

Hitachi Unified Storage Operations Guide

FASTFIND LINKS

Changes in this revision
Document Organization
Contents

© 2012-2015 Hitachi, Ltd. All rights reserved.

No part of this publication may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying and recording, or stored in a database or retrieval system for any purpose without the express written permission of Hitachi, Ltd. and Hitachi Data Systems Corporation (hereinafter referred to as "Hitachi").

Hitachi, Ltd. and Hitachi Data Systems reserve the right to make changes to this document at any time without notice and assume no responsibility for its use. Hitachi, Ltd. and Hitachi Data Systems products and services can only be ordered under the terms and conditions of Hitachi Data Systems' applicable agreements.

All of the features described in this document may not be currently available. Refer to the most recent product announcement or contact your local Hitachi Data Systems sales office for information on feature and product availability.

Notice: Hitachi Data Systems products and services can be ordered only under the terms and conditions of Hitachi Data Systems' applicable agreements. The use of Hitachi Data Systems products is governed by the terms of your agreements with Hitachi Data Systems.

Hitachi is a registered trademark of Hitachi, Ltd. in the United States and other countries. Hitachi Data Systems is a registered trademark and service mark of Hitachi in the United States and other countries.

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

France – Import pending completion of registration formalities

Hong Kong – Import pending completion of registration formalities

Israel – Import pending completion of registration formalities

Russia – Import pending completion of notification formalities

Distribution Centers – IDC, EDC and ADC cleared for exports

Table of Contents

1	Introduction	1-1
	Navigator 2 overview.	1-2
	Navigator 2 features	1-2
	Security features.	1-2
	Monitoring features	1-2
	Configuration management features	1-2
	Data migration features	1-2
	Capacity features	1-3
	General features	1-3
	Navigator 2 benefits	1-3
	Navigator 2 task flow.	1-4
	Navigator 2 functions	1-5
	Browser security	1-7
	Using the Navigator 2 online help	1-7
	Use with Hitachi Command Suite	1-9
2	System theory of operation	2-1
	Network standard and functions which the array supports	2-2
	RAID features.	2-2
	RAID technology task flow	2-3
	RAID levels	2-3
	RAID chunks and stripes	2-5
	Host volumes	2-6
	Number of volumes per RAID group	2-7
	Volume management and controller I/O management	2-7
	About the HUS Series of storage systems	2-7
	Recent features.	2-7
	Major controller features	2-8
	Understanding Navigator 2 key terms	2-8
	Navigator 2 operating environment.	2-9
	Firewall considerations.	2-10
	Anti-virus software considerations.	2-10
	Hitachi Storage Command Suite common components	2-10
3	Installation.	3-1
	Connecting Hitachi Storage Navigator Modular 2 to the Host.	3-2
	Installing Hitachi Storage Navigator Modular 2.	3-2
	Preparation.	3-2
	Setting Linux kernel parameters	3-5

Setting Solaris 8 or Solaris 9 kernel parameters	3-7
Setting Solaris 10 kernel parameters.	3-8
Operating or stopping Hitachi Command Suite (Device Manager, Tiered Storage Manager and Replication Manager) version 8.0.0 or later and navigator in the common server environment	3-9
Types of installations	3-12
Installing Navigator 2	3-13
Getting started (all users)	3-13
Installing Navigator 2 on a Windows operating system	3-14
If the installation fails on a Windows operating system.	3-15
Installing Navigator 2 on a Sun Solaris operating system	3-15
Installing Navigator 2 on a Red Hat Linux operating system	3-16
Updating Navigator 2	3-17
Setting the server certificate and private key.	3-18
Preinstallation information for Storage Features.	3-20
Environments.	3-20
Storage feature requirements	3-21
Requirements for installing and enabling features	3-21
Account Authentication	3-21
Audit Logging requirements.	3-21
Cache Partition Manager requirements	3-22
Data Retention requirements	3-22
LUN Manager requirements	3-22
Password Protection	3-22
SNMP Agent requirements.	3-22
Modular Volume Migration requirements	3-23
Installing storage features.	3-23
Enabling storage features	3-24
Disabling storage features	3-24
Uninstalling storage features	3-25
Uninstalling Navigator 2	3-25
Uninstalling when HCS V8.0.0 or greater is installed.	3-27
Starting Navigator 2 host and client configuration	3-27
Host side	3-27
Client side	3-28
For Windows	3-28
Changing JRE.	3-28
Changing JDK	3-30
For Linux and Solaris	3-30
Changing the Port Number for Applet Screen of Navigator 2	3-30
Changing the Port Number for Starting Navigator 2	3-31
Starting Navigator 2	3-33
Operations.	3-35
Setting an attribute.	3-37
Additional guidelines.	3-37

Help.	3-38
Understanding the Navigator 2 interface	3-40
Menu Panel	3-40
Explorer Panel	3-40
Button panel.	3-41
Page panel	3-41
Performing Navigator 2 activities	3-41
Description of Navigator 2 activities	3-43
Stopping the Hitachi Storage Command Suite Common Components	3-46
Starting the Hitachi Storage Command Suite Common Components	3-47

4 Provisioning 4-1

Provisioning overview	4-2
Provisioning wizards	4-2
Provisioning task flow	4-3
Hardware considerations	4-3
Verifying your hardware installation	4-3
Connecting the management console	4-3
Logging in to Navigator 2.	4-4
Selecting a storage system for the first time	4-6
Running the Add Array wizard	4-6
Running the Initial (Array) Setup wizard	4-8
Registering the Array in the Hitachi Storage Navigator Modular 2.	4-8
Initial Array (Setup) wizard — configuring email alerts	4-9
Initial Array (Setup) wizard — configuring management ports	4-11
Initial Array (Setup) wizard — configuring host ports	4-12
Initial Array (Setup) wizard — configuring spare drives	4-14
Initial Array (Setup) wizard — configuring the system date and time . .	4-14
Initial Array (Setup) wizard — confirming your settings.	4-14
Running the Create & Map Volume wizard	4-15
Manually creating a RAID group	4-15
Using the Create & Map Volume Wizard to create a RAID group	4-17
Create & Map Volume wizard — defining volumes	4-18
Create & Map Volume wizard — defining host groups or iSCSI targets .	4-19
Create & Map Volume wizard — connecting to a host	4-20
Create & Map Volume wizard — confirming your settings	4-21
Provisioning concepts and environments	4-21
About DP-Vols.	4-21
Changing DP-Vol Capacity	4-21
About volume numbers	4-22
About Host Groups	4-23
Creating Host Groups	4-23
Displaying Host Group Properties	4-24
About array management and provisioning	4-24

About array discovery	4-24
Understanding the Arrays screen	4-24
Add Array screen	4-25
Adding a Specific Array	4-25
Adding Arrays Within a Range of IP Addresses	4-25
Using IPv6 Addresses	4-26
Notes on Using Hitachi Storage Navigator Modular 2	4-26

