Series

7

HDS Confidential: For distribution only to authorized parties.

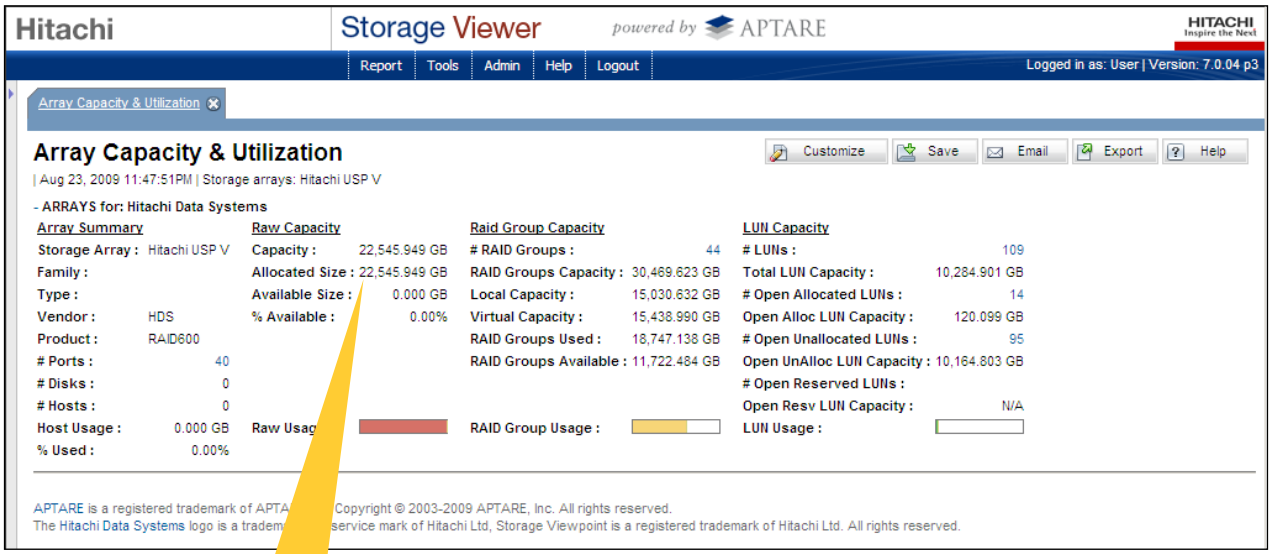
Architecture



Hitachi Storage Capacity Reporter (HSCR) utilizes operating systems API, such as WMI for Windows, and application specific API to gather the respective host and application capacity information. Even though no host agent is required, Storage Capacity Reporter will still need to log in to each host to gather the desired capacity information. HSCR uses OS commands on host to collect the required information. For example, df on Linux and Solaris is used to collect capacity information.

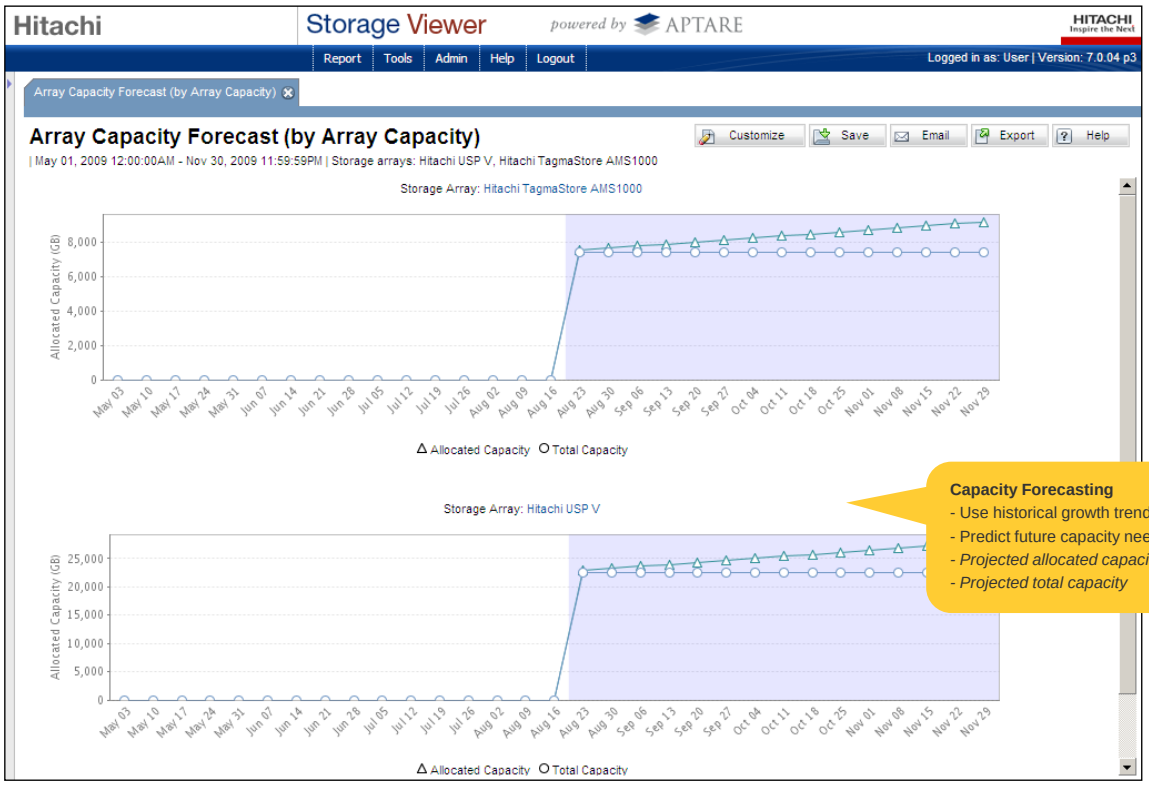
If Hitachi Storage Capacity Reporter is installed in a Windows Server then the WMI service can be deployed on the same server. Otherwise, the WMI service needs to be deployed on a Windows Server within the enterprise network.

End-to-End Storage System Capacity



HDS Confidential: For distribution only to authorized parties.

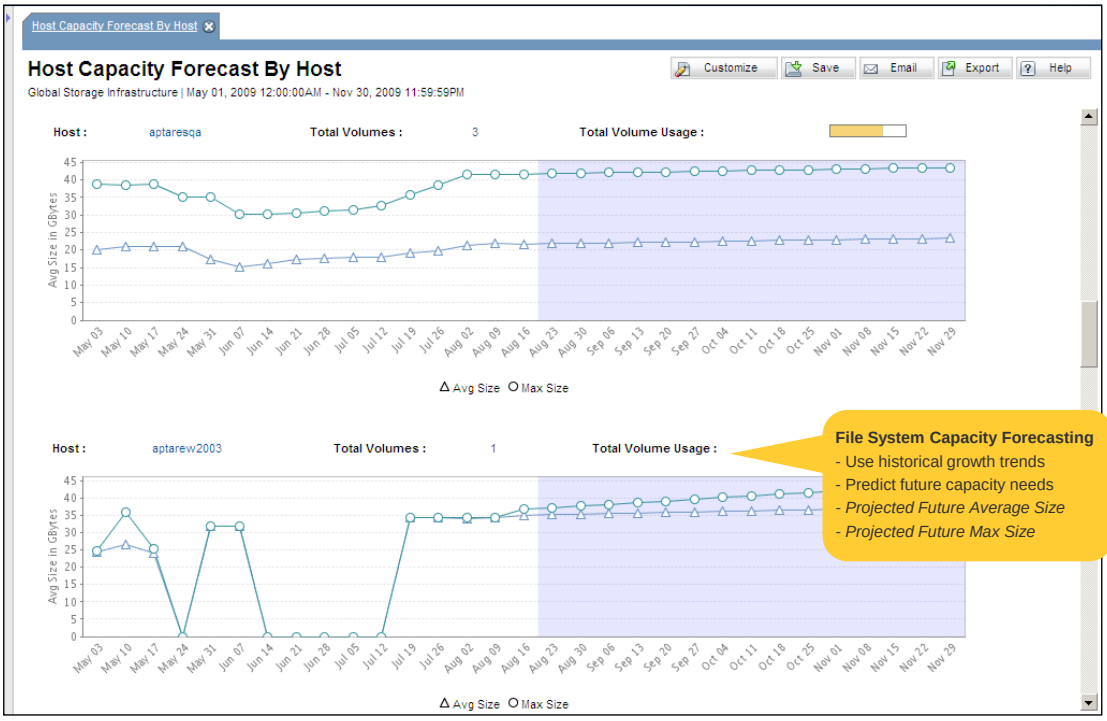
Storage System Capacity Planning



10

Features

- Host and File System - Capacity Planning



11

HDS Confidential: For distribution only to authorized parties.

Review Questions

1. Hitachi Storage Capacity Reporter depends on agents to collect capacity information.
   a) True
   b) False

2. Hitachi Backup Services Manager depends on Hitachi Capacity Reporter to collect backup capacity information.
   a) True
   b) False

12

HDS Confidential: For distribution only to authorized parties.

19. Hitachi Data Protection Suite

Module Objectives

- Upon completion of this module, the learner should be able to:
  - Discuss the components and architecture of the Hitachi Data Protection Suite
  - Identify the core components
  - Identify the differences between data archiving and data protection

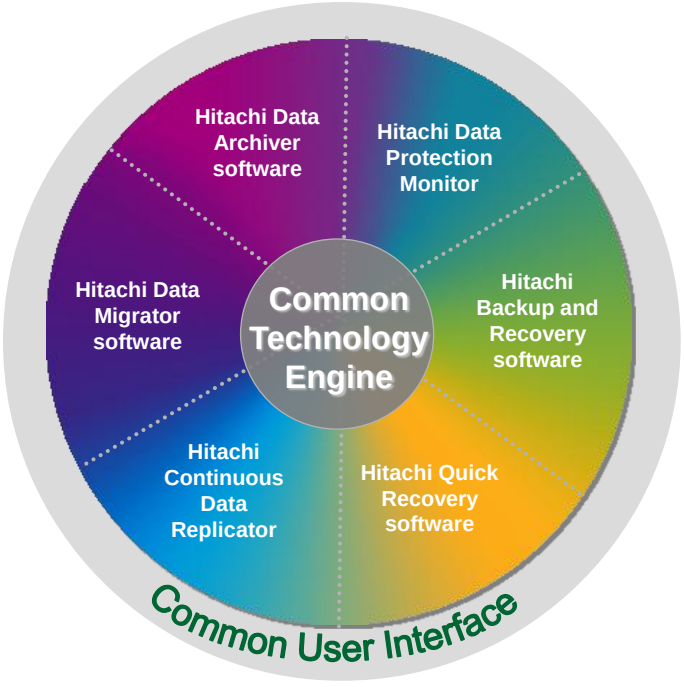
2

Hitachi Data Protection Suite, powered by CommVault®

What is Data Protection Suite?

