



| Hitachi Data Systems Enterprise Storage Boot Camp

THI0430

Book 1 of 2

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

Hi-Track, ShadowImage, TrueCopy. Essential NAS Platform and Universal Storage Platform

Hitachi Data Systems Trademarks

HiCard, HiPass, Hi-PER Architecture, HiReturn, Hi-Star, iLAB, NanoCopy, Resource Manager, SplitSecond, TrueNorth, 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.

©2011, Hitachi Data Systems Corporation. All Rights Reserved

HDS Academy 1071

Contact Hitachi Data Systems at www.hds.com.

Product Names

Enterprise Storage Systems

- Hitachi Virtual Storage Platform
- Hitachi Universal Storage Platform® V
- Hitachi Universal Storage Platform® VM

Legacy Products:

- Hitachi Universal Storage Platform®
- Hitachi Network Storage Controller
- Hitachi Lightning 9900™ V Series enterprise storage systems
- Hitachi Lightning 9900™ Series enterprise storage systems

Modular Storage Systems

- Hitachi Adaptable Modular Storage
- Hitachi Adaptable Modular Storage 2000 Family

Legacy Products:

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

Management Tools

- 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
- Hitachi TrueCopy® Heterogeneous Remote Replication bundle
- Hitachi TrueCopy® Remote Replication bundle (for modular systems)
- Hitachi TrueCopy® Synchronous
- Hitachi TrueCopy® Asynchronous

- Hitachi TrueCopy® Extended Distance

In-System Replication:

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

Hitachi Command Suite

- Hitachi Basic Operating System
- Hitachi Basic Operating System V
- Hitachi Command Director
- Hitachi Device Manager
- Hitachi Dynamic Link Manager Advanced
- Hitachi Dynamic Provisioning
- Hitachi Dynamic Tiering
- Hitachi Global Link Manager
- Hitachi Global Reporter
- Hitachi Replication Manager
- Hitachi Storage Capacity Reporter, powered by APTARE®
- Hitachi Storage Services Manager
- Hitachi Tiered Storage Manager
- Hitachi Tuning Manager
- Hitachi Virtual Server Reporter, powered by APTARE®

Other Software

- Hitachi Backup and Recovery, powered by CommVault®
- Hitachi Backup Services Manager , powered by APTARE®
- Hitachi Business Continuity Manager
- Hitachi Command Control Interface (CCI)
- Hitachi Volume Migration
- Hi-Track® Remote Monitoring system

Other Solutions and Terms

- Hitachi Content Platform
- Hitachi Compute Systems Manager
- Hitachi NAS Platform
- Hi-Star™ crossbar switch architecture
- Hitachi Universal Star Network V

Contents

Book 1 of 2

INTRODUCTION.....	IX
Welcome	ix
Course Description	x
Prerequisites	xi
Course Objectives.....	xii
Course Topics.....	xv
Learning Paths.....	xvii
HDS Academy Is on Twitter and LinkedIn	xviii
 1. ENTERPRISE STORAGE ARCHITECTURE.....	 1-1
Module Objectives	1-1
Virtual Storage Platform.....	1-2
VSP Full Configuration — 6 Rack.....	1-3
VSP Two Module System	1-5
Control and Drive Chassis Structure	1-6
19inch Industry Standard Width Rack — DKC	1-7
DKU Frames	1-8
DKC Components	1-9
Dual Cluster Structure of the DKC.....	1-10
DKC Components	1-11
Back End Director and Front End Director Pairs	1-13
DKC Components	1-14
Cache PCB and Component Specifications	1-15
DKC Components	1-16
VSP Hardware Architecture	1-19
Architecture Layout Dual Node	1-20
Service Processor (SVP)	1-21
DKU Box Structure.....	1-22
Expanded Support for Multiple HDD Types.....	1-23
Consistent or Mixed HDU boxes Supported.....	1-24
VSP Back End Cabling	1-25
VSP DKU and HDU Numbering — Front View.....	1-26
VSP DKU and HDU Numbering — Rear View	1-27
VSP B4 Layout — Front View.....	1-28
VSP B4 Layout — Rear View	1-29
VSP SAS Back End Paths.....	1-30
Encryption Support	1-31
Disk Sparing Operations.....	1-32
Hi-Track Monitor	1-33
Architecture — Storage	1-35
Data Redundancy	1-36
Supported RAID Configurations	1-37
RAID Configurations	1-38
Module Summary.....	1-42
Module Review	1-43
 2. ENTERPRISE STORAGE SOFTWARE.....	 2-1
Module Objectives	2-1
Software Tools for Configuring VSP	2-2
Maintenance and Administration GUI Comparison	2-3
SVP Application	2-4
Storage Navigator 2 Web Console	2-5
Storage Navigator 2 Login Screen.....	2-6
Command Line Interface	2-7

Hitachi Command Suite 7	2-8
Module Summary	2-9
Module Review	2-10
3. HITACHI STORAGE NAVIGATOR 2	3-1
Module Objectives	3-1
Interface	3-2
Setup	3-3
Security	3-4
GUI	3-5
Management Tasks	3-6
Provisioning Tasks	3-7
Viewing Storage System Information	3-8
System Alerts	3-9
Licensing	3-10
GUI Feature Highlights	3-13
GUI Overview	3-16
GUI Actions	3-17
Resource Tree and Summaries	3-19
Hardware Summary	3-20
Allocation Summary	3-21
Viewing Components	3-22
Viewing Parity Groups	3-23
Viewing Parity Group LDEVs	3-24
Viewing Parity Group LDEV Detail	3-25
Viewing Logical Devices	3-26
Viewing Storage Pools	3-28
Viewing Ports and Host Groups	3-29
Using Ports and Alert Link	3-30
Creating and Viewing Reports	3-31
Create Configuration Report Screen	3-32
HTML Table Reports	3-33
HTML Graphical Reports	3-34
CSV Text Reports	3-36
Module Summary	3-39
Module Review	3-40
4. HITACHI DEVICE MANAGER	4-1
Module Objectives	4-1
Device Manager	4-2
Overview	4-3
Benefits	4-4
Hitachi Storage Management Tools	4-5
Unified Data Center Management Tools	4-6
Overview	4-7
Architecture Overview	4-9
Server Components	4-10
Agent	4-11
Graphical User Interface	4-12
Configuration Overview	4-16
Command Line Interface	4-17
Register Storage System	4-18
Data Collection Tasks	4-20
Managing Hosts	4-22
Host Discovery	4-24
Allocate Storage	4-25
Instructor Demonstration	4-29
Module Summary	4-30
Module Review	4-31

5. HITACHI GLOBAL LINK MANAGER ADVANCED.....	5-1
Module Objectives	5-1
Hitachi Dynamic Link Manager (HDLM)	5-2
Overview	5-2
Benefits	5-3
Features	5-4
HDLM Features	5-5
Path Failover and Failback	5-6
Load Balancing	5-7
Round Robin — Extended Round Robin	5-11
Graphical User Interface	5-12
Optional Parameters	5-12
Operations.....	5-15
Instructor Demonstration.....	5-16
Hitachi Global Link Manager (HGLM).....	5-17
Multipathing — Problems and Solutions	5-17
Problems and Solutions	5-18
Configuration Overview.....	5-19
Features	5-20
Graphical User Interface	5-21
Software Path Management	5-22
Alert Management.....	5-23
Host Management.....	5-25
Instructor Demonstration.....	5-26
Module Summary.....	5-27
Module Review	5-28
 6. HITACHI DYNAMIC PROVISIONING.....	 6-1
Module Objectives	6-1
Virtualization	6-2
Capacity Virtualization and Data Mobility	6-3
Dynamic Provisioning	6-4
What is Dynamic Provisioning?	6-6
Benefits	6-9
Overview	6-10
Components.....	6-11
How Does It Work?	6-12
Setup and Monitoring.....	6-13
Thresholds	6-14
DP Operations	6-15
DP Volumes	6-18
Module Summary.....	6-24
Module Review	6-25
 7. HITACHI DYNAMIC TIERING.....	 7-1
Module Objectives	7-1
Tiered Storage and Information Lifecycle Management.....	7-2
Classic Provisioning Volumes From Storage Tiers	7-3
Classic Data Mobility.....	7-4
Challenges with Implementing Tiered Storage	7-5
Hitachi Dynamic Tiering	7-7
Hitachi Dynamic Tiering Builds on Dynamic Provisioning	7-9
HDT — Page Level Tiering.....	7-10
Hitachi Dynamic Tiering Benefits.....	7-12
How Hitachi Dynamic Tiering Fits into Tiered Storage	7-14
Hitachi Dynamic Tiering Specifications.....	7-15
HDP versus HDT	7-17

Hitachi Dynamic Tiering Operations.....	7-19
Hitachi Command Suite Operations for HDT	7-19
HDT Operations.....	7-20
Module Summary	7-30
Module Review	7-31
8. HITACHI TIERED STORAGE MANAGER.....	8-1
Module Objectives	8-1
Customer Storage Management Needs.....	8-2
Overview.....	8-3
Benefits.....	8-4
Requirements	8-5
Software Options	8-6
Components	8-7
Volumes Used in HTSM	8-8
Concepts	8-9
Operations.....	8-14
Tiers.....	8-15
Migration Groups.....	8-25
Migration Tasks	8-38
Migration Task.....	8-41
Migration Task Status.....	8-48
Module Summary	8-51
Module Review	8-52
9. EXTERNAL STORAGE MANAGER	9-1
Module Objectives	9-1
Hitachi Virtualized Tiered Storage Architecture	9-2
Virtualization	9-3
Capacity Virtualization and Data Mobility	9-4
Universal Volume Manager	9-5
Features	9-6
Configuring External Storage	9-7
Change the Port Attribute to External	9-8
Prepare the External System	9-9
Configure Physical Links	9-10
Externally Attached Storage Connectivity	9-11
Discover External Storage.....	9-12
Map External Storage.....	9-13
Concepts for Using External Volumes	9-14
External Volumes and External Volume Group.....	9-15
External Volume Groups	9-16
External Volumes and External Volume Group.....	9-17
Cache Mode Usage.....	9-19
Operations.....	9-20
Virtualize Volumes.....	9-21
Virtualize Storage	9-22
Virtualize Volumes.....	9-25
Data Migration	9-34
Module Summary	9-35
Module Review	9-36
10. HITACHI VIRTUAL PARTITION MANAGER	10-1
Module Objectives	10-1
Virtual Partition Manager Overview.....	10-2
Overview.....	10-3
Cache Logical Partition	10-4

External Storage Virtualization and Cache	10-5
External Storage — How Much Cache?	10-6
How Much Cache to Allocate in Internal Storage?	10-7
A Sensible Suggestion — Rule of Thumb	10-8
Guidelines	10-9
Specifications	10-10
Operations	10-11
Module Summary	10-19
Module Review	10-20

Contents

Book 2 of 2

11. HITACHI HIGH-PERFORMANCE NAS PLATFORM	11-1
12. ARCHIVING AND BACKUP	12-1
13. BUSINESS CONTINUITY	13-1
14. HITACHI REPLICATION MANAGER.....	14-1
15. HITACHI SHADOWIMAGE HETEROGENEOUS REPLICATION.....	15-1
16. HITACHI COPY-ON-WRITE SNAPSHOT	16-1
17. HITACHI TRUECOPY[®] HETEROGENEOUS REMOTE REPLICATION.....	17-1
18. HITACHI UNIVERSAL REPLICATOR.....	18-1
19. COMMAND CONTROL INTERFACE	19-1
20. HITACHI PERFORMANCE MONITOR	20-1
21. HITACHI TUNING MANAGER	21-1
22. HITACHI IT OPERATIONS SUITE	22-1
23. HITACHI COMMAND DIRECTOR.....	23-1
NEXT STEPS	N-1
GLOSSARY	G-1
EVALUATING THIS COURSE	E-1

Introduction

Welcome

- Student Introductions
 - Name
 - Position
 - Experience
 - What you expect from the course?

2

Course Description

- This ten day instructor-led course provides an introduction to the Hitachi Virtual Storage Platform, associated software products, and SAN and NAS implementation.
- The learner performs hands-on lab activities designed to build the skills necessary to operate the key software utilities. These activities provide an opportunity to work with UNIX and Microsoft® Windows® integrated with the Hitachi enterprise storage systems.
- The course is designed for pre-sales technical personnel, but can benefit any audience that requires deep technical training on enterprise products, making it appropriate for Hitachi TrueNorth® Partners and customers.
- This course is part of the Hitachi Certified Storage Professional Program and supports the Hitachi Data Systems Storage Foundations – Enterprise Exam (HH0-110).

3

Enterprise storage systems include:

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

Prerequisites

- A solid understanding of enterprise storage concepts and functionality

4

Course Objectives

- Upon completion of this course, you should be able to:
 - Describe the architecture, essential components, and features of Hitachi Enterprise storage systems
 - Describe the tools available for the management of Hitachi enterprise storage systems
 - Describe Hitachi Device Manager features and functions
 - Describe the Hitachi Dynamic Link Manager features
 - Describe how Dynamic Link Manager and Hitachi Global Link Manager help solve business continuity problems
 - Define dynamic provisioning and thin provisioning
 - Explain how to perform dynamic provisioning operations

5

- Upon completion of this course, you should be able to:
 - Describe how Hitachi Dynamic Tiering is used to optimize tiered storage utilization
 - Describe the customer business value of Hitachi Tiered Storage Manager
 - Demonstrate how Hitachi Command Suite is used to manage tiered storage
 - State the purpose of the Hitachi Virtual Partition Manager
 - Describe the features and usage of cache logical partitions
 - State the purpose and benefits of using Hitachi High-performance NAS Platform, powered by BlueArc®
 - Describe the Hitachi High-performance NAS Platform architecture

6

- Upon completion of this course, you should be able to:
 - Identify the hardware and software components of the Hitachi High-performance NAS Platform
 - Describe Hitachi Content Platform features and functions
 - Describe Hitachi Data Protection Suite, powered by CommVault® features and functions
 - Describe Hitachi Business Continuity solutions and the available software
 - Describe Hitachi Replication Manager features and functions
 - Describe the key features and operations of Hitachi ShadowImage® Heterogeneous Replication
 - Identify ShadowImage Heterogeneous Replication RAID Manager and Command Control Interface components

7

- Upon completion of this course, you should be able to:
 - Describe the purpose of Hitachi Copy-on-Write Snapshot
 - Compare the functionality of Copy-on-Write Snapshot to ShadowImage Heterogeneous Replication
 - Describe typical Copy-on-Write Snapshot operations
 - Describe the key benefits and features of Hitachi TrueCopy® Heterogeneous Remote Replication bundle
 - List the steps to configure replication with Hitachi TrueCopy® Synchronous
 - Identify the purpose of Hitachi Universal Replicator
 - Describe how Hitachi Tuning Manager enables storage performance optimization

8

- Upon completion of this course, you should be able to:
 - Describe the features and functions of Hitachi IT Operations Analyzer
 - Describe the features and functions of Hitachi IT Operations Director
 - Describe the purpose of Hitachi Command Director
 - Describe how Hitachi Command Director helps solve customer issues

9

Course Topics

Week One

Content	Labs
1. Enterprise Storage Architecture	1. Hitachi Storage Navigator 2 GUI
2. Enterprise Storage Software	
3. Hitachi Storage Navigator 2	
4. Hitachi Device Manager	2. Agent Installation and Configuration
5. Hitachi Global Link Manager Advanced	3. Hitachi Dynamic Link Manager Advanced
6. Hitachi Dynamic Provisioning	4. Hitachi Dynamic Provisioning
7. Hitachi Dynamic Tiering	5. Hitachi Dynamic Tiering Operations
8. Hitachi Tiered Storage Manager	6. Hitachi Tiered Storage Manager
9. External Storage Management	7. External Storage
10. Hitachi Virtual Partition Manager	8. Hitachi Virtual Partition Manager

10

Week Two

Content	Labs
11. Hitachi High-performance NAS Platform 12. Archiving and Backup 13. Business Continuity 14. Hitachi Replication Manager	9. Hitachi Replication Manager Initial Setup
15. Hitachi ShadowImage Heterogeneous Replication 16. Hitachi Copy-on-Write Snapshot	10. Hitachi <u>ShadowImage</u> Operations
17. Hitachi TrueCopy Heterogeneous Remote Replication 18. Hitachi Universal Replicator 19. Command Control Interface	11. Hitachi <u>TrueCopy</u> Operations
20. Hitachi Performance Monitor 21. Hitachi Tuning Manager	12. Hitachi Tuning Manager GUI
22. Hitachi IT Operations Suite 23. Hitachi Command Director	13. Hitachi Command Director Reports

11

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



12

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

Partner Xchange: <https://portal.hds.com/>

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.

13

These are the URLs for Twitter and LinkedIn:

<http://twitter.com/HDSAcademy>

http://www.linkedin.com/groups?gid=2972602&trk=hb_side_g

1. Enterprise Storage Architecture

Module Objectives

- Upon completion of this module, you should be able to:
 - Describe the architecture, essential components, and features of Hitachi Enterprise storage systems

2

Virtual Storage Platform

Innovative Virtualization

Centralized, secured and simplified storage management with multi-vendor storage support

Seamless Migration and 100% availability

90% reduction on USP V and USP VM migration effort with no application outage

Unique 3D Scaling Technology

High performance grid scale-up and scale-out storage architecture with advanced virtualization capabilities and connectivity

Dynamic Mobility

Enables fluid storage architecture, simplifies QoS and lowers storage TCO for block and file data

Sustainable IT

40% less power and half the footprint versus competition

Data Resilience and Protection

Proven replication and continuous data protection options to simplify backup and business continuity needs



Hitachi Virtual Storage Platform

3

USP V = Universal Storage Platform V

USP VM = Universal Storage Platform VM

QoS = Quality of Service

TCO = Total Cost of Ownership

The Hitachi Virtual Storage Platform scales in three dimensions.


It can **scale up** to provide improved performance and throughput as its host servers consolidate and scale up .

It can **scale out** to deliver capacity, performance and bandwidth from a shared pool of resources as more and more host servers are connected to the Hitachi Virtual Storage Platform.

It can **scale deep** by connecting external storage, enabling existing multivendor storage or lower-cost modular storage to add their resources to the scale up and scale out demands of host servers. This extends the life of existing storage assets and increases the return on assets.

VSP Full Configuration — 6 Rack

The maximum number of frames is six, two DKC boxes and 16 DKU boxes.



	1module	2module
HDD(2.5")	1,024	2,048
HDD(3.5")	640	1,280
CHA ports	80 (96 ¹)	176 (192 ¹)
Cache	256GB (512GB) ²	512GB (1,024B) ²

¹ ALL CHA configuration (Diskless)

² Enhanced (planned for V02)

4

A fully-configured VSP system contains 2 DKC boxes, in each of 2 separate racks, and 16 HDU boxes.


A fully-configured VSP system requires six racks.

- Each 19" rack is 60 cm wide, outside edge to outside edge
- Rack is 110 cm deep including rear door
- Total width of six racks is 278 mm or 11 inches wider, compared to the five cabinet, fully-configured USP V


The table on this page also shows cache capacity. VSP V01 will support 256GB cache in one module or 512GB cache in a 2 module configuration. These capacities are expected to double when V02 is released.

The rack naming convention is slightly different from the RAID600 USP V. Each rack has a two digit identifying number. Each of the DKC racks will be RK-00 and RK-01 respectively. The HDU racks associated with its DKC rack uses the same left digit followed by a 1 or 2 depending upon its physical position relative to the DKC rack.

Diskless configuration option is supported for the VSP. In the case of a diskless configuration, more CHA ports are possible.




VSP Two Module System



6

Control and Drive Chassis Structure


- Control chassis includes virtual storage directors (VSD), cache, FED and BED adapters, power supplies and service
- Drive chassis includes drives and SAS links
- Fan assembly can be opened so that a drive may be exchanged online
- 80 x 3.5 in drives or 128 x 2.5 in drives



7

- Looking at the control chassis, the design is a more modular, blade style structure.
- Virtual Storage Directors and Cache adapters are added to the front.
- FEDs, BEDs and Grid Switch adapters are added to the back.
- Services Processors are accessed from the back of the system as well.
- Two control chassis (14 rack units high) can be combined to operate as a single unit.
- The drive chassis (13 rack units high) contain either 2.5" or 3.5" drives.
- Fan doors are moved aside in order to service drives online.
- Fans on the opposite side run faster to move air when the other side is open.

19inch Industry Standard Width Rack — DKC




8

The VSP system is assembled in a rack frame. The rack is a Hitachi-custom rack that conforms to the industry standard 19in width. The rack is 42U high and 1100 mm (43.3 inches) deep.

The DKC box and DKU box containers are used to hold the VSP components. The DKC box is 14U high and the DKC box is 13U high. The first rack will contain one of DKU box, one front and one rear. This rack may also contain one or two DKU boxes.

DKU Frames



The VSP DKU frame is also a 19in industry standard width rack. All racks in a VSP system have the same outside dimensions.

A rack that is used as a DKU frame can hold three DKU boxes, each 13U high.

An HDU box contains HDDs in both the front and rear.


The weight of the DKU is 80kg.

The VSP supports two different internal HDU box structures — one that holds Large Form Factor (LFF) 3.5" disk drives and one that holds Small Form Factor (SFF) 2.5" disk drives. One HDU box can be either for LFF or SFF disk drives but not mixed. A mix of HDU boxes can be configured in one VSP system. LFF and SFF HDU boxes can be mixed in any configuration in the VSP.

An LFF HDU box can contain a maximum of 80 HDDs. An SFF HDU box can contain a maximum of 128 HDDs.

DKC Components

- Controller chassis (DKC) consists of:
 - Channel adapter (CHA) or front end directors (FED)
 - Disk adapter (DKA) or back end directors (BED)
 - Cache memory adapter or data cache adapter (DCA)
 - Grid switch adapter (GSW)
 - Virtual storage directors (VSD)
 - Service processor (SVP)
 - Cooling fan
 - AC-DC power supply



10


The battery and the Cache Flash Memory are also installed in the CMA to prevent data loss from a power outage or other event.

The storage system continues to operate when a single point of failure occurs, by adopting a duplexes configuration for each control board (CHA, DKA, CPC, GSW and VSD), a redundant configuration for the AC-DC power supply and the cooling fan.

The addition and the replacement of the components and the upgrade of the microcode is an online operation.

The SVP allows the engineers to set and modify the system configuration information, and also can be used for checking the system status. The SVP can also be configured to report system status and errors to Service Center and enables the remote maintenance of the storage system.

Dual Cluster Structure of the DKC



The components on VSP are installed in a two-cluster structure. The diagrams on this page show the front view and rear view of DKC#0. In both the front and rear of the DKC box, Cluster 1 components are found in the lower slots and Cluster 2 components are found in the upper slots of the DKC box.

When a VSP system includes two modules, both DKC#0 and DKC#1, the DKC slots in the second module have different identification codes. The main PCB types are found in the same slot locations and the cluster boundaries are the same in both DKC modules.

PCB types:

1. Data Cache Adapter (CACHE)
2. Virtual Storage Director (MP)
3. Grid Switch (ESW)
4. Back End Director (DKA)
5. Front End Director (CHA)

DKC Components

- Front End Director (FED)/Channel Adapter (CHA)
 - Controls data transfer between the hosts and the cache memory
 - Two types
 - Mainframes — FICON
 - Open Systems — FC
 - FC Specs
 - Max FED only Options — 8
 - In addition 4 CHA options can be installed instead of DKA options
 - 8 port or 16 port options (4 or 8 ports per CHA PCB)
 - 8Gbps (Auto negotiate 2/4/8Gbps)
 - Max Ports
 - FED Slots only — 64 (8 port option), 128 (16 port option)
 - FED+DKA Slots — 96 (8 port option), 192 (16 port option)

12

Note: One Option consists of two PCBs. One gets installed in Cluster 1 (CL1) and the second in Cluster 2 (CL2).

- **Back End Director (BED)/Disk Adapter (DKA)**
 - Controls data transfer between the disks and the cache memory
 - DKA Specs
 - Max BED Options — 4
 - Max BED options are reduced if FED installed in BED slots
 - 8 SAS ports per option
 - 32 SAS ports maximum
 - 6Gbps
 - Max Disk Drives per SAS port
 - 128 (2.5" Drives)
 - 80 (3.5" Drives)
 - Disk Types supported
 - SSD, SAS, SATA
 - Diskless VSP – Zero (0) BED options

13

The Back End Director (BED) boards execute all I/O jobs received from the processor boards and control all reading or writing to disks. There are one or two features (two or four BED boards) per chassis.


BED functions include the following:

- Execute jobs received from a VSD board
- Use DMA to move data in or out of data cache
- Create RAID-5 and RAID-6 parity with an embedded XOR processor
- Encrypt data on disk (if desired)
- Manage all reads or writes to the attached disks

Each BED board has eight 6Gbps SAS links. There are up to 640 LFF disks or 1024 SFF disks per chassis attached to the 16 or 32 6Gbps SAS links from these 2 or 4 BED boards.

Back End Director and Front End Director Pairs

- Operational control data is distributed to L2 memory
 - Data accelerator processors:
 - Transfer data to and from cache
 - Execute host command requests
 - Dual core, special purpose I/O processors built into a unique ASIC package
 - Designed to accelerate I/O and operational performance
 - Offload latency sensitive processing tasks directly onto the BEDs and FEDs




14

DKC Components

■ Cache Memory Adapter

- Caches user data blocks from drives via BED during a read
- Caches data from FED as part of a data write operation
- All write data is mirrored in cache
- Only one copy of cached data when performing read operations
- Up to 64GB of data cache and a total of 256GB in a single module
 - Two modules can contain 512GB of cache memory.
- Non-volatile data protection
 - Backup cached user data
 - Also used to protect the state
- Allows for three copies
 - One operational state
 - Two backup copies (mirrors)
 - Backed up to an onboard SSD



15

The Data Cache Adapter (DCA) boards are the memory boards that hold all user data and the master copy of Control Memory (metadata).

There are up to 8 DCAs installed per chassis, with 8GB to 32GB of cache per board (32GB to 256GB per chassis) when using the current 4GB DIMM (as part of the 16GB feature set). The two boards of a feature must have the same RAM configuration, but each DCA feature can be different.

The first two DCA boards in the base chassis (but not in the expansion chassis) have a region of up to 48GB (24GB per board) used for the master copy of Control Memory. Each DCA board also has a 500MB region reserved for a Cache Directory.

Each DCA board also has one or two on-board SSD drives (31.5GB each) for use in backing up the entire memory space in the event of an array shutdown due to power failure. If the full 32GB of RAMM is installed on a DCA, it must have two 31.5GB SSDs installed. On-board batteries power each DCA board long enough to complete several such shutdown operations back-to-back in the event of repeated power failures before the batteries have had a chance to charge back up.

Cache PCB and Component Specifications

No.	Item	Specification	
1	Board model name	DKC-F710I-CPC	A pair of CM Boards (included a 32GB-SSD and a battery)
2	DIMM model name	DKC-F710I-C16G DKC-F710I-C32G*	4GB-DIMM x 4 8GB-DIMM x 4*
3	SSD model name	DKC-F710I-BM64 DKC-F710I-BM128*	32GB-SSD x 2 64GB-SSD x 2*
5	Spare parts (FRU)	CM board (without SSD, battery, and DIMM) Cache memory (4/8*GB DIMM) Battery (12v) SSD (32/64*GB)	
6	Firmware on CM board	1) Backup and battery firm: Memory backup and Battery charge/discharge control Online firm update is available 2) SSD Firm: SSD internal control Online firm update will be available from 2 nd GA	

* Enhanced model


16

The table on this page identifies the component codes for the Cache PCB, the cache memory DIMMs, and the SSDs.

This information also indicates that the firmware for the battery management and memory backup is nondisruptively upgradeable starting with V01. Online upgrade of the SSD internal control firmware will be available in V02.

DKC Components

- Virtual Storage Director
 - Main processors and operational control data memory
 - Stores and manages internal operational metadata and state
 - Array groups, LDEVs, external LDEVs, runtime tables, mapping data for various software products
 - Overall state of system stored, referenced, and executed
 - Distributed to appropriate I/O offload processors on FED/BED



17

The VSD board is the VSP I/O processing board. There are two or four of these installed per chassis. Each board includes 1 Intel 2.33GHz Core Duo Xeon CPU with 4 processor cores and 12MB of L2 cache. There is 4GB of local DDR2 RAM on each board (2 DIMMs). This local RAM space is partitioned into five regions, with one region used for each core's private execution space, plus a shared Control Memory region used by all four cores.