5 Security 5-1

Security overview	5-2
Security features	5-2
Account Authentication	5-2
Audit Logging	5-3
Data Retention Utility	5-3
Security benefits.	5-3
Account Authentication overview	5-4
Account Authentication features.	5-4
Account Authentication benefits	5-4
Account Authentication caveats	5-5
Account Authentication task flow	5-5
Account Authentication specifications	5-8
Accounts	5-8
Account types.	5-9
Roles	5-9
Resources	5-10
Session	5-12
Session types for operating resources.	5-12
Advanced Security Mode	5-14
Changing Advanced Security Mode	5-14
Account Authentication procedures	5-15
Initial settings	5-15
Managing accounts	5-15
Displaying accounts	5-15
Adding accounts	5-17
Changing the Advanced Security Mode	5-18
Modifying accounts	5-19
Deleting accounts	5-21
Changing session timeout length	5-22
Forcibly logging out	5-23
Setting and deleting a warning banner	5-23
Troubleshooting	5-25
Audit Logging overview.	5-27
Audit Logging features	5-27
Audit Logging benefits	5-27

Audit Logging task flow	5-28
Audit Logging specifications	5-29
What to log?	5-30
Security of logs	5-30
Pulling it all together	5-30
Summary	5-31
Audit Logging procedures	5-32
Initial settings	5-32
Optional operations	5-32
Enabling Audit Log data transfers	5-32
Viewing Audit Log data	5-34
Initializing logs	5-35
Configuring Audit Logging to an external Syslog server	5-35
Data Retention Utility overview	5-36
Data Retention Utility features	5-36
Data Retention Utility benefits	5-36
Data Retention Utility specifications	5-37
Data Retention Utility task flow	5-39
Assigning access attribute to volumes	5-40
Read/Write	5-40
Read Only	5-40
Protect	5-41
Report Zero Read Cap. (Mode)	5-41
Invisible (Mode)	5-41
Retention terms	5-41
Protecting volumes from copy operations	5-42
Usage	5-43
Volume access attributes	5-43
Unified volumes	5-43
SnapShot and TCE	5-43
SYNCHRONIZE CACHE command	5-43
Host Side application example	5-44
Operating System (OS) restrictions	5-44
Volume attributes set from the operating system	5-44
Notes on usage	5-45
Notes about unified LU	5-45
Notes About SnapShot and TCE	5-45
Notes and restrictions for each operating system	5-46
Operations example	5-47
Initial settings	5-47
Configuring and modifying key settings	5-47
Data Retention Utility procedures	5-48
Optional procedures	5-48
Opening the Data Retention dialog box	5-48
Setting S-VOLs	5-50

Setting expiration locks	5-50
Setting an attribute	5-51
Changing the retention term	5-52
Setting the expiration lock	5-52
Setting S-VOL Disable	5-53

6 Provisioning volumes 6-1

LUN Manager overview	6-2
LUN Manager features.	6-2
LUN Manager benefits.	6-3
LUN Manager task flow	6-3
For Fibre Channel	6-3
For iSCSI	6-4
Understanding preconfigured volumes	6-5
LUN Manager specifications for Fibre Channel	6-5
About iSCSI	6-6
Design configurations and best practices.	6-8
Fibre Channel configuration	6-9
Conditions for Using LUN Manager for Fibre Channel	6-9
Fibre Channel design considerations	6-11
Fibre system configuration.	6-11
iSCSI system design considerations	6-11
iSCSI network port and switch considerations	6-12
Additional system design considerations	6-13
System topology examples	6-14
Assigning iSCSI targets and volumes to hosts	6-18
Preventing unauthorized SAN access	6-20
Avoiding RAID Group Conflicts	6-21
SAN queue depth setting	6-22
Increasing queue depth and port sharing	6-23
Increasing queue depth through path switching.	6-23
LUN Manager procedures	6-24
Using Fibre Channel	6-25
Using iSCSI	6-26
Fibre Channel operations using LUN Manager	6-29
About Host Groups	6-29
Adding host groups.	6-30
Enabling and disabling host group security	6-30
Creating and editing host groups	6-31
Initializing Host Group 000	6-35
Deleting host groups.	6-35
Changing nicknames	6-37
Deleting World Wide Names	6-37
Copy settings to other ports.	6-38

iSCSI operations using LUN Manager	6-39
Creating an iSCSI target.	6-40
Using the iSCSI Target Tabs.	6-40
Setting the iSCSI target security	6-43
Editing iSCSI target nicknames.	6-44
Adding and deleting targets	6-44
About iSCSI target numbers, aliases, and names	6-49
Editing target information	6-50
Editing authentication properties.	6-51
Initializing Target 000	6-52
Changing a nickname	6-52
CHAP users.	6-53
Adding a CHAP user	6-53
Changing the CHAP user	6-54
Setting Copy to the Other Ports	6-55
Setting Information for Copying	6-55
Copying when iSCSI Target Creation.	6-55
Copying when iSCSI Target Editing	6-55

7 Capacity 7-1

Capacity overview	7-2
Cache Partition Manager overview	7-2
Cache Partition Manager features	7-2
Cache Partition Manager benefits	7-2
Cache Partition Manager feature specifications.	7-3
Confirming Environments.	7-3
Cache Partition Manager task flow	7-3
Operation task flow	7-4
Stopping Cache Partition Manager	7-4
Pair cache partition	7-5
Partition capacity	7-6
Supported partition capacities	7-7
Segment and stripe size restrictions	7-9
Specifying partition capacity	7-10
Using a large segment.	7-10
Using load balancing	7-11
Using ShadowImage, Dynamic Provisioning, or TCE	7-11
Installing Dynamic Provisioning/Dynamic Tiering when Cache Partition Manager is Used.	7-11
Adding or reducing cache memory	7-13
Cache Partition Manager procedures.	7-14
Confirming Environments.	7-14
Initial settings.	7-14
Stopping Cache Partition Manager	7-15

Working with cache partitions	7-15
Adding cache partitions	7-16
Deleting cache partitions	7-17
Assigning cache partitions	7-17
Setting a pair cache partition	7-19
Changing cache partitions	7-20
Changing cache partitions owner controller	7-21
Installing SnapShot or TCE or Dynamic	7-22
VMWare and Cache Partition Manager	7-23
Cache Residency Manager overview	7-23
Cache Residency Manager features	7-23
Cache Residency Benefits	7-24
Cache Residency Manager task flow	7-24
Cache Residency Manager Specifications	7-25
Termination Conditions	7-26
Disabling Conditions	7-26
Equipment	7-27
Volume Capacity	7-27
Supported Cache Residency capacities	7-29
Restrictions	7-32
Cache Residency Manager procedures	7-33
Confirming environments	7-33
Initial settings	7-33
Stopping Cache Residency Manager	7-34
Setting and canceling residency volumes	7-34
NAS Unit Considerations	7-35
VMware and Cache Residency Manager	7-36

8 Performance Monitor 8-1

Performance Monitor overview	8-2
Monitoring features	8-2
Monitoring benefits	8-3
Monitoring task flow	8-3
Monitoring feature specifications	8-4
Analysis bottlenecks of performance	8-5
Launching Performance Monitor	8-6
Performance Monitor procedures	8-7
Initial settings	8-7
Optional operations	8-7
Optimizing system performance	8-8
Obtaining information	8-8
Using graphic displays	8-8
.	8-10
Working with the Performance Monitor Tree View	8-10

More about Tree View items in Performance Monitor	8-12
Using Performance Monitor with Dynamic Provisioning	8-16
Working with Graphing and Dynamic Provisioning	8-17
Displayed Items	8-24
Determining the ordinate axis.	8-26
Saving monitored data.	8-29
Exporting Performance Monitor information	8-30
Enabling performance measuring items	8-33
Working with port information	8-35
Working with RAID Group, DP Pool and volume information	8-35
Working with cache information	8-35
Working with processor information	8-35
Troubleshooting performance.	8-36
Performance imbalance and solutions	8-36
Dirty Data Flush	8-37