- Unified platform comprising solutions for data backup and recovery, migration, archiving, and replication
  - Designed for disk-to-disk storage area network (SAN) -based data protection.
  - Provides unified cataloging, indexing, and movement of data at various levels of granularity via a common user interface and policy engine.
- Hitachi Backup and Recovery
  - **Intelligent Media Agents**
    - Integrates with OS and applications to provide granular recovery:
      - Win 2000/2003 system state
      - Oracle, DB2, SQL Server tablespace or files
      - Exchange message, note or mailbox

- Backup and recovery software
- Data migrator software
- Data archiver software
- Quick recovery software
- Data protection monitor
- Continuous data replication



4

HDS Confidential: For distribution only to authorized parties.

The Data Protection Suite platform was designed to automate management of data by changing the approach from static to dynamic. It combines traditionally separate products for data movement with products for data management. Built on the proven CommVault Common Technology Engine (CTE) that moves, manages and catalogues data, the Data Protection Suite platform integrates the following storage management functions:

- Backup and Recovery Software
- Data Migration Software and Data Archiving Software
- Provide high data availability
- Enterprise data management
- Remote office data protection

The resulting product is a single automated platform with a global catalogue, indexing, data movement, and media management.

Backup and Recovery Software

Backup and Recovery

Backup and Recovery software offers you fast, granular backup and recovery of distributed heterogeneous environments. With Backup and Recovery software, policy-based actions can be enforced across all your operating environments, applications and storage types. This simplifies the management of complex storage environments. The Data Protection Suite Storage Environment is comprised of all data storage resources. Unlike other solutions, Backup and Recovery software addresses storage resources and data as logical entities rather than physical resources. Client data is logically mapped to storage resources (called *Storage Policies*), which eliminates an administrator's need to identify the physical location of data before attempting a restore.

Data Migrator Software

Hierarchical Storage Management (HSM)

Data Migrator software allows you to use existing storage resources more efficiently by overseeing the migration of your data to less expensive secondary storage. This includes file system data from a Microsoft® Windows® or UNIX host and also Microsoft Exchange Email messages and attachments. By shrinking the size of your primary storage, Data Migrator software reduces the backup window and prevents over-capacity, thereby easing IT administrative burden and improving system performance.

Data Archiver Software

Archiving for Exchange

Data Archiver software delivers the data archival, search and retrieval capabilities necessary to meet the mounting demand on organizations for consistent data retention and retrieval capabilities. This demand requires companies to meet data

retention regulations and standards set by the government such as SEC 17a-4 or Health Insurance Portability and Accountability Act (HIPAA) as well as establish policies to comply with laws, like the Sarbanes-Oxley (SOX) Act of 2002.

Quick Recovery Software

Point-in-Time Replication

Quick Recovery software offers a simple, unified interface that creates, administers, and tracks point-in-time copies of production data for testing, recovery, or business continuance reasons. Via centralized policy management and scheduling, administrators can define point-in-time Recovery Volume retention periods and recover data when it is required.

Data Protection Monitor Software

Service Level Management

Data Protection Monitor software can change the way IT departments manage data storage. It does so by providing the administrator with tools to create, define, monitor and run storage and data management services. By mapping these services to business units, Data Protection Monitor software gives the administrator a better understanding of storage resource utilization. This will ensure successful delivery of data protection, migration, archiving and replication services. Data Protection Monitor software consolidates physical, logical and application views of storage resources from across the datacenter or across the globe.

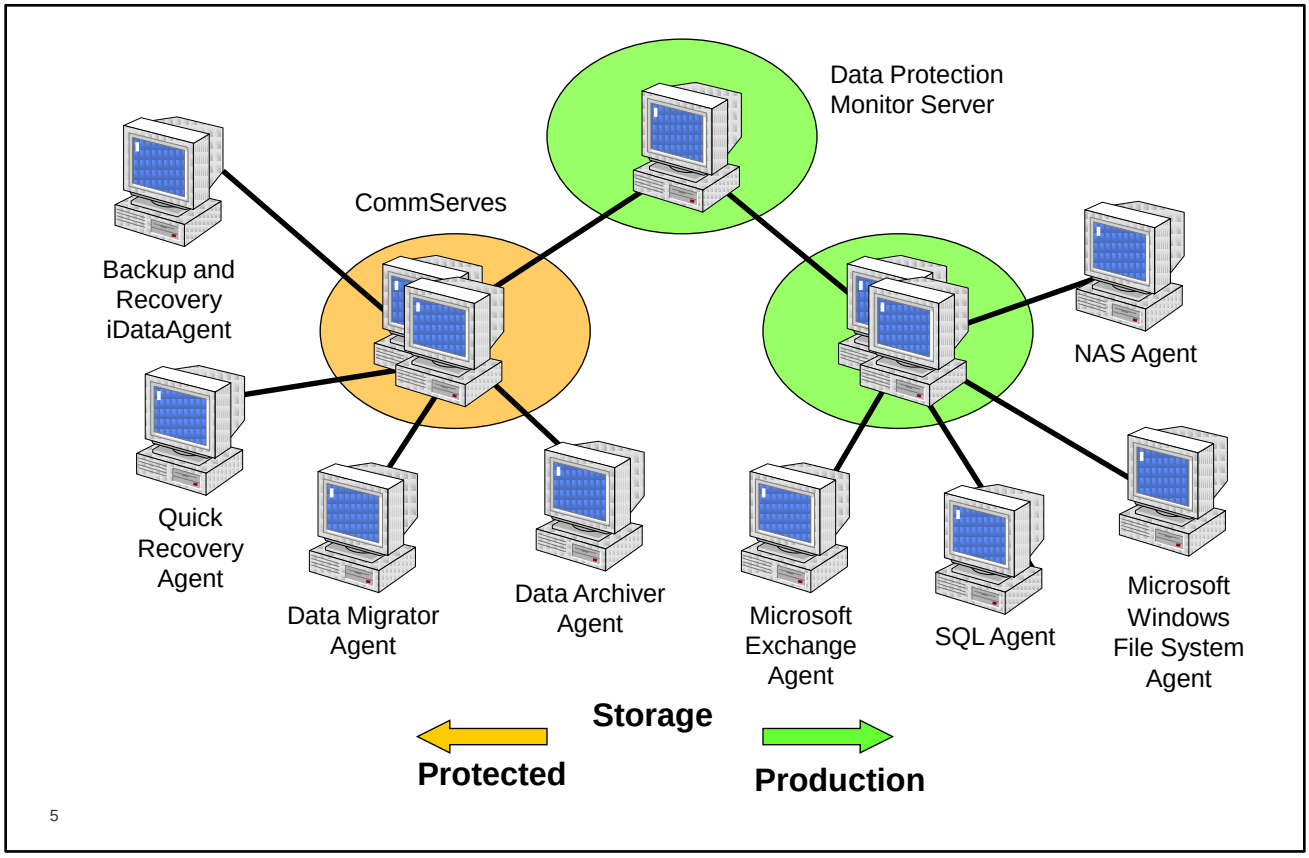
Continuous Data Replicator

Remote Office Data Protection

Continuous Data Replicator is a solution for remote office data protection over a WAN or LAN.

HDS Confidential: For distribution only to authorized parties.

Architecture



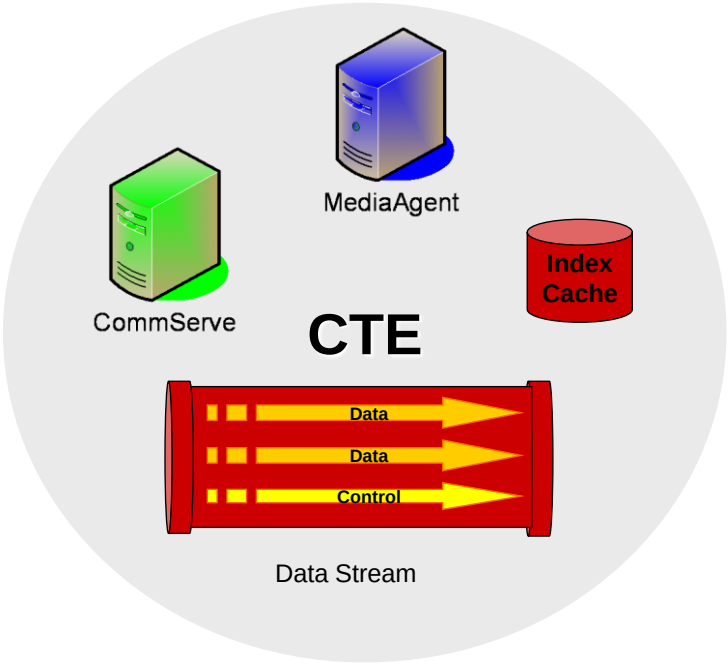
The Data Protection Suite architecture allows for flexible as well as scalable scenarios that can span enterprise data centers including all remote offices. Data Protection Monitor software aggregates and synthesizes information from multiple CommCells. It also supports traditional LAN environments, SAN, and NAS architectures. With its GUID (Globally Unique Identifier) capabilities, Data Protection Monitor software always provides a single instance view of your storage resources – assuring administrators of a single, accurate point of view.

From a logical perspective, the Data Protection Suite solution consists of a Data Protection Monitor server module and Data Protection Monitor agents for the CommCell. Data is collected from all CommCells as well as the resource levels within any registered CommCell.

From a functional perspective, the Data Protection Monitor software interacts with the CommCell to provide a unified, comprehensive view of all data movement and/or management aspects of the enterprise storage environment.

Common Technology Engine

- Core components for all Data Protection Suite modules:
  - CommServe
  - MediaAgent
  - Dual Level Index
  - Data Stream



The Common Technology Engine (CTE) is the starting point for understanding all activity of the CommCell. All other Data Protection Suite products interface with the CTE. It provides the necessary tools to manage and administer the client agents and manage the storage media associated with the CommCell.

It consists of four major components:

- The CommServe
- MediaAgent
- Dual Level Indexing
- Data Stream

A single **CommServe** provides administrative and management oversight of all tasks and data movement within a CommCell. The CommServe functions as the command and control center. All data protection and recovery activity is managed there. In addition, user security and all Data Protection Suite module licensing is handled by the CommServe.

The CommServe also houses the meta-database catalog. Loosely translated, the meta-database contains the "*data about the data*." That would include component configuration information, history, and location of all data backed up within the CommCell.

HDS Confidential: For distribution only to authorized parties.

Three major sub-components within the CommServe are the Job Manager, Scheduler, and the Event Manager. The Job Manager monitors and controls all the scheduled activity of the software and provides Backup and Recovery software with its restart capabilities. The Scheduler allows users to schedule tasks which are then initiated by the Job Manager. The Event Manager logs all events, providing unified notification of important events.

The **MediaAgent** is the workhorse for all data movement. It facilitates the transfer of data between the data source, the client computer, and the destination storage media. It is supported on Microsoft Windows, UNIX, Linux, and Netware systems. Each MediaAgent manages the reading/writing activity of all attached storage devices. The devices can be locally attached. For example, Small Computer System Interface (SCSI)), or remotely attached via Transmission Control Protocol/Internet Protocol (TCP/IP), or a Storage Area Network (SAN, either FC or Internet Small Computer System Interface (iSCSI).