Each VSD board executes all I/O requests for the LDEVs that are assigned to that board. No other VSD board can process I/Os for these LDEVs.

No user data is processed within the VSD itself.

The firmware loaded onto the VSD board contains five types of code and each Xeon core will schedule a process that depends upon the nature of the job it is executing. A process will be of one of the following types or their mainframe equivalents:

- Target process – Manages host requests to and from a FED board for a particular LDEV.
- External (virtualization) process – Manages requests to or from a FED port used in virtualization mode (external storage). Here the FED port is operated as if it were a host to operate the external array.
- BED process – Manages the staging or de-staging of data between cache blocks and internal disks via a BED board.


- HUR Initiator (MCU) process – Manages the *respond* side of a Hitachi Universal Replicator connection on a FED port.
- RCU Target (RCU) process – Manages the *pull* side of a remote copy connection on a FED port.

In addition there will be system housekeeping type processes.

- Grid Switch

- Provides multiple high performance interconnect paths to transfer user data between FED/BED and DCA
- Total of 96 ports @ 24 ports per GSW board

- The GSWs are the hub of the HiStar-E Network




18

The core of the VSP HiStar-E Network architecture is the Express Switch which is identified by the acronym GSW. The internal HiStar-E Network interconnections are different due to the separation of the microprocessors on the new MP PCBs. The GSW provides the highly redundant, high performance interconnection paths among the other main components of the HiStar-E Network as shown by the schematic diagram on this page.

- The Grid Switch boards supply the high performance interconnect paths that cross connect all of the other boards.
- There are two or four GSW boards per chassis.
- Each GSW board has 24 unidirectional ports, each port having a send path and a receive path, with each operating at 1024MBps.
- As such, the GSW supports an aggregate peak load of 24GBps send and 24GBps receive (or 48GBps overall).
- Eight of the GSW paths connect to the FED and BED boards for a total of 192GBps full duplex.
- Eight other paths attach to the cache boards.
- Four more paths connect to VSD boards, and the final four paths cross connect to the matching GSW board in the second chassis (if used).
- There are no connections among the two or four GSWs within a chassis. Every board in a chassis attaches to two or four GSWs.


VSP Hardware Architecture



19

The diagram on this page shows how the DKC components in a single DKC system are connected. A single DKC VSP system can support a maximum of 4 DKA features for a total of 16 back end data loops.

Architecture Layout Dual Node



20

When the VSP configuration needs to expand to more back end capacity and/or access, the second DKC, module 1 must be added. The two DKCs are connected between their respective GSWs.

Service Processor (SVP)

- Blade server that runs the application for performing hardware and software maintenance functions
- SVP PC provides three main functions:
 - Human interfaces
 - Storage system health monitoring and reporting
 - Performance monitor capabilities
- SVP application, web console, and Storage Navigator applications run on SVP PC
- If SVP PC fails or is unavailable, movement of I/O is not affected
- Support for High Reliability Kit — Second SVP
- Maintenance PC is used to connect to SVP for storage system management and administration

21


The purpose of the SVP blade PC is to provide the *human interface* to the Virtual Storage Platform (VSP).

Functions of the SVP include:

- Internal communication with the Micro Processors (MPs)
- Interface to perform microprogram exchange
- Platform where the GUIs run including:
 - ♦ Web Console
 - ♦ Storage Navigator 2 (SN2) (web services)
 - ♦ SVP Application
- Connection point for the Virtual Storage Platform (VSP) to the customer LAN
- Connection point for the CE Laptop
- Connection point for Hi-Track Monitor reporting connection
- Collects and reports internal errors (SIMs) and alarms
- Collects and reports workload and performance information through the SVP Monitor
- Interface through which to download dump information
- Interface through which to download or backup configuration information

DKU Box Structure

- HDD box contains HDDs and SSW(SAS Switch)-PKs
- Fan-assembly is hinged on outside edge of HDU box
- Fan assembly door is opened to provide access to HDD slots for replacement, insertion, and removal operations



LFF 3.5HDD box – maximum 80 HDD

22


A closer view of the HDU box structure is shown on this page. The example shown is the LFF structure which holds 3.5" HDDs. Each HDU box has two fan door assemblies on the front and two on the rear. From this diagram, you can see how the fan door assembly blocks access to the HDD slots when the fan doors are in their normal operating position.

The fan assembly latch mechanism is located from front-to-back between the upper sets of HDDs. The fan door latch is pushed or pulled, depending on which side you are standing on and on which side you need to open a fan door.

Note: Technicians who have worked with early versions of the VSP system have found that the fan door latch mechanisms are a bit touchy. You may have to walk to the other side of the system if the fan door latch gets stuck.

The SSW components are installed along the sides of the HDU box in both the front and rear.


Expanded Support for Multiple HDD Types




23

Consistent or Mixed HDU boxes Supported


When all drives are 3.5", maximum number is 1280



When all drives are 2.5", maximum number is 2048



2.5" box and 3.5" box can be intermixed



24


The VSP supports two different internal HDU box structures: one that holds Large Form Factor (LFF) 3.5" disk drives and one that holds Small Form Factor (SFF) 2.5" disk drives. One HDU box can be either for LFF or SFF disk drives but not mixed. A mix of HDU boxes can be configured in one VSP system. LFF and SFF HDU boxes can be mixed in any configuration in the VSP.

An LFF HDU box can contain a maximum of 80 HDDs. An SFF HDU box can contain a maximum of 128 HDDs.


The part number for an SFF HDU box is DKC-F710I-SBX.

The part number for an LFF HDU box is DKC-F710I-UBX. (It not known why *U* is used in the LFF number.)

VSP Back End Cabling



VSP DKU and HDU Numbering — Front View




In the VSP, each DKU is identified by a two digit number which includes the number of the DKC, 0 or 1, and a number for the DKU within the module.

Remember that the DKC#0 module can be configured on either the right or the left as compared to its DKC#1 partner.


This will add complexity to component identification at the customer site. In a two-module system, you will need to know or be able to determine the position of DKC#0 and DKC#1 relative to each other.

VSP DKU and HDU Numbering — Rear View




27

VSP B4 Layout — Front View




VSP B4 Layout — Rear View




29

VSP SAS Back End Paths



Encryption Support

- All Virtual Storage Platforms are encryption capable
 - Every BED has built in encryption capability
 - Encryption needs to be enabled by software license key
- New encryption mode of operations
 - XTS-AES 256 bit encryption
- Expanded key support
 - 32 keys per array
 - Encryption as access control



31

Every VSP is *encryption capable* and requires a license key to be installed to be activated.

Each array supports up to 32 encryption keys per platform allowing for encryption to be used as an access control mechanism within the array. This allows for different classifications of data to be stored on the same array with encryption providing a data leakage prevention mechanism.

Disk Sparing Operations

- **Dynamic Sparing (Pre-Emptive Copy)**
 - Means that ORM (online read margin) diagnostics have determined a drive to be suspect, or drive read/write error thresholds have been exceeded
 - Storage system spares out the drive even though the drive has not completely failed
 - Data is copied to spare drive (not recreated)
- **Correction Copy (Disk Failure)**
 - Occurs when a drive fails
 - If a spare is available, lost data is re-created on the spare, which logically becomes part of the array group
 - This mode invokes DRR chip, where pre-emptive does not
 - If a spare is not available, array group is at risk for a longer period of time than normal, plus that group will continuously run in degraded mode until bad drive is replaced

32

Hi-Track Monitor

- Installed on the Service Processor
- Service and remote maintenance tool
- Monitors the operation of the VSP-V 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 Support Center

33

Hi-Track Monitor is a software utility program that is optionally installed on the SVP PC.

The Hi-Track Monitor software monitors the operation of the VSP 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. 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. This feature provides a remote SVP mode for the large RAID systems that enables the specialist to operate the SVP as if they were at the site. This allows support specialists to provide immediate, remote troubleshooting and assistance to any Hi-Track location.

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

- Hi-Track SVP Agent for VSP is a Microsoft Windows® application that runs on the SVP
- Performs Hi-Track error monitoring, configuration data acquisition, and remote SVP mode functions
- Transport to Hi-Track Center can be either
 - HTTPS or FTP (SSL or standard)
 - Via public internet or dialup modem
- Can send email alerts to customer
 - User definable destinations
- Optionally, the Hi-Track FTP Relay application can be used to perform FTP transport across LAN segments in case the SVP does not have direct Internet connectivity.

Architecture — Storage

Storage Overview



35

Parity Groups are created from the physical disks.

A **RAID Level** and **emulation** is applied to the group.

The emulation creates equal sized stripes called **LDEVs** (Logical Devices).

LDEVs are mapped into a **LDKC, Control Unit** matrix (LDKC#:CU#:LDEV#).


Control Unit (CU): A Control Unit is a logical entity. All the Logical Devices (LDEVs) that have been carved out of a RAID Group have to be a part of a Control Unit (up to 256 on the Universal Storage Platform V). There can be a maximum of 256 LDEVs in each CU.

Logical DKC (LDKC): An LDKC is a set of Control Unit tables. Each LDKC contains 256 (hex FF) control unit tables. Currently LDKC can be set as 0 or 1. LDKC 0 contains open LDEVs only. LDKC 0&1 can be used in mainframe environments.

Emulations: When you wish to carve out LDEVs from a RAID group, you must specify the size of the LDEVs. The storage system supports various emulation modes which specify the size of each LDEV in a RAID group. Each RAID group can have only one emulation type. Storage system can have multiple RAID groups with different emulations, such as, Open-V.

Data Redundancy

- RAID Implementation



Four or eight physical HDDs are configured into a RAID group (also called a *Parity* group).

- Groups of 4 or 8 (16 or 32 with concatenation) HDDs are set up using one of three parity options:
 - Supported RAID Levels
 - RAID-1 (2D+2D)/(4D+4D)
 - RAID-5 (3D+1P)/(7D+1P)
 - RAID-6 (6D+2P)


36

The **Concatenated Array Group** feature allows you to configure all of the space from either two or four RAID-5 (7d+1p) Array Groups into an association of 16 or 32 drives whereby all LDEVs created on these Array Groups are actually striped across all of the elements. Recall that a slice (or partition) created on a standard Array Group is an **LDEV** (Logical Device), becoming a **LUN** (Logical Unit) once it has been given a name and mapped to a host port.


Supported RAID Configurations

■ RAID-1

2D + 2D Configuration



4D + 4D Configuration



37

Description:

- Also called *mirroring*
- Two copies of the data
- Requires twice the number of disk drives
- For writes, a copy must be written to both disk drives
- Two parity group disk drive writes for every host write
- Do not care about what the previous data was, just over-write with new data
- For reads, the data can be read from either disk drive

Read activity distributed over both copies reduces disk drive busy (due to reads) to ½ of what it would be to read from a single (non-RAID) disk drive.


Advantages: Best performance and fault-tolerance

Disadvantages: Uses more *raw* disks to implement which means a more expensive solution.

RAID Configurations

■ RAID-5

3D + 1P Configuration



7D + 1P Configuration



38

- For sequential reads and writes, RAID-5 is very good.
 - ♦ It's very space efficient (smallest space for parity), and sequential reads and writes are efficient, because they operate on whole stripes.
- For low access density (light activity), RAID-5 is very good.
 - ♦ The 4x RAID-5 write penalty is (nearly) invisible to the host, because it is asynchronous.
- For workloads with higher access density and more random writes, RAID-5 can be throughput-limited due to all the extra parity group I/O operations to handle the RAID-5 *write penalty*.

In the RAID-5 design data are written to the first three disks and the fourth disk has an error correction data set that would allow any one failing block to be reconstructed from the other three. This method has the advantage that only, effectively one disk is used out of the four for error correction, (parity), information.

Small-sized records are intensively read and written randomly in transaction processing. This type of processing generates many I/O requests for transferring small amounts of data. In such a situation, greater importance is placed on increased I/O performance (parallel I/O processing) than on increase in the rate of transferring large-volume data. RAID-5 has been introduced to be suitable for this type of transaction processing.


- Parity calculated by XOR-ing bits of data on the stripe
- Overlapping I/O requests allowed

Recovery from failure:

- Missing data recalculated from parity and stored on spare

■ RAID-6 – Striping With Dual Parity Drives

6D + 2P Configuration



39

Description:

- Sometimes called RAID-5DP
- Two parity schemes used to store parity on different drives
- An extension of the RAID-5 concept that uses two separate parity-type fields usually called *P* and *Q*
- Allows data to be reconstructed from the remaining drives in a parity group when any one or two drives have failed.

Note: The math is the same as for ECC used to correct errors in DRAM memory or on the surface of disk drives.

- Each host random write turns into 6 parity group I/O operations
 - ♦ Read old data, read old *P*, read old *Q*
 - ♦ (Compute new *P*, *Q*)
 - ♦ Write new data, write new *P*, write new *Q*
- Parity group sizes usually start at 6+2.
 - ♦ This has the same space efficiency as RAID-5 3+1.

Recovery from failure:

- Missing data recalculated from parity and stored on spare

Advantages:

- Very high fault-tolerance. Duplicate parity provides redundancy during correction copy.

Disadvantages:

- Uses additional space for second parity
- Slower than RAID-5 due to second parity calculation

Module Summary

- Described the architecture, essential components, and features of the Hitachi Enterprise storage systems
- Listed the distinct hardware components in the Virtual Storage Platform (VSP)
- Identified the minimum configuration for VSP
- Described the internal controller chassis architecture
- Listed the five main VSP internal controller chassis (DKC) components
- Identified the DKU, HDU, and B4 layouts

40

Module Review

1. List the key PCB's in the Logic Box of DKC.
2. The Virtual Storage Director is responsible for moving user data in and out of Data Cache Adapters? (True/False)
3. The configuration information for the Virtual Storage Platform is stored on the _____ PCB.
4. The Virtual Storage Platform uses the _____ interface for interfacing with SATA drives.

42


2. Enterprise Storage Software

Module Objectives

- Upon completion of this module, you should be able to:
 - Describe the tools available for the management of Hitachi Enterprise storage systems

2

Software Tools for Configuring VSP



CLI = Command Line Interface

Maintenance and Administration GUI Comparison

GUI Interface	Runs on SVP	Customer Use ?	Functions	Used By
SVP Application	yes	no	Hardware maintenance Hardware configuration Microcode exchange LUN mapping, LUSE, VLL	HDS CE and authorized partners only
Web Console	yes	no	License keys	HDS CE and authorized partners only
Storage Navigator	yes	yes	Provisioning LUN mapping, LUSE, VLL Replication (SI, TC and UR) Performance Monitor Partitioning	Customer Storage administrators and Storage Partition administrators
Device Manager	no	yes	Provisioning LUN mapping, LUSE, VLL SI Replication Performance Monitor	Customer Storage, Storage Partition and System administrators


4

This table compares the four main GUI applications that are used to view and manage the Virtual Storage Platform (VSP) storage systems. Three of the GUI applications run on the SVP PC: The SVP Application, Web Console, and Storage Navigator.

The fourth GUI interface is Hitachi Command Suite Device Manager software, which is installed and runs on a Microsoft Windows® or Sun® Solaris host other than the SVP PC.


The SVP Application and the Web Console applications are used primarily by the maintenance engineer.

SVP Application




The SVP Application is used by the Engineers for doing hardware and software maintenance. The Application is launched by accessing the Web Console application.

Storage Navigator 2 Web Console



The **Web Console** is Storage Navigator 2 that is accessed on the SVP as a user with the **maintenance** user account.

Storage Navigator 2 Login Screen



Storage Navigator 2 GUI is accessed from an end user PC, via the *public* IP LAN, using a supported web browser. In the customer environment, this *public* LAN may be a secured management LAN within the customer's network environment. We use the term *public* LAN to differentiate from the internal LAN within the VSP storage system. Storage Navigator 2 should *never* be accessed and used on the VSP internal LAN.

Command Line Interface

```
C:\Windows\system32\cmd.exe
raidcfg : User authentication has failed on Serial#(53030).
raidcfg: [EX_ENAUTH] Authentication failed with User
Refer to the command log(C:\HORCM\log0\horoc_R7007E8D7E0F_log.txt) for details.

c:\HORCM\etc>raidcom -login maintenance raid-mainte

c:\HORCM\etc>raidcom get ldev -ldev_id 0
Serial# : 53030
LDEV : 0
SL : 0
CL : 0
VOL_TYPE : OPEN-V-CVS
VOL_Capacity(BLK) : 112575360
NUM_LDEV : 1
LDEVs : 0
NUM_PORT : 4
PORTs : CL1-A-1 0 1dot30 : CL4-A-1 0 1dot30 : CL2-A-1 0 1dot30 : CL3-A-1 0 1dot30
F_POOLID : NONE
VOL_ATTR : CVS
RAID_LEVEL : RAID1
RAID_TYPE : 2D+2D
NUM_GROUP : 1
RAID_GROUPS : 01-01
DRIVE_TYPE : DKS5B-J300SS
DRIVE_Capa : 571679368
LDEV_NAMING :
STS : NML
OPE_TYPE : NONE
OPE_RATE : 100
MP# : 3

c:\HORCM\etc>
```

- RAIDCOM
 - In-band
 - Out-of-band
 - CCI

8


The Virtual Storage Platform (VSP) is the first enterprise storage system to include a unified, fully compatible command line interface. The VSP CLI supports all storage provisioning and configuration operations that can be performed through SN2.

The CLI is implemented through the **raidcom** command. The example on this page shows the **raidcom** command that retrieves the configuration information about an LDEV.

For **in-band** CCI operations, the command device is used, which is a user-selected and dedicated logical volume on the storage system that functions as the interface to the storage system for the UNIX/PC host. The dedicated logical volume is called command device and accepts commands that are executed by the storage system.

For **out-of-band** CCI operations, a virtual command device is used. The virtual command device is defined by specifying the IP address for the SVP. CCI commands are issued from the host and transferred via LAN to the virtual command device (SVP) and the requested operations are then performed by the storage system.

Hitachi Command Suite 7



The customer can also use the Device Manager component of the Hitachi Command Suite 7 storage management software products to view and administer the VSP storage system as well as any other HDS storage system.

Module Summary

- Described the tools available for the management of Hitachi Enterprise storage systems

10

Module Review

1. List the customer usable GUI applications for managing the Virtual Storage Platform.

3. Hitachi Storage Navigator 2

Module Objectives

- Upon completion of this module, you should be able to:
 - Identify the key features and functions of Storage Navigator 2
 - Understand Storage Navigator 2 components on the service processor

2

Interface



3

Setup

- SN2 is included with the service processor (SVP) during storage system installation
- If your client meets hardware, browser, flash and Java requirements, should be able to login with a simple URL:
 - `http:\\<ip-address | hostname>`
- Options include:
 - IPv6 communication
 - SSL encryption
- Setup Login Message for Login Page
 - Requires Security Administrator Modify permission

4

Security

- External authentication of privileged users
 - RADIUS and LDAP support
 - Allows organizations to leverage existing authentication data
 - Authentication support for:
 - Active directory
 - LDAP



5

External authentication of privileged users (storage admins) allows for the storage array to integrate into the existing customer security and compliance infrastructure (for example, existing authentication data), generally directory services environments such as Active Directory. This allow customers to provision and de-provision storage management access with existing infrastructure.

GUI

Hitachi Storage Navigator

File Actions Reports Settings View Help

Alert Audit Log Resource Locked

Logged in as: maintenance Logout

Storage Systems

VSP(S/N:53032)

Tasks Reports Components Parity Groups Logical Devices Pools Ports/Host Groups External Storages

Performance Monitor Administration General Tasks Create Host Groups Create LDEVs Add LUN Paths Create Pools Add External Volumes

<< Storage Systems

VSP(S/N:53032)

VSP(S/N:53032)

Hardware Summary

Storage System Name	VSP	Microcode Version	Main	70-01-26-00/00
Storage System Type	VSP		SVP	70-01-24/00
Serial Number	53032		RMI Server	06_01_00
IP Address	10.129.2.137	Total Cache Size		54.00 GB
Contact				
Location				

Edit Storage System

Allocation Summary

Internal/External: Total System Type: Total

Drive Capacity

Allocated (Excluding V-VOL)

Unallocated (Excluding V-VOL)

Reserved

Free

Total

Used DP Pool

Unused DP Pool

Other

970.67 GB [4%]

1.52 TB [8%]

1.54 TB [7%]

2.57 TB [13%]

28.99 GB [1%]

13.21 TB [67%]

19.82 TB

Number of LDEVs

Allocated (Excluding V-VOL)

Unallocated (Excluding V-VOL)

Reserved

V-VOLs

Total

52

76

16

118

262

Management Tasks


- With Storage Navigator, you can perform the following functions:
 - Provision the storage system
 - View and manage the storage system configuration
 - View system Alerts
 - Monitor and tune performance
 - Run Reports
 - Acquire logs for actions and commands performed on the storage system

7

The above list represents some of the common tasks that can be performed as part of day-to-day operations.

Provisioning Tasks

- Supported through SN2 include:
 - LDEVs (Create, Delete, Shred, Edit)
 - Pools (Create, Delete, Expand, Shrink)
 - External Volumes (Add, Delete)
 - Storage Allocation
 - Manage Host Groups (Create, Edit, Delete)
 - Add/Delete LUN, WWN
 - Edit Port Configuration
 - Speed, Topology, Fabric



8

You can connect multiple server hosts of different platforms to one port of your storage system. When configuring your system, you must group server hosts connected to the storage system by *host groups*.

For example, if HP-UX hosts and Windows hosts are connected to a port, you must create one host group for HP-UX hosts and create another host group for Windows hosts. Next, you must register HP-UX hosts to the corresponding host group and also register Windows hosts to the other host group.

Viewing Storage System Information

- Hardware and Storage Allocation Summaries
- Components
- Parity Groups
- LDEVs
- Storage Pools
- Host/Port Groups
- External Storage
- Create and View HTML Reports or CSV file to:
 - Examine a system or verify changes to a system
- Perform Task Management
- License Key Installation



9

Storage Navigator 2 can be used for viewing storage configuration information and downloaded as HTML or CSV reports.

System Alerts

- SN2 can be used to view Service Information Messages (SIM) generated on the storage
- Alerts are related to:
 - Hardware component failure
 - Software configuration or operations issues:
 - Pools capacity thresholds exceeded
 - LDEVs blocked
 - License related issues
 - External storage issues
 - Replication issues

The screenshot shows the 'Alerts' window in Hitachi Storage Navigator. It contains a table of alerts and a 'Detail' pane for the selected alert.

Ref.Code	Error Level	Status	Date
180100	Moderate	Not Complete	2010/11/23 12:41:58
451003	Service	Complete	2010/11/09 11:38:49
EFD000	Serious	Complete	2010/11/04 10:55:54
21D002	Moderate	Complete	2010/11/04 10:55:53
21D042	Moderate	Complete	2010/11/04 10:55:53
DF9003			17:51:51
622006			18:36:37
621006			18:35:13
620006			18:34:48
EFD000			18:43:15
21D002			18:43:14
21D042			18:43:14
307301			18:22:42

Detail	
Ref.Code	21D042
Error Level	Moderate
Status	Complete
Date	2010/11/04 10:55:53
Error Section	Port error
Error Detail	EXTERNAL DEVICE PATH CLOSE
Location	-

Total Rows : 13 Previous Next Close

10

Licensing

■ License Key Attributes

- Delivered in text format in a file with the .plk extension
- License keys have four dimensions:
 - Program Product (PP)
 - Serial Number of the storage system
 - Time duration
 - Storage capacity permitted to be managed in GB or TB

```
*PP=P7018,53030,05012:34U3S655B69D6,"Hitachi Storage Navigator software"  
*PP=P7019,53030,MRMNO)PQGPESRRXSAC8,"Hitachi SNMP Agent software"  
*PP=P7020.53030.W1WXY'Z0QZ021172QDZ,"Java API".....
```

- The license key is actually a 75 character string (it is truncated in this view).
- The storage capacity and time duration are not provided in human-readable format.

11

Storage Navigator 2 must be license key enabled. When the customer or engineer accesses Storage Navigator 2 with no license keys installed, Storage Navigator 2 will open the License Key interface by default. No other Storage Navigator 2 functions will be possible until the license keys have been installed.

The time duration for license keys can be one of three values:

- Permanent – Normal, agreed and long-term licensed usage.
- Temporary – Up to 120 days; used for trial and evaluation projects.
- Emergency – 7 to 30 days; used when a key is needed quickly. This requires a special agreement reached with the customer.

Capacity Definitions and Key Types

Capacity
Definitions:

- Usable Capacity
- Used Capacity
- Unlimited Capacity

License
Time
Duration
Types:

- Permanent
- Emergency – 30 days
- Temporary – 120 days
- Term – 365 Days

12

Two new storage capacity definitions have been created: usable capacity and used capacity.

Usable capacity is calculated based on the real, internal storage and the connected external storage capacity of a VSP. The program products of the BOS are licensed based on *usable capacity*.


Used capacity is the total allocated capacity including any and all replicated copies. That is to say the total capacity of all P-VOLs and all S-VOLs are added together to determine the *used capacity*. Used capacity is the basis for the replication program product licenses .

Some program products will be *capacity free*. When the customer buys the license for a *capacity free* program product, the customer is entitled to use that product functionality against an unlimited capacity.

This information is provided so that you are aware that there may be differences in licensed capacity calculation for different program products within the VSP system and also as compared to older enterprise storage systems.

Whenever a license key expires, the current configuration is retained but no new configuration changes are allowed with related the program product.


GUI



GUI Feature Highlights

Flash-based Interface

- Deploys use of Flex application features and Flash display qualities¹
- Continues some use of legacy Java components²



14

¹ Storage Navigator 2 is a Flex application served by **httpd** on the SVP, and rendered by clients supporting Flash 10 or greater.

² The use of Java is supported by the RMI server on the SVP.

■ DKC Status — Top Window

Hardware Summary
DKC information summary is in this table: FW version, S/N, IP address, etc.

Allocation Summary
DKC capacity allocation information is described in this table.

Hardware Summary	
Storage System Name	ss
Storage System Type	RAID600
Full Serial Number	315
IP Address	10.213.74.143
Contact	S
Location	S
Firmware Version	Main
SVP	70-00-05/00
RMI Server	06_01_00
Total Cache Size	88,000B

Allocation Summary					
Internal/External : Total	System Type : Total				
Drive Capacity	Allocated (Excluding V-VOL)	207.65GB [1%]	Number of LDEVs	Allocated (Excluding V-VOL)	2068
	Unallocated (Excluding V-VOL)	7.19TB [67%]		Unallocated (Excluding V-VOL)	1210
	Reserved (Excluding V-VOL)	- [-%]		Reserved (Excluding V-VOL)	-
	Used THP Pool	- [-%]		V-VOLs	3278
	Unused THP Pool	- [-%]		Total	
	Free	3.35TB [32%]			
	Total	10.60TB			

Other resources belonging to the storage system will have a summary at the top and detail in the form of a table below.

■ Reporting

- Detail Configuration Report, CSV file
- Configuration Report, HTML file
 - Graphic Reports
 - Table Reports

Create Configuration Report

1. Confirm

Enter a name for the task. Confirm the settings in the list and click Apply to add task in Tasks queue for execution.