9 SNMP Agent Support 9-1

SNMP overview.	9-2
SNMP features	9-2
SNMP benefits	9-3
Environments and requirements	9-4
SNMP task flow.	9-4
SNMP versions	9-5
SNMP managers and agents.	9-6
Management Information Base (MIB)	9-6
Object identifiers (OIDs)	9-6
SNMP command messages.	9-7
SNMP traps.	9-9
Supported configurations.	9-12
Frame types	9-13
License key.	9-13
Installing Hitachi SNMP Agent Support	9-13
Hitachi SNMP Agent Support procedures	9-15
Preparing the SNMP manager.	9-15
Preparing the Hitachi modular storage array	9-15
Creating an operating environment file	9-16
Creating a storage array name file	9-22
Registering the SNMP environment information	9-22
Registering the SNMP environment information	9-24
Confirming your setup	9-25
Operational guidelines.	9-26
MIBs	9-29
Supported MIBs	9-29
MIB access mode	9-29

OID assignment system	9-29
Supported traps and extended traps	9-32
MIB installation	9-35
MIB II	9-35
system group	9-36
.	9-38
at group.	9-40
ip group.	9-41
icmp group.	9-46
tcp group.	9-46
udp group	9-46
egp group	9-46
snmp group	9-47
Extended MIBs	9-50
dfSystemParameter group	9-50
dfWarningCondition group	9-51
dfCommandExecutionCondition group.	9-54
dfPort group.	9-56
dfCommandExecutionInternalCondition group	9-60
Additional resources	9-62
10 Virtualization	10-1
Virtualization overview	10-2
Virtualization features	10-2
Virtualization task flow	10-2
Virtualization benefits	10-3
Virtualization and applications	10-4
Storage Options	10-5
A sample approach to virtualization	10-5
Hitachi Dynamic Provisioning software	10-7
Storage configuration	10-8
Redundancy	10-8
Zone configuration	10-8
Host Group configuration	10-10
One Host Group per ESX host, standalone host configuration	10-10
One Host Group per cluster, cluster host configuration	10-10
Host Group options	10-10
Virtual Disk and Dynamic Provisioning performance	10-11
Virtual disks on standard volumes	10-11
11 Special functions.	11-1
Modular Volume Migration overview	11-2
Modular Volume Migration Manager features.	11-2

Modular Volume Migration Manager benefits	11-2
Modular Volume Migration task flow	11-2
Modular Volume Migration Manager specifications	11-3
Requirements	11-6
Supported capacity	11-6
Setting up Volume Migration	11-7
Setting volumes to be recognized by the host	11-7
Volume Migration components	11-7
Volume Migration pairs (P-VOLs and S-VOLs).	11-8
Reserved Volume	11-8
DMLU.	11-8
DMLU precautions	11-9
VxVM	11-11
MSCS	11-11
AIX	11-11
Windows 2000/Window Server	11-12
Linux and LVM	11-12
Windows 2000/Windows Server and Dynamic Disk	11-12
UNMAP Short Length Mode	11-12
Performance	11-12
Using unified volumes	11-13
Using with the Data Retention Utility.	11-14
Using with ShadowImage.	11-14
Using with Cache Partition Manager	11-15
Concurrent Use of Dynamic Provisioning	11-16
Concurrent Use of Dynamic Tiering	11-19
Dirty Data Flush Limit number	11-19
Load Balancing function.	11-19
Contents related to the connection with the host	11-19
Modular Volume Migration operations	11-21
Managing Modular Volume Migration	11-22
Pair Status of Volume Migration	11-22
Setting the DMLU	11-22
Removing the designated DMLU	11-23
Adding the designated DMLU	11-23
Adding reserved volumes.	11-24
Deleting reserved volumes.	11-26
Migrating volumes.	11-26
Changing copy pace	11-28
Confirming Volume Migration Pairs	11-29
Releasing Volume Migration pairs	11-30
Canceling Volume Migration pairs	11-31
Volume Expansion (Growth not LUSE) overview.	11-32
Volume Expansion features	11-32
Volume Expansion benefits	11-32

Volume Expansion task flow	11-32
Displaying Unified Volume Properties	11-33
Selecting new capacity	11-33
Modifying a unified volume	11-33
Add Volumes	11-34
Separate Last Volume	11-34
Separate All Volumes	11-35
Power Savings overview	11-36
Power Saving features.	11-36
Power Saving benefits.	11-37
Power Saving task flow	11-37
Power Saving specifications	11-39
Estimated Spin-Up time.	11-44
Power down best practices	11-44
Power saving procedures.	11-45
Power down.	11-45
Power up.	11-45
Power saving requirements	11-45
Start of the power down operation.	11-45
RAID groups that cannot power down	11-46
Things that can hinder power down or command monitoring	11-46
Number of times the same RAID group can be powered down	11-47
Extended power down (health check).	11-47
Turning off of the array.	11-47
Time required for powering up.	11-47
Operating system notes	11-48
Advanced Interactive eXecutive (AIX).	11-48
Linux.	11-48
Hewlett Packard UNIX (HP-UX)	11-48
Windows	11-48
Solaris	11-48
Viewing Power Saving status	11-50
Powering down	11-53
Notes	11-54
Powering up.	11-55
Notes	11-56
Viewing volume information in a RAID group	11-56
Failure notes	11-57
Overview	11-59
Security	11-59
HDPS AUX-Copy plus aging and retention policies	11-60
HDPS Power Saving vaulting	11-61
HDPS sample scripts	11-63
Windows scripts	11-64
Power down and power up	11-64
Using a Windows power up and power down script	11-64

Powering down	11-66
Powering up	11-66
UNIX scripts	11-66
Power down	11-66
Power up	11-67
Using a UNIX power down and power up script	11-68
Powering down	11-69
Powering up	11-69
Power Saving Plus	11-70
Example of Power Saving Plus	11-70
Preparing to Use Power Saving Plus	11-70
Power Saving Plus Specifications	11-71
About Power Savings Plus	11-76
Effect of Power Saving	11-77
Drive Layout in RAID Group	11-77
DBS/DBL/DBX	11-77
General Power Saving Plus details	11-79
Power Saving details by operating system	11-81
Notes on Failures	11-82
Notes on Hosts	11-82
Operations Example	11-83
Example of operation in non I/O link mode	11-83
Example of operation in I/O link mode	11-83
Removing Power Saving	11-84
Operations of Power Saving Plus	11-84
Displaying Power Saving Information	11-84
Requesting Non I/O-linked Spin Down	11-86
Requesting I/O-linked Spin Down	11-87
Requesting I/O-linked Drive Power OFF	11-87
Requesting I/O-linked Spin Down with Drive Power OFF	11-88
Requesting Remove Power Saving (Spin Up)	11-89

A Specifications A-1

Navigator 2 specifications	A-2
System requirements	A-2
Windows server	A-2
Virtual OS	A-3
Windows client settings	A-3
Solaris™ (SPARC)	A-4
CPU: SPARC minimum 1 GHz (2 GHz or more is recommended)	A-4
Client	A-5
HP-UX	A-5
AIX	A-6
Linux	A-7

Host	A-7
Client.	A-8
IPv6 supported platforms	A-11
Considerations at Time of Operation.	A-13
Volume formatting	A-13
Constitute array	A-14
B Recording Navigator 2 Settings	B-1



Preface

Welcome to the Hitachi Unified Storage Navigator Modular 2 (HSNM2) Operations Guide. This document describes how to use the Hitachi Unified Storage Navigator Modular storage system provisioning software.