Index Cache
The major function of the index cache is to provide quick access to indexed data on the media during a browse\restore operation. Whenever a new data protection job starts, the Data Protection Suite will generate a list of all the data objects (for example, files/subdirectories, database objects, mailbox objects) that will be backed up from the source server. This list, along with the path to the actual archive file that stores the data on the media, is called the ***index***.

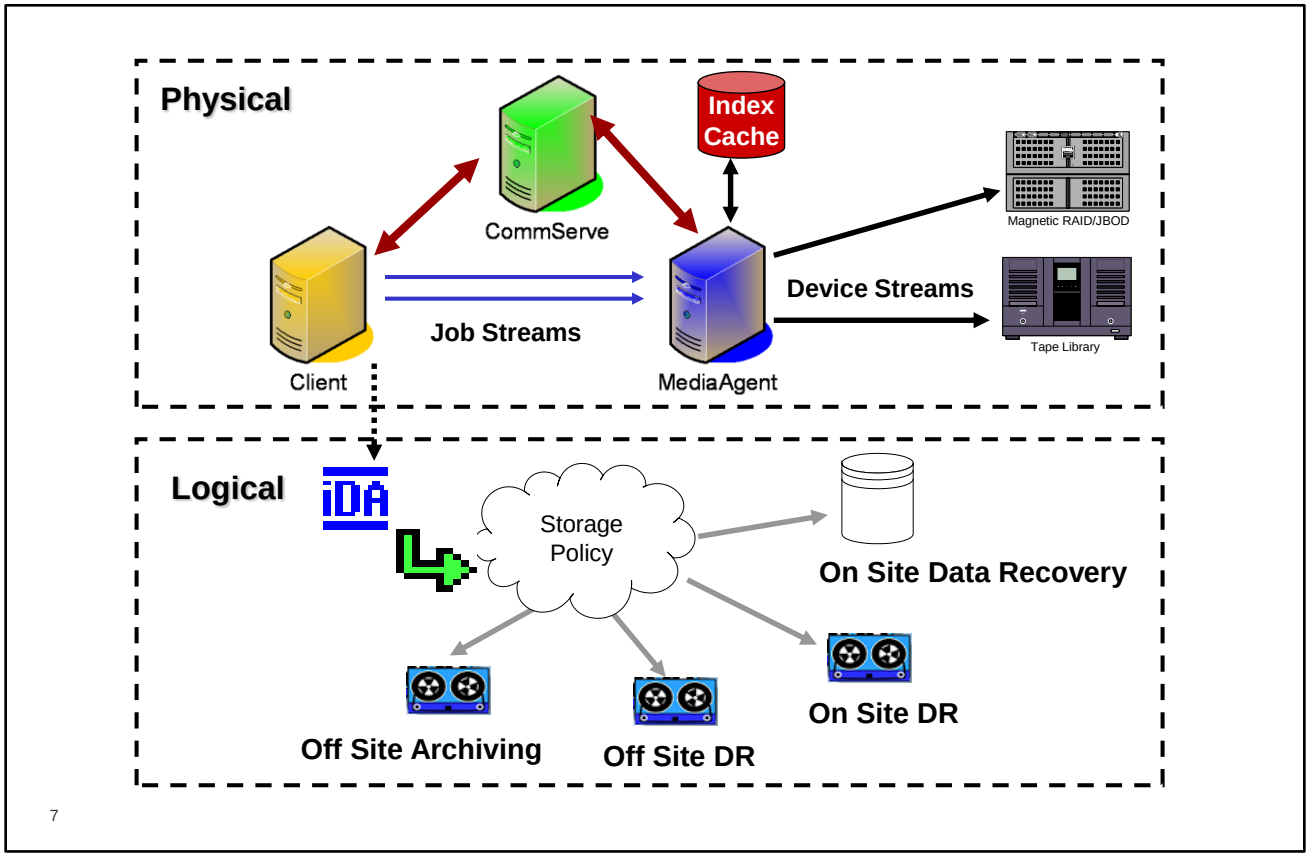
During a data protection operation, or *backup*, the source server sends data to the MediaAgent for archiving on the media. During a data recovery operation, or *restore*, the MediaAgent accesses data on the storage media and sends it to the destination server. The connection set up between the MediaAgent, source or destination server, is called a *Data Pipe* or what is more commonly referred to as a *Data Stream.*

A **Data Stream** is CommVault's patented high performance data mover used to move data back and forth between a data source and a MediaAgent or between two MediaAgents. It consists of a shared memory component on each system for multi-thread processing with multiple connections between the components. It handles all data encryption, compression, and disk I/O.

The data stream takes advantage of the shared memory speed and efficiency to double up on data movement functions. Each stream opens multiple connections between the data source and MediaAgent. One of these streams is used for control while the others move data. Whether performing local backups to a MediaAgent installed on the same system, or over the network to a MediaAgent installed on a different system, the data stream processes remain the same.

The CommServe can only be installed on a Microsoft Windows server platform. It may reside on its own dedicated system (recommended), or a system that also contains a MediaAgent and/or client modules. The system requirements are found in the Data Protection Suite Online Documentation or on the CommVault Systems Website.

Data Flow



To get a better understanding of how all of the CommCell components work in conjunction, let us examine the basic data flow in a data protection operation as it relates to the physical layout and logical layout of Backup and Recovery software.

When a backup operation is conducted, the CommServe will act as the central point of communication. This means that the CommServe must always be running and accessible via TCP/IP in order for jobs to commence. The CommServe will communicate with the Client and the MediaAgent to coordinate operations and ensure necessary resources are available.

Each client machine will have *i*DataAgents installed. These agents are logical components which interface between Backup and Recovery software and the application or system being backed up. This allows Backup and Recovery software to backup and restore data directly to and from an application in a format the application can interpret. At a minimum, a file system *i*DA specific to the operating system will be installed on the client.

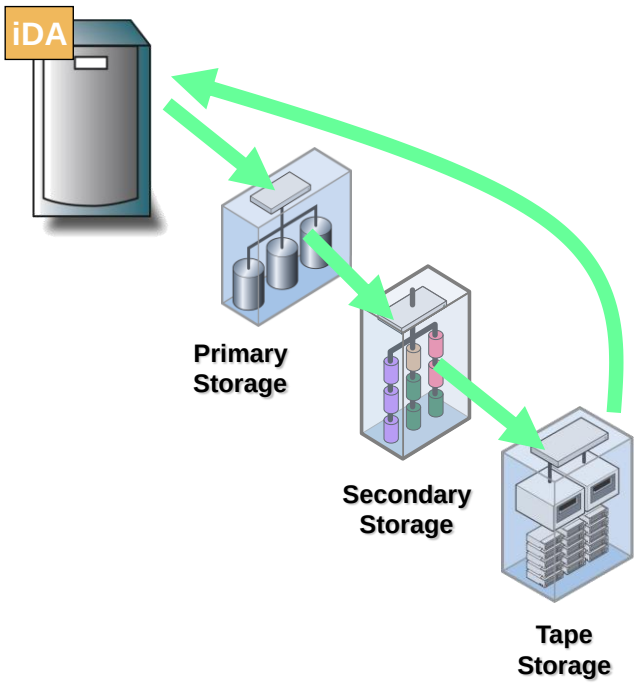
*i*DataAgents are logically divided into backup sets. A backup set is a complete collection of data being protected within the *i*DA. Each *i*DataAgent will have at least one backup set called the Default Backup Set.

HDS Confidential: For distribution only to authorized parties.

A backup set can be divided into logical containers called subclients. This is the smallest collection unit for data being protected. The use of subclients facilitates the ability to manage data uniquely within a backup environment. Each subclient will be associated with a storage policy. When creating subclients, data contained within the subclient will be exclusive to that subclient. This prevents data from being backed up more than once within a backup set.

Operational Concept – Data Migrator Software

- Tiered Storage
- Immediate reclamation of disk space
- Transparent recall of data



8

Introduction to Data Migrator software

Data Migrator software is a hierarchical data management solution that allows an organization to reclaim critical primary storage space by moving and automatically recalling data from secondary storage for active use. By applying policies to select, move and retain fixed content, organizations gain significant economic and productivity benefits. It allows companies to address aggressive data growth by redistributing older data across tiered-storage architecture. Data Migrator software's transparent in-place recall mechanism maximizes data accessibility for users, without imposing the need for extensive training, physical staging or catalog maintenance over the data's lifecycle. With this, administrators maintain more efficient primary storage architectures.

As part of the Data Protection Suite, Data Migrator software fully complements data protection strategies for Microsoft Internet Exchange, Microsoft Windows, Solaris and Network Appliance file system environments. The impact from migrating data is immediate: the substantial reduction in the production environment's physical storage space directly translates into reclaimed disk capacity and improvements in the speed and size of backup and recovery operations.

HDS Confidential: For distribution only to authorized parties.

Data Migrator Software for Microsoft Windows

Data Migrator software for Microsoft Windows is a policy based migration of files from Volumes, Directories or Shares to secondary storage locations. These policies contain rules definable on thresholds, age, size, filters, and or folders. This process is transparent to end-user and integrated through the file system. Migration also retains the ACL security of the files that are moved. A pointer or Stub replaces the actual file. To the end-user this stub looks like the original file and gives the user or application, the ability to recall the original file for use.

All content is containerized and optionally compressed and encrypted. Each migration job is equivalent to an Archive file that contains chunks of data. Each chunk is 2GB by default and each file is registered in the index by the offset within the chunk.

Data Migrator Software for Solaris/Linux

Data Migrator software for UNIX supports the ability to migrate and recover data from Network File System (NFS) mount points, through the use of CXFS (which is a special file system type for Data Migrator software on UNIX) and the "cxfstab" file. Once you have configured a subclient and run a migration, the directories and mount points specified in the subclient content become CXFS partitions or mount points. Like any other locally mounted file system, CXFS must be explicitly exported in order to become available to remote NFS clients.

Data Migrator Software for Netware

Data Migrator software for Netware also removes data from production storage to secondary storage leaving stubs for the files. The option for a common internet file system (CIFS) share or Netware Client is supported.

Data Migrator Software for Network Appliance

This agent unlocks a content area that has been traditionally blocked as a content source for Data Migrator software or Hitachi Systems Manager. Data files stored on a NAS device in Microsoft Windows/CIFS shares are eligible for Migration. A driver is distributed to all Microsoft Windows 2000/2003/XP clients that may mount the CIFS shares and initiate automatic recalls of migrated files. This driver loads as a local Microsoft Windows service and monitors all file I/O from that client machine. If a stub is identified within the read/write request, the driver will read the recovery data key wrapped within a digital signature in the stub file and contact the migration server to restore the item to the stub location.

Due to the restrictive, proprietary nature of NAS device lacking Hitachi Systems Manager application programming interfaces (APIs), the iDataAgent only supports Microsoft Windows data on the Filer. Support for NFS/Unix data will be added in future releases and incorporate the evolving NetApp Hitachi Systems Manager API.