Task Name: 101006-CreateConfigurationReport (Max. 32 Characters)

Selected Reports	
Report Type	Format
Configuration Reports	HTML
Detail Configuration Reports	CSV
Total: 2	

Go to tasks window for status | Back | Next | Apply | Cancel | Help

Report Configuration

Report Types

- Table Configuration Report
- Table Details
- Table Summary
- Table Trend
- Table Usage
- Table Volume
- Table Capacity
- Table Performance
- Table Health
- Table Security
- Table Backup
- Table Recovery
- Table Migration
- Table Replication
- Table Archiving
- Table Deletion
- Table Purge
- Table Retention
- Table Compliance
- Table Audit
- Table Monitoring
- Table Alerting
- Table Notification
- Table Logging
- Table Troubleshooting
- Table Support
- Table Training
- Table Documentation
- Table Help

Cache Details

Cache Group	Cache Type	Cache Size (MB)	Cache Usage (MB)	Cache Free (MB)	Cache Full (MB)	Cache Error (MB)	Cache Warning (MB)	Cache Critical (MB)	Cache Fatal (MB)
Cache Group 1	Cache Type 1	Cache Size (MB)	Cache Usage (MB)	Cache Free (MB)	Cache Full (MB)	Cache Error (MB)	Cache Warning (MB)	Cache Critical (MB)	Cache Fatal (MB)
Cache Group 2	Cache Type 2	Cache Size (MB)	Cache Usage (MB)	Cache Free (MB)	Cache Full (MB)	Cache Error (MB)	Cache Warning (MB)	Cache Critical (MB)	Cache Fatal (MB)
Cache Group 3	Cache Type 3	Cache Size (MB)	Cache Usage (MB)	Cache Free (MB)	Cache Full (MB)	Cache Error (MB)	Cache Warning (MB)	Cache Critical (MB)	Cache Fatal (MB)
Cache Group 4	Cache Type 4	Cache Size (MB)	Cache Usage (MB)	Cache Free (MB)	Cache Full (MB)	Cache Error (MB)	Cache Warning (MB)	Cache Critical (MB)	Cache Fatal (MB)
Cache Group 5	Cache Type 5	Cache Size (MB)	Cache Usage (MB)	Cache Free (MB)	Cache Full (MB)	Cache Error (MB)	Cache Warning (MB)	Cache Critical (MB)	Cache Fatal (MB)


<<LDEV>>

ECC Group	HDEV#	LDEV Emulation	LDEV Type	Volume Size(Cyl)	Volume Size(MB)	Volume Size(Blocks)
ECC Group 3-2	0	OPEN-V	Regular	62675	58757.81	1.2E+08
ECC Group 3-2	1	OPEN-V	Regular	62675	58757.81	1.2E+08
ECC Group 3-2	2	OPEN-V	Regular	62675	58757.81	1.2E+08
ECC Group 3-2	3	OPEN-V	Regular	62675	58757.81	1.2E+08
ECC Group 3-2	4	OPEN-V	Regular	62675	58757.81	1.2E+08
ECC Group 3-2	5	OPEN-V	Regular	62675	58757.81	1.2E+08

16

GUI Overview

■ GUI Display Components



Drop Down Menus – File, Actions, Reports, Settings, View (magnification of text ... 100%, 150%, or 200%), Help.

Full Help – Full help resource, navigable by topic, and searchable.

Status Bar – Provides status of alerts, audit log space, and resource lock status.

Context Help – Context-sensitive help, typically brings up a quick-reference table.

Resource Tree – Per storage system resources, resource relations and details. All actions can be started from here. You can also access the Performance Monitor and Administration items, for example, user accounts.

Accordion Items – Within *Explorer*, accordion items can be selected to efficiently use screen space.

Resource Summary - Summary of current resource tree object, typically statistics.

Resource Detail – Details of current resource tree object, varies by object.

Resource Actions – Common tasks related to displayed resource are readily available.

GUI Actions

- SN2 Wizards lead users through use case actions

Set or choose options and/or attribute for setting items

Selected items listed in this table

Concatenated actions are described in **Next Task Option**.
Next: Go to next wizard action.
Finish: Go to task confirm window to finish action.

LDEV ID	LDEV Name	Parity Group ID	Drive Type/RPM	RAID Level	Emulation Type
00:00:00	nickname	P01-9	FC/10K	1(2D+2D)	OPEN-V
00:00:64	nickname	P01-9	FC/10K	1(2D+2D)	OPEN-V
00:00:65	nickname	P01-9	FC/10K	1(2D+2D)	OPEN-V

18

Wizards, based on use-case analysis, combine subsequent steps for common tasks. This results in fewer clicks and in intuitive use.

Note: You can exit out of a Wizard at any step.

■ Task Management

- View execution condition, status, and order

Task summary statistics for all submitted tasks, are broken out by status, that is 2 items are Completed.

Task list

Task Name	Status	Type	User Name	Submission Time	Start Time	Completion Time	Auto Delete
100492-Crea...	In Progress (2%)	Create LDEVs	mainten...	2010/04/02 02:09:59	2010/04/02 02:10:00		Enabled
100491-Crea...	Completed	Create User	mainten...	2010/04/01 21:01:06	2010/04/01 21:01:06	2010/04/01 21:01:11	Enabled
100325-Edt...	Completed	Edit Storage...	mainten...	2010/03/25 23:01:52	2010/03/25 23:01:53	2010/03/25 23:01:59	Enabled

The **Task** list is new and reflects the ability to queue tasks and run them when able. This allows one or more users to potentially queue up a number of tasks that can run when able, and do so unattended. Users can confirm task completions at a convenient time.

The task list will be discussed more in Module 4, “Viewing and Managing the Storage System.”

Resource Tree and Summaries

The screenshot displays the Hitachi Storage Navigator interface. On the left is the 'Resource tree' containing 'Storage Systems' and 'VSP(S/N:53005)'. Below this are sections for 'Performance Monitor', 'Administration', and 'General Tasks'. The main area shows the 'Storage Systems' section with 'VSP(S/N:53005)' selected. This displays two summary tables: 'Hardware Summary' and 'Allocation Summary'. The 'Allocation Summary' includes a pie chart and a table of storage capacity usage.

Hardware Summary

Storage System Name	VSP	Microcode Version	Main SVP
Storage System Type	VSP		RMI
Serial Number	53005		
IP Address	10.4.9.2	Total Cache Size	
Contact	VSP Administrator		
Location	VSP Lab		

Allocation Summary

Internal/External : Total System Type: Total

Drive Capacity	Allocated (Excluding V-VOL)	10.00 GB [1%]	Number of LDEV
Unallocated (Excluding V-VOL)	16.54 TB [70%]		
Reserved	Used DP Pool	0.00 MB [0%]	
	Unused DP Pool	2.61 TB [11%]	
	Other	813.27 GB [3%]	
Free	3.63 TB [15%]		
Total	23.59 TB		

20

The Virtual Storage Platform (VSP) Summaries (Hardware and Allocation) are displayed by default upon login.

Selecting the VSP (S/N 53nnn) object displays the summaries on demand.

Hardware Summary

Hardware Summary				
Storage System Name	VSP	Microcode Version	Main	70-01-04-00/00
Storage System Type	VSP		SVP	70-01-01/00
Serial Number	53005		RMI Server	06_01_00
IP Address	10.4.9.2	Total Cache Size		54.00 GB
Contact	VSP Administrator			
Location	VSP Lab			
				Edit Storage System

Item	Description
Storage System Name	Displays the device name of the storage system.
Storage System Type	Displays the model type of the storage system.
Serial Number	Displays the serial number of the storage system.
IP Address	Displays the IP address of the SVP.
Contact	Displays the contact information such as personnel and telephone number where you can inquire about the storage system.
Location	Displays the location of the connected storage system.
Microcode Version	• <i>Main</i>: Displays the version of the microprogram installed in the connected storage system. • <i>SVP</i>: Displays the version of the Storage Navigator installed in the SVP. • <i>RMI Server</i>: Displays the version of the RMI server installed in the SVP.
Total Cache Size	Displays the total size of the cache memory in the storage system.
Edit Storage System	Edits the displayed item in the window.

21

The **Edit Storage System** button allows changes to:


- Storage System Name
- Contact
- Location

Note: These are all user-defined text fields

SVP information that cannot be edited includes:

1. Storage System Type (factory)
2. Serial Number (factory)
3. IP Address (assigned by installation tech)
4. Versions (assigned by SVP/Microcode updates)
5. Total Cache Size (known to SVP)

Allocation Summary



Notes:

- The Internal/External drop-down list can be set to Total, Internal Only, or External Only.
- The System Type drop-down list can be set to Total, Open Only, or Mainframe Only.

The arrows above indicate the following:

The count of **Allocated LDEVs** provisioning **Allocated space**

The count of **Unallocated LDEVs** accounting for **Unallocated space**

The count of **Reserved LDEVs** supporting **Reserved space (V-VOLs)**. V-VOLs indicates the number of V-VOLs that are virtual volumes (LDEVs) and do not consume the actual disk capacity. V-VOLs are currently used for Copy-on-Write Snapshot and Dynamic Provisioning. This appears only when **Internal/External** is **Total**.

Viewing Components

The screenshot shows the Hitachi Storage Navigator 2 interface. The left sidebar contains the Explorer pane with 'Storage Systems' expanded, showing a tree view of VSP(S/N:53005) components including Tasks, Reports, Components, Parity Groups, Logical Devices, Pools, Ports/Host Groups, and External Storages. The main pane displays the 'Components' page for the selected storage system. A summary bar shows 'Number of Controller Chassis' as 1. Below this is a table with columns 'Chassis ID' and 'Chassis Type'. The first row shows 'DKC-0' as a 'Controller Chassis'. A yellow callout box points to this link, stating 'Link drills down to expose available MPs.' An arrow points from this callout to the 'Processor Blades' table, which lists MPB0 and MPB2. The 'Processor Blades' table has columns: MP Blade ID, MP Blade Name, Status, Cluster, and Auto Assignment. The status for both MPB0 and MPB2 is 'Normal'.

MP Blade ID	MP Blade Name	Status	Cluster	Auto Assignment
MPB0	MPB-1MA	Normal	1	Enabled
MPB2	MPB-2MC	Normal	2	Enabled

- Parity Groups and related LDEVs typically are either owned by one processor blade or the other.
- This balances I/O requests across both processor blades and still allows for a change of MP ownership in case of failure.

Viewing Parity Groups

Hitachi Storage Navigator

File Actions Reports Settings View Help

Alert Audit Log Resource Unlocked Logged in as: STompkins Logout

Explorer

Storage Systems

VSP(S/N:53005)

Tasks

Reports

Components

Parity Groups

Logical Devices

Pools

Ports/Host Groups

External Storages

Performance Monitor

Administration

General Tasks

Create Host Groups

Create LDEVs

Add LUN Paths

Create Pools

Add External Volumes

Create Configuration Report

Monitor Performance

more...

<< Storage Systems

Parity Groups

VSP(S/N:53005) > Parity Groups

Summary

Capacity	Internal	Free	3.63 TB
		Total	23.59 TB
	External	Free	0.00 MB
		Total	0.00 MB

Parity Groups

Filter ON OFF Select All Pages Options 1 / 1

Parity Group ID	LDEV Status	RAID Level	Base Emulation Type	Capacity	Number of LDEVs	Drive Type/RPM	Protection	Encryption Key
				Free Total	Unallocated Total			
1-1	Normal	5(7D+1P)	OPEN-V	0.00 GB 1878.82 GB	1 1	SAS/10k	Standard	Disable
1-2	Normal	5(7D+1P)	OPEN-V	0.00 GB 1878.82 GB	1 1	SAS/10k	Standard	Disable
1-3	Normal	5(7D+1P)	OPEN-V	0.00 GB 1878.82 GB	1 1	SAS/10k	Standard	Disable
1-4	Normal	5(7D+1P)	OPEN-V	0.00 GB 1878.82 GB	1 1	SAS/10k	Standard	Disable
1-5	Normal	5(7D+1P)	OPEN-V	0.11 GB 805.21 GB	2 3	SAS/10k	Standard	Disable
1-6	Normal	5(3D+1P)	OPEN-V	0.21 GB 805.21 GB	1 1	SAS/10k	Standard	Disable
1-7	Normal			0.21 GB 805.21 GB	1 1	SAS/10k	Standard	Disable
1-8	Normal			0.21 GB 805.21 GB	0 1	SAS/10k	Standard	Disable
1-9	Normal			0.80 GB 536.80 GB	1 1	SAS/10k	Standard	Disable
1-10	Normal			0.80 GB 536.80 GB	0 1	SAS/10k	Standard	Disable
1-11	Normal	1(2D+2D)	OPEN-V	0.80 GB 536.80 GB	0 1	SAS/10k	Standard	Disable
2-1	Normal	1(2D+2D)	OPEN-V	0.00 GB 536.80 GB	1 1	SAS/10k	Standard	Disable
2-2	Normal	1(2D+2D)	OPEN-V	0.00 GB 536.80 GB	1 1	SAS/10k	Standard	Disable
2-3	Normal	1(2D+2D)	OPEN-V	0.00 GB 536.80 GB	1 1	SAS/10k	Standard	Disable

Selected: 0 of 23

Create LDEVs Format LDEVs Edit Encryption Export

Link exposes related LDEVs.

Clicking the **Parity Group ID** link will display a list of LDEV members

Columns of interest at this point include:

- RAID Level
 - ♦ RAID-1: (2 + 2)
 - ♦ RAID-5: (3 + 1) or (7 + 1)
 - ♦ RAID-6: (6 + 2)
 - ♦ CoW Snapshot: CoW Snapshot volumes
- Protection
 - ♦ SATA-W/V: Write and Verify method on SATA drives
 - ♦ SATA-E: Enhancing method on SATA drives
 - ♦ Standard: Indicates SAS, SSD, or External
- Encryption Key
 - ♦ Key-ID: Shows key ID for the encrypted parity group
 - ♦ Disable: For the non-encrypted parity groups

Viewing Parity Group LDEVs

The screenshot displays the Hitachi Storage Navigator web interface. On the left, the 'Explorer' pane shows a tree view of storage systems. Under 'Storage Systems', 'VSP(S/N:53005)' is expanded, showing 'Parity Groups' and 'Internal'. The 'Internal' group is further expanded, showing a list of LDEVs from 1-1 to 1-10. LDEV 1-2 is selected. The main pane shows the details for LDEV 1-2, including a summary table and a table of LDEVs. A yellow callout box points to the '1-2' link in the Explorer pane, stating 'Link exposes LDEV properties.'

Hitachi Storage Navigator

File Actions Reports Settings View Help | Alert Audit Log Resource Unlocked Logged in as: STompkins Logout

Explorer

Storage Systems

- VSP(S/N:53005)
 - Tasks
 - Reports
 - Components
 - Parity Groups
 - Internal
 - 1-1
 - 1-2**
 - 1-3
 - 1-4
 - 1-5
 - 1-6
 - 1-7
 - 1-8
 - 1-9
 - 1-10

Performance Monitor Administration General Tasks

- Create Host Groups
- Create LDEVs
- Add LUN Paths
- Create Pools
- Add External Volumes
- Create Configuration Report
- Monitor Performance

1-2

VSP(S/N:53005) > Parity Groups > Internal > 1-2

Summary

LDEV Status	Normal	Drive Type/RPM	SAS/10k
RAID Level	5(7D+1P)	Interleaved Parity Groups	1-2
Capacity	Free: 1.75 MB	Number of LDEVs	Unallocated: 1
	Total: 1.83 TB		Total: 1

LDEVs

LDEV ID	LDEV Name	Status	Emulation Type	Individual Capacity	Attribute
00:00:00:C		Normal	OPEN-V CVS	1879.82 GB	-

Selected: 0 of 1

Create LDEVs Edit LDEVs Format LDEVs More Actions

Tip: When you list resources in Explorer, as shown above for Parity-Group LDEVs, clicking on the next Parity Group will display its LDEVs. You can sequentially click through the LDEV listings per Parity Group without as much navigation.

Viewing Parity Group LDEV Detail

LDEV Properties

LDEV Properties		
LDEV ID		00:00:02
LDEV Name		
Parity Group	ID	5-4
	Interleaved Parity Groups	-
	RAID Level	1(2D+2D)
	Drive Type/RPM	SAS/10k
	Protection	Standard
Pool	Encryption	Disabled
	Name (ID)	-
	RAID Level	-
Type		-
Emulation Type		OPEN-V CVS
Capacity		536.80 GB
Individual Capacity		536.80 GB
Provisioning Type		Basic
Status		Normal
Attribute		-
Command	Security	-
Device	User Authentication	-
Attribute	Device Group Definition	-
UUID		
CLPR		00:CLPR0
Access Attribute		Read/Write
SSID		0004
Cache Mode		-
Page Status		-
Current MP Blade ID		MPB0
Current MP Blade Name		MPB-1MA
Assigned MP Blade ID		MPB0
Assigned MP Blade Name		MPB-1MA

26

Use the LDEV Properties screen when you really, *really* want to know everything about an LDEV.

Viewing Logical Devices

Hitachi Storage Navigator

File Actions Reports Settings View Help | Alert Audit Log Resource Unlocked Logged in as: STompkins Logout

Explorer

- Storage Systems
 - VSP(S/N:53005)
 - Tasks
 - Reports
 - Components
 - Parity Groups
 - Logical Devices**
 - Pools
 - Ports/Host Groups
 - External Storages
- Performance Monitor
- Administration
 - General Tasks
 - Create Host Groups
 - Create LDEVs
 - Add LUN Paths
 - Create Pools
 - Add External Volumes
 - Create Configuration Report
 - Monitor Performance

Storage Systems

Logical Devices

VSP(S/N:53005) > Logical Devices

Summary

Number of LDEVs	Open Allocated	1	Number of LDEVs	Mainframe Allocated	0
	Open Unallocated	25		Mainframe Reserved	0
	Open Reserved	5	Total Number of LDEVs		34
	Open V-VOLs	3			

LDEVs

Filter: ON OFF Select All Pages Options 1 / 1

LDEV ID	LDEV Name	Status	Parity Group ID	Pool Name(ID)	RAID Level	Emulation Type	Capacity	Provisioning Type	Attribute	Number Paths
00:00:00		Normal	1-1	-	5(7D+1P)	OPEN-V CVS	1878.82...	Basic	-	
00:00:01	krispy k...	Normal	5-10	-	6(6D+2P)	OPEN-V CVS	10.00 GB	Basic	-	
00:00:02		Normal	5-4	-	1(2D+2D)	OPEN-V CVS	536.80 GB	Basic	-	
00:00:03		Normal	5-3	-	1(2D+2D)	OPEN-V CVS	536.80 GB	Basic	-	
00:00:04		Normal	5-2	-	1(2D+2D)	OPEN-V CVS	536.80 GB	Basic	-	
00:00:05		Normal	5-1	-	1(2D+2D)	OPEN-V CVS	536.80 GB	Basic	-	
00:00:06		Normal	5-12	-	6(6D+2P)	OPEN-V CVS	0.04 GB	Basic	-	
00:00:07	ST1	Normal	5-8	-	5(3D+1P)	OPEN-V CVS	0.21 GB	Basic	-	
00:00:08		Normal	5-9	-	6(6D+2P)	OPEN-V CVS	1610.42...	Basic	-	
00:00:0C		Normal	1-2	-	5(7D+1P)	OPEN-V CVS	1878.82...	Basic	-	
00:00:0D		Normal	1-3	-	5(7D+1P)	OPEN-V CVS	1878.82...	Basic	-	
00:00:0E		Normal	1-4	-	5(7D+1P)	OPEN-V CVS	1878.82...	Basic	-	
00:01:00		Normal	5-12	-	6(6D+2P)	OPEN-V CVS	20.00 GB	Basic	-	
00:01:01		Normal	5-12	-	6(6D+2P)	OPEN-V CVS	20.00 GB	Basic	-	
00:01:02		Normal	5-12	-	6(6D+2P)	OPEN-V CVS	20.00 GB	Basic	-	

Selected: 1 of 34

Create LDEVs Add LUN Paths Edit LDEVs More Actions

27

Clicking the **LDEV ID** link will display a panel of LDEV details.

Columns of interest at this point include:

- **LDEV Name** – User assigned during creation to indicate owner, application or some other good idea
- **Provisioning Type**
 - ♦ **Basic**: Basic volume (internal volume).
 - ♦ **DP**: V-VOLs of Dynamic Provisioning.
 - ♦ **External**: External volume.
 - ♦ **CoW Snapshot**: CoW Snapshot volumes

- Attribute
 - ♦ **Command Device:** Command Device.
 - ♦ **System Disk:** System Disk.
 - ♦ **Journal Volume:** Journal Volume
 - ♦ **Pool Volume:** Pool volume. The number in the parenthesis shows the pool ID.
 - ♦ **Reserved Volume:** Reserved Volume.
 - ♦ **Reserved Volume for BC:** Reserved volume of ShadowImage.
 - ♦ Volume other than the above ones.

This topic will be covered in more detail in an upcoming module.

Viewing Storage Pools

Hitachi Storage Navigator

File Actions Reports Settings View Help | Alert Audit Log Resource Unlocked Logged in as: STompkins Logout

Explorer

- Storage Systems
 - VSP(S/N:53005)
 - Tasks
 - Reports
 - Components
 - Parity Groups
 - Logical Devices
 - Pools**
 - Ports/Host Groups
 - External Storages

Performance Monitor

Administration

General Tasks

- Create Host Groups
- Create LDEVs
- Add LUN Paths
- Create Pools
- Add External Volumes
- Create Configuration Report
- Monitor Performance

Pools

VSP(S/N:53005) > Pools

Summary

DP	Pool Capacity (Used/Total)	0.00 MB / 2.61 TB [0 %]
	Remaining Configurable Pool Capacity	3986.04 TB
	DP Licensed Capacity	Unlimited
CoW Snapshot	V-VOL Capacity (Total/Remaining Configurable)	1.02 TB / 15006.50 TB [1 %]
	Pool Capacity (Used/Total)	0.00 MB / 805.00 GB [0 %]
	Remaining Configurable Pool Capacity	30.86 TB
	SI/CoW Snapshot Licensed Capacity	Unlimited
Number of Pools		3 (Max Allowed: 128)

Pools

Filter: ON OFF Select All Pages Options 1 / 1

Pool Name (ID)	Status	Number of Pool VOLS	Capacity			User-Defined Threshold (%)	Subscription (%)		Pool Type	Tier Management
			Total	Used	Used (%)		Current	Limit		
vsp_pool...	Normal	2	1605.84...	0.00 GB	0	70	65	Un...	DP	-
stone.br...	Normal	1	805.00 GB	0.00 GB	0	80	-	-	CoW Snaps...	-
STPool(2)	Normal	2	1067.88...	0.00 GB	0	70	-	Un...	DP	-

Selected: 0 of 3

Create Pools Create LDEVs Expand Pool More Actions

28

Clicking the **Pool Name** link will display a list of LDEV members and related details.

Columns of interest at this point include:

- Pool Type
 - ♦ **DP**: Indicates Dynamic Provisioning or Dynamic Tiering storage
 - ♦ **CoW Snapshot**: Indicates *Copy-on-Write Snapshot* storage
- Tier Management
 - ♦ **Auto** or **Manual**: In combination with *DP* above, indicates that Dynamic Tiering is enabled
 - ♦ In combination with *DP* above, hyphen indicates that Dynamic Tiering is *not* enabled.

This topic will be covered in more detail in a later module.

Viewing Ports and Host Groups

The screenshot shows the Hitachi Storage Navigator interface. The left sidebar contains the Explorer menu with options like Storage Systems, Tasks, Reports, Components, Parity Groups, Logical Devices, Pools, Ports/Host Groups (selected), and External Storages. The main content area is titled 'Ports/Host Groups' and shows a summary table and a detailed table of host groups.

Summary Table:

Summary				
Number of Ports	Target	12	Initiator	0
	RCU Target	2	External	2
	Total	16		

Host Groups Table:

Port ID	Host Group Name	Host Mode	Port Security
CL1-A	1A-G00 (00)	00[Standard]	Enabled
CL1-A	Sample HSD...	2C[Window...]	Enabled
CL3-A	3A-G00 (00)	00[Standard]	Disabled...
CL5-A	5A-G00 (00)	00[Standard]	Disabled...
CL7-A	7A-G00 (00)	00[Standard]	Disabled...
CL1-B	1B-G00 (00)	00[Standard]	Disabled...
CL3-B	3B-G00 (00)	00[Standard]	Disabled...
CL5-B	5B-G00 (00)	00[Standard]	Disabled...
CL7-B	7B-G00 (00)	00[Standard]	Disabled...
CL2-A	2A-G00 (00)	2C[Window...]	Disabled...
CL4-A	4A-G00 (00)	00[Standard]	Disabled...
CL6-A	6A-G00 (00)	00[Standard]	Disabled...
CL8-A	8A-G00 (00)	00[Standard]	Disabled...
CL2-B	2B-G00 (00)	00[Standard]	Disabled...
CL4-B	4B-G00 (00)	00[Standard]	Disabled...
CL6-B	6B-G00 (00)	00[Standard]	Disabled...

Selected: 0 of 17

Buttons: Create Host Groups, Add LUN Paths, Add Hosts, More Actions

29


Clicking the **Port ID** link will display a list of **Host Group Names** with access to all details.

Clicking the **Host Group Name** link will display **HBA WWN** addresses, **LUNs**, and **Host Mode** options.

This topic will be covered in more detail in an upcoming module.

Using Ports and Alert Link

- Alert link launches Java legacy component.



30

Alerts tab is displayed by default, but ports are also viewable.

Creating and Viewing Reports

Hitachi Storage Navigator

File Actions **Reports** Settings View Help

Alert Audit Log Resource Unlocked Logged in as: STompkins Logout

Explorer

Storage Systems

- VSP(S/N:53005)
 - Tasks
 - Reports**
 - Components
 - Parity Groups
 - Logical Devices
 - Pools
 - Ports/Host Groups
 - External Storages

Performance Monitor

Administration

General Tasks

- Create Host Groups
- Create LDEVs
- Add LUN Paths
- Create Pools
- Add External Volumes
- Create Configuration Report**
- Monitor Performance

<< Storage Systems

Reports

VSP(S/N:53005) > Reports

Summary

Number of Reports 1(Max Allowed:20)

Reports

Filter ON OFF Select All Pages Options 1 / 1

Report Name	User Name	Submission Time
101011-CreateConfigurationReport	STompk...	2010/10/11 12:03:50

Click link to list all HTML reports

You can generate a report from three locations.

Selected: 0 of 1 Create Configuration Report Download Reports Delete Reports Export

31

Creating a Configuration Report

Perform these steps:

1. Use the **Reports** (drop-down) menu.
2. From **General Tasks**, click **Create Configuration Report** wizard.
3. In **Reports** (resource tree), look for the button.

Up to 20 reports will be listed in the **Reports** tab; click the **Report Name** link to access all HTML sub-reports.

Tip: Sort or filter by username to locate your reports.

Create Configuration Report Screen

1. Confirm

Enter a name for the task. Confirm the settings in the list and click Apply to add task in Tasks queue for execution.

Task Name: (Max. 32 Characters)

Report Type	Format
Configuration Reports	HTML
Detail Configuration Reports	CSV

Total: 2


☐ Go to tasks window for status Back Next Apply Cancel Help

32

HTML can be viewed from Storage Navigator 2.

CSV requires saving a compressed file and accessing it when you want to find the CSV files.

HTML Table Reports



Configuration Reports

Report Types

- Storage System Summary**
- Physical View
- Cache Memories
- Channel Adapters
- Ports
- Host Groups
- Hosts
- LUNs
- Logical Devices
- Logical Device Expansions
- Parity Groups
- MP Blades
- MP Blade Details
- Disk Adapters
- Spare Drives

Storage System Summary

This report shows a summary of the storage system.

Storage System Type
VSP

Serial Number
53005

IP Address
126.109.101.15

Micro Program Versions

Main	7000530000
SVP	70005300
SSVP	70000800
HTP	700039
DKAF	700109
FCDG	700005
CUDG4	700023
RAM BOOT	700002
ROM BOOT	700003
CMBK	702155
Config	70006200
HDD	DKS5B-J300SS : 4F09
HDD(SVP)	DKS5B-J : 004F09
	DKS5C-J : 005F05
	DKS5B-K : 004F59
	DKS5A-H : 005F02
	SDT2A-S : 000044
	SDT5A-S : 000044
SSW	000063
SCHIP	DKS5B-J300SS : -
Battery	702155
Exp	H05476
CMSSD	
Printout Tool	70-00-53/00

33

← Graphical Report Icon

1. The **Storage System Summary** is displayed by default.
2. In a multi-tab browser, it may be possible to right click and open multiple reports in new tabs.
3. Multi-tab browsers include IE 8 and Firefox 3.6.10, and others.
4. The **Physical View**, **Cache Memories**, and **Channel Adapters** are graphical reports.

HTML Graphical Reports

Configuration Reports

Report Types

- Storage System Summary
- Physical View
- Cache Memories
- Channel Adapters
- Ports
- Host Groups
- Hosts
- LUNs
- Logical Devices
- Logical Device Expansions
- Parity Groups
- MP Blades
- MP Blade Details
- Disk Adapters
- Spare Drives

Physical View **Front View** Back View

This report shows disk controllers and disk units, and includes cache PCBs, MP blades, channel adapters, disk adapters, data drives, and spare drives.