Please read this document carefully to understand how to use this product, and maintain a copy for reference purposes.

This preface includes the following information:

- ☐ [Intended audience](#)
- ☐ [Product version](#)
- ☐ [Changes in this revision](#)
- ☐ [Changes in this revision](#)
- ☐ [Document Organization](#)
- ☐ [Related documents](#)
- ☐ [Document conventions](#)
- ☐ [Convention for storage capacity values](#)
- ☐ [Accessing product documentation](#)
- ☐ [Getting help](#)
- ☐ [Comments](#)

Intended audience

This document is intended for system administrators, Hitachi Data Systems representatives, and authorized service providers who install, configure, and operate Hitachi Unified Storage systems.

This document assumes the following:

- The user has a background in data processing and understands storage systems and their basic functions.
- The user has a background in data processing and understands Microsoft Windows and their basic functions.
- The user has a background in data processing and understands Web browsers and their basic functions.

Product version

This document applies to Hitachi Unified Storage firmware version 0981/D and to HSNM2 version 28.12 or later.

Changes in this revision

- Under [Environments \(page 3-20\)](#) updated firmware, Storage Navigator 2, and CCI versions.

Document Organization

Thumbnail descriptions of the chapters are provided in the following table. Click the [chapter title](#) in the first column to go to that chapter. The first page of every chapter or appendix contains links to the contents.

Chapter Title	Description
Chapter 1, Introduction	Describes an overview of the product
Chapter 2, System theory of operation	Describes how to install and enable Hitachi Unified Storage provisioning features.
Chapter 3, Installation	Describes the basic flow of tasks involved with setting up provisioning software for Hitachi Unified Storage systems.
Chapter 4, Provisioning	Describes how to provision the Hitachi Unified Storage systems.
Chapter 5, Security	Describes account authentication and audit log features that provide intruder filtering safety for Hitachi Unified Storage systems.
Chapter 6, Provisioning volumes	Describes how to configure volumes for your storage system.
Chapter 7, Capacity	Describes how to set up cache partitions and work with cache residency items.
Chapter 8, Performance Monitor	Describes how to monitor the Hitachi Unified Storage systems.

Chapter Title	Description
Chapter 9, SNMP Agent Support	Describes how to configure Simple Network Manage Protocol to manage a distributed network of storage systems from a single, centralized location.
Chapter 10, Virtualization	Describes how to create virtual sessions for storage system configuration.
Chapter 11, Special functions	Describes how to configure storage systems using Modular Volume Migration Manager, Data Retention, Power Savings, and Data Migration, Volume Expansion and Shrinking, RAID Group Expansion, DP VOL Expansion, Mega Volumes, USP, and VSP.
Appendix A, Specifications	Describes specifications.
Appendix B, Recording Navigator 2 Settings	Provides a worksheet for your network settings.



HSNM2 also provides a command-line interface that lets you perform operations by typing commands from a command line. For information about using the Dynamic Provisioning command line, refer to the *Hitachi Unified Storage Command Line Interface Reference Guide*.

Related documents

This documentation set consists of the following documents.

Hitachi Unified Storage Firmware Release Notes, RN-91DF8304

Contains late-breaking information about the storage system firmware.

Hitachi Storage Navigator Modular 2 Release Notes, RN-91DF8305

Contains late-breaking information about the Navigator 2 software.

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

Hitachi Unified Storage Getting Started Guide, MK-91DF8303

Describes how to get Hitachi Unified Storage systems up and running in the shortest period of time. For detailed installation and configuration information, refer to the Hitachi Unified Storage Hardware Installation and Configuration Guide.

Hitachi Unified Storage Hardware Installation and Configuration Guide, MK-91DF8273

Contains initial site planning and pre-installation information, along with step-by-step procedures for installing and configuring Hitachi Unified Storage systems.

Hitachi Unified Storage Hardware Service Guide, MK-91DF8302

Provides removal and replacement procedures for the components in Hitachi Unified Storage systems.

Hitachi Unified Storage Operations Guide, MK-91DF8275 — this document

Describes the following topics:

- Adopting virtualization with Hitachi Unified Storage systems
- Enforcing security with Account Authentication and Audit Logging.
- Creating DP-Vols, standard VOLs, Host Groups, provisioning storage, and utilizing spares
- Tuning storage systems by monitoring performance and using cache partitioning
- Monitoring storage systems using email notifications and Hi-Track
- Using SNMP Agent and advanced functions such as data retention and power savings
- Using functions such as data migration, VOL Expansion and VOL Shrink, RAID Group expansion, DP pool expansion, and Mega VOLs

Hitachi Unified Storage Replication User Guide, MK-91DF8274

Describes how to use the four types of Hitachi replication software to meet your needs for data recovery:

- ShadowImage In-system Replication
- Copy-on-Write SnapShot
- TrueCopy Remote Replication
- TrueCopy Extended Distance

Hitachi Unified Storage Command Control Interface Installation and Configuration Guide, MK-91DF8306

Describes Command Control Interface installation, operation, and troubleshooting.

Hitachi Unified Storage Dynamic Provisioning Configuration Guide, MK-91DF8277

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

Hitachi Unified Storage Command Line Interface Reference Guide, MK-91DF8276

Describes how to perform management and replication activities from a command line.

Document conventions

The following typographic conventions are used in this document.

Convention	Description
Bold	Indicates text on a window, other than the window title, including menus, menu options, buttons, fields, and labels. Example: Click OK .

Convention	Description
<i>Italic</i>	Indicates a variable, which is a placeholder for actual text provided by you or the system. Example: copy <i>source-file target-file</i> Angled brackets (< >) are also used to indicate variables.
screen or code	Indicates text that is displayed on screen or entered by you. Example: <code># pairedisplay -g oradb</code>
< > angled brackets	Indicates a variable, which is a placeholder for actual text provided by you or the system. Example: <code># pairedisplay -g <group></code> Italic font is also used to indicate variables.
[] square brackets	Indicates optional values. Example: [a b] indicates that you can choose a, b, or nothing.
{ } braces	Indicates required or expected values. Example: { a b } indicates that you must choose either a or b.
vertical bar	Indicates that you have a choice between two or more options or arguments. Examples: [a b] indicates that you can choose a, b, or nothing. { a b } indicates that you must choose either a or b.
underline	Indicates the default value. Example: [<u>a</u> b]

This document uses the following symbols to draw attention to important safety and operational information.

Symbol	Meaning	Description
	Tip	Tips provide helpful information, guidelines, or suggestions for performing tasks more effectively.
	Note	Notes emphasize or supplement important points of the main text.
	Caution	Cautions indicate that failure to take a specified action could result in damage to the software or hardware.

The following abbreviations for Hitachi Program Products are used in this document.