Data Migrator-Network Storage (DM-NS) supports Microsoft Windows data residing on NetApp, EMC or Blue Arc NAS devices. Our framework is based on

Note data movement using the CIFS file access protocol; which means it is not bound to specific restrictions on NAS/OS.

: If snapshots are managed locally on the Filer, a migration job will not produce immediate free space on the NAS server. The capacity freed by the migrated space will not be released to the disk scratch pool until the snapshots expire.

Data Migrator Software for Microsoft Internet Exchange

The size of Microsoft Internet Exchange Server Information Store will grow over time. To keep its size under control, administrators and users have to either delete or migrate unwanted messages periodically to release mailbox disk space. However, the migrated messages usually reside on users' workstations as .PST files, which is difficult to manage. So it is necessary to have centralized migration management to keep the size of mailboxes within a healthy range, as well as to provide easy access to the migrated messages.

The Data Migrator software for Microsoft Internet Exchange Agent is designed for the above purpose on Microsoft Exchange Servers. It will automatically migrate messages satisfying certain criteria, and replace them with stubs containing information for recovery. Users can double-click the message stub in Outlook or Outlook Web Access (OWA) to recover the original message. Alternatively, users can also ask the administrator to browse the CommCell Console to recover the message. The latter method is especially useful if the stub is deleted.

Data Migrator software for Microsoft Exchange honors Single Instance Storage by capturing all of the properties of the message and comparing each to the Microsoft Exchange Data Store.

So when a Data Migrator software job runs we group the entire data together and single instance both the message data and attachment data, which is then written to the archive file. This also includes any PST data migrated during that run with Data Migrator software.

While this preserves the single instance state of that data related to Microsoft Internet Exchange, it will not single instance across jobs or servers (as in the case of 1 PowerPoint file found in two exchange servers).

Data Migrator Software Web Proxy Agent for Microsoft Internet Exchange

The Data Migrator software Web Proxy Agent provides Outlook users with the option to route stub recovery through the Internet Information Service (IIS) server over hyper-text transfer protocol (HTTPS). Normally, two direct channels of communication are opened from the Outlook Add-in client, one to the Microsoft Exchange Server and one to the CommServe, over ports 8401 and 8402. When using the Data Migrator software Web Proxy Agent, a single channel of communication is opened through the IIS server. Stub recovery is performed through a single, encrypted and authenticated Secure Sockets Layer (SSL) port.

This solution takes advantage of the remote procedure call (RPC) over HTTP facilities of Microsoft Windows Server 2003; therefore, this solution only applies to Microsoft Internet Exchange 2003 accounts on a Microsoft Internet Exchange 2003 Server, running on Microsoft Windows Server 2003.

Data Archiving

- Compliance
- Complete collection of email
- Ensure accessibility
- Retrieve data for external audits
- Deep search and retrieval with audit controls
- Full-text and sampling

Archiving	Data Protection
Long term data retention	Short and medium term data retention
Removes data from source	Leaves data on source
Archiving is optimized for fixed content, specifically content that does not change over time such as an email message.	Data Protection designed to handle dynamic data depending on the Agent (multiple file versions for example)
Designed to retrieve data outside of the operational environment. Data is retrieved out-of-place.	Designed to recover data back into a production environment. Data can be restored in or out-of-place.

9

Introduction

The Data Archiver software agent is designed for long term storage and indexing of data and to meet security and compliance standards. The primary function of Data Archiver software is to preserve data outside of the operational environment. Data Archiver software removes the data from the source client once it has been archived and/or indexed. In this way, large amounts of data can be stored, and reviewed at a later time. The table contrasts archiving with data protection.

Data Archiver software for Microsoft Exchange

Data Archiver software for Microsoft Exchange works in conjunction with the message journaling feature of Microsoft Exchange Server software to archive all incoming and outgoing messages and attachments. This is done at the Microsoft Exchange Message Store level. All incoming messages and outgoing messages are captured in the Microsoft Exchange Journal Mailbox, which is then archived with the Data Archiver software agent according the schedules you set. Incoming messages are written to the journal as they come into the Microsoft Exchange Server. Therefore, all messages are recorded, regardless of whether the message recipient deletes the message in his or her mailbox.

HDS Confidential: For distribution only to authorized parties.

The archived data remains available for retrieval for the length of time set by the retention rules. During this period the data is readily available for access and retrieval by authorized users. They can retrieve duplicate copies from the archive, which remains under the management of the Storage Policy. All archive data is protected from any edit or alteration from all users within the context of the archival system (i.e., the archive data on the storage media cannot be written to or altered).

Module Review

1. The Data Protection Suite was designed to automate the management of data by combining traditionally separate products for data movement and data management.

   What are two major features of Hitachi Data Protection Suite? (Choose two)

   a. Fixed content performance reporting and tuning
   b. Backup and recovery software
   c. Data archiving
   d. Storage capacity configurator

10

HDS Confidential: For distribution only to authorized parties.

20. Hitachi Content Platform

Module Objectives

- Upon completion of this module, the learner should be able to:
  - Define active archive
  - Articulate the benefits of active archive
  - Describe the architecture of Hitachi Content Platform

2

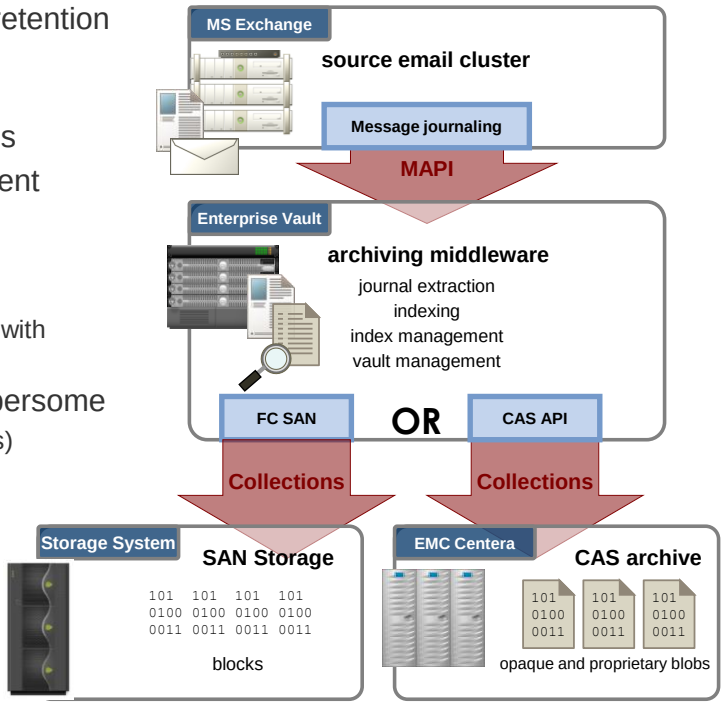
Email Archiving Challenges

Problem definition:

- SEC 17a-4 driving change in data retention practices (WORM)
- SOX driving need for audit and transparency of business operations
- Quantity of emails overloading current systems
  - 500K – 1M emails per day
  - 70% - 100% growth per year
  - Ingestion performance cannot keep up with volume
- Data discovery and access is cumbersome
  - Distributed environments (multiple sites)
  - Number of searchers are increasing

Previous solutions:

- Large enterprises: Centera/KVS
- SMB: NAS or modular storage



3

The email application (Microsoft Exchange) has two possible storage targets for storing email and attachments, SAN storage and archive. SAN storage is typically viewed as primary array-based storage and is typically "block-based". The archive (Content Archive Platform or Centera) is object based and has features to control the retention and authentication of the data and objects.

The middleware application (Enterprise Vault) makes a decision based on customer supplied criteria (age, size, and custom) on which target to use for storage of that data, either SAN Storage, or Archive. If it chooses the Archive, it can further specify retention time, and access control.

MAPI = Messaging Application Programming Interface

CAS = Content Addressed Storage

HDS Confidential: For distribution only to authorized parties.

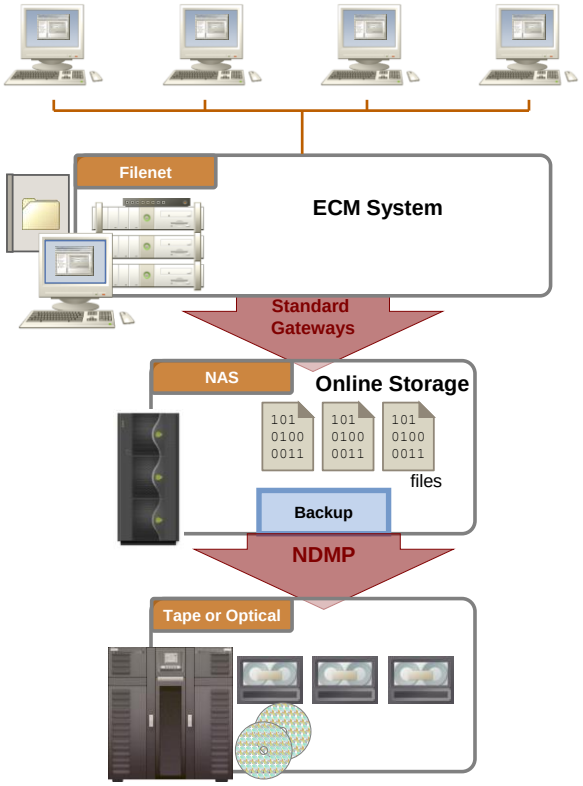
Document Archiving Storage Challenges

Problem definition:

- No assurance that information is secure, unchanged and protected.
- Need for security to restrict unauthorized users from accessing critical information.
- Inability to scale and manage massive volumes of regulated content.
- Inability to quickly find and retrieve information that has been archive.

Previous solutions:

- NAS, Automated Tape



4

What an Active Archive Solution Must Deliver

- Fixed content retention
  - Policies for preserving content for long periods
  - May require Write Once - Read Many (WORM)
  - Authentication
- Scalability
  - Capacity — hold thousands of terabytes in a single repository
  - Volume of files — ingest a growing volume of new files such as email
- Reliability
  - Be able to withstand simultaneous points of failure
  - Withstand site disasters
- Accessibility
  - Support and provision files to multiple applications from the same archive
  - Search and access files using standard and open methods



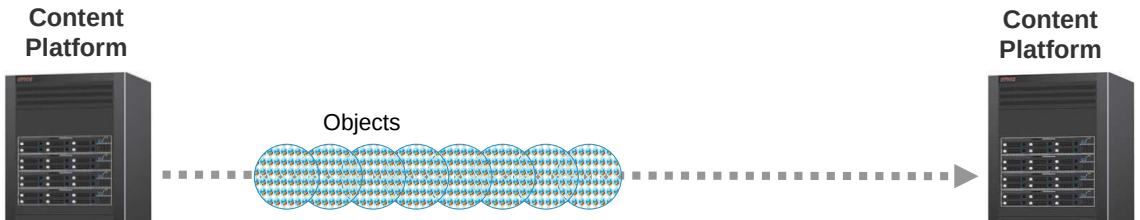
Fixed Content File Types

5

Customers demand a product that assures the retention of authentic, fixed content in an immutable form that provides the scalability needed to address ever-increasing volumes of new content and the associated growth in storage capacity. It must provide the reliability required to meet customer DR/BC policies as well as SLAs needed to ensure content is accessible when needed. The Hitachi Data Systems product delivers on all of these with the most robust platform for fixed content archiving.