< Front View > Storage System Type: VSP, Serial Number: 53005, Microcode Version = 7000530000

Module#0

DKU-01

HDU-015

HDU-014

HDU-011

HDU-010

HDU-005

HDU-004

HDU-001

HDU-000

DKC0

Not Installed

Not Installed

CACHE-2CB

CACHE-2CC

Not Installed

MPB-2MC

MPB-1MA

Not Installed

CACHE-1CA

CACHE-1CB

Not Installed

Not Installed

RK-00

DKU-04

Not Installed

DKU-07

Not Installed

DKU-03

Not Installed

DKU-06

Not Installed

DKU-02

Not Installed

DKU-05

Not Installed

RK-01

RK-02

Drive Type Legend

- SAS
- SATA
- SSD
- Spare
- Not Installed

34

The legend indicates drive type as follows:

- **SAS** stands for Serial Attached SCSI
- **SATA** stands for Serial ATA
- **SSD** stands for Solid State Device

Notes:

1. The DKC0 **Front View** displays installed Cache Memory and Processor (MP) blades.
2. There is no physical cabling in front.

Configuration Reports

Report Types

- [Storage System Summary](#)
- [Physical View](#)
- [Cache Memories](#)
- [Channel Adapters](#)
- [Ports](#)
- [Host Groups](#)
- [Hosts](#)
- [LUNs](#)
- [Logical Devices](#)
- [Logical Device Expansions](#)
- [Parity Groups](#)
- [MP Blades](#)
- [MP Blade Details](#)
- [Disk Adapters](#)
- [Spare Drives](#)

Physical View [Front View](#) [Back View](#)

This report shows disk controllers and disk units, and includes cache PCBs, MP blades, channel adapters, disk adapters, data drives, and spare drives.

< Back View > Storage System Type: VSP, Serial Number: 53005, Microcode Version = 7000530000

Module#0

The diagram shows the physical layout of the storage system. It includes six large gray blocks labeled DKU-07, DKU-04, DKU-01, DKU-06, DKU-03, and DKU-05, each with the text 'Not Installed'. Below these are two smaller gray blocks labeled RK-02 and RK-01, also with 'Not Installed'. To the right, there are several smaller blocks: HDU-017, HDU-016, HDU-013, HDU-012, HDU-007, HDU-006, HDU-003, HDU-002, and DKC0. The DKC0 block shows installed components: CHA-2QU, DKA-2MU, DKA-1AU, and CHA-1EU. A legend at the bottom indicates: SAS (blue), SATA (yellow), SSH (pink), Spare (green), and Not Installed (gray).


35

The DKC0 **Back View** indicates installed Channel Adapters (CHA) and Disk Adapters (DKA).

All physical cabling is in back (cabling is not shown in the report).

CSV Text Reports

- Generating a CSV Report
 - Select report
 - Click **Download Reports**
 - Save file




36

- You will get a **Ready to Download** message, click **OK**.
- When you save, you will see a **Downloading** message.
- When it disappears, your file is ready.

Also:

1. You may select only one report at a time to download.
2. For **Delete Reports** you *can* select multiple reports and delete them as a single task.

■ Saving a Compressed File




37

Save with default file name, or your own name.

■ Saved File

- Extract file to folder; drill down; open *AllConf.csv*
- Smaller individual files are also available



The diagram illustrates the process of saving and opening a CSV file. It starts with a 'Report_10101' icon, which is extracted to a folder named 'dkc200'. From this folder, the file 'AllConf.csv' is accessed. This file is then opened in Microsoft Excel, showing a spreadsheet with the following data:

	A	B	C	D	E	F	G	H	I	J	K	L	M	N
1	<<Micro Version>>													
2	DKCMAIN	HTP	DKAF	SVP	SSVP	FCDG	CUDG4	ROM BOO	RAM BOO	CMBK	Config	HDD	HDD(SVP)	SSW
3	7E+09	700102	701009	70010100	70010200	700006	700024	700003	700002	702255	70010300	DKS5B-J3C	DKS5B-J : I	63
4													DKS5C-J : 005F06	
5													DKS5B-K : 004F59	
6													DKR2C-H : 00A31H	
7													DKS5A-H : 005F03	
8													SDT2A-S : 000048	
9													SDT5A-S : 000048	
10	<<Subsystem ID>>													
11	DEV#	Star	DEV#	End	SSID									
12	0:00:00	00:00:3f			4									
13	0:00:40	00:00:7f			4									
14	0:00:926	00:00:bf			4									
15	00:00:c0	00:00:ff			4									
16	0:01:00	00:01:3f			5									
17	0:01:40	00:01:7f			5									
18	0:00162	00:01:bf			5									
19	00:01:c0	00:01:ff			5									

38

AllConf.csv concatenates all *.csv data into one file.

Extract to a folder to avoid polluting your current location.

Module Summary

- Identified the key features and functions of Storage Navigator 2
- Explained Storage Navigator 2 components on the SVP

Module Review

1. What are the common provisioning management tasks?
2. From which three locations can actions be started?
3. Configuration Reports can be generated to which file types?
4. Name the software components that a client must have installed to successfully access and use Storage Navigator 2.


4. Hitachi Device Manager

Module Objectives

- Upon completion of this module, you should be able to:
 - Describe Hitachi Device Manager software features and functions

2

Device Manager



Hitachi Storage Command Suite:

BOS (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 now uses the Device Manager configuration information rather than polling all the storage systems, or requiring the 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/or application.

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

Overview

- Provides a single platform for centrally managing, configuring, and monitoring Hitachi storage systems
- Helps raise storage management efficiency in these environments and reduce costs
- Presents a logical view of storage resources, while maintaining independent physical management capabilities
- Allows administrators to precisely control all managed storage systems
- Helps automate entire storage environments

4

Hitachi Device Manager software provides a single platform for centrally managing, configuring, and monitoring Hitachi storage systems. By significantly boosting the volume of storage that each administrator can manage, the single-point-of-control design of Device Manager software can help raise storage management efficiency in these environments as well as reduce costs. Easy-to-use Device Manager software logically views storage resources, while maintaining independent physical management capabilities. By offering a continuously available view of actual storage usage and configuration, Device Manager software allows administrators to precisely control all managed storage systems.


This results in highly efficient use of administrative time and storage assets. When combined with other Hitachi Command Suite products, Device Manager helps automate entire storage environments.

Benefits


- Reduced total cost of ownership made possible through consolidation and simplification
- Increased storage utilization
- Higher data availability delivered by reducing storage administrator errors and the time it takes to deploy new capacity for critical applications
- Real-time management and automation
- Investment protection in ISV management software (open API)
- Multiple levels of secure access to multiple storage administrators

5

Hitachi Storage Management Tools



Unified Data Center Management Tools



Designed for Hitachi Storage Systems and Sun StoreEdge 99xxV series, Hitachi Device Manager delivers centralized management of distributed Hitachi environments. With an intuitive easy-to-use graphical interface and a complete command line interface (CLI), Device Manager enables logical, physical, and host management views. This allows customers to facilitate provisioning and storage pooling for primary and secondary storage. It also provides an open standards device interface into Hitachi storage systems for the heterogeneous Hitachi Storage Services Manager and other third part storage management tools. Using it, administrators can integrate and manage various models and types of storage systems as a single virtual storage pool, allowing our customers to do more with less.

Overview

- Unified Data Center Management
 - No need to switch to native element managers for routine tasks
- Usability
 - Task oriented GUI
 - Intuitive operations
- Scalability and Performance
 - One management server can manage up to 5 million objects
 - Quick GUI response, CLI performance improvement
- Serviceability
 - Discovers hosts and virtual servers without deploying host agents
 - Integrated and simplified installation

8

Product Features:

- Easy-to-use and consistent intuitive interface used throughout Hitachi Command Suite software modules
- Multiple levels of security for storage administrators
- Enables capacity analysis and reporting based on physical and secondary storage
- Allows managing Hitachi storage via included SMI-S standard provider

Customer Benefits:


- Reduces costly errors associated with manual storage management procedures
- Improves IT productivity through consistent administrative operations across storage platforms anywhere in the enterprise
- Simplified storage centralized management of geographically dispersed storage
- Increased return on investments
- Maximizes performance, reliability, and availability

- **Functionality**


- Storage operations
 - Allocating volumes
 - UnAllocating volumes
 - Creating volumes
 - Virtualizing storage systems (virtualize volumes)
 - Virtualizing storage capacity (HDP pools)
- Managing storage resources
 - Group management of storage resources (logical groups, LDEV level)
 - Searching storage resources and outputting reports
- User management
- Security settings

9

Architecture Overview



Server Components



Device Manager Server

Program that interacts with storage subsystems and hosts based on requests from management clients (Web Client and CLI).

Common Component

A set of common programs/functions used by all Hitachi Command Suite products:

- Web Server
- Error and Log Management
- Database interactions
- Security and User Management

Database

Hitachi Relational Database (HiRDB) stores all the data collected by HDvM and other HCS products.

- Embedded Database – no management interface or functions

Agent

- Components and Functions
 - Includes support for:
 - Hitachi Provisioning Manager
 - Hitachi Replication Manager
 - Hitachi Global Link Manager
 - Optional – depending on required functionality
 - If installed reports the following:
 - Host information
 - Information about volumes allocated to the host
 - Information about file systems allocated to the host
 - Replication information

12

Optional

- Depends on required functionality

Reporting information

- Host
- Host names
- IP addresses
- HBA WWN
- iSCSI name

Volumes allocated to the host

- LDEV number
- Storage subsystem
- LUN
- LDEV type

File systems allocated to the host

- File system types
- Mount points
- Usage

Replication


- Pair types
- Statuses

Graphical User Interface

- Key Features
 - Completely new
 - Easy to learn, easy to use, task-oriented
 - Adobe Flex technology
 - Adobe Flash Player 10.1 required
 - Improved response time

Components — Login Screen

- URL is the same as current version:
[http://**hostname**:23015/DeviceManager/](http://hostname:23015/DeviceManager/)



Alternative content

[Get Adobe Flash player](#)



14

When Flash Player is not installed, page with the link to Adobe Flash Player download site is displayed

When an old version of Flash Player is installed, the message for upgrading to Flash Player 10 is displayed

Layout

The screenshot displays the Hitachi Command Suite 7 interface. The left sidebar shows the 'Resources' section with 'Storage Systems' selected. The main area shows a 'Summary' table and a 'Storage Systems' table.

Summary	
No. of Storage Systems	6
Open-Allocated Capacity	428.00 GB
Open-Reserved Capacity	1.09 TB
No. of Volumes	3782
Open-Unallocated Capacity	6.36 TB
Mainframe-Unspecified Capacity	7.01 TB

Storage System	No. of Volumes	Open-Allocated Capacity	Open-Unallocated Capacity	Open-Reserved Capacity	Mainframe-Unspecified Capacity	Last Refreshed
NSC55@10...	578	98.02 GB	780.69 GB	26.00 GB	1.16 TB	2010-10-01 02...
NSC55@10...	564	198.02 GB	582.69 GB	0 MB	1.16 TB	2010-10-01 00...
USP_VM@10...	663	26.00 GB	1.07 TB	401.67 GB	1.16 TB	2010-10-01 02...
USP_VM@10...	665	34.00 GB	1.06 TB	401.67 GB	1.16 TB	2010-10-01 00...
USP_VM@10...	658	4.00 GB	1.36 TB	200.83 GB	1.16 TB	2010-10-01 00...
USP_VM@10...	654	67.95 GB	1.52 TB	90.00 GB	1.16 TB	2010-10-01 00...

15

Task Oriented

- Settings categorized as *Basic* and *Advanced*
- Only necessary information displayed
- Minimizes administrator input
- Tasks
 - All operations for storage configuration are managed as *tasks*
 - Creation and execution can be separated
 - User can continue next operation without waiting for task completion
- Scheduler
 - Capability to set a schedule for the task

The 'Allocate Storage' dialog box is shown. It includes fields for Host, Volume Size, Number of Volumes, Storage System, Volume Type, and Drive Type. There are also sections for 'Advanced Option' and 'Path Option' with a table for storage paths.

Host	HBA	Port	CHA	Subsystem
HOST_IC5_1	11.22.33.44...	CL7-B	CHA...	USP_V@172...

16

Tasks can be scheduled.

Allocate Volumes
Verify the following plan for allocating volumes.

Plan Summary

Host	WLABA
Storage System	USP_VM@10.2.24.10
Volumes	1
Volume Capacity	1.00 GB
Volume type	Basic Volume
Volume Location	Internal
Drive Type	FC
RAID Level	RAID5(3D+1P)
Drive Speed (RPM)	15000
Parity Group/DP Pool	1-1-1

Plan Details

Task name: * Allocate-1

Description:

Schedule

☐ Now

☒ Later:

Date: 2010-10-01

Time: 14:15

☐ Display task status

* Required

Back Allocate Cancel Help

17

Configuration Overview

Configure Server

- **Install Hitachi Device Manager Server**
 - Install Device Manager Server
 - Install Device Manager Agents
- **Configure Server**
 - License Key Installation
- **Administration Operations**
 - Register Managed Resources
 - Set up Users, Permissions
 - Set up Security Policies
 - Set up SMTP

18


Provisioning Operations

- **Set Up Virtualized Storage Environment**
 - Configure Dynamic/Thin Provisioning Volumes
 - Configure External Volumes
 - Create Volumes
- **Allocate Storage**
 - Allocate Storage
 - Unallocate Storage
- **Group Storage**
 - Create Logical Groups
- **Monitor and Optimize Storage**
 - Alerts
 - Task Management

19

Command Line Interface


- Perform device manager operations
 - Issue commands from CLI client interactively
 - Use a script
- Communicates with and runs as client of device manager server
- In-Band CLI
 - Communicates to storage system through command device
 - Included with HDvM



20

Register Storage System

■ Add Storage System



21

Storage systems need to be registered to monitor and manage resources such as Volumes, Parity groups, HDP pools, and External storage systems from Hitachi Command Suite (HCS).

Before registering a storage system, verify that the target storage system is connected to the network.

In addition, determine, as appropriate:

- The target storage system type (device type)
- The IP address (IPv4 or IPv6) for the target storage system
- The user name and password of the storage system

You must have administrator permission to perform this task.

■ Add Storage System (cont')

Enterprise Storage

Add Storage System

To add a storage system to the device manager, a storage system's type, IP address, and the login username and password must be specified.

Storage System Type: VSP

IP Address:

User Name:

Password:

☐ Go to Data Collection Task Status

OK
Cancel
Help

Modular Storage

Add Storage System

To add a storage system to the device manager, a storage system's type, IP address (Controller 0) must be specified. All other fields are optional.

Storage System Type: AMS/WMS/SMS

IP Address (Controller 0): *

IP Address (Controller 1):

Protocol: Any

User Name:

Password:

☐ Go to Data Collection Task Status

OK
Cancel
Help

* Required field

22

Before registering a storage system, verify that the target storage system is connected to the network and the IP addresses for management interface.

Data Collection Tasks

- Check status of *Add Storage System* task

Type	Target	Target Type	User Name	Status	Start Time	Completion Time	Elapsed Time
Refresh Stor...	AMS2300@10...	AMS2300	system	✓ Completed	2010-08-31 ...	2010-08-31 ...	00:02:57
Add Hosts	10.2.25.40-1...	Windows	system	✓ Completed	2010-08-31 ...	2010-08-31 ...	00:00:09
Add Storage ...	10.2.24.30, ...	AMS/WMS/SMS	system	✓ Completed	2010-08-31 ...	2010-08-31 ...	00:03:29
Add Hosts	10.17.105.3-...	Windows	system	✓ Completed	2010-08-31 ...	2010-08-31 ...	00:00:08
Add Hosts	10.17.105.3-...	Windows	system	✓ Completed	2010-08-31 ...	2010-08-31 ...	00:00:10
Add Storage ...	10.17.104.57...	AMS/WMS/SMS	system	✗ Failed	2010-08-31 ...	2010-08-31 ...	00:02:33
Add Hosts	10.17.105.3-...	Windows	system	✓ Completed	2010-08-31 ...	2010-08-31 ...	00:00:08
Add Hosts	192.168.1.115	Windows	system	✓ Completed	2010-08-31 ...	2010-08-31 ...	00:00:09
Add Storage ...	10.129.2.137	RAID700	system	✓ Completed	2010-08-31 ...	2010-08-31 ...	00:05:23

Selected: 0 of 9

Delete Tasks

23

The **Data Collection Tasks** tab displays the task status for tasks associated with administration of Storage Systems and Hosts. The tasks include:

- Add Storage System
- Edit Storage System
- Remove Storage System
- Refresh Storage System
- Add Hosts
- Edit Hosts
- Remove Hosts
- Refresh Hosts

These tasks do not get registered in the **Tasks & Alerts** tab.

■ Check details for *Add Storage System* task

Data Collection Task Detail

Task Summary

Type	Refresh Storage Systems	User Name	system
Status	Completed	Progress	100 %
Target	AMS2300@10.2.24.30		
Target Type	AMS2300		
Start Time	2010-08-31 11:38:20		
Completion Time	2010-08-31 11:41:18	Elapsed Time	00:02:57
Message	n/a		

Storage Systems

Filter: ☐ ON ☒ OFF Select all Show/Hide Columns Rows/Page: 30 Page: 1 / 1

Name	IP Address	Storage System Model	Status	Message
AMS2300@10.2.24.30	10.2.24.30, 10.2.24.31	AMS2300	✓ Completed	n/a


Selected: 0 of 1

Close Help

24

Managing Hosts

- Add Hosts using the following:
 - Discover hosts with Host Data Collector (Agentless Discovery)
 - Manually
 - Install HDvM agent on the host



25

Administration > Managed Resources > Hosts screen provides operations for managing hosts

Register a host to monitor and manage it by using Hitachi Command Suite (HCS), along with the other registered hosts and storage systems. To register a host, you can specify the IP address and authentication information so HCS can discover it automatically, or you can add a host manually by specifying the WWN or iSCSI name.

Before registering a host, verify that the target host is connected to the network.

In addition, verify, as appropriate:

- The target host name, when manually registering a host
- The IP address (IPv4 or IPv6) for the target host
- The user name and password for the target host
- The unregistered WWN or iSCSI name, when manually registering a host

Supported host types

Host Type		Description	Required Program
Open hosts	Normal host	Environment in which virtualization software is not installed	• Device Manager server • Host Data Collector • Device Manager agent¹
	Virtualization server	Virtual environment created by using virtualization software	• Device Manager server • Host Data Collector • Device Manager agent¹
	Virtual machine	Physical environment in which virtualization software is installed ²	• Device Manager server • vMA
Mainframe hosts		Computers that use mainframe volumes	• Device Manager server • Mainframe Agent
¹ Required if you want to manage copy pairs or if you want to use Device Manager to manage hosts (OSs) whose volume usage cannot be acquired by Host Data Collector.			
² Physical environment in which Windows Server® (Hyper-V™) is installed is treated as a normal host by Device Manager.			

Host Discovery

- Host Data Collector (HDC) discovers and collects information from hosts using
 - IP address, user ID and password

Add Hosts

To discover hosts automatically, specify their IP addresses, OS type, and other parameters.

☒ Discover Hosts ☐ Add Hosts Manually

Specify Hosts By: ☒ IP Address Range ☐ List of IP Addresses

10.2.25.41 - 10.2.25.50

OS Type: Windows

User ID: Administrator

Password: ***

Host Data Collector: 10.2.25.52

[Advanced Settings](#)

Protocol: ☒ Windows Service Management ☐ Windows Management Instrumentation

☐ Go to Data Collection Task Status


OK Cancel Help

Specify Hosts By: ☐ IP Address Range ☒ List of IP Addresses

e.g. 172.17.77.162, 10.208.17.*, 10.208.15.1-10.208.15.255

Allocate Storage

- Multiple hosts can be connected on one port
- Add LUN security by implementing Host Groups on ports
- Host Groups collect Volumes and Hosts (WWPN) into logical grouping




28

You can connect multiple server hosts of different platforms to one port of your storage system. When configuring your system, you must group server hosts connected to the storage system by **host groups**.

For example, if HP-UX hosts and Windows hosts are connected to a port, you must create one host group for HP-UX hosts and also create another host group for Windows hosts. Next, you must register HP-UX hosts to the corresponding host group and also register Windows hosts to the other host group. The host group also includes the LUN Mappings to the volumes that can be accessed by the host.

When registering a host, you are requested to specify the WWN of the host bus adapter.

■ Launch **Allocate Volumes** from General Tasks



29

Launching **Allocate Storage** from General Tasks or from **Actions** option on the menu bar gives all the options for allocating storage. User will need to configure all the requested information.

You can reduce the amount of information requested by launching **Allocate Storage** from a specific target object, for example, from Volume or Parity Group.

■ Allocate Volumes

Allocate Volumes

To allocate volumes to a host, first select the host and specify the number of volumes, capacity, and characteristics. To define I/O connections between the host and storage system, edit the paths.

Host:

No. of Volumes:

Volume Capacity:

Required capacity: GB

Storage system:

Volume type:

Volume location:

[Advanced Options](#)
[Path Options](#)
[Host Group and LUN Settings](#)

The screenshot shows the 'Allocate Volumes' dialog box. Red arrows point to the dropdown menus for 'Storage system', 'Volume type', and 'Volume location'. The 'Storage system' dropdown shows a list of hosts including AMS500@172.17.42.184, RAID700@172.17.45.30, USP@172.17.41.165, USP_V@172.17.41.57, and USP_V@172.17.41.58. The 'Volume type' dropdown shows 'Any', 'Dynamic/Thin Provisioning', and 'Basic Volume'. The 'Volume location' dropdown shows 'Internal' and 'External'.

30


■ Advanced Options

31

Format Type – Does not appear if the volumes of specified sizes are found in the existing configuration.

Quick Format – Allows a volume to be presented to a host as soon as the volume is created. The format process is executed in the background. Quick Format requires you to configure a system disk (using the Element Manager) before using Quick Format.

Instructor Demonstration



- Register Storage System
- Add Host
- Allocate Storage



Module Summary

- Described Hitachi Device Manager software features and functions

Module Review

1. List three key features of Hitachi Command Suite products.
2. List the components of Device Manager.
3. List five provisioning tasks that are supported through Device Manager.
4. What are the task status displayed in the global monitoring area ?

35

5. Hitachi Global Link Manager Advanced

Module Objectives

- Upon completion of this module, the learner should be able to:
 - Describe Hitachi Dynamic Link Manager software features and functions
 - Describe Hitachi Global Link Manager software features and functions

Hitachi Dynamic Link Manager (HDLM)

Overview

- Server based path management software that provides path failover and load balancing capabilities to provide a high availability SAN environment.
- Provides:
 - Support for Fibre Channel connectivity
 - Automatic path discovery
 - Automatic path failover and failback
 - Two applications: Command Line Interface (CLI) and GUI
 - Integrates with Hitachi Storage Management software


4

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.

Benefits

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



5

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 SCSI 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

- **Enables Multipathing**
 - From host to devices
 - Share I/O workloads
- **Path Failover**
 - Automatically redirects I/O operations to alternate path
- **Failback**
 - Automatically places recovered path back online

6

Multipathing

- From host to devices
 - ♦ Allowing access to the device even if a specific path is unavailable
- Used to share I/O workloads
 - ♦ Improve performance

Path Failover

- Automatically redirects I/O operations to alternate paths if a failure occurs
 - ♦ Processing to continue without interruption
 - ♦ Removes threat of I/O bottlenecks
 - ♦ Protected data paths
 - ♦ Increases performance and reliability

Failback

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

HDLM Features

- **Load Balancing**
 - Intelligently allocates I/O requests across all available paths
 - Ensures continuous operation at optimum performance levels
- **Path health checking**
 - Automatically checks the path status at regular user-specified intervals

7


Load Balancing

- Intelligently allocates I/O requests across all available paths
 - ♦ Prevent a heavily loaded path from adversely affecting processing speed
 - ♦ Ensures continuous operation at optimum performance levels
 - Improved system and application performance

Path health checking

- HDLM automatically checks the path status at regular user-specified intervals
 - ♦ Eliminates the need to perform repeated manual path status checks
 - ♦ Proactive identification of issues


Path Failover and Failback



- 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

Load Balancing




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.

Note: *HDLM stands for Dynamic Link Manager software.

■ Clustered Environment

- Microsoft Windows
- Sun Solaris
- HP-UX
- AIX
- Linux



10

Load Balancing on Cluster

- The Dynamic Link Manager allows load balancing over the clustering system in a safe manner.
- The path failover and failback of the Dynamic Link Manager will work along with cluster's node failover and failback.

Microsoft Windows

- Microsoft Cluster Server
- Oracle RAC
- Veritas Cluster Server

Sun Solaris

- Sun Cluster
- VERITAS Cluster Server
- Oracle RAC

HP-UX

- MC/Serviceguard
- Oracle RAC

AIX

- HACMP
- Veritas Cluster Server
- Oracle RAC

Linux

- Redhat AS Bundle Cluster
- SuSE Linux Bundle Cluster
- Veritas Cluster Server
- Oracle RAC

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.

11

Load Balancing Algorithm	Description
Least I/Os	I/O operations are issued to path that has least number of I/Os being processed*
Extended least I/Os (Default)	I/O operations are issued to path that has least number of I/Os being processed* • For sequential access, a certain number of I/Os are issued to one path in succession – Next path is chosen according to least I/Os algorithm • For random access, I/Os are issued to multiple paths according to the least I/Os algorithm

12


* Regardless of I/O block size

Load Balancing Algorithm	Description
Least blocks	I/O operations are issued to path with least pending I/O block size*
Extended least blocks	I/O operations are issued to path with least pending I/O block size* • For sequential access, a certain number of I/Os are issued to one path in succession – Next path is chosen according to least blocks algorithm • For random access, I/Os are issued to multiple paths according to least blocks algorithm


13

* Regardless of I/O count

Round Robin — Extended Round Robin



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



Status	Path ID	Path Name	Subsystem	Emulation Type	LUN	SLPR	CHA	CLPR	Type
✓	000007	0004.0000.00...	HITACHI.9970/9980.200...	OPEN-V	0000	-	2A	-	Ow
✓	000008	0004.0000.00...	HITACHI.9970/9980.200...	OPEN-V	0001	-	2A	-	Ow
✓	000009	0004.0000.00...	HITACHI.9970/9980.200...	OPEN-V	0002	-	2A	-	Ow
✓	000010	0004.0000.00...	HITACHI.9970/9980.200...	OPEN-V	0003	-	2A	-	Ow
✓	000011	0004.0000.00...	HITACHI.9970/9980.200...	OPEN-V	0004	-	2A	-	Ow
✓	000012	0004.0000.00...	HITACHI.9970/9980.200...	OPEN-V	0000	-	1A	-	Ow
✓	000013	0004.0000.00...	HITACHI.9970/9980.200...	OPEN-V	0001	-	1A	-	Ow
✓	000014	0004.0000.00...	HITACHI.9970/9980.200...	OPEN-V	0002	-	1A	-	Ow
✓	000015	0004.0000.00...	HITACHI.9970/9980.200...	OPEN-V	0003	-	1A	-	Ow
✓	000016	0004.0000.00...	HITACHI.9970/9980.200...	OPEN-V	0004	-	1A	-	Ow
✓	000024	0005.0000.00...	HITACHI.9970/9980.200...	OPEN-V	0000	-	1A	-	Ow
✓	000025	0005.0000.00...	HITACHI.9970/9980.200...	OPEN-V	0001	-	1A	-	Ow
✓	000026	0005.0000.00...	HITACHI.9970/9980.200...	OPEN-V	0002	-	1A	-	Ow
✓	000027	0005.0000.00...	HITACHI.9970/9980.200...	OPEN-V	0003	-	1A	-	Ow
✓	000028	0005.0000.00...	HITACHI.9970/9980.200...	OPEN-V	0004	-	1A	-	Ow
✓	000029	0005.0000.00...	HITACHI.9970/9980.200...	OPEN-V	0000	-	2A	-	Ow
✓	000030	0005.0000.00...	HITACHI.9970/9980.200...	OPEN-V	0001	-	2A	-	Ow
✓	000031	0005.0000.00...	HITACHI.9970/9980.200...	OPEN-V	0002	-	2A	-	Ow
✓	000032	0005.0000.00...	HITACHI.9970/9980.200...	OPEN-V	0003	-	2A	-	Ow
✓	000033	0005.0000.00...	HITACHI.9970/9980.200...	OPEN-V	0004	-	2A	-	Ow

16

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.

For Hitachi modular storage systems, Dynamic Link Manager software selects the switching destination path from owner paths and then from non-owner paths.

For Hitachi enterprise storage systems, all paths that access the same LU are candidates for the switching destination path. All other paths that run through the same path are switched.