Abbreviation	Description
ShadowImage	ShadowImage In-system Replication
SnapShot	Copy-on-Write SnapShot
TrueCopy	A term used when the following terms do not need to be distinguished: • True Copy • True Copy Extended Distance • True Copy remote replication
TCE	TrueCopy Extended Distance

Abbreviation	Description
Volume Migration	Modular Volume Migration
Navigator 2	Hitachi Storage Navigator Modular 2

Convention for storage capacity values

Physical storage capacity values (for example, disk drive capacity) are calculated based on the following values:

Physical capacity unit	Value
1 KB	1,000 bytes
1 MB	1,000 KB or 1,000 ² bytes
1 GB	1,000 MB or 1,000 ³ bytes
1 TB	1,000 GB or 1,000 ⁴ bytes
1 PB	1,000 TB or 1,000 ⁵ bytes
1 EB	1,000 PB or 1,000 ⁶ bytes

Logical storage capacity values (for example, logical device capacity) are calculated based on the following values:

Logical capacity unit	Value
1 block	512 bytes
1 KB	1,024 (2 ¹⁰) bytes
1 MB	1,024 KB or 1024 ² bytes
1 GB	1,024 MB or 1024 ³ bytes
1 TB	1,024 GB or 1024 ⁴ bytes
1 PB	1,024 TB or 1024 ⁵ bytes
1 EB	1,024 PB or 1024 ⁶ bytes

Accessing product documentation

The Hitachi Unified Storage user documentation is available on the HDS Support Portal: <https://portal.hds.com>. Please check this site for the most current documentation, including important updates that may have been made after the release of the product.

Getting help

The Hitachi Data Systems customer support staff is available 24 hours a day, seven days a week. If you need technical support, please log on to the HDS Support Portal for contact information: <https://portal.hds.com>

Comments

Please send us your comments on this document: doc.comments@hds.com.
Include the document title, number, and revision, and refer to specific sections and paragraphs whenever possible.

Thank you!

Introduction

This chapter provides an introduction to the Storage Navigator Modular 2 (Navigator 2).

The topics covered in this chapter are:

- ❑ [Navigator 2 overview](#)
- ❑ [Navigator 2 functions](#)

Navigator 2 overview

The Hitachi Data Systems Navigator 2 empowers you to take advantage of the full power of your Hitachi storage systems. Using Navigator 2, you can configure and manage your storage assets from a local host and from a remote host across an Intranet or TCP/IP network to ensure maximum data reliability, network up-time, and system serviceability.

The role that the Navigator 2 management console plays is to provide views of feature settings on the storage system in addition to enabling you to configure and manage those features. The following section provides more detail about what features Navigator 2 provides to optimize your experience with the Hitachi Unified Storage system.

Navigator 2 features

Navigator 2 provides the features detailed in the following sections.

Security features

- **Account Authentication** - Account authentication and audit logging provide access control to management functions.
- **Audit Logging** - Records all system changes.
- **SAN Security** - SAN security software helps ensure security in open systems storage area networking environments through restricted server access.

Monitoring features

- **Performance Monitor** - Performance monitoring software allows you to see performance within the storage system.

Configuration management features

- **LUN Manager** - Software that manages volumes streamlines configuration management processes by allowing you to define, configure, add, delete, expand, revise and reassign VOLs to specific paths without having to reboot your storage system.
- **Replication Software** - Replication setup and management feature provides basic configuration and management of Hitachi ShadowImage® products, Hitachi Copy-on-Write Snapshot software and Hitachi TrueCopy® mirrored pairs.
- **System Maintenance** - System maintenance feature allows online controller microcode updates and other system maintenance functions.
- **SNMP** - Simple Network Management Protocol (SNMP) function agent support includes MIBs specific to Hitachi Data Systems and enables SNMP-based reporting on status and alerts for Hitachi storage systems.

Data migration features

- **Modular Volume Migration Manager** - Modular volume migration software enables dynamic data migration.

- **Cache Residency Manager** - This feature allows you to "lock" and "unlock" data into a cache in real time for optimal access to your most frequently accessed data.

Capacity features

- **Cache Partition Manager** - This feature allows the application to partition the cache for improved performance.
- **RAID Group Expansion** - Online RAID group expansion feature enables dynamic addition of HDDs to a RAID group.

General features

- **Point and click GUI** - Point-and-click graphical interface with initial set-up wizards that simplifies configuration, management, and visualization of Hitachi storage systems.
- **Real-time view of environment** - An immediate view of available storage and current usage.
- **Deployment efficiency** - Efficient deployment of storage resources to meet business and application needs, optimize storage productivity, and reduce the time required to configure storage systems and balance I/O workloads.
- **Access protection** - Protection of access to information by restricting storage access at the port level, requiring case-sensitive password logins, and providing secure domains for application-specific data.
- **Data redundancy** - Protection of the information itself by letting you configure data-redundancy and assign hot spares.
- **System management** - functions for Hitachi storage systems, such as storage system status, event logging, email alert notifications, and statistics.
- **Major platform compatibility** - Compatibility with Microsoft® Windows®, UNIX, and Linux environments.
- **Online help** - Online help to enable easy access to information about use of features.
- **Command Line Interface** - A full featured and scriptable command line interface. For more information, refer to the *Hitachi Unified Storage Command Line Interface Reference Guide*.

Navigator 2 benefits

Navigator 2 provides the following benefits:

- **Simplification** - Simplifies storage configuration and management for the HUS family of storage systems.
- **Access protection** - Protects access to information by allowing secure permission to assigned storage
- **Performance enhancement** - Enhances data access performance to key applications and protects data availability of mission-critical information

- **Optimization of data retrieval** - Optimizes storage administrator productivity by reducing the time required to configure storage systems and balance I/O workloads
- **Enables integration** - Facilitates integration of Hitachi storage systems with enterprise management products
- **Cost reduction** - Reduces storage costs.
- **Long-term planning enabler** - Improves the organization's long-term sustainable business strategy.
- **Establishment of metrics** - Identifies clear metrics with a full analysis of the payback period and savings potential.
- **Capacity provisioning** - Provisions content storage capacity to organizations and to post production end users.

Navigator 2 task flow

This section details the task flow associated with the Navigator 2 Management Console.

1. You install and provision a Hitachi Unified Storage system
2. You install Navigator 2.
3. Using Navigator 2, you access data on your host systems.
4. You store data in the HUS system.
5. You create volumes and assign pieces of the stored data to volumes.
6. You partition portions of the cache on the HUS and assign data to the partitions.
7. You set up Performance Monitor to view and monitor activity on the Hitachi Unified Storage system.
8. You set up SNMP Agent Function to generate traps when certain thresholds have been exceeded.
9. You set up the Audit Logging function to send logs to the syslog when certain events occur.

Figure 1-1 shows how Navigator 2 connects directly to the front-end controller of the HUS family storage system.




Figure 1-1: Navigator 2 task flow

The front-end controller communicates to the back-end controller of the storage system, which in turn, communicates with the Storage Area Network (SAN), often a Fibre Channel switch. Hosts or application servers contact the SAN to retrieve data from the storage system for use in applications, commonly databases and data processing programs.

Navigator 2 functions

Table 1-1 details the various functions.