HDS Confidential: For distribution only to authorized parties.

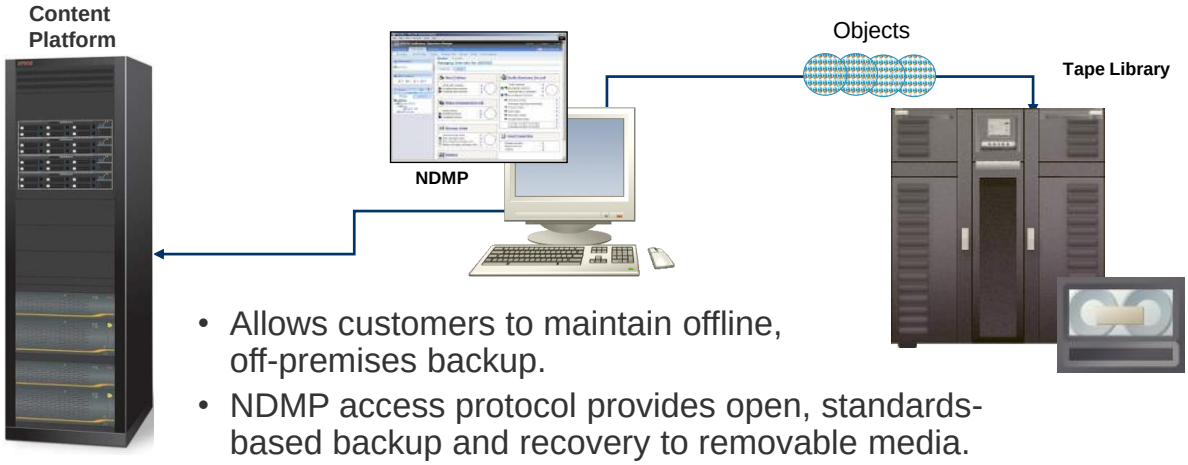
Archive Object Replication



- Cluster-to-cluster replication supported today with client tool and professional service delivery
- Supported use cases:
  - Primary archive to a stand-by replica (unidirectional) for traditional off-site disaster recovery
  - Remote office to central archive configuration
- Future release will include advanced replication (One to Many; Many to One; Bi-directional; Ring topology)

6

Archive Object Backup



- Allows customers to maintain offline, off-premises backup.
- NDMP access protocol provides open, standards-based backup and recovery to removable media.
- Freedom of choice among leading backup tools such as Data Protection Suite, CommVault Galaxy, NetBackup, and many others.
- Supports full, incremental and differential backups.
- Backup sets can be signed, encrypted and/or compressed and are backed up as complete objects.
- Recovery is vendor and platform independent.

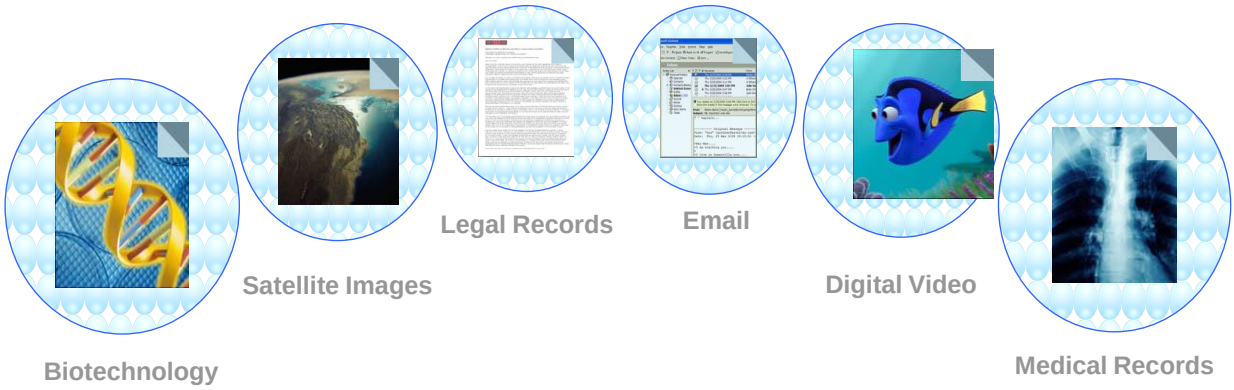
7

Diligent ProtecTier is an example of a virtual tape library that can be used to connect to the Content Archiving Platform. It is a software + storage implemented solution running on a Linux platform and uses advanced data similarity analysis for data reduction.

HDS Confidential: For distribution only to authorized parties.

Fixed Content

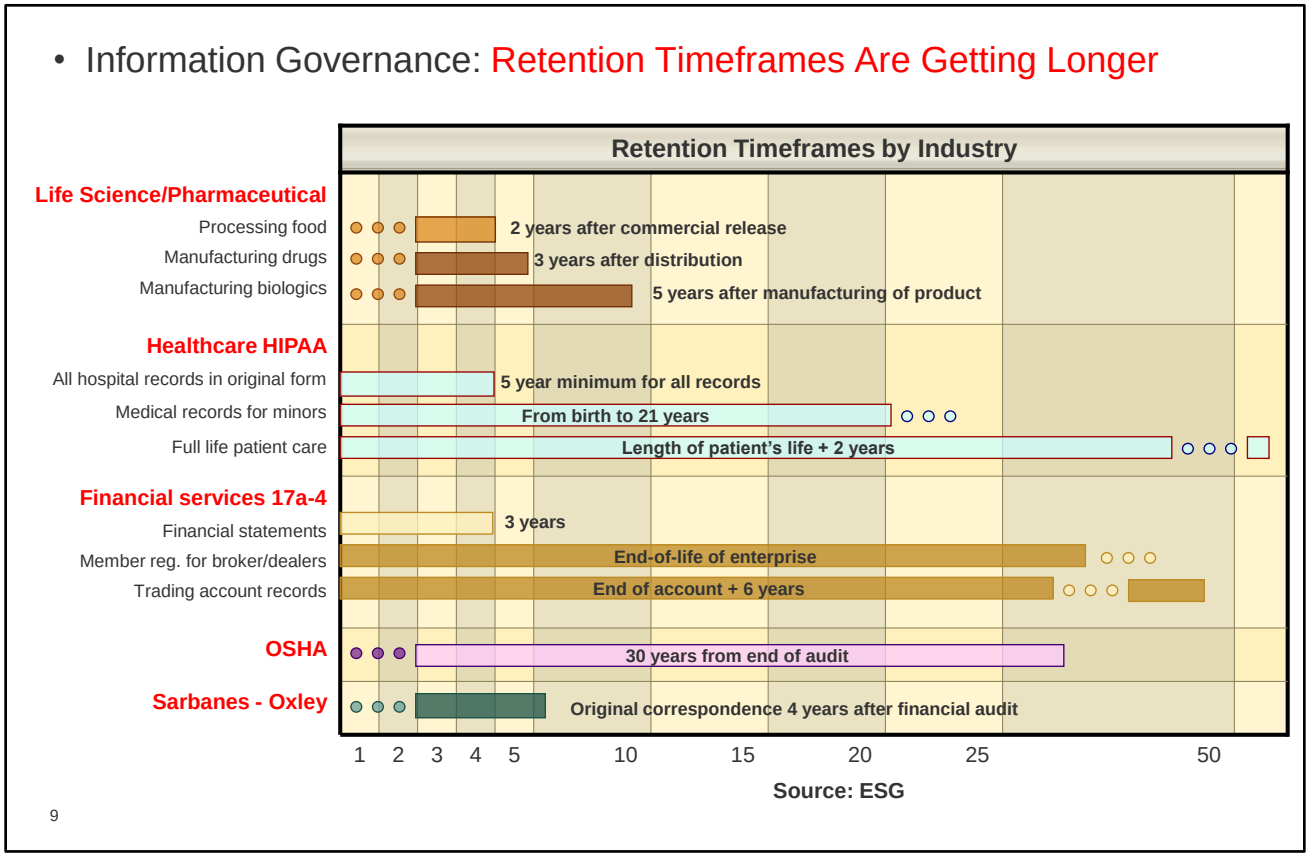
- What is **Fixed Content?**
  - Data objects that have a **long-term value**, **do not change over time**, and are easily accessible and secure



8

Organizations across all industries need to address the management of business-critical information assets over time and have them accessible in the future. Such types of content includes documents, images, graphics, technical data, and video where more and more of it is being created digitally or converted from physical form into digital form.

Retention Times



While government regulations have a significant impact on content archiving and preservation for prescribed periods, compliance does not necessarily require immutable or Write Once, Read Many (WORM)-like media. In many cases, the need for corporate governance of business operations and the information generated are related to the need to retain authentic records. This requirement ensures adherence to corporate records management policies as well as the transparency of business activities to regulatory bodies. As this chart illustrates, the retention periods for records are significant, from two years to near indefinite.

HIPPA – (**Health Insurance Portability and Accountability Act (HIPAA) of 1996)** HIPAA has enacted several mandates to improve the access and portability of patient health records while maintaining strict privacy and security. A critical aspect of the HIPAA privacy ruling is Data Protection, requiring compliant backup methodologies to ensure the security and confidentiality of patient records. Health care providers who engage in electronic transactions must observe privacy safeguards to restrict the use and disclosure of individually identifiable health information.

OSHA – (**Occupational Safety and Health Administration)**

A US government agency in the Department of Labor to maintain a safe and healthy work environment .

Sarbanes–Oxley

The Sarbanes-Oxley (SOX) Act, passed in the year 2002, outlines the procedures of storing financial records. All companies and business organizations must comply with the SOX procedures of financial records storage, ensuring that there are no accounting error related scandals or illegal financial activities. The Sarbanes-Oxley Act legislates the time period, during which the financial records of the company must be maintained, along with the manner in which they should be done so.

All Roads Lead To Hitachi Content Platform



- **Legal / Compliance**
  - External legal mandate to archive
  - Strict policies for retention and deletion
  - Examples: HIPAA, OSHA, Sarbanes-Oxley
- **Long Term Preservation**
  - Internal need to preserve data for later retrieval
  - Examples: NASA, Library of Congress, Oil and Gas
- **Offloading Primary Storage (Tiering)**
  - Utilize Archive to deduplicate files from all over the network and store a **gold** copy in the platform for access and control.
  - This creates space on their other tiers of storage for use and reduces their backup windows.

10

HDS Confidential: For distribution only to authorized parties.