Operations

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

```

C:\Program Files\Hitachi\DynamicLinkManager\bin>dlnkmgr view -path
Paths:000008 OnlinePaths:000008
PathStatus IO-Count IO-Errors
Online 205 0

PathID PathName ChaPort Status DskName Type IO-Count IO-Errors DNum HDevName
000000 0006.0000.0000000000000018.0000 HITACHI .OPEN-U Own 16 0 0 -
000001 0006.0000.0000000000000018.0001 HITACHI .OPEN-U Own 14 0 0 -
000002 0006.0000.0000000000000018.0002 HITACHI .OPEN-U Own 12 0 0 -
000003 0006.0000.0000000000000018.0003 HITACHI .OPEN-U Own 10 0 0 -
000004 0007.0000.0000000000000008.0000 HITACHI .OPEN-U Own 15 0 0 -
000005 0007.0000.0000000000000008.0001 HITACHI .OPEN-U Own 46 0 0 -
000006 0007.0000.0000000000000008.0002 HITACHI .OPEN-U Own 46 0 0 -
000007 0007.0000.0000000000000008.0003 HITACHI .OPEN-U Own 46 0 0 -
KAPL01001-I The HDLM command completed normally. Operation name = view
C:\Program Files\Hitachi\DynamicLinkManager\bin>_

```


17

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.

Instructor Demonstration



- GUI
 - View Path Status
 - Change Path Status
- CLI
 - View Path Status
 - Change Path Status



Hitachi Global Link Manager (HGLM)

Multipathing — Problems and Solutions

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


20

Problems and Solutions

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

21



Configuration Overview



22

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



23

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

Path Management — Global view of paths and Hidden Devices (HDEVs) for all Dynamic Link Manager 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 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.



25

Manage the entire **Dynamic Link Manager** software multipathing 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 multipathing environment.
- Adjust the online or offline path status for single or multiple hosts
- Adjust load balancing for individual LUNs

Alert Management

SNMP-based

The screenshot shows the 'Alerts' tab in the Global Link Manager software. The left sidebar contains a navigation menu with options like Resources, My Groups, Administration, Alerts, and Settings. The main area displays a table of alerts with columns for Severity, Mark, Date, Host Name, Category, Component, Message ID, and De. The table lists several critical alerts related to path failures and component errors.

Below the screenshot, a diagram illustrates the alert flow. A 'Management Server' (Global Link Manager Software) is connected to two 'Host' boxes (Dynamic Link Manager software). A 'Path failure' event on the first host triggers a 'Trap' to the server. The second host has 'SNMP ON' and also sends a 'Trap' to the server.

Events sending alert

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


26

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.

SNMP-based

- SNMP setting on each HDLM instance can be configured only through HGLM



27

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

- Provides ability to manage all aspects of HDLM host
 - Host Information View




28

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.

Instructor Demonstration



29

- Discover Servers
- View Status
- View Reports
- Configure LUN Level settings



Module Summary

- Described Hitachi Dynamic Link Manager software features and functions
- Described Hitachi Global Link Manager software features and functions

30

Module Review

1. List some of the path management features offered by HDLM.
2. Round-Robin is the default load balancing algorithm used by HDLM. (true/false)
3. What protocol does HDLM use for alerting HGLM about path errors ?


6. Hitachi Dynamic Provisioning

Module Objectives

- Upon completion of this module, you should be able to:
 - Describe the use and function of Hitachi Dynamic Provisioning (HDP)

2


Virtualization




Hitachi Enterprise Storage Systems provided multiple options for virtualizing storage and capacity.

Capacity Virtualization and Data Mobility

- Hitachi Dynamic Provisioning
 - Supports **capacity virtualization**



- Hitachi Tiered Storage Manager
 - Supports transparent, nondisruptive data volume mobility



4

Hitachi Dynamic Provisioning software uses thin provisioning to simplify and add agility to the application storage provisioning process and can save significant money on storage purchases. This is important for companies faced with ongoing rapid growth of their data storage requirements and escalating storage and storage management expenses. Dynamic Provisioning also can reclaim existing previously allocated storage within the virtualized pool, and on externally attached heterogeneous systems. Also, to improve performance and simplify performance management, Hitachi Dynamic Provisioning software automates optimization of spindle usage and provides dynamic load balancing.

Hitachi Tiered Storage Manager data management software enables IT administrators to easily and interactively match application quality of service requirements to heterogeneous storage assets. Tiered Storage Manager enables easy storage tier creation and the nondisruptive movement of data volumes to match application-driven price, performance and availability characteristics. As the cost, quality of service and functionality demands of application data change over time, data must be migrated between tiers. This is how a tiered storage architecture maintains an optimal match of storage system characteristics and data requirements.

Dynamic Provisioning

- Allows storage to be allocated to an application without actually being physically mapped until it is used.
 - This *as needed* method means storage allocations can exceed the amount of storage physically installed.
 - Physical storage capacity can be added without an application service interruption.
- Provides the following benefits:
 - Reducing initial installation costs because you purchase only the required physical disk capacity at the start
 - Decreasing management expenses and idle time caused by changing the configuration of both the storage system and the host
 - Improving performance by implementing wide striping

5

- Easier management
 - Simpler planning (reactive → strategic management)
 - Avoid or defer tough decisions about LDEV size
 - Control, change, and optimize: capacity, performance, reliability, cost tier
- Naturally balances performance
 - With large pools, there are more spindles than static provisioning.
 - User can scale performance by growing pool.
- Over provisioning capacity
 - Simplify capacity planning
 - Reduce need for *urgent change*
 - Reduce the need for this decision: “Disk X on server Y has run out. We’ll have to re-allocate the whole estate.”

6

- **Wide striping and storage performance**
 - DP software eliminates the need for outside experts to fine tune application I/O performance.
 - Using wide striping techniques, the storage system automatically spreads the I/O load of all applications accessing the common pool of storage across the available spindles.
 - This process eliminates hot spots and optimizes I/O response times, leading to consistently high application performance.


7

Improved performance is one of the key benefits of using Dynamic Provisioned volumes.

What is Dynamic Provisioning?

- Hitachi Dynamic Provisioning
 - 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 using fewer physical disks than formerly would be required.

- 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.




9

Fat Provisioning occurs on traditional storage arrays where large pools of storage capacity are allocated to individual applications, but remain unused (that is, not written to) with storage utilization often as low as 50%.

Thin Provisioning is a mechanism that applies to large scale centralized computer disk storage systems. Thin Provisioning allows space to be easily allocated to servers, on a *just enough* and *just in time* basis.

Over Allocation is a mechanism that allows server applications to allocate more storage capacity than has been physically reserved on the storage array itself. This allows leeway in growth of application storage volumes, without having to accurately predict which volumes will grow by how much. Physical storage capacity on the array is only dedicated when data is actually written by the application, not when the storage volume is initially allocated.

- Provision virtual capacity to hosts and applications
 - Virtual maximum capacity is specified and provisioned.
 - Real capacity is provisioned from dynamic provisioning pool as host writes are received in **42MB Pages**.



10

The illustration shows the allocation, and timing of allocation, of the 42MB pages. Page is a Hitachi Dynamic Provisioning (DP) construct. It is the unit of allocation and is allocated upon receipt of a write for either the first unit of allocation for the DP V-Volume or a write that requires additional allocation due to filling up previously allocated pages.

So the light gray shades represent the allocation, but the data is not written until the first write is received and settled, at which point the gray turns darker gray.

The point is that subsequent 42MB pages are not required to be contiguous and, in fact, will be spread randomly over the real LDEVs within the DP Pool.

A 42 MB page will be striped over a single LDEV within a single parity group that is in the HDP Pool. (That is, the 42 MB stripe will not be spread over all parity groups within the DP Pool unless Parity Group Concatenation is used.) In this case the LDEV is striped over up to four parity groups.

Benefits

- Application Storage Provisioning simpler and faster for administrator
 - Draw from the dynamic provisioning pool without adding physical disks
 - Define an initial virtual volume, often without an increase in volume capacity or change of configuration
 - Add physical storage separately, nondisruptively
- Reduces application outages, saves time, and keeps costs down
- Performance or throughput improvements from pooling resources
- Ease of provisioning
- Ease of management


11

With Hitachi Dynamic Provisioning software, application storage provisioning is much simpler, faster, and less demanding on the administrator.


- To configure additional storage for an application, the administrator can draw from the Dynamic Provisioning pool without immediately adding physical disks.
- When a (virtual) volume of the maximum anticipated capacity is defined initially, the administrator does not have to increase the volume capacity and change the configuration as often.
- Additionally, when more physical storage is needed, the administrator is required only to install additional physical disks to the Dynamic Provisioning disk pool without stopping any host or applications during the process.

This decoupling of physical resource provisioning from application provisioning simplifies storage management, reduces application outages, saves time, and keeps costs down.

Overview




Components



13

How Does It Work?

- The Storage Controller presents virtual volumes to hosts into which it maps storage from a single physical pool as needed.
- Technical Implementation:
 - Logical volume: Virtual volume – similar to a snapshot V-VOL. Logical volumes are assigned Pool Volume storage as needed.
 - Pool Volume: Physical volumes to store the actual data for logical volumes. Uses technology similar to Snapshot Pool volumes.
 - Dynamic Mapping Table: Dynamic access table between logical volumes and pool volumes.



14

Setup and Monitoring


- Dynamic Mapping Table and Shared Memory – SVP
- Configure Pool and HDP Volumes – Storage Navigator program and Hitachi Device Manager
- Monitoring Pool – Storage Navigator program, RAID Manager CCI, Device Manager and Hitachi Tuning Manager
 - Pool information available
 - HDP Pool ID
 - Capacity
 - Amount of free space
 - Thresholds (2 values)
 - LDEV# of HDP Pool volumes
 - RAID level, drive type
 - Virtual volume information available
 - Virtual capacity, capacity consumed from pool
 - Threshold
 - HDP Pool ID

15

RAID Manager CCI stands for Hitachi Command Control Interface (CCI) software.

Thresholds

- Two pool thresholds are set for the pool capacity. One is fixed at 80% and settable between 5% to 95%.
- One V-VOL threshold is for the DP volume. It is variable from 5% to 300%.
- An alert (SIM/SNMP Trap) is reported if the used capacity exceeds the thresholds.



16

For each HDP volume, a single variable threshold is provided.

- Variable threshold to be set between 5% and 300% for each DP V-Volume.
- Represents x% times the current unallocated capacity.
- If free capacity in the DP Pool for this DP V-Volume falls below this threshold (in capacity), then an alert (SIM/SNMP Trap) is reported.

DP Operations

- Steps for implementing Dynamic Provisioning
 - Create DP Pool
 - Add normal LDEVs
 - Open-V, 8GB – 4TB
 - LDEVs can be added dynamically at later point to expand the pool:
 - If DP Pool is running out of free space
 - If you are adding new servers with new capacity requirements
 - Create DP VOLs (Virtual Volumes)
 - 46MB – 4TB each
 - Format is fast.
 - Associate DP VOLs with DP Pool
 - DP VOLs get capacity allocated as 42MB pages from the associated DP Pool.
 - Assign DP VOL to Host
 - Host sees the full capacity.

17


- Impact of replication operations
 - Replicating a normal LDEV to a DP VOL results in full capacity being allocated for the DP VOL from the pool.
 - Replicating a DP VOL to another DP VOL results in used capacity (not DP VOL capacity) being utilized in the destination DP Pool.

18

Zero Page Reclaim — Recovering Empty Space

OS or other utility writes all binary zeroes to data blocks.

DP Zero Page Reclaim recognizes pages as *empty* and returns capacity to the pool as available.



19


It is possible to set the option **Zero Page Reclaim (ZPR)** for each DP VOL. When ZPR is set, the storage system scans the DP Pool for the DP VOL and will discard any 42 MB blocks that contain only binary zeros. Reclaimed space is returned to the DP Pool as available capacity that can be used to satisfy future writes.

This can be a big benefit when *real* LDEVs are migrated to DP VOLs or to a DP Pool. If the *real* volume contained a lot of *non-data* space, then the Zero Page Reclaim option can reclaim that space.

Zero Page Reclaim can help the customer recover significant amounts of storage capacity that are otherwise thought to be *used* when, in fact, no data is recorded.

You can migrate or replicate a *thick* LDEV into a DP Pool. Then you can apply the Zero Page Reclaim logic against those volumes to reclaim the space. The customer will recognize the recovery of physical storage capacity and can delay the next storage purchase. This also delivers the improved back end performance by causing the LUNs to be distributed among more spindles of all the parity groups that make up the DP Pool.

- Creating Dynamic/Thin Provisioning Volumes — Steps
 - Create Pools
 - Create Dynamic/Thin Provisioning Volumes (Create Volumes)



DP Volumes

■ Prerequisites

- Hitachi Dynamic Provisioning license on storage system
- Volumes to be added in the pool should not be in use (by a host or any other software)
- Volume Size: 8GB to 4TB
- Volumes per pool: Max 1024
- Number of Pools: Max 128
- Up to 63,232 volumes per pool
- Page allocation unit: 42MB

21

■ Launch Create Pool from General Tasks

The screenshot displays the Hitachi Dynamic Provisioning (DP) web interface. The left sidebar shows the 'General Tasks' menu with 'Create Pool' highlighted. The main window shows the 'Create Pool' dialog box. The 'Storage system' is set to 'USP_VM@10.2.24.10' and the 'Pool name' is 'GenPool'. The 'Physical storage for pool' section shows a table of 'Selected Parity Groups'.

Parity Group	Drive Type	Drive Speed (RPM)	RAID Level	Total Capacity	Unallocated Capacity	Free Capacity	No. of Available Volumes	Ext. Storage System	Ext. Storage Model
1-6-1	FC	15000	RAID5(...)	401.67 GB	401.67 GB	401.67 GB	1	-	-

Below the table, the 'No. of Parity Groups' is set to 1, and the 'DP Pool capacity' is 397.56 GB. The 'Advanced Options' section includes 'Pool ID' (1), 'Subscription setting' (Disable), and 'Usage rate threshold' (70%).

22

■ Create Pool — Specify Pool Attributes, Parity Group

Create Pool

To create a pool, first specify its name and physical storage.

Storage system:

Pool name:

Physical storage for pool

Parity Group	Drive Type	Drive Speed (RPM)	RAID Level	Total Capacity	Unallocated Capacity	Free Capacity	No. of Available Volumes	Ext. Storage System	Ext. Storage Model
1-6-1	FC	15000	RAID5(...)	401.67 GB	401.67 GB	401.67 GB	1	-	-

Select

No. of Parity Groups: 1

DP Pool capacity: 397.56 GB

Advanced Options

Pool ID: [Show Pool ID Usage](#)

Thresholds

Subscription setting: ☒ Disable ☐ Enable

Usage rate threshold (5% - 95%): Warning Limit

Subscription threshold (10% - 1000%):

Show Plan Cancel Help

23

■ Create Pool — Add Parity Groups to Pool

Select Parity Groups

Select Parity Groups to use as the physical storage for the pool. Note Parity Groups with already allocated volumes are not shown.

Parity Group	Drive Type	Drive RPM	RAID Level	Total Capacity	Unallocated Capacity	Free Capacity	Number of Usable Volumes	CLPR	External Storage System Name	External Storage System Model	Cache Mode
1-5	SAS	10000	RAID6(6D...	1.57 TB	1.57 TB	1.57 TB	1	0	n/a	n/a	n/a

Selected: 0 of 1

Parity Group	Drive Type	Drive RPM	RAID Level	Total Capacity	Unallocated Capacity	Free Capacity	Number of Usable Volumes	CLPR	External Storage System Name	External Storage System Model	Cache Mode
1-6	SAS	10000	RAID6(6D...	1.57 TB	1.57 TB	1.57 TB	1	0	n/a	n/a	n/a

Selected: 0 of 1

Number of Parity Groups: 1

Estimated Pool Size: 1.57 TB

OK Cancel Help

Show Plan Cancel Help

24

■ Show Plan

Create Pool

Verify the following plan for creating a pool.

Plan Summary

Storage System	USP_VM@10.2.24.10	Pool Name	GenPool
No. of Parity Groups	1	DP Pool Capacity	397.56 GB
Drive Type	FC	RAID Level	RAID5(3D+1P)
Drive Speed (RPM)	15000	Volume Location	Internal
Used % threshold (user defined)	70 %	System Threshold	80 %
Subscription Warning	Disabled	Subscription Limit	Disabled
Pool ID	1		

Plan Details

Parity Group Information

Parity Group	Drive Type	Drive Speed (RPM)	RAID Level	Total Capacity	Unallocated Capacity	Free Capacity	Usable Volumes	Ext. Storage System	Ext. Storage Model
1-6-1	FC	15000	RAID5(3...	401.67 GB	401.67 GB	401.67 GB	1	-	-

DP Pool Volume Information

Volume	Capacity	Parity Group
00:05:00	401.67 GB	1-6-1

Task name: * Create-Pool-1

Description:

[Schedule](#)

* Required

Back Create Cancel Help

25

■ Post pool creation tasks

- Use **Create Volume** to create DP volumes with required capacity.
- Use **Allocate Storage** to allocate volumes to hosts.

26

- Expand a pool
 - Allows you to add parity groups to the pool
 - All unallocated capacity on the Parity Group is added to the pool.

Expand Pool
Verify the plan for expanding a pool.

Plan Summary

Storage System	USP_VM@10.2.24.12	Pool Name	
Parity Groups	0 -> 1	DP Pool Capacity	196.34 GB -> 396.70 GB
Drive Type	FC	RAID Level	RAID5(3D+1P)
Drive Speed (RPM)	15000	Volume Location	Internal
Used Threshold 1	5 %	System Threshold	80 %
Subscription Warning	Disabled	Subscription Limit	Disabled
Pool ID	DP 10		

Parity Group Information

Parity Group	Drive Type	Drive Speed (RPM)	RAID Level	Total Capacity	Unallocated Capacity	Free Capacity	Usable Volumes	Ext. Storage System	Ext. Storage Model
1-3-1	FC	15000	RAID5(3D+1P)	401.66 GB	200.83 GB	401.66 GB	5	-	-

DP Pool Volume Information

Volume	Capacity	Parity Group
00:02:06	40.16 GB	1-3-1
00:02:09	40.16 GB	1-3-1
00:02:07	40.16 GB	1-3-1
00:02:05	40.16 GB	1-3-1
00:02:08	40.16 GB	1-3-1

Task name: * Expand-Pool-1
Description:

* Required

Back Expand Cancel Help

27

- Always look at the final plan to check how much capacity will be added.

The screenshot shows the 'Expand Pool' window with the following sections:

Plan Summary

Storage System	USP_VM@10.2.24.12	Pool Name	
Parity Groups	0 -> 2	DP Pool Capacity	196.34 GB -> 433.57 GB
Drive Type	FC	RAID Level	RAID5(3D+1P)
Drive Speed (RPM)	15000	Volume Location	Internal
Used Threshold 1	5 %	System Threshold	80 %
Subscription Warning	Disabled	Subscription Limit	Disabled
Pool ID	DP 10		

Plan Details

Parity Group Information

Parity Group	Drive Type	Drive Speed (RPM)	RAID Level	Total Capacity	Unallocated Capacity	Free Capacity	Usable Volumes	Ext. Storage System	Ext. Storage Model
1-1-1	FC	15000	RAID5(3D+1P)	351.21 GB	347.21 GB	347.21 GB	24	-	-

DP Pool Volume Information

Volume	Capacity	Parity Group
00:09:16	10.00 GB	1-1-1
00:09:08	10.00 GB	1-1-1
00:09:06	10.00 GB	1-1-1
00:09:17	10.00 GB	1-1-1
00:09:01	10.00 GB	1-1-1
00:09:0C	10.00 GB	1-1-1
00:09:0A	10.00 GB	1-1-1
00:09:04	10.00 GB	1-1-1

Task name: * Expand-Pool-1

* Required

Buttons: Back, Expand, Cancel, Help

In the above screenshot we have 347.25GB available as Unallocated capacity of the parity group we have selected to add to the pool, but the actual capacity added is only 240GB. Only 24 Volumes x 10GB of the total unallocated volumes are above the minimum volume size (8GB) to be added in a pool.

■ Using the pools

- Create Thin/Dynamic Provisioning Volumes using **Create Volume** task.

Create Volumes

To create volumes, first specify the number of volumes, capacity, and characteristics.

No. of Volumes:

Volume Capacity: GB

Required capacity: 2.00 GB

Storage System: USP_VM@10.2.24.12

Volume Type: Dynamic Provisioning

Internal/External: Internal

[Advanced Options](#)

Drive Type: FC

Drive Speed (RPM): 15000

RAID Level: RAID5(3D+1P)

Pool: DP 10 [Select Pool](#)

LDEV ID: ☐ auto ☒ manual

Initial LDEV ID: LDCK CU DEV

00 : :

Show Plan Cancel Help

29

Module Summary

- Described the use and function of Hitachi Dynamic Provisioning software

30

Module Review

1. List the advantages of using Hitachi Dynamic Provisioning.
2. The page size used for storage allocation from the pool is _____.
3. Explain the steps for using DP Volumes.
4. What monitoring options are available for monitoring the pool usage?

7. Hitachi Dynamic Tiering


Module Objectives

- Upon completion of this module, you should be able to:
 - Describe how Hitachi Dynamic Tiering is used to optimize tiered storage utilization

2

Tiered Storage and Information Lifecycle Management

- Once created, most data is rarely or never accessed again.
- Economics dictate that such data should be moved from expensive storage tiers to less expensive lower tiers.
- This is the premise and promise of information and data lifecycle management methodologies and tools.
- But what can be moved and how do you move it efficiently?
- The answers have been complex and challenging.




3

Once created, most data is rarely or never accessed again. It should be moved off of your most expensive storage tier to lower, less expensive lower tiers. This is the premise and promise of data lifecycle management methodologies and tools. But how do you determine what can be moved and efficiently move it? The answers to this question often have been complex and problematic, until now.

Classic Provisioning Volumes From Storage Tiers

- Start with drive technology
 - A simple criterion that gives us a sense of cost and performance
- Define tiers
 - SSD = “tier 1”
 - SAS or FC drives = “tier 2”
 - SATA drives = “tier 3”
- Advanced: Use Dynamic Provisioning and define separate pools
 - Ease of provisioning
 - CAPEX deferral and avoidance
 - Performance improvements




4

This animated slide introduces how storage tiers are currently constructed.

Classic Data Mobility

- Two layers of data mobility
 - Infrastructure layer or block level
 - Object layer or file level
- Use virtualization and data mobility tools to move (or re-provision) the volume without host disruption.
- Target pool or tier based on
 - Promotion or demotion
 - Consolidation or migration
 - SLO, performance or cost change
- Automated with policy based management
 - Based upon pre-set SLAs



5

This animated slide shows how data mobility is accomplished in the current tiering environments at the block, then file/content level.

The problem is often the data can only be moved in large pieces (volumes or files) and not at a level that is fine grained enough to get the most out of DLM and tiered storage.

Challenges with Implementing Tiered Storage

- Operational and management issues
 - Planning and design before moving of volumes
 - Movement of volumes to ensure SLAs
 - Manual and intensive process
 - Managing and moving data to take advantage of tiers is operationally complex, time consuming and hard to optimize.
 - Difficulty or inability to classify and segregate data to take advantage of tiering
 - Need to improve storage performance
 - Need to reduce storage media costs

6

Operational Issues


Change Management Logistics

- ♦ Coordination
- ♦ 3-5 weeks of planning
- ♦ 75-100 hours of labor per instance
- ♦ Can be scheduled on weekends only
- ♦ Manually intensive
- ♦ High IT risk

In order to take advantage of Data Lifecycle Management and the economics offered by tiered storage a data classification exercise is required. This is often time consuming and expensive. And even after, large applications often cannot be subdivided and must remain completely on Tier 1 storage.

- Customer has a very large database and wants to:
 - Improve the performance
 - Re-architect the database environment to leverage tiered storage for cost savings
 - Improve storage performance while retaining the cost benefits of a *tiered environment*
- Challenges
 - Professional services associated data classification is expensive and time consuming.
 - Managing tiered storage is operationally expensive.
 - Moving volumes around requires customer to have idle capacity which is expensive.

Hitachi Dynamic Tiering



8

With **Hitachi Dynamic Tiering** software the complexities and overhead of implementing data lifecycle management and optimizing use of tiered storage are solved. Dynamic Tiering software simplifies storage administration by eliminating the need for time consuming manual data classification and movement of data to optimize usage of tiered storage.


Hitachi *Dynamic Tiering* automatically moves data on fine grain pages within Dynamic Tiering virtual volumes to the most appropriate media according to workload to maximize service levels and minimize TCO of storage.

For example, a database index that is frequently read and written will migrate to high performance flash technology while older data that has not been touched for a while will move to slower, cheaper disks.

No elaborate decision criteria are needed; data is automatically moved according to simple rules. One, two or three tiers of storage can be defined and used within a single virtual volume using any of the storage media types available for the Hitachi Virtual Storage Platform. Tier creation is automatic based on user configuration policies, including media type and speed, RAID level and sustained I/O level requirements. Using ongoing embedded performance monitoring and periodic analysis, the data is moved at a fine grain sub-LUN level to the most appropriate tier.

The most active data moves to the highest tier. During the process, the system automatically maximizes the use of storage keeping the higher tiers fully utilized. Hitachi Dynamic Tiering is available on Virtual Storage Platform only.


Hitachi Dynamic Tiering Builds on Dynamic Provisioning



Dynamic Tiering uses Hitachi Dynamic Provisioning technology and thus inherits its advantages of simplified provisioning, capital savings and self-optimizing performance maximization.

HDT — Page Level Tiering

- Different tiers of storage are now in one pool.
- If data becomes less active, it migrates to lower level tiers.
- If activity increases, data will be promoted to a higher tier.
- Since 20% of data accounts for 80% of the activity, only the active part of a volume will reside on the higher performance tiers.



10

The pool contains multiple tiers (not the other way around, as in USP V/HDP or USP/HTSM).

The logical volumes have pages mapped to the pool (same as USP V/HDP). Those pages can be anywhere in the pool on any tier in that pool.

The pages can move (migrate) within the pool for performance optimization purposes (move up/down between tiers).

HDT will try to use as much of the higher tiers as possible. (T1 and T2 will be used as much as possible while T3 will have more spare capacity).

You can add capacity to any tier at any time. You can also remove capacity dynamically. So, sizing a tier for a pool is a lot easier.

Quantity added/removed should be in ARRAY Group quantities.

The first version of HDT (with VSP at GA):


Up to a maximum 3 tiers in a pool

We will start with managing resources in a 3 tier approach. That may mean:

- 1-Flash drives, 2-SAS, 3-SATA or
- 1-SAS(15k), 2-SAS(10k), 3-SATA (or something else)
- The Pool's tiers are defined by HDD type.
- No support for RAID 10
- No external storage supported (in v01)
- No mainframe support in v01

Hitachi Dynamic Tiering Benefits

Automate and Eliminate Complexities of Efficient Tiered Storage Use



■ Solution Capabilities

- Automate data placement for higher performance and lower costs
- Simplified ability to manage all storage tiers as single entity
- Self-optimized for high performance and space efficiency
- Page-based granular data movement for highest efficiency and throughput

■ Business Value

- Significant savings by moving data to lower cost tiers
- Increase storage utilization up to 50%
- Easily align business application needs to the right cost infrastructure

11

With Hitachi Dynamic Tiering software the complexities and overhead of implementing data lifecycle management and optimizing use of tiered storage are solved. Dynamic Tiering software simplifies storage administration by eliminating the need for time consuming manual data classification and movement of data to optimize usage of tiered storage.

Hitachi *Dynamic Tiering* automatically moves data on fine-grain pages within Dynamic Tiering virtual volumes to the most appropriate media according to workload to maximize service levels and minimize TCO of storage.

For example, a database index that is frequently read and written will migrate to high performance flash technology while older data that has not been touched for a while will move to slower, cheaper disks.

No elaborate decision criteria are needed; data is automatically moved according to simple rules. One, two, or three tiers of storage can be defined and used within a single virtual volume using any of the storage media types available for the Hitachi Virtual Storage Platform. Tier creation is automatic based on user configuration policies, including media type and speed, RAID level and sustained I/O level requirements. Using ongoing embedded performance monitoring and periodic analysis, the data is moved at a fine grain page level to the most appropriate tier.

The most active data moves to the highest tier. During the process, the system automatically maximizes the use of storage keeping the higher tiers fully utilized.