Table 1-1: Function details

Category	Function Name	Description	Notes	Online Usage
Components	Component status display	Displays the status of a component such as tray.		Yes
Groups	RAID Groups	RAID group: Creates, deletes, or displays a RAID group.		Yes
		VOL creation: Used to add a volume. A new volume is added by specifying its capacity.		Yes
		VOL deletion: Deletes the defined volume. User data is deleted.		Yes
		VOL formatting: Required to make a defined volume accessible by the host. Writes null data to the specified volume, and deletes user data.		Yes

Table 1-1: Function details

Category	Function Name	Description	Notes	Online Usage
	Host Groups	Review, operate, and set host groups.		Yes
	iSCSI Targets	Review, operate, and set iSCSI targets.		Yes
Settings	iSCSI Settings	View and configure iSCSI ports.		Yes
	FC Settings	View and configure FC ports.		Yes
	Port Options	View and configure port options.		Yes
	Spare Drives	View, add, or remove spare drives.		Yes
	Licenses	View, install, or de-install licensed storage features.		Yes
	Command devices	View and configure command devices.		Yes
	DMLU	View and configure the Differential management volumes for replication/migration.		Yes
	SNMP Agent	View and configure SNMP Agent Support Function		Yes
	LAN	View and configure LAN.		Yes
	Drive Recovery	View and configure options to recovery drives.		Yes
	Constitute Array	Input and output constitute array parameters.		Yes
	System Parameters	View and configure system parameters		Yes
	Verification Settings	View and configure verification for the drive and cache.		Yes
	Parity Correction	Recovery parity status of the volumes.		Yes
	Mapping Guard	View and configure Mapping Guard for the volumes		Yes
	Mapping Mode	View and configure mapping mode.		Yes
	Boot Options	View and configure boot options	Array must be restarted to enable the settings.	Yes
	Format Mode	View and configure format mode for the volume.		Yes
	Firmware	Refer/update firmware.	Array must be restarted to enable the settings.	Yes
	E-mail Alert	View and configure E-mail Alert function in the array.		Yes
	Date & Time	View and configure the Date & Time in the array.		Yes
	Advanced Settings	View and configure advanced settings.		Yes

Table 1-1: Function details

Category	Function Name	Description	Notes	Online Usage
Security	Secure LAN	Set the SSL certificate and validity/invalidity of the normal port.		Yes
Performance	Monitoring	View and output the monitored performance in the array.		Yes
	Tuning Parameter	Configure the parameter to performance in the array.		Yes
Alerts & Events	-	Displays the alerts and events.		Yes
Error Monitoring	Report when a failure occurs and controller status display	Polls the array and displays the status. If an error is detected, it is output into log.	Contact your maintenance personal.	Yes

Browser security

Depending on the setting of the browser security, the Hitachi Storage Navigator Modular 2 screen might not appear or might not be displayed correctly. In addition, the message "DMEG000220: Cannot display because JRE is not installed or it is disabled. Please install the JRE or enable it and try again." will appear.

If using Internet Explorer 6 or later, register the IP address or host name of the server running Hitachi Storage Navigator Modular 2 in the browser's Trusted sites setting. Also, use the default setting for "Trusted sites" security. If it is not default setting, check that the following items are enabled in "Trusted sites" security (they are enabled by default).

- Run ActiveX controls and plug-ins
- Script ActiveX controls marked safe for scripting
- Active scripting
- Submit nonencrypted form data

Using the Navigator 2 online help

This document covers many, but not all, of the features in Navigator 2 software. Therefore, if you need information about a Navigator 2 function that is not included in this document, please refer to the Navigator 2 online help in the Navigator GUI. To access the help, click the **Help** button on the Navigator 2 GUI and select **Help**. For convenience, the **Help** button is available regardless of the window displayed in Navigator 2.

The online help provides several layers of assistance.

- The **Contents** tab shows how the help topics are organized. You can "drill down" the topics to quickly find the support topic you are looking for, and then click a topic to view it.

- The **Index** tab lets you search for information related to a keyword. Type the keyword in the field labeled **Type in the keyword to find:** and the nearest match in the Index is highlighted. Click an index entry to see the topics related to the word. Click a topic to view it. If only one topic is related to an index entry, it appears automatically when you click the entry.
- The **Search** tab lets you scan through every help topic quickly for the word or words you are looking for. Type what you are looking for in the field labeled **Type in the word(s) to search for:** and click **Go**. All topics that contain that text are displayed. Click a topic to view it. To highlight your search results, check **Highlight search results**.




Figure 1-2: Help menu




Figure 1-3: Home page of the Navigator 2 online help

Use with Hitachi Command Suite

When operating the navigator on the common server that is identical to Hitachi Command Suite version 8.0.0 and integrating the user authentication at the time of starting up the navigator with Hitachi Command Suite, the collaboration setting is required. To uninstall Hitachi Command Suite or the navigator, canceling the collaboration setting is required. See [Operating or stopping Hitachi Command Suite \(Device Manager, Tiered Storage Manager and Replication Manager\) version 8.0.0 or later and navigator in the common server environment on page 3-9](#).

Note that Hitachi Storage Command Suite and Hitachi Command Suite are treated in the same way in this manual.

System theory of operation

This chapter describes the Navigator 2 theory of operation.

The topics covered in the chapter are:

- ❑ Network standard and functions which the array supports
- ❑ RAID features
- ❑ RAID levels
- ❑ Host volumes
- ❑ About the HUS Series of storage systems
- ❑ Major controller features
- ❑ Understanding Navigator 2 key terms
- ❑ Navigator 2 operating environment

Network standard and functions which the array supports

The user LAN port of the array supports the network standard and functions detailed in [Table 2-1](#).

Table 2-1: Network standards and functions

Item	Standard and Functions
Standard	IEEE 802.3 10BASE-T IEEE 802.3u 100BASE-TX IEEE 802.3ab 100BASE-T
Protocol	ARP, ICMP, ICMPv6, IPv4, IPv6, NDP, TCP, UDP
Routing	RIPv1, RIPv2, RIPv6
IP Address Resolution	DHCPv4 Router advertisement
Standard and function not affecting the use of the array	Port VLAN IEEE 802.1Q : Tag VLAN IEEE 802.1D : STP (Spanning Tree Protocol) IEEE 802.1w : Rapid STP (RSTP) IEEE 802.1s : Multiple Instances Spanning Tree Protocol (MISTP) IEEE 802.3ad : Link Aggregation
Communication Port	2000/tcp (Non Secure) 28355/tcp(Secure) The array uses the above TCP port for Hitachi Storage Navigator Modular 2 communication. Hi-track communication port 80

RAID features

To put RAID to practical use, some techniques such as striping, mirroring, and parity disk are used.

- **Striping** - To store data spreading it on several Disk Drives. The technique segments logically sequential files, in a way that accesses sequential segments for different physical storage devices. Striping is useful when a processing device requests access to data more quickly than a storage device can provide access. By performing segment accesses on multiple devices, multiple segments can be accessed concurrently. This provides more data access throughput, which avoids causing the processor to idly wait for data accesses.
- **Disk Drives** - The time required to access each Disk Drive is shortened and thus, time required for reading or writing is shortened.
- **Mirroring** - It means to copy all the contents of one Disk Drive to one or more Disk Drives at the same time in order to enhance reliability.
- **Parity disk** - It is a data writing method used when configure RAID with three or more Disk Drives. Parity of data in the corresponding positions of two or more Disk Drives is generated and stored on another Disk Drive.

RAID technology task flow

1. When I/O processing spans multiple Disk Drives (when the stripe size is too small) during transaction processing in RAID 5, the system does not perform optimally. So several events occur.
2. The stripe size of 256 k bytes is set as a default value in this subsystem.
3. When the Cache Partition Manager function of the priced option is used, the stripe size can be changed to 256 k bytes or 512 k bytes for each VOL.
4. Lump writing of data on the Disk Drive and pre-reading of old data are performed by use of the cache memory so as prevent occurrence of write penalty as far as possible.
5. A Write penalty may occur for various reasons.
6. In the RAID 5 configuration, 3 to 16 Disk Drives compose one parity group (2D+1P to 15D+1P);
7. In the RAID 6 configuration, 4 to 30 Disk Drives compose one parity group (2D+2P to 28D+2P). Since parity data is generated from 2 to 15 data disks in the group, when partial writing of one stripe in the group occurs in the transaction processing, it is necessary to generate the corresponding parity data in the group once again.
8. For RAID 5, since parity data is calculated by the following calculation formula, data before update parity before update and data after update are necessary to create the parity.