Product Lines



Hitachi Content Platform 300

- Fully integrated archive appliance with **RAIN** architecture
- Integrated software and storage on each node
- RAID-5 Protection with DPL2
- Available in 2 – 42.5 TB configurations

OR



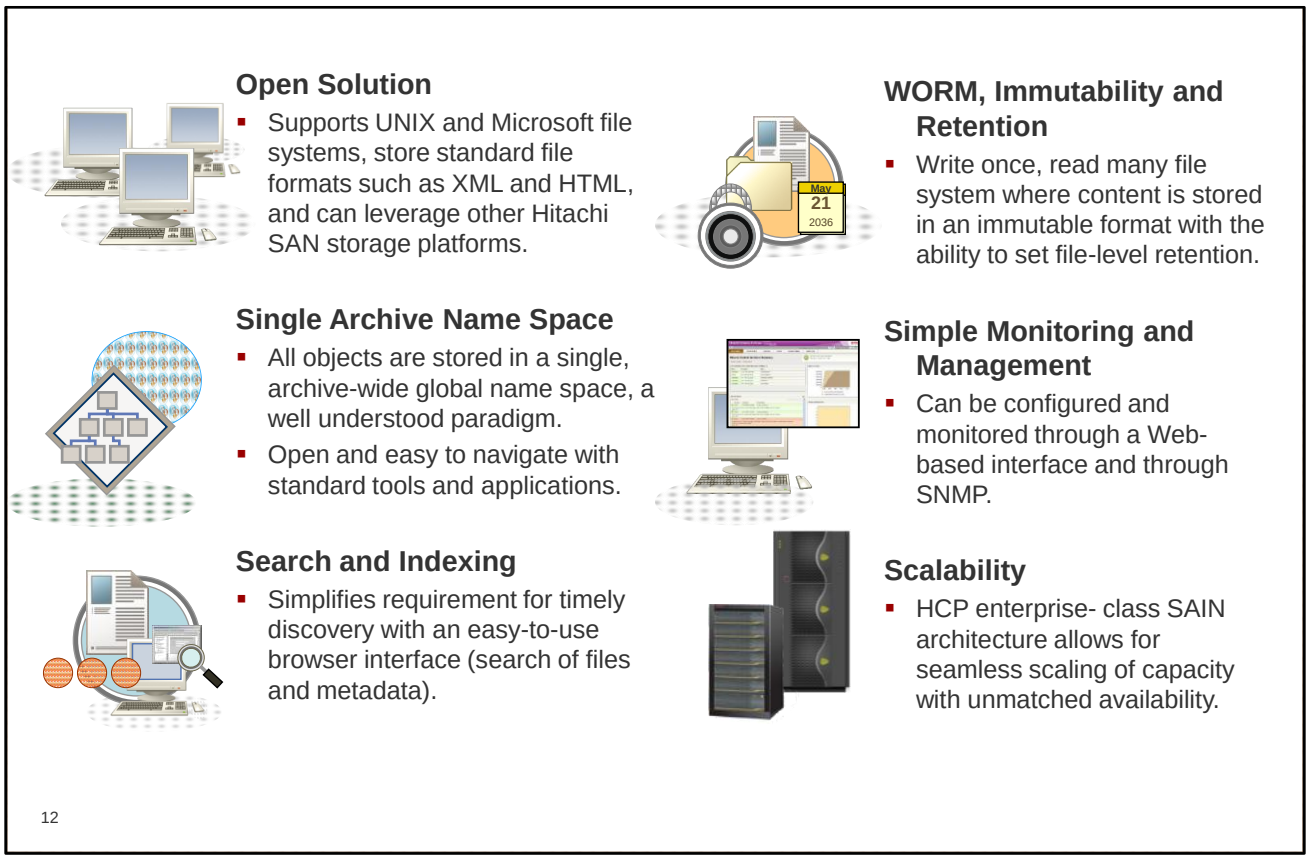
Hitachi Content Platform 500

- **SAIN** configuration – SAN-attached Array of Independent Nodes
- Leverages world-class Hitachi storage platforms to provide scalability high availability and performance
- RAID-6 Protection
- Provides the highest-level of security
- Available in 9 TB – 40 PB configurations

11

RAIN – Redundant Array of Independent Nodes

Customer View of the Archive



Although the platform is complex, its presentation is wonderfully simple. It presents itself as a big directory, just like what you would see on your PC or MAC. It operates in multiple protocols, such as NFS for UNIX/Linux users, CIFS for Microsoft Windows, and it has WebDAV and HTTP interfaces as well. There are no barriers for simple access. The Hitachi Content Platform software automatically provides for WORM storage, and lets the users assign retention in various ways: across the archive, across a directory, or by each object. This guarantees that the file and record will not change, and is immutable, which can be defended in a court of law. Simple user interfaces make the cluster easy to manage.

HDS Confidential: For distribution only to authorized parties.

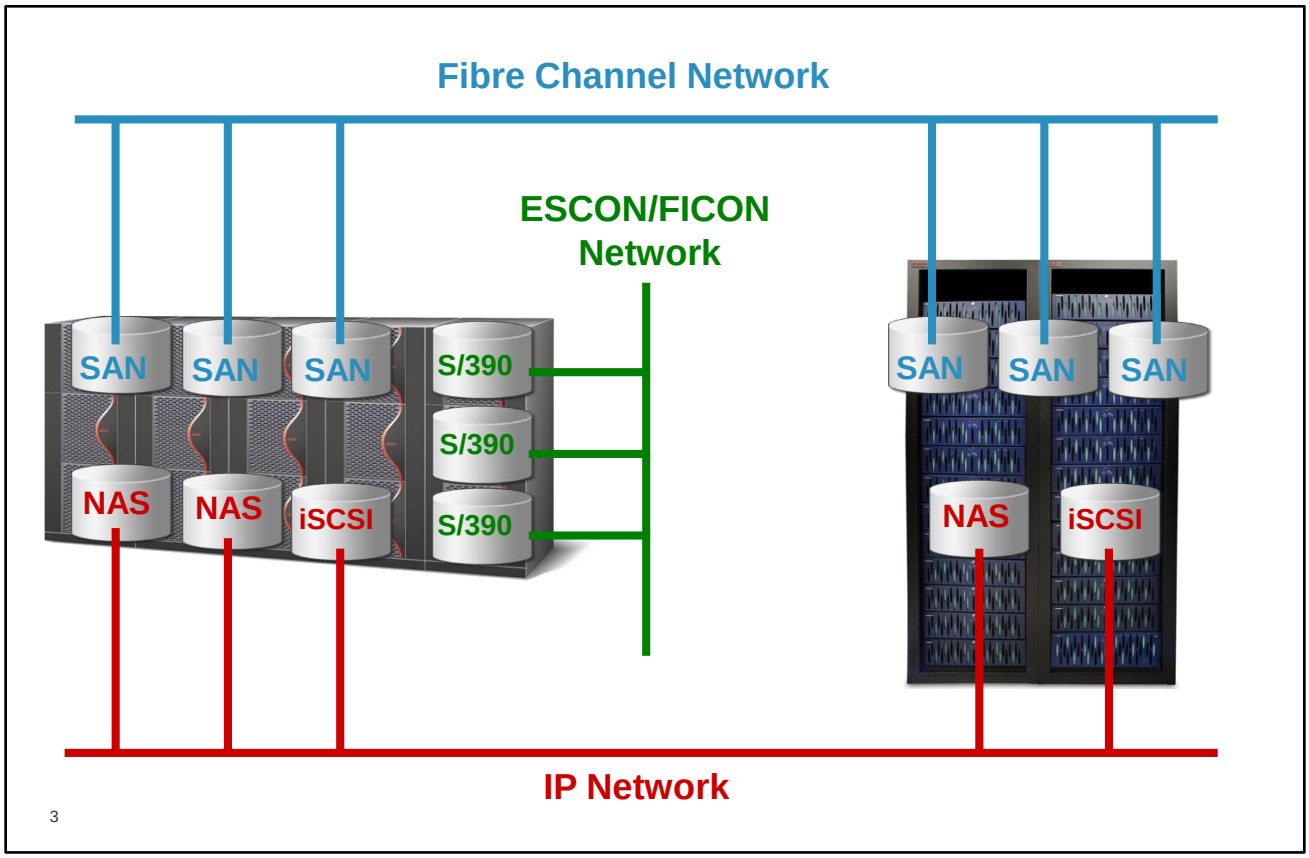
21. Hitachi Essential NAS Platform Overview

Module Objectives

- Upon completion of this module, the learner should be able to:
  - State the purpose and position of the Essential NAS Platform
  - Recognize the history and idea behind NAS technologies
  - Identify the hardware components
  - Identify the software components
  - List the data protection features

2

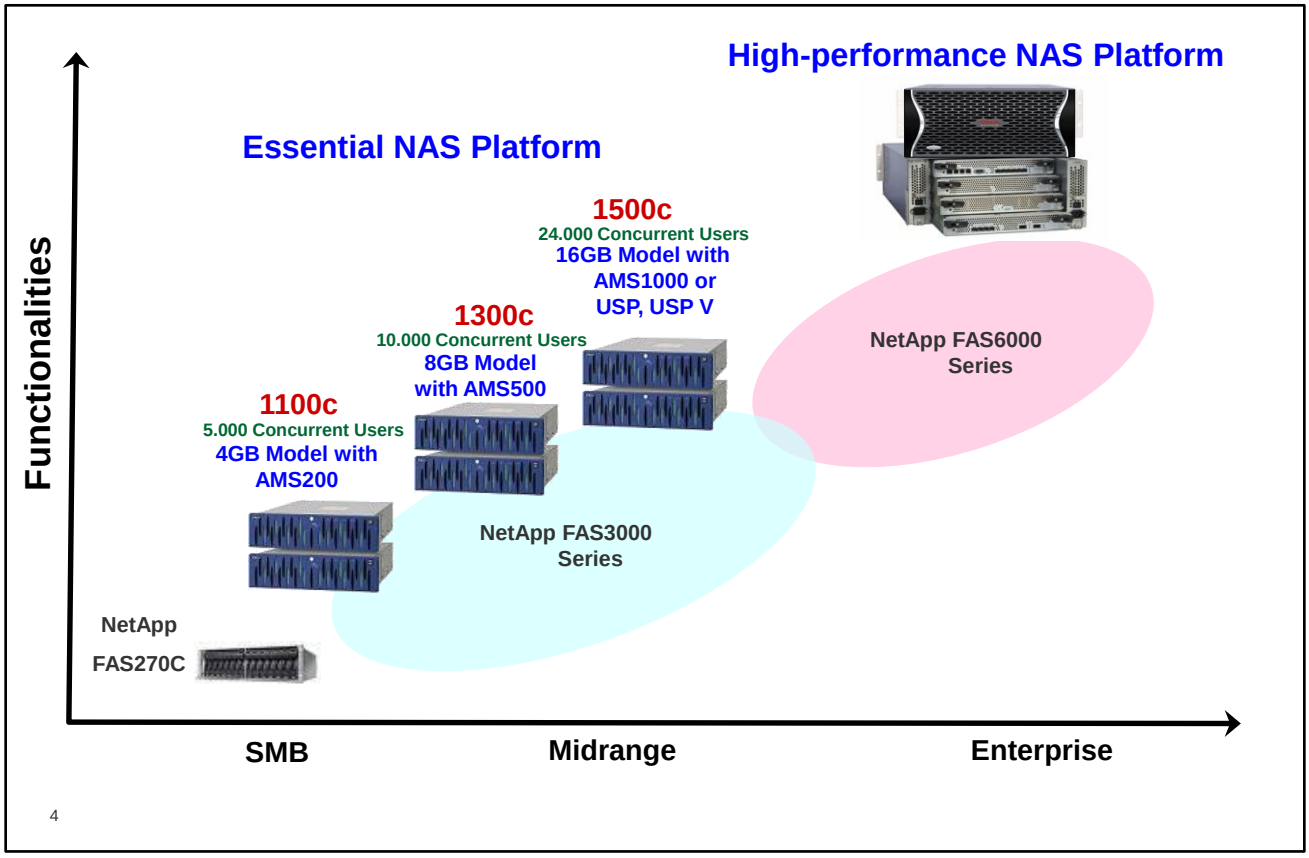
Hitachi Storage System Connectivity History



Storage Area Network (SAN)

HDS Confidential: For distribution only to authorized parties.