How Hitachi Dynamic Tiering Fits into Tiered Storage

- Just like media type, RAID configurations, and speeds in volume design, Hitachi Dynamic Tiering multitier volumes are another way of delivering tailored storage cost and performance service levels.
- Depending on requirements, Hitachi Dynamic Tiering pools and volumes can be configured to optimize for:
 - Maximum performance
 - Balanced performance and cost
 - Minimum cost for lower tiers

Hitachi Dynamic Tiering volumes are another kind of volume in a tiered storage architecture

		Attributes by Tier				Recoverable	Retain
		Avail.	Perf.	Cost (relative)			
Data Tier	Primary Tier	99.999 %	5	\$\$\$\$\$	A		1 mo
	Lower Cost Tier For Business Data	99.99 %	4	\$\$\$\$	A		1-6 mo
CRITICAL	Low Critical	99.9 %	3	\$\$	B		6-36 mo
ARCHIVAL	Archive Tier	99. %	1	\$	C		7-100 yrs

Hitachi Dynamic Tiering volumes deliver superior service levels at lower cost

12

With or without Hitachi Dynamic Tiering, a proper analysis and design of individual customer requirements is needed

The graphic is to illustrate the concept of tiering. The detailed labels and numbers do not mean anything specific in this context.

Hitachi Dynamic Tiering Specifications

- Available on Virtual Storage Platform (VSP)
- HDP and HDT license key installed on storage system
- First version
 - Up to a maximum 3 tiers in a pool
 - Most often a SSD-SAS-SATA hierarchical model
 - The pool's tiers are defined by HDD type.
 - No external storage supported (in v1)
 - Capacity of any tier can be added to the pool. Capacity can also be removed from a pool.
 - Tier management
 - Fills top tiers as much as possible. Monitors I/O references. Adjusts page placement according to trailing 24 hour heat map cyclically (adjustable 1 to 24 hours)
 - Automatic or manual controls are available.
 - Tier management (*migration* up and down tier) is automatic and built into the system's firmware.
 - Hitachi Dynamic Tiering is a unique product (HDP add-on).

13

This applies to the V1 specifications. Same for the next slide.

We will start with managing resources in a 3 tier approach. That may mean:

- 1-SSD, 2-SAS, 3-SATA or
- 1-SAS15k, 2-SAS10k, 3-SATA (or something else)

Note: External is not supported at this point. Tiers are based on HDD type.

Tier management (automatic!) is important. Future enhancements will continue to optimize this area. In V1 we monitor I/O references and *balance* (or migrate) the data across tiers in a cycle. Monitoring essentially goes on 24/7. Following a *cycle*, we begin a background operation to move or relocate pages to better fit the I/O pattern.

- Each tier in a pool has a calculated sustainable workload level.
 - Collectively the average I/O level measurements for pages located in the tier is targeted to not exceed the calculated sustainable I/O level.
 - Some pages may not be moved up-tier if the predicted workload level is to high.
 - Normally the tier's workload level would not be a factor. Most used pages would be located in the highest tiers.
- Page size is not adjustable; it remains 42MB.
- External storage and mainframe support is on the roadmap for a future release.
 - Pools will not share open and mainframe formatted volumes.
 - Page size for mainframe will be unique as well.

14

We also calculate the average sustainable I/O rate of the tier. We avoid over driving a tier as well. In some cases a Tier may not be able to handle all the heaviest used pages, so we may elect to keep some in a lower tier. This is unlikely to be a real problem in most cases, but nevertheless we watch for this outlier case.

We will have reporting in Device Manager and Tuning Manager with utilization performance data that can be used to help:

- Size tiers in pools
- Size pools
- Help chargeback (this is a topic that can get overly complex. Basic representation is to chargeback on quality of service not on specific utilization levels)
- Determine DP VOL spread across the tiers
- Determine IO amounts per tier
- And so on

Finally note that page size stays at 42MB.

HDP versus HDT

Item		Specifications	
		HDP (single tier/pool)	Hitachi Dynamic Tiering (multi-tier/pool)
Functions	Rebalance	Per pool Run when pool capacity is expanded or reduced (new)	Per tier in pool Run only in the concerned tier when capacity is expanded or reduced
	Zero data reclamation	Request by HDP volume Started from GUI and HTSM task Perform rebalance after ZPR	Same Same After ZPR perform rebalance in each tier
Pool	Number	128 (HDP plus Snapshot pools)	Same (shared with HDP and Snapshot)
	Tier	No tier differentiation within a pool	Max 3 Tiers in v1. More planned after v1.
	Max capacity	1.1PB	Same. Increase planned to 3.3 PB as future enhancement.
	Supported media	SSD, SAS15K, SAS10K, SATA Media cannot be mixed	Same HDD supported. Up to a mix of 3 media types per pool. Optional automatic pool expansion with automatic storage selection in future enhancement.
	RAID level	RAID-1, RAID-5, RAID-6 (cannot be mixed)	RAID-5 and RAID-6 (can be mixed).
	External Storage	Possible	External support as a future enhancement.

15

Note that we have Pool capacity reduction in the HDP version. (this is removing a Pool Volume from a pool).

Note: No RAID-1 support exists in Hitachi Dynamic Tiering V1. Also no external pool Hitachi Dynamic Tiering V1 support.

ZPR = Zero Page Reclaim

ZPR examines volumes of physical capacity on or on storage connected to the storage system and returns unused storage blocks back to the storage pool and reclaims storage space

Item		Specifications	
		HDP	Hitachi Dynamic Tiering
HDP Volume	Number	64k	Same
	Capacity	46MB to 60TB (new)	Same. Expandable and shrinkable (new). Concurrent grow/shrink across replication pairs in future enhancement.
	Emulation Type	OPEN-V Mainframe support planned for V2 MF-HDP requires unique pool, Page size is 38MB. ZPR will be automatic.	OPEN-V. Mainframe support as a future enhancement
Conversion between HDP pool and Hitachi Dynamic Tiering pool		Bidirectional (but see restrictions): HDP → Hitachi Dynamic Tiering: Not allowed if HDP pool is RAID-1 or contains external volumes. Hitachi Dynamic Tiering → HDP: Not allowed if Hitachi Dynamic Tiering pool contains multiple media types or a mix of RAID-5 and RAID-6.	

16

Note that we have a way to convert a HDP pool to Hitachi Dynamic Tiering (and back) given the configuration restrictions already mentioned.

Hitachi Dynamic Tiering Operations

Hitachi Command Suite Operations for HDT

- Device Manager allows you to:
 - Create and Delete a HDT Pool
 - Expand an HDT Pool
 - Create and Delete Virtual Volumes
 - View Tier Definition and Usage

18

HDT Operations

- Create HDT Pool

- Device Manager allows you to enable HDT while creating a new pool.

Hitachi Command Suite 7

To create a pool, first specify its name and physical storage.

Storage system: VSP@172.17.45.30

Pool name: HDTPool

Hitachi Dynamic Tiering: ☐ Disable ☒ Enable

Physical storage for pool

Available Tiers

Tier No	Drive Type	Drive Speed (RPM)	
☒	1	SSD	-
☒	2	SAS	15000
☒	3	SAS	10000

Selected Parity Groups

Tier 1	Tier 2	Tier 3							
Parity Group	Drive Type	Drive Speed (RPM)	RAID Level	Total Capacity	Unallocated Capacity	Free Capacity	No. of Available...	Ext. Storage System	Ext. Storage Model
No Data									

No. of Parity Groups: 0 Selected available capacity: 0 MB

DP Pool capacity: 0 MB

Advanced Options

more... Selected: 0 of 6

Allocate Volumes Create Volumes **Create Pool** Expand Pool Delete Pools

If HDT enables the interface, this allows you to select the parity groups to be added to each tier. The **Select Parity Group** button allows you to select the parity groups to include in the tier.

- Create Pool
- Advanced Options
 - Page reallocation
 - Automatic
 - Manual
 - Reallocation cycle
 - [1,2,4,8,24] hrs
 - Monitoring interval
 - Start time
 - End Time

20

Reallocation

- If you select **Auto**, performance monitoring and tier relocation are automatically performed.
- If you select **Manual**, you can manually perform performance monitoring and tier relocation with the CLI commands.

Reallocation Cycle

- When you select **24 Hours** (default value), performance monitoring and tier relocation is performed once a day.
- You can also select 1 or 2 or 4 or 8 hr intervals.

Monitoring Period field

- Specify the start and end time of performance monitoring in 00:00 to 23:59 (default value).
 - ♦ One or more times must be taken between the starting time and the ending time, and the starting time must be before the ending time.
 - ♦ You can view the information gathered by performance monitoring with Storage Navigator/Tuning Manager.

When you select any of **1 Hour**, **2 Hours**, **4 Hours** or **8 Hours**:

- Performance monitoring is performed every hour you selected, starting at 00:00.

Note: When Auto is set and if all the V-VOL pages cannot be completed migrating in one cycle, at the next cycle, the last processed V-VOL will start being migrated with the updated information.

Expand Pool — Add capacity to a Tier

Expand Pool

To expand the pool, first specify the physical storage to add.

Storage System VSP@172.17.45.35

Pool Name DP 0

Pool type HDT

Physical Storage for Pool

Available Tiers

Tier No	Drive Type	Drive Speed (RPM)	Volume Location	RAID Level	Parity Groups	Total Capacity
1	SAS	10000	Internal	RAID5(7D+1P)	4	3.14 TB

Selected Parity Groups

Tier 1 | Tier 2 | Tier 3

Parity Group	Drive Type	Drive Speed (RPM)	RAID Level	Total Capacity	Unallocated Capacity	Free Capacity	No. of Available Volumes	Ext. Storage System	Ext. Storage Model
1-12	SAS	10000	RAID5(3D+1P)	805.21 GB	805.21 GB	805.21 GB	1	-	-

No. of Parity Groups: 1 Selected available capacity: 805.21 GB **Select Parity Group**

No. of Parity Groups to add: 4->5

DP Pool capacity: 3.14 TB->3.92 TB

Used %: 5%->4% (Used Threshold 1: 70%, Used Threshold 2: 80%)

Subscription: 42%->33% (Subscription Warning: -%, Subscription Limit: -%)

Selected: 1 of 4

Allocate Volumes Create Volumes Create Pool **Expand Pool** Delete Pools

Expand Pool allows you to add more capacity to a particular tier in a HDT pool. The Expand Pool dialog allows you to select parity groups that need to be added to the pool.

■ Expand Pool — Add new Tier in the HDT pool

The screenshot displays the 'Expand Pool' dialog box in the Hitachi Dynamic Tiering (HDT) interface. The dialog is titled 'Expand Pool' and contains the following sections:

- Summary:** Shows 'No. of DP Pools', 'No. of DP Vols', and 'No. of DP Pool Vols'.
- DP Pools:** A list of DP Pools (DP 0, DP 1, DP 2, DP 3) with a filter and status indicators.
- Available Tiers:** A table with columns: Tier No, Drive Type, Drive Speed (RPM), Volume Location, RAID Level, Parity Groups, and Total Capacity. Tier 1 is selected.
- Selected Parity Groups:** A table with columns: Parity Group, Drive Type, Drive Speed (RPM), RAID Level, Total Capacity, Unallocated Capacity, Free Capacity, No. of Available Volumes, Ext. Storage System, and Ext. Storage Model. It shows 'No Data'.
- Summary:** Shows 'No. of Parity Groups to add' (4->4), 'DP Pool capacity' (3.14 TB->3.14 TB), 'Used %' (5%->5%), and 'Subscription' (42%->42%).
- Buttons:** 'Allocate Volumes', 'Create Volumes', 'Create Pool', 'Expand Pool' (highlighted in red), and 'Delete Pools'.

Expand Pool also allows you to add new tiers in the HDT Pool.

Whenever new drives are added, Expand Pool will present the option of creating a new tier if the drive configuration is not part of any existing tier in the HDT Pool.

■ Create Volumes

The screenshot displays the 'Create Volumes' dialog box overlaid on the 'DP Pools' table. The dialog box has a title bar 'Create Volumes' and a close button. Below the title bar, it says 'To create volumes, first specify the number of volumes and capacity.' The 'No. of Volumes' is set to 1. The 'Volume Capacity' is set to 10 GB. The 'Required capacity' is 10.00 GB. There is an 'Advanced Options' section with 'LDEV ID' set to 'auto'. The 'Show Plan', 'Cancel', and 'Help' buttons are at the bottom of the dialog.

The 'DP Pools' table shows the following data:

Pool	Name	No. of DP Vols	No. of DP Pools	Capacity	Used Capacity	Free Capacity	Used %	Free %	Pool Type
DP 0		77	4	3.14 TB	176.90 GB	5 %	1.32 TB	42 %	HDT
DP 1		16	4	1.56 TB	1.14 GB	1 %	147.24 GB	9 %	HDT
DP 2		8	2	1.56 TB	0 MB	0 %	100.07 GB	6 %	HDT
DP 3		0	2	507.85 GB	0 MB	0 %	0 MB	0 %	HDP

The 'Selected: 1 of 4' status is shown at the bottom left. The 'Create Volumes' button is highlighted with a red box in the bottom right corner of the interface.

Once the HDT pool is created, you can select the HDT Pool and click **Create Volumes** button to create new volumes in the HDT Pool. These volumes can then be presented to the host, using the **Allocate Volumes** button.

■ Delete Pools

Delete Pool

Error: The pool "[DP 0]" cannot be deleted because a path is set to a volume. Verify the volumes and try again. (KAIC15223-E)

Do you want to delete these pools?

Pool	Pool Name	Total Capacity	Free Capacity	Used Threshold 1	Used Threshold 2	Subscription Warning	Subscription Limit	No. of DP...	Status	Pool Type
DP 0		4.46 TB	2.80 TB	70 %	80 %	Disabled	Disabled	143	Normal	HDT

Task name: * Delete-Pool-3

Description:

☐ Display task status

* Required


Delete Cancel Help

Selected: 1 of 4

Allocate Volumes Create Volumes Create Pool Expand Pool **Delete Pools**

HDT Pools can be deleted if:

- DP Volumes from the Pools are not mapped to a host
- DP Volumes from the Pools are not Used/reserved by any other program product

 HDS Confidential: For distribution only to authorized parties. **Page 7-27**

- HDT versus HDP Pool Detail
 - DP0 — HDT

All Storage Systems > VSP@172.17.45.30 > DP Pools > DP 0

Help

Summary

Actions

Pool	DP 0	Name		Capacity	4.46 TB
No. of DP Vols	143	Subscribed Capacity	2.39 TB	Used Capacity	1.52 TB
No. of DP Pool Vols	7	Subscription	53 %	Used %	34 %

DP Vols

H/W Tier

DP Pool Vols

Parity Groups

Filter

On

Off

Select All

Column Settings

Rows/page: 30

Page: 1 / 5

All Storage Systems > VSP@172.17.45.30 > DP Pools > DP 2

Help

Summary

Actions

Pool	DP 2	Name		Capacity	3.14 TB
No. of DP Vols	10	Subscribed Capacity	12.67 GB	Used Capacity	0 MB
No. of DP Pool Vols	4	Subscription	0 %	Used %	0 %

DP Vols

DP Pool Vols

Parity Groups

Filter

On

Off

Select All

Column Settings

Rows/page: 30

⏪

⏩

Page: 1 / 1

⏪

⏩

Volume	Label	Host	WWN	Port	Host Group	LUN	Capacity	Volume Attribute
--------	-------	------	-----	------	------------	-----	----------	------------------

Note the additional H/W Tier tab on HDT pool.

- Tier Definition
- Tier Utilization (Used %)

27

The information can be used for checking the spread of data on the HDT Pool Tiers. Also the **Used%** can be used for identifying tiers that need additional capacity.

Module Summary

- Described how Hitachi Dynamic Tiering is used to optimize tiered storage utilization

Module Review

1. List the benefits of HDT.
2. What is the page size in an HDT pool?
3. In an HDT pool, page size can be changed. (True or False)
4. Can an HDP pool be converted to an HDT pool?

30

8. Hitachi Tiered Storage Manager

Module Objectives

- Upon completion of this module, you should be able to:
 - Describe Hitachi Tiered Storage Manager features and functions
 - Demonstrate how Tiered Storage Manager is used to move data across storage tiers
 - Create tiers
 - Describe the criteria available to create tiers
 - Create migration groups
 - Describe the usage of rules in defining migration groups
 - Perform migrations by creating and executing migration tasks
 - Check the status of migration tasks

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 the 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

Overview

- Hitachi Tiered Storage Manager (HTSM) enables you to **nondisruptively migrate** data within the VSP.
 - Provides safe and efficient data storage by transferring data to the most suitable storage resource in response to changing storage requirements for applications
- HTSM 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 also are supported.

Benefits

- Easy to use service-level based volume migration
- Provides non-interruptive, completely transparent volume movement
- Safely moves data when application requirements increase or decrease
- Schedules 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
- Has ability to use LDEVs that are larger than the source LDEV in a migration as a migration target¹

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.

Requirements

- HTSM relies on other Hitachi products for its operations:
 - Hitachi Device Manager
 - Communicates with the storage controller system
 - Hitachi Volume Migration V2 (license installed on Storage)
 - Contains the migration engine embedded in the Hitachi Virtual Storage Platform, Hitachi Universal Storage Platform™ (USP V), Hitachi Universal Storage Platform™ VM (USP VM), Hitachi Universal Storage Platform™ (USP), or Hitachi Network Storage Controller (NSC) microcode
 - License automatically included with HTSM license

Software Options

- HTSM features can be expanded using optional Hitachi products for its operations.
 - Linkage with Hitachi Tuning Manager enables additional search fields to select volumes (for Tiers):
 - Array Group Busy Rate
 - Array Group Max Busy Rate
 - Hitachi Volume Shredder is required to enable data shredding.
 - The volume shredding engine is embedded in the VSP 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 to provide data protection for volumes involved in migrations using Hitachi Tiered Storage Manager.

Components

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

8

Volumes Used in HTSM

- 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 of source and target volumes must be the same.
 - Sizes of the source and target volumes must be the same.
 - You can migrate to bigger volumes but the bigger volume will be automatically carved to a smaller volume before migration.
 - 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 the source volume, you can now perform migration. If the capacity of a target volume is greater than that of the source volume, HTSM 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`.

Concepts

- HTSM introduces several concepts:
 - Storage Domain
 - Migration Group
 - Storage Tier
 - Migration Task

10

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

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 enterprise storage system as the domain controller (USP Family onwards).

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.

- 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 = an application's volumes used for logging.
 - When an actual list of volumes is created, supplemental volumes can be added as required; for example, volumes can be added to a given application.

12

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 must 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.

- 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.

For 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.

- **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: All LDEV sizes are not 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, a great challenge occurs when exact LDEV size matching is required to perform operations such as LDEV migrations.

Operations

- Steps:
 - Create Migration Groups
 - Create and Execute Migration Tasks
 - Check Task Status for completion

Tiers

- By specifying filter conditions based on a volume's characteristics, or by manually specifying a volume, you can group volumes to be used as migration target candidates. Each group is called a tier.
- You can add filter conditions or volumes until the necessary tier has been created.
- Volume list for a Tier is dynamic in nature
 - New volumes appear in the tier as and when they get created
 - Volumes deleted from storage system will not be listed in the Tier.

16

- Creating Tiers based on Internal/External Storage/Volumes attributes
 - Create a storage tier by specifying an externally connected storage subsystem
 - Create a tier by specifying an array group
 - Create a tier by specifying disk attributes.

17

Create a tier by specifying an externally connected storage subsystem

- To migrate infrequently used data from a domain controller to an externally connected storage subsystem, specify **Subsystem** or **Subsystem Serial Number** in the filter conditions for creating the tier.

Create a tier by specifying an array group

- When an array group contains multiple volumes that are frequently accessed, resulting in degraded response to hosts, you might want to migrate the data into volumes in another array group. In this case, specify **Array Group** in the filter conditions for creating the tier.

Create a tier by specifying disk attributes

- To migrate data to a faster disk for quicker response to hosts, specify **Disk Type** or **Disk RPM** in the filter conditions for creating the storage tier.

- Creating Tiers based on Dynamic Provisioning attributes
 - DP Pool as a migration target
 - DP Volume created automatically
 - Existing DP Volume can also be the target for a migration

18

If you use Dynamic Provisioning Volumes, you can also create the following tier:

Create a tier by specifying a DP pool volume or DP pool

- When you are using the Dynamic Provisioning feature you can specify a DP pool volume or DP pool as a migration target. When you specify a DP pool as the migration target, a DP volume that has the same size and emulation type as the source volume is created in the DP pool when migration is performed. To dynamically create a DP pool volume in the DP pool when performing migration, use either of the following methods to create the tier:

Specify **Volume(s)** as the **filter type**, **Dynamic Provisioning** as the filter condition, and **DP-Pool-VOL** as the value.

Specify **Pool** as the filter type and **Free Capacity** of the pool as the filter condition.

Create Tiers

■ **Resources tab > Storage Tiers**

Hitachi Command Suite 7

File Actions Tools Help Logged in as: System Log Out

Dashboard Resources Tasks & Alerts Search Administration Quick Find

Resources

Storage Systems

Hosts

Logical Groups

Tiers

All Tiers > USP_V@10.74.64.20 Help

Summary

Storage System	USP_V@10.74.64.20	No. of Storage Systems	USP_V
Total Capacity	1.22 TB		

Tiers

Select All Column Settings Rows/page: 30 Page: 1 / 1

Tier	Description	Total Capacity	No. of Volumes	% of Domain	Total Subscribed Capacity	Unallocated Capacity	Allocated Capacity	Consumed Capacity	No. of Subscribed Volumes
No Data									

Selected: 0 of 0 Create Tier Edit Tier Delete Tier

Waiting: 0 In Progress: 0 Completed: 5 Failed: 0

19

Resources tab > Tiers

Create Tier

Specify the name, description, and resources for the tier.

Storage system: USP_V@10.74.64.20

Name: * Internal Volumes

Description: All Internal Volumes

Select Resources: ☐ By Search Criteria ☒ Manually

Resource type: ☐ DP Pool ☒ Volumes

Volumes

Filter ☐ On ☒ Off Select All Column Settings Rows/page: 30 Page: 1 / 1

Volume	Storage System	Host	Capacity	Used Capacity	Drive Type	Drive Speed (RPM)	RAID Level	Parity Group	Pool	Ext. Storage System	Ext. Storage Vendor
00:00:02	USP_V...	-	2.00 GB	2.00 GB	FC	15000	RAID5(...	1-1-1	-	-	-
00:00:03	USP_V...	-	2.00 GB	2.00 GB	FC	15000	RAID5(...	1-1-1	-	-	-
00:00:04	USP_V...	-	2.00 GB	2.00 GB	FC	15000	RAID5(...	1-1-1	-	-	-

Selected: 0 of 3 Add Volumes Remove Volumes

* Required OK Cancel Help


more... Selected: 0 of 0 Create Tier Edit Tier Delete Tier

Waiting: 0 In Progress: 0 Completed: 5 Failed: 0


20

The volumes that should appear as part of a storage tier can be specified manually or through a search criteria.

Tiers created based on manual selection are a static list of volumes.



Tiers can also be created based on a search criteria. The tiers created based on search criteria present a dynamic list of volumes.



Create Tier with DP Pool

Tier with Volumes

The screenshot shows the Hitachi Tiered Storage Manager interface. The top navigation bar includes Dashboard, Resources, Tasks & Alerts, Search, and Administration. The left sidebar shows a tree view with Resources, Storage Systems, Hosts, Logical Groups, and Tiers. The main content area is titled 'Tier with Volumes' and shows a summary table for the tier 'USP_V@10.74.64.20'.

Tier Name	Internal Volumes	Description	All Internal Volumes	Total Capacity	1.05 TB
No. of Volumes	225	% of Domain	85.7 %	Total Subscribed Capacity	958.84 GB
Unallocated Capacity	924.84 GB	Allocated Capacity	34.00 GB	Consumed Capacity	958.84 GB
No. of Subscribed Volumes	213				


Below the summary table is a table of internal volumes. The table has columns: Volume, Label, Host, WWN, Port, Host Group, LUN, Capacity, Parity Group, Pool, RAID Level, and Volume Attribute. The table shows 13 rows of data, all with a capacity of 2.00 GB and a RAID level of RAID...

Volume	Label	Host	WWN	Port	Host Group	LUN	Capacity	Parity Group	Pool	RAID Level	Volume Attribute
00:00:00		WLAB0A	10.00...	CL1-A...	WLAB0A	0000	2.00 GB	1-1-1	-	RAID...	CVS
00:00:01		WLAB0A	10.00...	CL1-A...	WLAB0A	0001	2.00 GB	1-1-1	-	RAID...	CVS
00:00:02		-	-	-	-	-	2.00 GB	1-1-1	-	RAID...	CVS
00:00:03		-	-	-	-	-	2.00 GB	1-1-1	-	RAID...	CVS
00:00:04		-	-	-	-	-	2.00 GB	1-1-1	-	RAID...	CVS
00:00:05		-	-	-	-	-	2.00 GB	1-1-1	-	RAID...	CVS
00:00:06		-	-	-	-	-	2.00 GB	1-1-1	-	RAID...	CVS
00:00:07		-	-	-	-	-	2.00 GB	1-1-1	-	RAID...	CVS
00:00:08		-	-	-	-	-	2.00 GB	1-1-1	-	RAID...	CVS
00:00:09		-	-	-	-	-	2.00 GB	1-1-1	-	RAID...	CVS
00:00:0A		-	-	-	-	-	2.00 GB	1-1-1	-	RAID...	CVS
00:00:0B		-	-	-	-	-	2.00 GB	1-1-1	-	RAID...	CVS
00:00:0C		-	-	-	-	-	2.00 GB	1-1-1	-	RAID...	CVS
00:00:0D		-	-	-	-	-	2.00 GB	1-1-1	-	RAID...	CVS

23

Whenever a user/process accesses the tier; a list of volumes is generated based on the search criteria.


Any change in volume configurations on the storage system can result in fewer or more volumes being listed for a tier.



24

Clicking **Edit Tier** allows a user to change the definition of a tier.

Delete Tier



Clicking **Delete Tier** allows a user to delete the definition of a tier.

Migration Groups

- Volumes that are to be migrated, must be grouped together. This group is called a *migration group*.
- Although multiple migration groups can be created in a single storage domain, each volume can belong to only one migration group
 - A volume can be moved between migration groups during operations.
- Migration group is a static list of volumes.
 - When the migration group is created, volumes can be manually added or removed from the migration group.
- One volume can appear in only one migration group.

26

Migration group consists of source volumes for a migration operation.

- Protecting volumes against migrations
 - Performance-sensitive volumes should remain on internal storage.
 - Backup volumes with no paths should stay on external storage.
 - Do not allow system to select that volume in any other migration group.
 - Do not migrate attribute of a migration group.
 - Volumes in the group cannot be migrated.

27

To prevent operational errors, you can use the characteristics of migration groups for the following operations:

Grouping volumes you do not want to be selectable as migration sources

- You can use the group attributes to specify whether a group can be migrated. By adding the volumes you do not want to be selectable as migration sources to a group that cannot be migrated, you can prevent those volumes from being migrated by mistake.

Hitachi Tiered Storage Manager selects volumes without a path as migration target candidates. Therefore, if volumes without a path contain data that you do not want to be overwritten, register these volumes in a migration group.

■ Rules

- Specify the criteria to select or avoid the target array groups
- Parity Group Selection
 - The volumes from array groups for the Migration Group volumes will be selected as the target volumes.
- Parity Group Avoidance
 - The volumes from array groups for the Migration Group volumes will not be selected as the target volumes.

Create MG

General **Rule** Notification

Use this form to set Rules for making out Migration Plan.

Parity Group Selection: Balance Capacity

Avoid Parity Groups used by the following Migration Groups:

Migration Group List:

Avoidance List: Performance_Avoid

Add > < Remove

28

For each migration group, set the LDEV selection rules which define the parity groups from which Hitachi Tiered Storage Manager selects the migration target volumes.

There are two types of LDEV selection rules: the parity group avoidance rule and the parity group selection rule.

If you do not want any parity group containing LDEVs that belongs to another migration group to be selectable as a target parity group, specify the **parity group avoidance rule**.

To ensure that the migration target is selected based on how parity groups are used, specify the **parity group selection rule**.

You can specify array group avoidance rules and the array group selection rule at the same time. Tiered Storage Manager excludes the array groups related to the migration groups specified by array group avoidance rules, and then selects target LDEV candidates from the selected array groups based on the array group selection rule.

Parity Group Selection Rule

- Use this rule to specify how the array groups from which Tiered Storage Manager selects volumes are chosen. The default method is Balance Capacity.