RAID 5 - [New parity] = ([Data before update] EOR [Data after update]) EOR [Parity before update]

RAID 6 - [New P parity] = ([Data before update] EOR [Data after update]) EOR [P parity before update]
[New Q parity] = [Coefficient parity] AND ([Data before update] EOR [Data after update]) EOR [Q parity before update]

RAID levels

Your Hitachi storage system supports various RAID configurations. Review the information in this section to determine the best RAID configuration for your requirements.

The Hitachi Unified Storage systems support RAID 0 (2D to 16D), RAID 1, RAID 5 (2D+1P to 15D+1P), RAID 6 (2D+2P to 28D+2P) and RAID 1+0 (2D+2D to 8D+8D).

[Table 2-3](#) describes RAID levels supported by the HUS systems.

Table 2-2: HDS supported RAID levels

Item	Description	Advantage/Disadvantage
RAID 0	RAID 0 stripes data across disk drives to attain higher throughput.	Advantage: Because disk drives having redundant data is not needed, disk drives can be used efficiently. Disadvantage: Data is lost in any failure of the disk drive.
RAID 1	RAID 1 provides data redundancy by copying all the contents of two disk drives to another (mirroring). Read/write performance is a little better than the individual disk drive.	Advantage: Data is not lost, even if a failure occurs in any disk drive. Performance is not lowered, even if a disk drive fails. Disadvantage: RAID 1 is expensive because it requires twice the disk capacity.
RAID 5	RAID 5 consists of three or more disk drives. It uses one of them as a parity disk and writes divided data on the other disk drives. Recovery from a failure of a data is possible by utilizing the parity data. Since the parity data is stored on all the disk drives, a bottleneck of the parity disk does not occur.	Advantage: When reading data, RAID 5 stripes data across disk drives in the same way as that in RAID 0 to attain higher throughput. Disadvantage: When writing data, since parity data is required to be updated, performance of writing small random data is lowered although there is no problem regarding writing of continuous data. The performance is also lowered when a disk drive fails.

Table 2-2: HDS supported RAID levels

Item	Description	Advantage / Disadvantage
RAID 6	Based on RAID 5, with another level of parity. making data significantly safer than RAID 5.	Advantage: Can withstand the loss of two disks in an array. While there is no performance penalty for read operations, there is a performance penalty on write operations due to overhead associated with parity calculations. Disadvantage: Imposes a write penalty, write operations require disks to read the data, read the first parity, read the second parity, write the data, write the first parity, and then finally write the second parity. This is equivalent to a six-times write penalty.
RAID 1+0	Combines two RAID levels: mirroring (RAID 1) and striping (RAID 0). Mirroring delivers the highest level of availability RAID offers, with the fastest rebuild times when a disk fails. Striping uses the appropriate chunk size for high-performance I/O.	Advantage: Like RAID 6, can withstand the loss of two disks in a RAID group, as long as both failed disks are not in the same mirror. Imposes only a 2x write performance hit. Rebuilding a failed disk is much faster than with RAID 5. Disadvantage: High space for cost — 50% of total array capacity, regardless of number of disks.

RAID chunks and stripes

A RAID Group is a logical mechanism that has two basic elements: a virtual block size from each disk (a **chunk**) and a row of chunks across the group (the RAID **stripe**). The chunk size is typically set to 64KB on midrange systems, but is adjustable on Hitachi HUS systems. In the HUS series, the RAID chunk size defaults to **256KB**, and is adjustable to either **64KB or 512KB** (per volume) via the Storage Navigator Modular management tool. This does not require the installation of the CPM package as used to be the case. Note that the Dynamic Provisioning software always uses a 256K chunk size, and this is not changeable.

The **stripe size** is the sum of the chunk sizes across a RAID Group. This only counts the “data” chunks and not any mirror or parity space. Therefore, on a RAID-6 group created as 8D+2P (ten disks), the stripe size would be 512KB (8 * 64KB chunk) or 2MB (8 * default 256KB chunk).

Note that some usage replaces *chunk* with “stripe size,” “stripe depth” or “interleave factor,” and *stripe size* with “stripe width,” “row width” or “row size.” The chunk is the primary unit of protection management for either parity or mirror RAID mechanisms.

Physical I/O is not performed on a chunk basis as is commonly thought. On Open Systems, the entire space presented by a volume is a continuous span of 512 byte blocks, known as the **Logical Block Address** range (LBA). The host application makes I/O requests using some native request size (such as a file system block size), and this is passed down to the storage as a unique I/O request. The request has the starting address (of a 512 byte block) and a length (such as the file system 8KB block size).

The storage system will locate that address within that volume to a particular disk sector address, and then proceed to read or write only that amount of data — not that entire chunk. Also note that this request could require physical I/O to two disks if the host 8KB logical block spans two chunks. It could have 2KB at the end of one chunk and 6KB on the beginning of the next chunk in that stripe.

Because of the variations of file system formatting and such, there is no way to determine where a particular block may lie on the raw space presented by a volume. Each file system will create a unique variety of metadata in a quantity and distribution pattern that is related to the size of that volume.

Most file systems also typically scatter writes around within the LBA range — an outdated holdover from long ago when file systems wanted to avoid a common problem of the appearance of bad sectors or tracks on disks. What this means is that attempts to align application block sizes with RAID chunk sizes is a pointless exercise.

These also have a native “stripe size” that is selectable when creating a logical volume from several physical storage volumes. In this case, the LVM stripe size should be a multiple of the RAID chunk size due to various interactions between the LVM and the volumes.

One example is the case of large block sequential I/O. If the LVM stripe size is equal to the RAID chunk size, then a series of requests will be issued to different volumes for that same I/O, making the request appear to be several random I/O operations to the storage system. This can defeat the system’s sequential detect mechanisms and turn off sequential prefetch, slowing down these types of operations.

Host volumes

On a midrange system, when space is carved out of a RAID Group and made into a volume. Once that volume is mapped to a host port for use by a server, it is known as a volume and is assigned a certain World Wide Name if using Fibre Channel interfaces on the system. On an iSCSI configuration, the volume gets a name that is associated with an NFS mount point.

Number of volumes per RAID group

When configuring a midrange storage system, one or more volumes can be created per RAID Group, but the goal should be to clearly understand what percentage of that group's overall capacity will contain active data. In the case where multiple hosts attempt to simultaneously use volumes that share the same physical disks in an attempt to fully utilize capacity, seek and rotational latency may become performance limiting factors. In attempting to maximize utilization, RAID Groups should contain both active and less frequently used volumes. This is true of all physical disks regardless of size, RAID level, and physical characteristics.

It is also true that, if many small volumes are carved out of a single RAID Group, their simultaneous use will create maximum seek times on each disk, reducing the maximum sustainable small block random IOPS rate to the disk's minimum.