Essential NAS Platform Positioning



AMS stands for Hitachi Adaptable Modular Storage

USP stands for Universal Storage Platform

Product Overview – Speeds and Feeds

Item		Description
Processor		Dual CPU, Woodcrest (2.33GHz, Dual Core)
Memory		4GB, 8GB, and 16GB depending on model
Ethernet Ports		Default: 1000Base-T Optional (4ports, Including Management port): 2 Slots optional: 1000Base-T(8ports) and/or 1000Base-SX (8ports)
Storage Ports		Default: Primary FC Dual port HBA (1, 2 and 4Gbps autosensing) Optional: Second FC Dual port HBA (1, 2 and 4Gbps autosensing)
Chassis		3U 19” rack mountable 430mm x 656mm x 3U
Power supply	AC input	Single phase AC100~120V Or AC200~240V 50/60±1Hz
	AC-DC power supply	640W Basic: AC-DC power supply Optional: AC-DC power supply (Backup)
	Temporary blackout	Endure until 20ms (in case of dual power supply)
	Battery	None

5

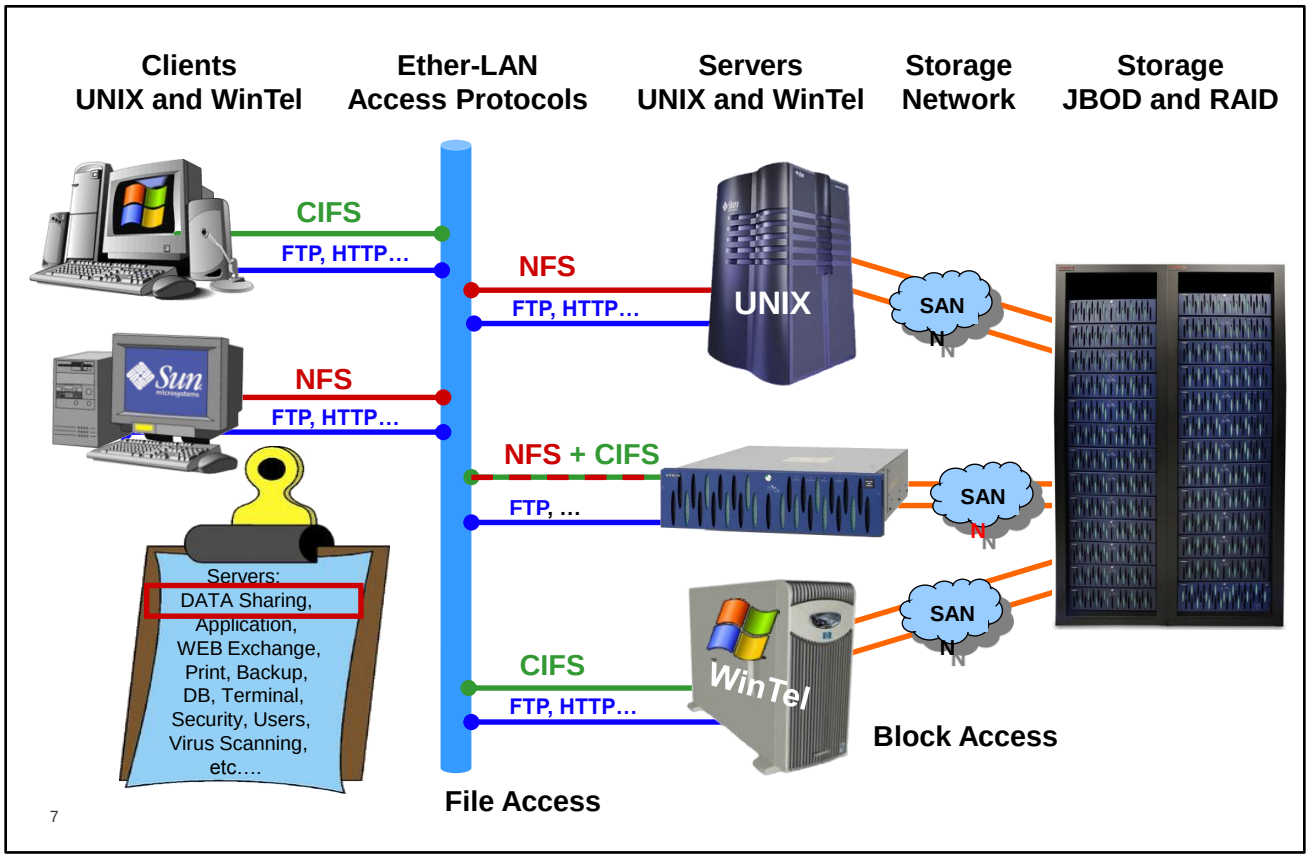
HDS Confidential: For distribution only to authorized parties.

Hardware Comparison

Item		Current NNC	Essential NAS
Processor	Type	**Single Core** CPU (2.8GHz)	**Dual Core** CPU (2.33GHz)
	Quantum	**1 / node**	**2 / node**
Memory	Amount	**4GB: H, M** **2GB: S, XS**	**4GB, 8GB, 16GB: H, M, S, NSC55, USP and newer products**
	Unit of replacement	2GB	2GB
Extension slot	For data ports	**None**	**2**
	For storage ports	**None**	**2**
Data ports	Type	**1000Base-T**	**1000Base-T** **1000Base-SX**
	Quantum	**4**	**Max. 20** Basic: 4 (1000Base-T) Optional: 16 (8 ports 1000Base-T or 8 ports 1000Base-SX)
Storage ports	Type	**PCIe (10Gbps)** **FC (4Gbps)**	**FC HBA (4Gbps)**
	Quantum	**PCIe: 1** **FC: 1**	**Standard 2 ports, with optional HBA 4 ports**
Management port	Type	1000Base-T	1000Base-T
	Quantum	1	1
Maintenance port	Type	100Base-T	100Base-T
	Quantum	3	2
Serial port	-	None	1 (Only for Diagnostics! – Should stay covered!)
Chassis	Form	**1U** 19" rack	**3U** 19" rack
	Dimension (WxDxH)	483mm x 653mm x **43mm**	430mm x 656mm x **127mm**
Power supply	AC input	Single phase AC100~120V Or AC200~240V 50/60±1Hz	Single phase AC100~120V Or AC200~240V 50/60±1Hz
	AC-DC power supply	640W **Single**	640W (Peak value) **Standard: Single** **With option: Dual**
	Temporary blackout	Endure until 15ms	Endure until 20ms (in case of dual power supply)
	Battery	None	None

6

Essential NAS Server



HDS Confidential: For distribution only to authorized parties.

Platform Introduction

- Two form factors with identical features and functionalities:
  - NAS Gateway attached via FC to Adaptable Modular Storage and Universal Storage Platform Family.
  - NAS Filer bundled and pre-installed with Adaptable Modular Storage, takes about 15 minutes to install and configure with the NAS Wizard.



- CIFS Capability:
  - Full NTFS ACL enables Microsoft Windows customers to have granular security configuration.
  - Up to 24K concurrent sessions.
  - Expands number of CIFS shares up to 7,500 for all models.
- IP and Block Replication (CLI with first release)
  - Provides IP based remote copy utilizing the technology of Sync Image differential snapshots.
  - Provides sync replication utilizing the TrueCopy Replication software technology.
- User friendly and intuitive NAS Manager GUI.
- Allows upgrade from NAS Blade and/or Adaptable Modular Storage system and Workgroup Modular Storage system with NAS Option.

8

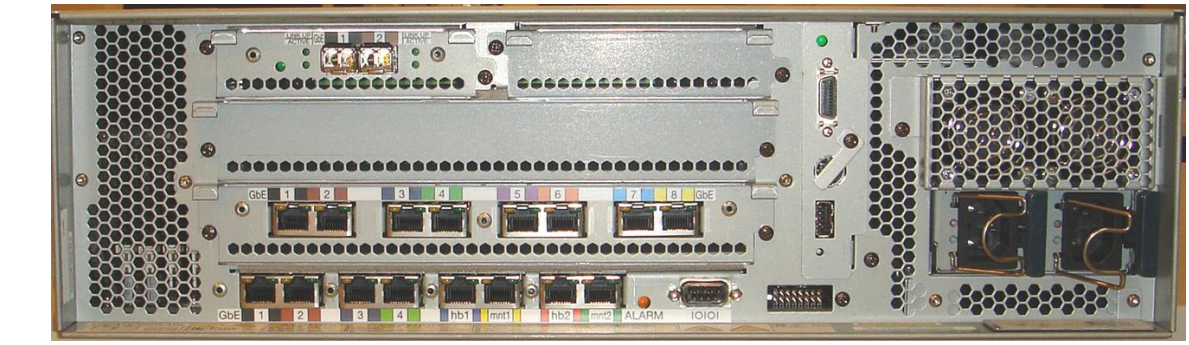
Data Protection Features

- Essential NAS Platform supports the following backup functions:
  - Backup/Restore using Network Data Management Protocol
    - File System Backup (through LAN or SAN)
  - Snapshot Sync Image
    - Multigenerations of Snapshot
    - One “Differential Volume” for all changes (“copy on write”)
  - Replication Utility for Sync Image (RUS)
    - Replication over IP using Sync Image technology
  - Antivirus Agent
    - A proprietary agent that provides virus check function through external virus scan server

9

HDS Confidential: For distribution only to authorized parties.

Rear and Front Views

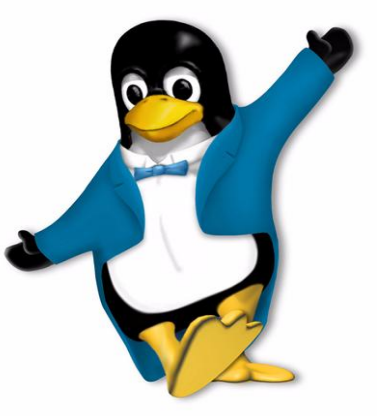




10

Software Components

- The NAS OS is equipped with the following NAS features and functions required for NAS Management operation:
  - CIFS Server
  - NFS Server
  - FTP Server
  - Hitachi Tuned xFS
  - LVM
  - Failover
  - Installer
  - Web server
  - SSH Server
  - NTP client
  - SNMP agent
  - DNS client
  - NIS client
  - and more…



11

Optional:

Hitachi Network Attached Storage Sync Image Differential Snapshot: (Base Suite)

The NAS Sync Image program provides functionality needed to create snapshots, which can be used to restore data in a file system shared in a NAS system. The data can be restored to a previous state, even after a client has mistakenly modified the data.