The screenshot shows the 'Create MG' form with the 'Rule' tab selected. The 'Parity Group Selection' dropdown menu is open, showing three options: 'Balance Capacity' (selected), 'Maximum Coverage', and 'Minimum Coverage'. The dropdown is highlighted with a red rectangle. Below the dropdown, there are fields for 'Avoid Parity Groups used in Migration Groups:' and 'Migration Group List:'.

29


Parity Group Selection Rule

■ Balance Capacity

Storage tier (before migration)



Status after migration with BalanceCapacity specified



30

Balance Capacity


Specify this method when the equal use of array groups has priority over the location of the data. Tiered Storage Manager selects volumes so that the remaining capacity of each parity group that can be selected as a migration target is equal, or nearly so.

For example, you can use this method in the following case: When the storage subsystem has limited remaining storage capacity, you can reserve more groups for use as migration target candidates for other migration groups.


Parity Group Selection Rule

■ Maximum Coverage

Storage tier (before migration)



Status after migration with MaximumCoverage specified



31

Maximum Coverage

Specify this method when you want to distribute data among as many parity groups as possible. Tiered Storage Manager selects volumes from as many parity groups as possible.

For example, you can use this method in the following cases:


When using an application that requires fast responses If all of the data is placed in the same parity group, the required performance might not be achieved due to I/O conflicts. Distributing the data over multiple array groups would improve the response time.

When using an application in which errors can be isolated by using partial blockage, distribute the data among multiple parity groups to prevent any disruption to system operations. If an error occurs in any parity group, the data in the other parity groups can still be accessed.


Parity Group Selection Rule

■ Minimum Coverage

Storage tier (before migration)



Status after migration with MinimumCoverage specified



32

Minimum Coverage

Specify this method when you want to consolidate as much data as possible in a specific parity group. Tiered Storage Manager selects LDEVs from as few parity groups as possible.

For example, you can employ this method when using an application that does not support error isolation by using partial blockage. Distributing the data among different parity groups increases the risk that an error in any one parity group could halt system operation. Placing all the data in a highly reliable parity group reduces the risk of hardware errors affecting the applications.

Avoidance Rules Use Case

- Use these rules to create a migration task in which the parity groups containing the migration target volumes are different from the parity groups containing the migration source volumes.

The screenshot shows the 'Create MG' form with the 'Rule' tab selected. The 'Parity Group Selection' is set to 'Balance Capacity'. Below this, the 'Avoidance Rules' section is highlighted with a red box. It contains a 'Migration Group List' on the left, an 'Add >' button, a '< Remove' button, and an 'Avoidance List' on the right containing the text 'Performance_Avoid'.

33

Specifying a parity group avoidance rule for a migration group is optional.

- **Avoidance Rules use case**

- Scenario — You should place a ShadowImage P-VOL and S-VOL in different array groups. If the P-VOL and S-VOL are placed in the same array group and an error occurs in that array group, the data on both the P-VOL and the S-VOL is no longer accessible. In that case creating volume copies becomes meaningless.
- Registering the P-VOL and S-VOL in different migration groups and then setting the array group avoidance rule prevents the P-VOL and S-VOL from being placed in the same array group by mistake.

34

When a migration task is created for a migration group that an avoidance rule has been specified for, Tiered Storage Manager will also exclude volumes in the avoided parity groups from the migration target candidates.

Hitachi Command Suite 7

File Actions Tools Help

Dashboard Resources

Resources

Storage Systems

Hosts

Logical Groups

Tiers

All Tiers

HDS9970V@10.74.64.20

USP_V@10.74.64.20

DP_Pool

External Volumes

Internal Volumes

General Tasks

Allocate Volumes

Create Pool

Create Volumes

Manage Replication

Migrate Volumes

Unallocate Volumes

Virtualize Volumes

more...

Waiting: 0 In Progress: 0

Hitachi Command Suite 7

File Help

Logged in as: system Close

Storage Domains

Storage Domains

TSM_USP_V@10.74.64.20

Storage Domains > TSM_USP_V@10.74.64.20

Migration Groups

Storage Tiers

Detailed Information

Summary Information

NameTSM_USP_V@10.74.64.20

Refresh Status(Success)

Total Capacity1.22 TB

No. of MGs0

No. of Tiers3


Rows/Page25

Page1 of 1

Action	Name	Actual Capacity		Subscribed Capacity		
		Total	% of Domain	Total	Unallocated (%)	Allocated (%)
	DP_Pool	-	-	-	-	-
	External Volumes	180.00 GB	14%	180.00 GB	120.00 GB (66%)	60.00 GB (34%)
	Internal Volumes	1.05 TB	85%	958.84 GB	924.84 GB (96%)	34.00 GB (4%)


Refresh Tree

Resources tab > Migrate Volume (General Tasks)



36

Specify the General, Rules and Notification details.



37

Specify a reminder date related to the migration group.

The notification could be used by the administrator for reminding about a pending migration for the volumes in the migration group.

Other Operations

Storage Domains

Storage Domains > TSM_USP_V@10.74.64.20 > Migration Groups > WLAB0A

Summary Information

Name	WLAB0A		
Migration Status	Can Migrate	Yes	
Task Type	No. of Volumes	9	
Task Status	Total Capacity	270.00 GB	
	No. of Subscribed Volumes	9	
	Total Subscribed Capacity	270.00 GB	
	Consumed Capacity	270.00 GB	

Volume List Detailed Information Rule

Rows/ Pages: 25 of 1 Page 1 of 1

LDEV ^	Label	Capacity	Status	Locked	Can Migrate	I/O Consumer	Port/HSD	Subs
00:05:01	Internal	30.00 GB	Used	Unlocked	Yes			USP_V
00:05:02	Internal	30.00 GB	Used	Unlocked	Yes	WLAB0A	CL1-A/WLAB0A CL2-A/WLAB0A	USP_V
00:05:04	Internal	30.00 GB	Used	Unlocked	Yes			USP_V
00:06:00	Cache_Enabled	30.00 GB	Used	Unlocked	Yes			HDS9
00:06:01	Cache_Enabled	30.00 GB	Used	Unlocked	Yes	WLAB0A	CL1-A/WLAB0A CL2-A/WLAB0A	HDS9
00:06:02	Cache_Enabled	30.00 GB	Used	Unlocked	Yes			HDS9
00:06:50	Cache_Disabled	30.00 GB	Used	Unlocked	Yes			HDS9
00:06:51	Cache_Disabled	30.00 GB	Used	Unlocked	Yes			HDS9
00:06:52	Cache_Disabled	30.00 GB	Used	Unlocked	Yes	WLAB0A	CL1-A/WLAB0A	HDS9

Refresh Tree

38

Selecting a migration group allows you to edit or delete the migration group.

Migration Tasks

- Can be created/launched from migration group or storage tiers
- System suggests the possible migration targets from the tiers.
- Migration targets are selected based on following criteria:
 - Source is same size as target volume.
 - Target volume is larger than source volume.
 - Target volume size is reduced to match source.
 - Remaining space set as free space.
- User can modify the suggested list.
- If a task cannot be performed for a volume in a migration group, **Out Of Execution** is displayed as the status of that volume.
 - You can choose to remove a volume from the migration list.

39

After you have defined migration groups and storage tiers, you can create migration tasks. Hitachi Tiered Storage Manager displays all of the free volumes suitable as migration target volume candidates. Select migration target volumes from the displayed target volume candidates. In addition to the migration target volumes you select, Tiered Storage Manager displays volumes that can be used as migration target candidates.

To perform migration, you can use the migration target volumes selected by Tiered Storage Manager or you can re-select migration target volumes from the volumes displayed as migration target candidates.

Create

Storage Domains

- Storage Domains
 - TSM_USP_V@10.74.64.20
 - Migration Groups**
 - Storage Tiers

Migration Groups
Storage Domains > TSM_USP_V@10.74.64.20 > Migration Groups

Summary Information			
No. of MGs	3	Total Capacity	284
No. of Volumes	16	Capacity per Domain	224
No. of Subscribed Volumes	16	Total Subscribed Capacity	284
Consumed Capacity	284.00 GB		

Migration Groups					
Rows/Page: 25 14					
☐	Action	Name	Migration Status	Last Migration Time	Can Migrate
☐	  	OrcL_SVR1			Yes
☐	  	Performance_Avoid			Yes
☐	  	Small_Volumes			Yes

Storage Domains

- Storage Domains
 - TSM_USP_V@10.74.64.20
 - Migration Groups**
 - Storage Tiers

OrcL_SVR1
Storage Domains > TSM_USP_V@10.74.64.20 > Migration Groups > OrcL_SVR1

Summary Information			
Name	OrcL_SVR1		
Migration Status	Can Migrate	Yes	
Task Type	No. of Volumes	9	
Task Status	Total Capacity	270.00 GB	
	No. of Subscribed Volumes	9	

40

- Select Source and Target
 - Check compatibility of volumes

Migrate MG Help

Properties

Please select a Migration Group and a Storage Tier to create a Migration Task.

Storage Domain: TSM_USP_VM@10.2.24.13

Migration Group(Source): Orcl_SVR1

Storage Tier(Target): All_Volumes

Information : Orcl_SVR1

Information : All_Volumes

Check Compatibility

Migration Group(Source): Orcl_SVR1 → Storage Tier(Target): All_Volumes

Time Estimate 0:00:00 (Shredding : + 0:00:00 , Zero data discard : + 0:00:00)

Compatibilities:

41

Select the migration group and storage tier to which you want to migrate the volumes. Click **check compatibility** to continue.

Migration Task

- Check Compatibility compares the source and target, and presents the Migration Profile and also the suggested migration plan.

Profile for Migration			
System could not automatically match all volumes to target volumes.			
Configuration	No. of Unallocated Volumes in Tier		Difference
	Needed	Available	
OPEN-V, 2.00 GB (4195200 Blocks)	2	0 (0)	-2 (-2)

42

A check is performed to make sure that the number of migration target volumes is sufficient for the migration source. If there is no problem, the migration source volumes and migration target volumes are automatically paired, and the result is displayed in the **Migration Pairs** table.

If there is a problem during the compatibility check, a message will be output to the **Profile for Migration** list, and the Migration Pairs table will display a list of the migration pairs that have been set at the migration target.

- Based upon the results of the compatibility check, Hitachi Tiered Storage Manager can present multiple options (migration plans) for migrating the volumes.
- **Edit Pairs** button allows you to edit the default source and target combinations suggested by Tiered Storage Manager.

Migration Pairs:


☐	Source		Target		Status	Subsystem		Ctrl. Parity Group		P
	Vol. ID	LDEV	LDEV			Source	Target	Source	Target	S
☐		00:00:01	00:00:01	00:00:01		USP_VM@10.2.24.13	USP_VM@10.2.24.13	1-1-1	1-1-1	1-
☐		00:02:00	00:02:00	00:02:00		USP_VM@10.2.24.13	USP_VM@10.2.24.13	X1-1	X1-1	X
☐		00:02:01	00:02:01	00:02:01		USP_VM@10.2.24.13	USP_VM@10.2.24.13	X1-1	X1-1	X
☐		00:05:00	00:05:00	00:05:00		USP_VM@10.2.24.13	USP_VM@10.2.24.13	X3-1	X3-1	X


|

- If no compatible volumes are found, the status column indicates with a question mark (?).
 - User may need to provision new volumes (with compatible attributes) through Device Manager.

Migration Pairs:

Source		Target		Status	Subsystem		Ctrl. Parity Group		Parity Group		S
Vol. ID	LDEV	LDEV			Source	Target	Source	Target	Source	Target	
 00:00:00	00:00:00				USP_V@10.74.64.20		1-1-1		1-1-1		
 00:00:01	00:00:01				USP_V@10.74.64.20		1-1-1		1-1-1		



- *Note:* The **Edit Pairs** button is disabled.

44

You will need to go to the Hitachi Device Manager to create some volumes with the same size as source volumes.

Replace Pair/Do Not Migrate

- Alternate target options available

Edit Migration Pairs

Help

Migration Group(Source): SAP_Volumes

Storage Tier(Target): All_Volumes

Current Pair(s): select a pair to replace or uncouple.

	Source	Target	Status	Subsystem	Ctrl. Parity Group	Parity Gro
	Vol. ID	LDEV	LDEV	Source	Target	Source
☐	00:00:03	00:00:03	00:09:0C	USP_V@10.74.64.20	USP_V@10.74.64.20	1-1-1
☐	00:00:07	00:00:07	00:00:13	USP_V@10.74.64.20	USP_V@10.74.64.20	1-1-1
☒	00:00:08	00:00:08	-	USP_V@10.74.64.20		1-1-1

Available Pair(s): select an alternate pair to replace the current pair.

Replace Pair

Analyze Target PG

	Source	Target	Assigned	Subsystem	Ctrl. Parity Group	Parit
	Vol. ID	LDEV	LDEV	Source	Target	Source
☐	00:00:08	00:00:08	00:00:08	USP_V@10.74.64.20	USP_V@10.74.64.20	1-1-1
☐	00:00:08	00:00:08	00:00:0D	USP_V@10.74.64.20	USP_V@10.74.64.20	1-1-1
☐	00:00:08	00:00:08	00:00:0E	USP_V@10.74.64.20	USP_V@10.74.64.20	1-1-1
☐	00:00:08	00:00:08	00:00:0F	USP_V@10.74.64.20	USP_V@10.74.64.20	1-1-1
☐	00:00:08	00:00:08	00:00:12	USP_V@10.74.64.20	USP_V@10.74.64.20	1-1-1

Do Not Migrate

OKCancel

Clicking the **Edit Pair** button brings up the Edit Migration Pairs dialog box.

Replace Pair/Do Not Migrate

- No alternate target options available

Edit Migration Pairs Help

Migration Group(Source): Orcl_SVR1 Storage Tier(Target): All_Volumes

Current Pair(s): select a pair to replace or uncouple.

	Source		Target	Status	Subsystem		Ctrl. Parity Group		Parity Gr
	Vol. ID	LDEV	LDEV		Source	Target	Source	Target	Source
	00:00:01	00:00:01	00:00:04		USP_VM@10.2.24.13	USP_VM@10.2.24.13	1-1-1	1-4-1	1-1-1
	00:02:00	00:02:00	00:02:00		USP_VM@10.2.24.13	USP_VM@10.2.24.13	X1-1	X1-1	X1-1
	00:02:01	00:02:01	00:02:01		USP_VM@10.2.24.13	USP_VM@10.2.24.13	X1-1	X1-1	X1-1
	00:05:00	00:05:00	00:05:00		USP_VM@10.2.24.13	USP_VM@10.2.24.13	X3-1	X3-1	X3-1

Do Not Migrate

Available Pair(s): select an alternate pair to replace the current pair. Replace Pair | Analyze Target PG

	Source		Target	Assigned	Subsystem		Ctrl. Parity Group		Parity
	Vol. ID	LDEV	LDEV		Source	Target	Source	Target	Source
	00:02:00	00:02:00	00:02:00	00:02:00	USP_VM@10.2.24.13	USP_VM@10.2.24.13	X1-1	X1-1	X1-1

OK Cancel

46

■ Migration to Larger Volume

Profile for Migration
 The candidate migration plan contains at least one target volume that has a larger capacity than its corresponding source volume. This target volume(s) will be resized to match its corresponding source volume(s), freeing the excess capacity.
 The migration plan contains volumes that are not targets for migration. Check the migration plan, and then create the migration task.

Configuration	No. of Unallocated Volumes in Tier		Difference
	Needed	Available	
OPEN-V, 2.00 GB (4195200 Blocks)	0	191 (111)	+191 (+111)

Migration Pairs:

☐	Source		Target		Status	Subsystem		Ctrl. Parity Group		Pa
	Vol. ID	LDEV	LDEV			Source	Target	Source	Target	
☐	 00:00:03	00:00:03	00:09:0C			USP_V@10.74.64.20	USP_V@10.74.64.20	1-1-1	2-4-1	1-
☐	 00:00:07	00:00:07	00:00:13			USP_V@10.74.64.20	USP_V@10.74.64.20	1-1-1	1-4-1	1-
☐	 00:00:08	00:00:08	-			USP_V@10.74.64.20		1-1-1		1-

47


If there are multiple migration target candidate volumes that have a larger capacity than the migration source volume, and all of these volumes are displayed, you must select the migration target volumes from the list of volumes.

■ Settings

Migration Task Settings:

Task Description: 	Notification Email: ☐ Execute immediately after confirmation. ☐ Erase remaining data on source volumes. ☐ Discard zero data after migration.
---	--

OK Cancel



48


Enter the task description and notification email address as necessary, and then select when the task you created is to be executed. Next, choose whether the data at the migration source is to be deleted after the task is executed.

By default, a task goes into standby mode when it is created if **Execute immediately after confirmation** is not checked. The task can be executed from the HCS interface.

By default, the data on the migration source volumes will not be erased after migration. In order to prevent inappropriate disclosure of information, Hitachi recommends that you shred the data on the migration source volumes after migration.

Migration Task Status

■ Check Task Status



In the main HCS window, switch to the **Tasks and Alerts** tab. Select **HTSM Tasks**. Specify the search criteria for the Tasks and click **Search**.

The tasks are listed in the Search results panel. Click on a Task ID to see the task details and status.

■ Task Details

HITACHI

Task - TKIndeeyng

Help

Summary Information

Task ID	TKIndeeyng	Progress	
Task Owner	System	Time Estimate	0:00:52
Task Type	Migration	Creation Time	2010-12-15 01:11:19
Domain	TSM_USP_V@10.74.64.20	Execution Request Time	
Migration Group	SAP_Volumes	Start Time	
Status	Standby	Time Elapsed	
Capacity	4.00 GB	End Time	
Notification Email			
Task Description			

Detailed Information

Component Status

Rows/Page: 25

Page: 1 of 1

Source	Target	Status	Task Execution	Subsystem
Vol. ID	LDEV	LDEV	Progress	Time Elapsed
00:00:03	00:00:03	00:09:0C	Standby	0%
00:00:07	00:00:07	00:00:13	Standby	0%
00:00:08	00:00:08	00:00:08	Out Of Execution	0%

Refresh

Close

50

■ Execute a Task

Hitachi Command Suite 7

File Help Logged in as: system Close

Search

Find all Task(s)

Creation Date(Year-Month-Day)
From 2010-12-8
To 2010-12-15

Show entries matching
All of the conditions below:
Same Attribute cannot be specified two or more times.

Attribute: Value:

Task List (Total:1)

Rows/ Page: 25 Page: 1 of 1

☒	Action	Task ID ^	Task Type	Status	Task Execution	Progress	Time Elapsed	Time
☒		TKIndeeyng	Migration	Standby				

Task status icon	Description
	The task is pending execution. To execute a task that has this icon, select it, and then click the Execute Tasks button.
	The task is pending execution or is being executed. This icon is displayed if creation, deletion, or re-creation of a volume in a migration task is pending or is currently being performed. Wait until processing is completed.
	The requested processing has completed.
	Stop or immediate stop processing is being performed for the task.
	The task has been canceled or execution failed.

Execute Tasks Stop Tasks Cancel Tasks Delete Tasks

Tasks can be executed using the Web client or CLI client, and can be executed either immediately after they are created or at some later time. You cannot create an additional task for a migration group while a task for that migration group remains uncompleted.

Module Summary

- Described Hitachi Tiered Storage Manager features and functions
- Demonstrated how Tiered Storage Manager is used to move data across storage tiers and how to create tiers
- Described the criteria available to create tiers and how to create migration groups
- Described the usage of rules in defining migration groups and:
 - How to perform migrations by creating and executing migration tasks
 - How to check the status of migration tasks

Module Review

1. What is the purpose of HTSM?
2. What is the pre-requisite software for HTSM?
3. What are the steps to perform migration?
4. Target volume can be larger than the source volume.
(True or False)


9. External Storage Manager

Module Objectives

- Upon completion of this module, you should be able to:
 - Demonstrate how Universal Volume Manager is used to manage tiered storage

2


Hitachi Virtualized Tiered Storage Architecture



The Hitachi virtualized tiered storage architecture consolidates heterogeneous storage solutions in a single managed pool. Within this pool, different storage tiers are established. Hitachi virtualization technologies enable the separation of logical views from physical assets, enabling an architecture that provides the right cost, performance, reliability and availability characteristics of storage, as needed, matched with application requirements. It provides thin provisioning, advanced replication, logical partitioning, virtual connections, support for heterogeneous storage and the sheer performance and scalability needed to truly deliver storage services on demand. At the same time, it eliminates vendor lock-in. Tiered storage can be scaled out using hardware appropriate to the tier.

This means organizations can strategically align storage infrastructure with business requirements.


Virtualization




Hitachi Enterprise Storage Systems provided multiple options for virtualizing storage and capacity.

Capacity Virtualization and Data Mobility

- Hitachi Dynamic Provisioning
 - Supports **capacity virtualization**



- Hitachi Tiered Storage Manager
 - Supports transparent, nondisruptive data volume mobility




5

Hitachi Dynamic Provisioning uses thin provisioning to simplify and add agility to the application storage provisioning process and can save significant money on storage purchases. This is important for companies faced with ongoing rapid growth of their data storage requirements and escalating storage and storage management expenses. Dynamic Provisioning also can reclaim existing previously allocated storage within the virtualized pool, and on externally attached heterogeneous systems. Also, to improve performance and simplify performance management, Hitachi Dynamic Provisioning automates optimization of spindle usage and provides dynamic load balancing.

Hitachi Tiered Storage Manager data management enables IT administrators to easily and interactively match application quality of service requirements to heterogeneous storage assets. Tiered Storage Manager enables easy storage tier creation and the nondisruptive movement of data volumes to match application-driven price, performance and availability characteristics. As the cost, quality of service and functionality demands of application data change over time, data must be migrated between tiers. This is how a tiered storage architecture maintains an optimal match of storage system characteristics and data requirements.

Universal Volume Manager


- Universal Volume Manager features and benefits
 - Maps heterogeneous external storage resources as internal volumes
 - Enables a dynamic tiered storage infrastructure
 - Creates storage pools independent of physical location
 - Simplifies the management of heterogeneous storage devices
 - Integrates competitive and heterogeneous storage resources into a single storage domain
 - All enterprise features and program products that are available to internal volumes are also available to external volumes.



6

Features

- Operate multiple storage systems as if they are all part of one storage system



7

Universal Volume Manager provides the virtualization of a multi-tiered storage area network comprised of heterogeneous storage systems. It enables the operation of multiple storage systems connected to enterprise storage systems as if they were all in one storage system and provides common management tools and software.


The shared storage pool comprised of external storage volumes can be used with storage system-based software for data migration and replication, as well as any host-based application. Combined with Hitachi Tiered Storage Manager, Universal Volume Manager provides an automated data lifecycle management solution, across multiple tiers of storage.

Configuring External Storage

- External storage is configured using the following procedures:
 - Change the port attribute to External in VSP
 - Prepare a volume in the external system
 - Set up physical links
 - Discover the external system volumes
 - Map an external volume
 - Define the host path

Change the Port Attribute to External

- VSP storage systems Channel front end directors (FEDs) are used to connect to external storage systems.
- Change the port attribute to External.




9

You set the port used for Universal Volume Manager as the external port. External port is used for discovering volumes presented from the External Storage system. The external storage system cannot be connected to the ports other than the external port. In order to set the port attribute to external, you need to release the paths set to the port. The attribute of the port where the paths are set cannot be changed to external. Before starting the Universal Volume Manager operations, you need to know the ports whose attributes can be changed to external.

Note: The ports whose attributes are set for remote copy software (for example: RCU target and initiator) or the other features cannot be used as external ports for External Storage Management. Change the port attribute to external if the port attribute is set to other than external.

Prepare the External System

- Create LUNs on external storage system.
 - If needed
- Set Host Mode or equivalent to Windows.
 - Virtual Storage Platform storage systems must see the external device as a Microsoft Windows System.
- Configure LUN masking and security on external storage system.
 - Enable LUN access for the VSP storage systems.




10

To use an external volume with the VSP storage systems, you first prepare the volumes in the external system that is (are) connected to the VSP systems. The volume in the external system should be between 46MB and 4TB. However, if the external volume is more than 4TB, you can access the data stored up to 4TB only. You cannot map external volumes of 46MB or smaller.

Note: If LUN Security is enabled on the port, you need to add the WWN of the external port on the VSP storage systems to the Host Group which has the volumes that need to be set as the external volumes.

Configure Physical Links


- Physically cable the systems.
- Zone the external ports to the external storage system.



11


Cable the external storage system and the VSP External Ports to the SAN. Configure zoning to enable the External Ports to connect to the External Storage System. Fibre Channel Arbitrated Loop (FC-AL) and direct connections also are supported.

Externally Attached Storage Connectivity



Externally attached storage can be attached directly via FC, attached through a FC switch, or attached via FC extender. Data can be copied or moved between internal and external storage over FC distances.

Discover External Storage




Discover external storage. This is similar to rescanning the SCSI bus on a host. If the previous steps were performed correctly, the VSP storage systems will see the LUNs on the external storage system after the external storage discovery.

You cannot use the external volume by just discovering the external storage system. You next need to perform the *Add LU* operation to make the external volume available to the VSP storage systems .

After you discover the external storage systems and the LUNs available, you can map volumes in the external storage system as volumes in the local storage system. However, you need to identify which volumes in the external storage system you intend to use.


Map External Storage



Configure the external volumes as Open-V, which is mapped as an internal volume, to define a path between the host and the internal storage system. Use the LUN Manager feature to define a path. External volumes mapped as OPEN-V volumes within the VSP storage systems allow you to access existing data on the external volumes.

Concepts for Using External Volumes

- Define the Host Path (map LDEVs to Host Group LUNS).




15

If you set the Open emulation type for the external volume, which is mapped as an internal volume, define a path between the host and the internal storage system. Use appropriate interface to define a path. External volumes mapped as OPEN-V volumes within the VSP storage systems allow you to access existing data on the external volumes.

External Volumes and External Volume Group


- An external volume is **registered** as a **Virtual Device** (VDEV).
- You apply the **OPEN-V** emulation to the volume.
- Create custom-sized **VDEVs** from the original volume.



16

External Volume Groups

- External volumes need to be mapped to external volume groups.
 - Create an External Volume Group (a Virtual Parity Group).
 - Set the external volume attributes.
 - Map external volumes to the External Volume Group.



17

When you map an external volume as an internal volume, you need to register the external volume to an external volume group. The user can classify the external volumes into the groups according to the use. The group is called external volume group (ExG). For instance, you can register multiple volumes in one external storage system to one external volume group. You must assign numbers from 1 to 16,384 to external LU groups. A maximum of 16,384 external volume groups can be created. The maximum number of volumes that can be registered in one external group is 256.

External Volumes and External Volume Group

- **Map** an external volume to VSP
 1. **Register** to an external volume group (ExG, a virtual parity group).
 2. Set the external volume **attributes**:
 - **Cache Mode**

Disabled: The local storage system signals the host that an I/O operation has completed only after the local storage system has synchronously written the data to the external storage system.

Enabled: The local storage system signals the host that an I/O operation has completed and then asynchronously destages the data to the external storage system.
 - **Inflow Control**

Disabled: I/O from the host during the retry operation is written to the cache memory even after the writing operation to the external volume is impossible.

Enabled: The writing operation to cache is stopped and the I/O from the host is not accepted when the writing operation to the external volume is impossible.

18

External Volume Attributes

Cache Mode

- Cache mode specifies whether the write data from the host to the external storage system is propagated synchronously or asynchronously. All I/O to and from the local storage system (either Enable or Disable) always uses cache. Write operations are always backed up in duplex cache.
- If you select **Enable**, after receiving the data into the local storage system cache memory, the local storage system signals the host that an I/O operation has completed and then asynchronously destages the data to the external storage system.
- If you select **Disable**, the local storage system signals the host that an I/O operation has completed only after the local storage system has synchronously written the data to the external storage system.


Inflow Control

- Inflow control specifies whether the writing operation to the cache memory is stopped (Enable) or continued (Disable) when the writing operation to the external volume is impossible. By default, inflow control is set to Disable.
- If you select Enable, the writing operation to cache is stopped and the I/O from the host is not accepted when the writing operation to the external volume is impossible.
- If you select Disable, the I/O from the host during the retry operation is written to the cache memory even after the writing operation to the external volume is impossible. Once the writing operation to the external volume becomes normal, all the data in the cache memory is written to the external volume (all the data is destaged).

Cache Mode Usage

■ Write to virtual volume

- Cache mode=**On**: Write from Host is written to cache memory in VSP, and completion of the write operation is reported to Host.
- Cache mode=**Off**: Write from Host is written to the external storage, and completion of the write operation is reported to Host.