A *snapshot* is a duplicate copy of the data in a file system at a particular time. If the system administrator creates a file share for the created snapshot, the client can view the data at the time the snapshot was created, and thus can restore data from the snapshot even if the client has changed the data by mistake.

Hitachi Network Attached Storage Backup Restore: (Base Suite) –

- A proprietary product that provides and manages backup/restore functions.
- Provides browser-based GUI that a NAS administrator accesses via HTTPS.
- NDMP server collaborates with other vendor's NDMP backup software.

Hitachi Network Attached Storage Anti Virus Agent: (Base Suite)

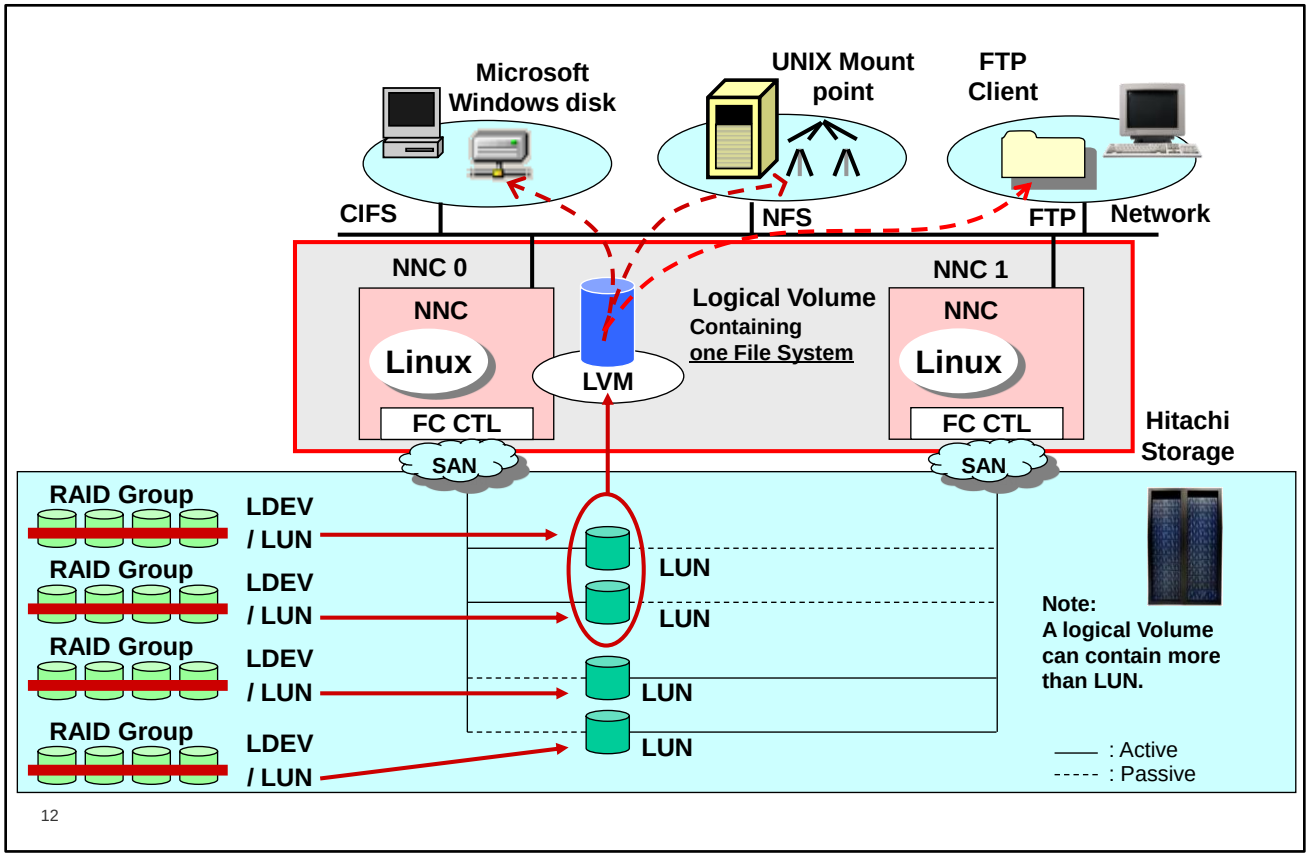
HDS Confidential: For distribution only to authorized parties.

A proprietary agent that provides virus check function via collaboration with external virus scan server. Customer purchased Scan Engine is sold and supported by Symantec, not HDS. The initial release will support Symantec's enterprise-class Anti-Virus software solution.

NAS Replication Utility for Sync Image (RUS): (Option)

A program for remote copy/remote backup snapshots from one site to other site. RUS is based on the technology of Sync Image Differential Snapshot.

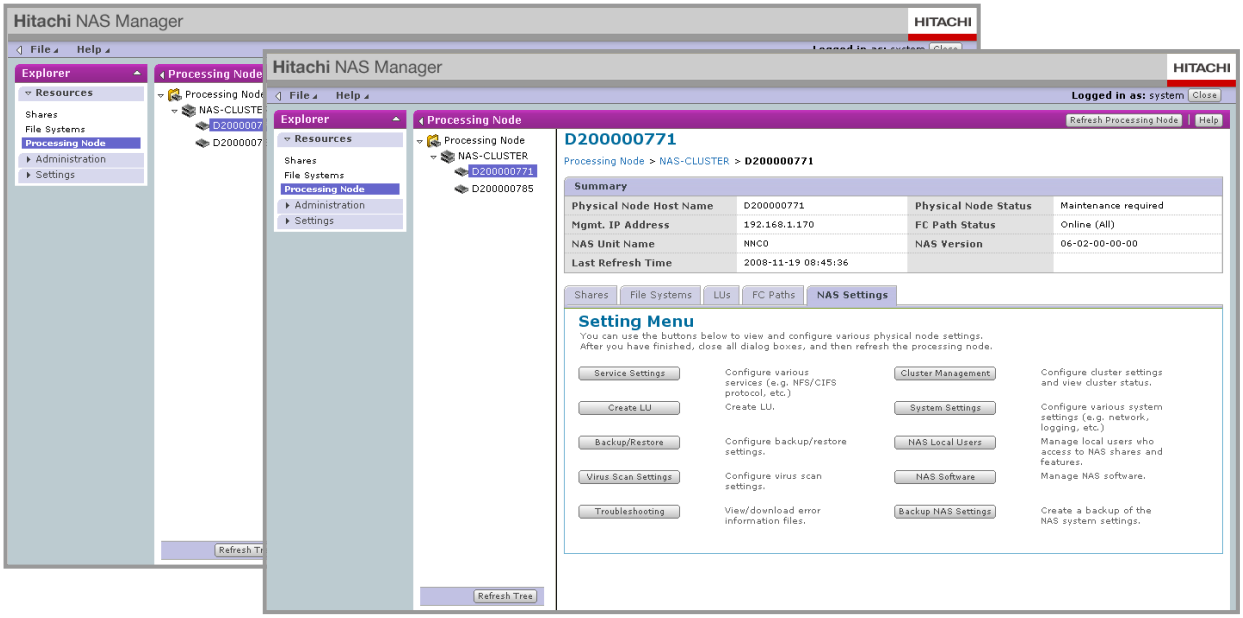
Data Storage Overview



HDS Confidential: For distribution only to authorized parties.

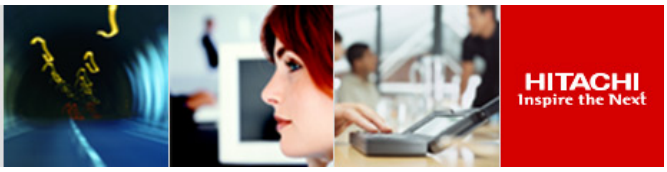
Management

- NAS Manager GUI
  - User friendly and intuitive NAS Manager GUI which satisfies both inexperienced users and savvy users



13

HDS Confidential: For distribution only to authorized parties.



Training Course 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 — **G**roup **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 ***I**nternet **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

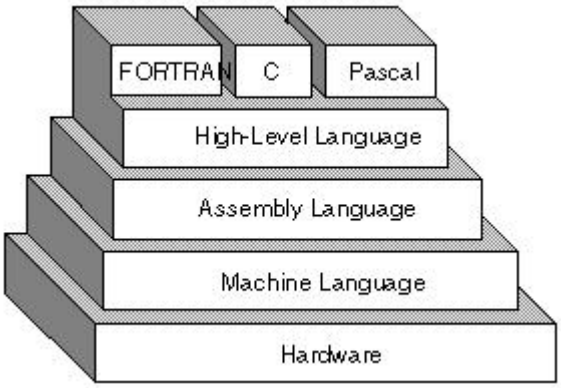
MCU — Main Disk Control Unit; the local CU of a remote copy pair.

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

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—

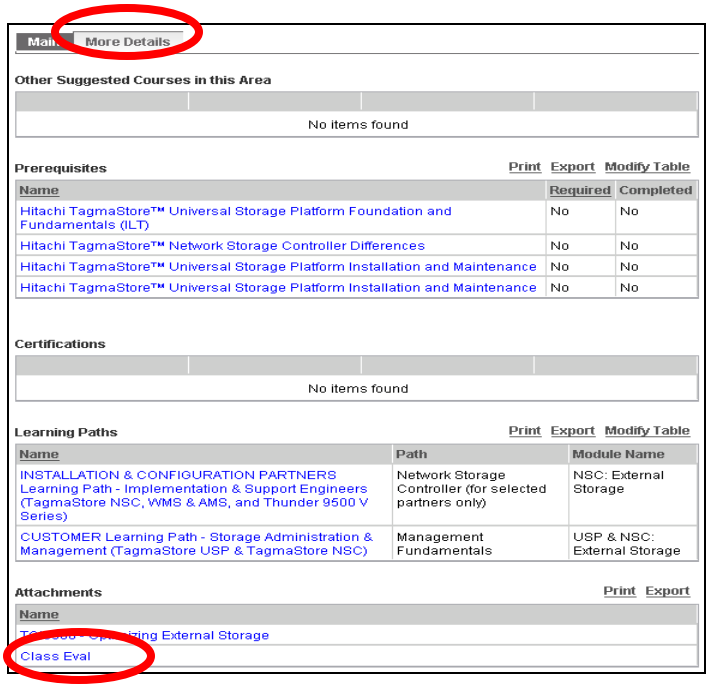
Zone — A collection of Fibre Channel Ports that are permitted to communicate with each other via the fabric

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.



6. Under Attachments, click the **Class Eval** link. The **Class Evaluation** form opens.

Complete the form and submit.