19

■ Cache Mode impact

- In an improperly designed configuration, cache mode enabled can have a negative impact on performance.
- Example: Accessing a busy external LDEV on SATA drives on external storage.
 - If the SATA drives on external storage cannot write data fast enough, the write pending data can back up in VSP cache and cause performance issues for hosts attached to the VSP.
- Use Hitachi Virtual Partition Manager to employ partitioned cache for external volumes or storage with lower performance attributes.

20

Operations

- Universal Volume Manager is resident in the microcode. It is a comprehensive part of the microcode.
- Universal Volume Manager is licensed on capacity.
- Operational specifications
 - External volume should be less than or equal to 4TB.
 - Specify OPEN-V for the emulation type for external volumes.
 - You must set the LUN attributes for the external attributes on VSP.
 - All LDEVs created in the external volume group will have the same attributes.
 - LUSE volumes cannot consist of external volumes with different attributes.
 - When the Cache Mode is set as Disable for the external volume, the volume cannot be used for cache residency.
 - Any capacity above the licensed capacity will not be discovered.


21

Important Operational Specifications for Universal Volume Manager Operations

- If the external volume is more than 4TB, you can use up to 4TB. You cannot access data over 4TB.
- Specify OPEN-V for the emulation type for the external volume during mapping process for the external volume on the VSP. You can access all existing data on the external volume after the mapping.
- The external volume attributes set on the external subsystem (such as Port security, LUN security, LDEV Guard attribute) are not applied to the mapped external volumes. Set the attributes to the mapped external volumes on VSP as necessary.

Virtualize Volumes

■ Virtualize volumes — Steps



22

Virtualize Storage

- Steps for virtualizing storage
 - Prepare environment
 - Add External Storage system to Device Manager
 - Supported HDS storage systems will be listed in dropdown for **Storage System Type** in the **Add Storage System**.
 - Supported other vendor storage systems will need to be added as **SMI-S enabled** from **Storage System Type** in the **Add Storage System**.
 - Any other system will need to be virtualized using the Storage Navigator interface.
 - The virtualized volumes will appear in HSC as external volumes after they are virtualized.

- Steps for virtualizing storage (*cont'*)
 - Set up connectivity
 - Identify FC ports that will link the two systems
 - The ports can be linked as:
 - Direct attached or
 - Through a fabric (recommended)
 - Set up internal storage
 - Configure external ports (using Element Manager)
 - Identify volumes to be virtualized
 - Execute Virtualize Volumes task from Device Manager


24

Currently Device Manager can virtualize volumes that are in the Unallocated-Open pool.

Use Element Manager for the Enterprise Storage to virtualize any volume already allocated and in use on the external.

Port Settings

■ Using Element Manager to change the Port Settings



25

Note: The attribute of the port where the LU paths are already set cannot be changed to external. Therefore, you must identify ports whose attributes can be changed to external before starting the virtualize operations.

Virtualize Volumes

■ Launch Virtualize Volumes

Hitachi Command Suite 7

Virtualize Volumes

To virtualize volumes, first select the external and internal storage systems, specify the external volumes. To define I/O connections between the storage systems, edit the paths.

Select the external and internal storage systems.

External: NSC55@10.2.24.20

Internal: USP_VM@10.2.24.10

Specify the Volumes to virtualize

No. of Volumes: 3

Volume Capacity: 2 GB

Required capacity: 6.00 GB

Advanced Options

No. of Paths: 2

Storage Port	Controller	Ext. Storage System	Int. Storage System	Int. Port WWN
CL2-A	MIX-2F	NSC55@10.2.24.20	USP_VM@10.2.24.10	50.06.0E.80.0...
CL1-A	MIX-1A	NSC55@10.2.24.20	USP_VM@10.2.24.10	50.06.0E.80.0...

[Edit Paths](#)

[Path Priority Settings](#)

[Host Group and LUN Settings](#)

[External Volume Configuration](#)

Show Plan Cancel Help

26

- External
 - Virtualized Storage
 - Any supported external storage
 - Registered in Device Manager
- Internal
 - Virtualizing Storage
 - VSP
 - USP V/USP VM
 - USP/NSC55

Virtualize Volumes

To virtualize volumes, first select the external and internal storage systems, specify the external volumes. To define I/O connections between the storage systems, edit the paths.

Select the external and internal storage systems.

External: NSC55@10.2.24.20

Internal: USP_VM@10.2.24.10

Specify the Volumes to virtualize

No. of Volumes: 3

Volume Capacity: 2 GB

Required capacity: 6.00 GB

Advanced Options

No. of Paths: 2

Storage Port	Controller	Ext. Storage System	Int. Storage System	Int. Port WWN
CL2-A	MIX-2F	NSC55@10.2.24.20	USP_VM@10.2.24.10	50.06.0E.80.0...
CL1-A	MIX-1A	NSC55@10.2.24.20	USP_VM@10.2.24.10	50.06.0E.80.0...

[Edit Paths](#)

[Path Priority Settings](#)

[Host Group and LUN Settings](#)

[External Volume Configuration](#)

Show Plan Cancel Help

27

■ External Storage

- Specifies volume parameter for volumes to be created and virtualized
- Advanced Options
 - Specify filter for choosing parity group for volumes to virtualize

Specify the storage system, source storage, and their paths.

Specify Storage System

Source: USP@172.17.41.165

Target: RAID700@172.17.45.30

Source Storage

Volume Size: 50 GB

Number of Volumes: 1

Required Capacity: 50.00 GB

Drive Type: Any

Advanced Option

Drive RPM: 10000

RAID Level: Any

Parity Group: 1-4-1

Number of ports: 2

Storage Port	Controller	Source Storage System	Target Storage System	Target Port WWN
CL1-A	CHA-1E	USP@172.17.4...	RAID700@172....	50.06.0E.80.0...
CL1-E	CHA-1F	USP@172.17.4...	RAID700@172....	50.06.0E.80.0...


Path Priority Setting

Host Group, LUN Number Setting

External Volume Configuration

Show Plan Cancel Help

■ Specify **Paths** linking the external storage and internal storage



29

Always verify the paths when discovering external storage for the first time. When the paths have been set and the virtualization process is complete, the current settings become the default for next virtualize storage task.

- Specify **Paths** linking the external storage and internal storage

Edit Paths

You can edit the paths graphically or by using the selection table. To create paths, choose ports to connect. First, select storage ports on one end, then connect them to the ports on the other end. You can also select paths to remove them.

Ext. Storage System: NSC55@10.2.24.21
Internal Storage System: USP_VM@10.2.24.11
No. of Paths: 2
Edit Path By: ☐ Topological Graph ☒ Selection Table

Ext. Storage Port: CL1-A

Internal Storage Ports	
Storage System	Port
USP_VM@10.2...	50.06.0E.80.05.64.2...
USP_VM@10.2...	50.06.0E.80.05.64.2...
USP_VM@10.2...	50.06.0E.80.05.64.2...
USP_VM@10.2...	50.06.0E.80.05.64.2...

Add > < Remove

Selected Paths			
Ext. Port	Ext. Controller	Internal Storage System	Internal Storage Port
CL7-B	MIX-1A	USP_VM@10.2...	50.06.0E.80.05.64.24.44
CL8-A	MIX-2F	USP_VM@10.2...	50.06.0E.80.05.64.24.54

OK Cancel Help

30

- Specify **Host Group, LUN Number Setting** for host group to be created on external storage

Virtualize Volumes

CL8-A MIX-2F NSC55@10.2... USP_VM@10.2... 50.06.0E.80.0...

Edit Paths

Path Priority Settings

Storage Port	Internal Port WWN
CL7-B	50.06.0E.80.05.64.24.44
CL8-A	50.06.0E.80.05.64.24.54

Increase Priority Decrease Priority

Host Group and LUN Settings

Host group name: R6010000
Host Mode Options: ☐ VERITAS Database Edition/Advanced Cluster ☐ TPRLO ☐ iSCSI Delayed ACK improvement
LU Number (hex): Start from 0000

External Volume Configuration

Ext. Parity Group: Start from E 1 (1-16384)
CLPR: 0
Inflow control: ☐ Enable ☒ Disable
Cache mode: ☒ Enable ☐ Disable

Show Plan Cancel Help

31

These settings represent the host group creation on the external storage.

■ Specify External Volume Configuration

The screenshot shows the 'Virtualize Volumes' dialog box. The 'External Volume Configuration' section is highlighted with a red rectangle. It contains the following fields and options:

- Ext. Parity Group:** Start from E 1 (1-16384)
- CLPR:** 0
- Inflow control:** ☐ Enable ☒ Disable
- Cache mode:** ☒ Enable ☐ Disable

Other visible sections include 'Path Priority Settings' with a table of storage ports and internal port WWNs, and 'Host Group and LUN Settings' with fields for host group name, host mode options, and LU number.

32

Just the way a Parity Group is the container for Internal Volumes, External Parity Group is the container for External Volumes.

Always have a strategy or method for assigning the External Parity Group number. For example, you can use the first or last four digits of the serial number of external storage as the External Parity Group number. This you can map back volumes in a specific External Parity Group to a external storage system.

■ External Volume Configuration

– External Parity Group



External Volume Configuration

Ext. Parity Group: Start from E (1-16384)

CLPR:

Inflow control: ☐ Enable ☒ Disable

Cache mode: ☒ Enable ☐ Disable

– Inflow Control

- Indicates if writes should be accepted in case of external volume cannot be accessed:
 - Enabled: Writes will be rejected
 - **Disabled: Writes will be accepted**

– Cache Mode

- Indicates if write data for the external volume is destaged to external storage:
 - Synchronously: Cache Mode – Disabled
 - **Asynchronously: Cache Mode – Enabled**

33

Inflow control

- If you select **Enable**, the writing operation to cache is stopped and the I/O from the host is not accepted when the writing operation to the external volume is impossible.
- If you select **Disable**, the I/O from the host during the retry operation is written to the cache memory even after the writing operation to the external volume is impossible. When the writing operation to the external volume becomes normal, all the data in the cache memory is written to the external volume (all the data is destaged).

Cache Mode

- If you select **Enable**, after receiving the data into the local storage system cache memory, the local storage system signals the host that an I/O operation has completed and then asynchronously destages the data to the external storage system.
- If you select **Disable**, the local storage system signals the host that an I/O operation has completed only after the local storage system has synchronously written the data to the external storage system.
- *Note:* Depending on the I/O profile from host and also the external storage configuration, incorrect settings of these options can have a significant impact on cache utilization. This may result in performance issues for other hosts connected to the internal storage.
- You can use Virtual Partition Manager to partition the cache on enterprise storage and use a dedicated cache partition to isolate I/O going to external volumes.

■ Plan Summary

Virtualize Volumes

Verify the following plan for virtualizing volumes.

Plan Summary		
Internal	Storage System	USP_VM@10.2.24.11
External	Storage System	NSC55@10.2.24.21
	No. of Volumes	1
	Volume Capacity	2.00 GB
	Drive Type	Unknown
	Drive Speed (RPM)	10000
	RAID Level	RAID5(3D+1P)
	Parity Group	1-1-1

Plan Details

Task name: * Virtualize-3

Description:

Schedule

☐ Display task status

* Required

Back Virtualize Cancel Help

■ Plan Details

Virtualize Volumes

[Plan Details](#)

Parity Group Information

Parity Group	Storage System	RAID Level	Drive Type	Drive Speed (RPM)	Total Capacity	Allocated Capacity	Unallocated Capacity	Free Capacity
1-1-1	NSC55@10...	RAID5(3D+...	Unknown	10000	199.01 GB	102.01 GB	97.00 GB	97.00 GB

Volume Information

Volume	Capacity	Parity Group
00:14	2.00 GB	1-1-1

External Volume Information

Ext. Volume	Ext. Parity Group	Capacity	Inflow Control	Cache Mode	Int. Storage System
auto	E1	2.00 GB	Disabled	Enabled	USP_VM@10.2.24.11

Path Information

Host Group	Host Mode	Host Mode Options	Ext. Storage Port	Int. Storage WWN	Volume	Label	LUN	Path Group
R6010000	Windows E...	2TB or Lar...	CL7-B	50.06.0E.8...	00:14		0000	new
R6010000	Windows E...	2TB or Lar...	CL8-A	50.06.0E.8...	00:14		0000	new

Task name: * Virtualize-3

Description:

[Schedule](#)

☐ Display task status

* Required

Back Virtualize Cancel Help

35

■ External storage — Internal Volume view

– Virtualized volumes and their internal LDEV IDs

Hitachi Command Suite 7

File Actions Tools Help Logged in as: System Log Out

Dashboard Resources Tasks & Alerts Search Administration

All Storage Systems > USP_VM@10.2.24.10 > External Storage

Summary

No. of Internal Vols: 6 Total Internal Capacity: 12.00 GB

Internal Volumes External Volumes Available External Volumes

Filter On Off Select All Column Settings Rows/page: 30 Page: 1 / 1


Volume	Label	Host	WWN	Port	Host Group	LUN	Capacity	Parity	Pool	RAID Level	Volume Attribute
00:00:...		-	-	-	-	-	2.00 GB	E100-1	-	RAID5(3D...	External, ...
00:00:...		-	-	-	-	-	2.00 GB	E100-2	-	RAID5(3D...	External, ...
00:00:...		-	-	-	-	-	2.00 GB	E100-3	-	RAID5(3D...	External, ...
00:00:...		-	-	-	-	-	2.00 GB	E100-4	-	RAID5(3D...	External, ...
00:00:...		-	-	-	-	-	2.00 GB	E100-5	-	RAID5(3D...	External, ...
00:00:...		-	-	-	-	-	2.00 GB	E100-6	-	RAID5(3D...	External, ...

Selected: 0 of 6

Waiting: 0 In Progress: 0 Completed: 3 Failed: 1

36

- External storage — External Volumes view
 - Virtualized Volume with their external LDEV IDs and configuration on external storage



Hitachi Command Suite 7

File Actions Tools Help Logged in as: System Log Out

Dashboard Resources Tasks & Alerts Search Administration Quick Find

All Storage Systems > USP_VM@10.2.24.10 > External Storage

Summary

No. of Internal Vols: 6 Total Internal Capacity: 12.00 GB

Internal Volumes External Volumes Available External Volumes

Filter On Off Select All Column Settings Rows/page: 30 Page: 1 / 1


Volume	Label	Ext. Storage System	Serial No.	Host	WWN	Port	Host Group	LUN	Capacity	Parity Group	RAID Level	Volume Attribute
00:27		NCS...	80061	EXSP_RA...	50.06.0E...	CL1-A, C...	ExtGrp1	000A	2.00 GB	1-1-1	RAID5(3...	CVS
00:28		NCS...	80061	EXSP_RA...	50.06.0E...	CL1-A, C...	ExtGrp1	000B	2.00 GB	1-1-1	RAID5(3...	CVS
00:29		NCS...	80061	EXSP_RA...	50.06.0E...	CL1-A, C...	ExtGrp1	000C	2.00 GB	1-1-1	RAID5(3...	CVS
01:21		NCS...	80061	EXSP_RA...	50.06.0E...	CL1-A, C...	ExtGrp1	0007	2.00 GB	1-1-1	RAID5(3...	CVS
01:22		NCS...	80061	EXSP_RA...	50.06.0E...	CL1-A, C...	ExtGrp1	0008	2.00 GB	1-1-1	RAID5(3...	CVS
01:23		NCS...	80061	EXSP_RA...	50.06.0E...	CL1-A, C...	ExtGrp1	0009	2.00 GB	1-1-1	RAID5(3...	CVS

Selected: 0 of 6

Waiting: 0 In Progress: 0 Completed: 3 Failed: 1

37

- External storage — Available External Volumes
 - Volumes presented for virtualization (on external storage) but not virtualized as yet



Hitachi Command Suite 7

File Actions Tools Help Logged in as: System Log Out

Dashboard Resources Tasks & Alerts Search Administration Quick Find

All Storage Systems > USP_VM@10.2.24.10 > External Storage

Summary

No. of Internal Vols: 6 Total Internal Capacity: 12.00 GB

Internal Volumes External Volumes Available External Volumes

Filter On Off Select All Column Settings Rows/page: 30 Page: 1 / 1

Volume	Label	Ext. Storage System	Serial No.	Host	WWN	Port	Host Group	LUN	Capacity	Parity Group	RAID Level	Volume Attribute
00:20		NCS...	80061	EXSP_RA...	50.06.0E...	CL1-A, C...	ExtGrp1	0000	2.00 GB	1-1-1	RAID5(3...	CVS
00:21		NCS...	80061	EXSP_RA...	50.06.0E...	CL1-A, C...	ExtGrp1	0001	2.00 GB	1-1-1	RAID5(3...	CVS
00:22		NCS...	80061	EXSP_RA...	50.06.0E...	CL1-A, C...	ExtGrp1	0002	2.00 GB	1-1-1	RAID5(3...	CVS
00:23		NCS...	80061	EXSP_RA...	50.06.0E...	CL1-A, C...	ExtGrp1	0003	2.00 GB	1-1-1	RAID5(3...	CVS
00:24		NCS...	80061	EXSP_RA...	50.06.0E...	CL1-A, C...	ExtGrp1	0004	2.00 GB	1-1-1	RAID5(3...	CVS
00:25		NCS...	80061	EXSP_RA...	50.06.0E...	CL1-A, C...	ExtGrp1	0005	2.00 GB	1-1-1	RAID5(3...	CVS
00:26		NCS...	80061	EXSP_RA...	50.06.0E...	CL1-A, C...	ExtGrp1	0006	2.00 GB	1-1-1	RAID5(3...	CVS

Selected: 0 of 7

Waiting: 0 In Progress: 0 Completed: 3 Failed: 1

38

Data Migration

- Universal Volume Manager does **not** migrate any data between volumes; it just maps the external volumes within VSP.
- When the mapping of external volumes is complete, all Hitachi software can use the external volumes.
- For migrating data from internal to external or vice versa, the following options are available:
 - Hitachi Tiered Storage Manager
 - Nondisruptive data migration
 - Recommended tool for migration with external storage
 - ShadowImage In-system replication
 - Application downtime required for data migration
 - Primary purpose data replication

39

Module Summary

- Demonstrated how Universal Volume Manager is used to manage tiered storage

40

Module Review

1. What is the procedure to set up the physical link for a Universal Volume Manager operation?
2. What is the size of an external volume that can be recognized by Universal Volume Manager?
3. What is an external volume group?
4. What does the Inflow Control attribute define?
5. What does the I/O Cache mode attribute define?

10. Hitachi Virtual Partition Manager

Module Objectives

- Upon completion of this module, you should be able to:
 - Describe how Virtual Partition Manager enables the logical partitioning on Enterprise storage systems
 - List the operations supported through Hitachi Virtual Partition Manager
 - Create new Cache Logical Partition (CLPR)
 - Assign resources to CLPR
 - Delete cache partitions
 - Use cache partitions in other program products

2

Virtual Partition Manager Overview

- Hitachi Virtual Partition Manager enables data center administrators to perform logical partitioning of cache and disk capacity, including external storage, on the VSP to create independently managed cache partitions.
- These logical partitions help maintain quality of service by acting as dedicated cache resources that are independently managed and reserved for specific applications.
- Benefits:
 - Assures quality of service
 - Storage resources optimized to application and business requirements

3

Assures Quality of Service

Virtual Partition Manager dedicates resources (for example, cache, disk) for exclusive use by specific applications to maintain priority and quality of service for business-critical applications. You can secure and restrict access to storage resources to ensure confidentiality for specific applications. You can also use Virtual Partition Manager to adjust data storage resources dynamically to satisfy changing business requirements.

Storage Resources Optimized to Application/Business Requirements

Virtual Partition Manager allows customer to match data to appropriate storage resources based on availability, performance, capacity and cost. It improves flexibility by allowing dynamic changes to cache partitions while in use.

Overview


- Cache Logical Partition (CLPR)
 - The cache logical partition function allows you to divide the cache into multiple virtual cache memories, to lessen I/O contention.
 - CLPRs can be defined by using Storage Navigator program.

Cache Logical Partition


- Created to guarantee cache resources to LDEVs on a specific parity groups
- Created due to performance reasons
- Consists of:
 - Cache (not shared with any other CLPR)
 - Parity groups (that can use the cache in the CLPR)
- Cache in a CLPR is not shared with any other CLPR.
- Minimum cache size is 4GB for a CLPR.
 - Can be increased dynamically, in increments of 2GB
- By default, all cache is allocated to CLPR0.
- Can create 31 new CLPRs
- Creating a new CLPR involves:
 - Reassigning cache from CLPR0 to the new CLPR
 - Reassigning parity groups from CLPR0 to new CLPR
- Deleting a CLPR moves the resources (cache and parity groups) back to CLPR0.

5

■ Use example



If Host B load is low, Host A can use a large portion of cache.



Although Host B load is low, Host A cannot use the cache assigned to CLPR of Host B.

6

External Storage Virtualization and Cache

Cache Mode Disabled


For random reads, data still passes through the internal storage cache. The use of the cache management MRU/LRU (Most Recently Used/ Least Recently Used) queues and sequential prefetch are disabled in "Cache Mode Disabled", so that these data blocks do not remain in cache for further reuse. In the case of random writes, data blocks are written to internal cache and then out to external storage, just as with internal disks, but the write operation complete is not reported to the host until the data has been written to the external storage device cache.

Cache Mode Enabled

For random reads, the use of cache management MRU/LRU queues and sequential prefetch are enabled and can benefit older storage technology. In the case of random writes, data blocks are written to internal cache, a write operation complete is returned to the host, and then the data is written to external storage, just as with internal disks.

Sequential Operations

For sequential read or write patterns, I/Os are never subject to the MRU/LRU system (whether stand-alone storage or with UVM). These are handled in bursts, with a number of cache slots being allocated for the duration of the operation and placed in the free queue when complete.



7

Cache Mode Disabled

For random reads, data still passes through the internal storage cache. The use of the cache management MRU/LRU (Most Recently Used/ Least Recently Used) queues and sequential prefetch are disabled in "Cache Mode Disabled", so that these data blocks do not remain in cache for further reuse. In the case of random writes, data blocks are written to internal cache and then out to external storage, just as with internal disks, but the write operation complete is not reported to the host until the data has been written to the external storage device cache.

Cache Mode Enabled


For random reads, the use of cache management MRU/LRU queues and sequential prefetch are enabled and can benefit older storage technology. In the case of random writes, data blocks are written to internal cache, a write operation complete is returned to the host, and then the data is written to external storage, just as with internal disks.

Sequential Operations

For sequential read or write patterns, I/Os are never subject to the MRU/LRU system (whether stand-alone storage or with UVM). These are handled in bursts, with a number of cache slots being allocated for the duration of the operation and placed in the free queue when complete.

External Storage — How Much Cache?

- If you add the same cache to the internal storage system, then:
 - It will get expensive \$\$.
 - The cache in the external storage will be underused.




How Much Cache to Allocate in Internal Storage?

- Why not use the external cache for read hits?
- Limit the CLPR size to force high level of read miss
- Configure enough CLPR space to accommodate the normal write 'working set'



**Applications with Write
component need to be Cached**

A Sensible Suggestion — Rule of Thumb



10

Guidelines

- For scenarios where workload analysis data is not available, the Cache Mode should initially be set to **Disable**. Then after workload analysis, it can be changed dynamically to **Enable**.
- For SATA or FC disk used for archive or backup storage, Cache Mode should be set to **Disable**.
- The cache on the internal storage has a major impact on the performance of the external storage, especially the write operations that take advantage of cache in the internal storage.

11

Specifications

#	Item	Content
1	Maximum number of CLPRs	32
2	Minimum unit of CLPR	Parity group
3	Change unit of CLPR	Increase size by 2GB
4	CLPR Capacity	4GB – 512GB
5	Support emulation type	All Open types supported by Hitachi storage systems
6	Max number of CLPRs	32
7	LUSE	Supported
8	RAID Level	All Open types supported by Hitachi storage systems


12

Operations

- Can be performed from Hitachi Storage Navigator
- The operations include the following tasks:
 - Create new CLPR
 - Migrate (cut) resources (cache/parity groups) to new CLPR
 - Move resources from existing CLPR to another CLPR
 - Delete CLPR

13


■ Validate the License Key



14

Before you start VPM operations, ensure that the license key is installed.


■ Launch Virtual Partition Manager



15

Storage Navigator > Settings > Partition Definition

■ Filter Parity Groups based on CUs



16

Specify how you want to view information in the CLPR resource list. In the Select CU dialog box, choose one of the following:

- **All CUs:** Shows the information about all CUs on the CLPR resource list.
- **Specific CU:** Choose specific LDKC and CU. This shows only CLPRs from the selected CU.
- **Unallocated:** Shows information about only the CUs unallocated to CLPR on the CLPR resource list.

■ Create CLPR

The screenshot shows the 'Create CLPR' dialog box in the Hitachi Virtual Partition Manager. The 'Partition Definition' tab is selected, and the 'Cache Logical Partition' section is active. The 'Subsystem' tree shows 'CLPR' and '00'. The 'Detail For CLPR in Subsystem' section contains the following fields:

- CLPR Name :
- Cache Size : GB (4 - 106GB)
- Cache Residency Size : GB (0.0 - 0.0GB)
- Num. of Cache Residency Areas : (0 - 16384)

The table below the fields is empty:

No.	Resource Type	Address	Properties
-----	---------------	---------	------------

Page: / 1 Previous Next


Apply Cancel

17

Specify the following information for the new CLPR:

- **CLPR Name** - Type the name of the cache logical partition, in up to 16 alphanumeric characters.
- **Cache Size** - Select the cache capacity. You may select from 4GB to 508GB, in 2GB increments. The default value is 4GB. The size of the cache is allocated from CLPR 0, but you must leave at least 8 GB remaining in CLPR 0.
- **Cache Residency Size** - Select the cache capacity. You may select from 4GB to 508GB, in 2GB increments. The default value is 4GB. The size of the cache is allocated from CLPR0, but you must leave at least 8GB remaining in CLPR0.
- **Num of Cache Residency Areas** - Type the desired capacity for the Cache Residency Manager area. The range of values is 0 to 16384 and the default value is 0.

■ Move resources to new CLPR



18

Select the Parity Groups from the source CLPR and paste in to the new CLPR.

When moving resources to and from cache logical partitions:

- You can move resources only within the same CU.
- All interleaved parity groups must be in the same CLPR.
- LUSE volumes cannot be set across more than one CLPR.
- If a parity group contains one or more LDEVs that have defined Cache Residency Manager extents, you cannot migrate that parity group to another CLPR.

■ New CLPR with Parity Groups

License Key | SNMP Information | E-mail Information | Partition Definition

Partition Definition

Cache Logical Partition CU: All CUs Select CU

Subsystem

- CLPR
 - 00: CLPR0
 - 01: Dev_Volumes

No.	Resource Type	Address	Properties
00001	Parity Group	1-1	805.21 GB RAID5G
00002	Parity Group	1-2	805.21 GB RAID5G
00003	Parity Group	1-3	805.21 GB RAID5G
00004	Parity Group	1-4	805.21 GB RAID5G
00005	Parity Group	1-5	805.21 GB RAID5G

Page: 1 / 1 Previous Next

Detail For CLPR in Subsystem

CLPR Name : Dev_Volumes

Cache Size : 16 GB (4 - 106GB)

Cache Residency Size : 0.0 GB (0.0 - 0.0GB)


Num. of Cache Residency Areas : 0 (0 - 16384)

Apply Cancel

19

Click **Apply** to create the new CLPR.

- Delete a CLPR



20

Before deleting a CLPR, you must first migrate resources from that CLPR to CLPRs that will not be deleted.

CLPR0 cannot be deleted.

Module Summary

- Described how Virtual Partition Manager enables the logical partitioning on Enterprise storage systems
- Listed the operations supported through Hitachi Virtual Partition Manager
- Described steps to create new Cache Logical Partition (CLPR)
 - How to assign resources to CLPR
- Described steps to delete cache partitions
- Explained how to use cache partitions in other program products

21

Module Review

1. Define SLPR and CLPR.
2. The Virtual Partition Manager is a performance enhancing tool. (True or False)
3. What is the difference between the Storage Administrator and Storage Partition Administrator?
4. The SLPRs function independently from each other including cache destaging. (True or False)
5. SLPRs and CLPRs can affect performance of each other. (True or False)
6. What is the maximum number of SLPRs that can be defined?
7. What is the maximum amount of cache that can be assigned to a CLPR?