Volume management and controller I/O management

On nearly every midrange storage system from any vendor, the individual volumes are tightly bound to an "owning" controller. This is because there is no global sharing between the controllers of either the data or its metadata. Each controller is independently responsible for managing these two objects. On enterprise storage systems, there is no concept of either a "controller" or "volume ownership." All data and metadata on most enterprise systems are globally shared by all front-end processors.

About the HUS Series of storage systems

A short discussion about the evolution of the HUS out of the AMS 2000 family of modular storage systems may be helpful to your understanding of features and concepts in the system.

the HUS is the successor to the AMS 2000, the midrange Hitachi modular storage systems that were the current price list modular family during the past three years. The HUS family systems have much higher performance and introduced features and designs from the HUS systems.

The HUS family systems have still higher performance and incorporate some significant features that were present in the AMS 2000 family, and introduce new features that were not present in the previous generation of modular devices. The HUS 110, 130, and 150 models comprise the current generation.

Recent features

A major shift in the approach to implementing storage has occurred with more instances of automatic provisioning. The DF systems achieve this approach with the following features:

Load Balancing - The HUS family uses the Hitachi Dynamic Load Balancing Controller. These are proprietary purpose-built Hitachi designs, not (like so many others) generic Intel OEM small server boards with a Windows/Linux operating system, generic Fibre Channel disk adapters, and a storage software package.

Dynamic I/O Servicing - The ability to dynamically manage I/O request execution between the controllers on a per volume basis is a significant departure from all other current midrange architectures. The back-end engine is a Serial Attached SCSI (SAS) design that allows 3.5" SAS or SSD drives to be freely intermixed in the same 15-disk trays. There is also a 24-disk 2.5" SAS tray, and a 3.5" high density drawer option which uses a pullout tray and vertically inserted disks. It holds either 38 SAS disks or 7200RPM SAS disks with no intermixing, and no SSDs.

Major controller features

- **Active / Active Symmetric front-end design** - Allows use of any port to dynamically access any
- **SAS Matrix Engine back-end architecture** - Provides SAS controllers, more paths and a dynamic, fault-tolerant matrix connection from the SAS controllers to the individual drives
- **Hardware I/O Load Balancing** - Maintains a more even distribution of backend I/O workloads between the SAS Matrix Engines of the two controllers over time
- **Hitachi Dynamic Provisioning** - Separates the logical from the physical allocation (thus delaying the storage purchasing decision), and to spread the I/Os across every RAID Group in the HDP pool (wide striping)
- **A standard 15"-disk tray** - Used for 3.5" SAS and SSD drives

Understanding Navigator 2 key terms

Before you install the Navigator 2 software, it is important to understand a few key terms associated with Navigator 2. [Table 2-3](#) defines a few key terms associated with Navigator 2.

Table 2-3: Understanding Navigator 2 key terms

Term	Explanation
Host group	A group that virtualizes access to the same port by multiple hosts since host settings for a volume are not made at the physical port level but at a virtual port level.
Profile	A set of attributes that are used to create a storage pool. The system has a predefined set of storage profiles. You can choose a profile suitable for the application that is using the storage, or you can create a custom profile.
Pool	A collection of volumes with the same configuration. A storage pool is associated with a storage profile, which defines the storage properties and performance characteristics of a volume.

Table 2-3: Understanding Navigator 2 key terms

Term	Explanation
Snapshot	A point-in-time copy of a primary volume. The snapshot can be mounted by an application and used for backup, application testing, or data mining without requiring you to take the primary volume offline.
Storage domain	A logical entity used to partition storage.
Volume	A container into which applications, databases, and file systems store data. Volumes are created from virtual disks, based on the characteristics of a storage pool. You map a volume to a host or host group.
RAID	Redundant Array of Independent Disks (RAID) — A disk array in which part of the physical storage capacity is used to store redundant information about user data stored on the remainder of the storage capacity. The redundant information enables regeneration.
Parity Disk	A RAID-3 disk that provides redundancy. RAID-3 distributes the data in stripes across all but one of the disks in the array. It then writes the parity in the corresponding stripe on the remaining disk. This disk is the parity disk.
Volume (formerly called LUN)	Logical unit number (LUN) — An address for an individual disk drive, and by extension, the disk device itself. Used in the SCSI protocol as a way to differentiate individual disk drives within a common SCSI target device, like a disk array. Volumes are normal.
iSCSI	Internet-Small Computer Systems Interface (iSCSI) — A TCP/IP protocol for carrying SCSI commands over IP networks.
iSCSI Target	A system component that receives an iSCSI I/O command. The command is sent to the iSCSI bus address of the target device or controller.
iSCSI Initiator	The component that transmits an iSCSI I/O command to the iSCSI bus address of the target device or controller.

Navigator 2 operating environment

You install Navigator 2 on a management platform (a PC, a Linux workstation, or a laptop) that acts as a console for managing your HUS storage system. This PC management console connects to the management ports on the HUS storage system controllers, and uses Navigator 2 to manage your storage assets and resources. The management console can connect directly to the management ports on the HUS storage system or via a network hub or switch.

Before installing Navigator 2 on the management console, confirm that the console meets the requirements in the following sections. For an optimum Navigator 2 experience, the management console should be a new or dedicated PC.



TIP: To obtain the latest compatibility information about supported operating systems, NICs, and various devices, see the Hitachi Data Systems interoperability matrix at <http://www.hds.com/products/interoperability/>.

Firewall considerations

A firewall's main purpose is to block incoming unsolicited connection attempts to your network. If the HUS storage system is used within an environment that uses a firewall, there will be times when the storage system's outbound connections will need to traverse the firewall.

The storage system's incoming indication ports are ephemeral, with the system randomly selecting the first available open port that is not being used by another Transmission Control Protocol (TCP) application. To permit outbound connections from the storage system, you must either disable the firewall or create or revise a source-based firewall rule (not a port-based rule), so that items coming from the storage system are allowed to traverse the firewall.

Firewalls should be disabled when installing Navigator 2 (refer to the documentation for your firewall). After the installation completes, you can turn on your firewall.



NOTE: For outgoing traffic from the storage system's management port, there are no fixed port numbers (ports are ephemeral), so all ports should be open for traffic from the storage system management port.

If you use Windows firewall, the Navigator 2 installer automatically registers the Navigator 2 file and Command Suite Common Components as exceptions to the firewall. Therefore, before you install Navigator 2, confirm that no security violations exist.

Anti-virus software considerations

Anti-virus programs, except Microsoft Windows' built-in firewall, must be disabled before installing Navigator 2. In addition, Navigator 2 cannot operate with firewalls that can terminate local host socket connections. As a result, configure your anti-virus software to prevent socket connections from being terminated at the local host (refer to the documentation for your anti-virus software).

Hitachi Storage Command Suite common components

Before installing Navigator 2, be sure no products other than Hitachi Storage Command Suite Common Component are using port numbers 1099, 23015 to 23018, 23032, and 45001 to 49000. If other products are using these ports, you cannot start Navigator 2, even if the Navigator 2 installation completes without errors.

If other Hitachi Storage Command products are running:

- Stop the services or daemon process for those products.
- Be sure any installed Hitachi Storage Command Suite Common Components are not operating in a cluster configuration. If the host is in the cluster configuration, configure it for a stand-alone configuration according to the manual.
- Back up the Hitachi Storage Command database before installing Navigator 2.

Installation

This chapter provides information on installing and enabling features.

After ensuring that your configuration meets the system requirements described in the previous chapter, use the instructions in this chapter to install the Navigator 2 software on your management console PC.

The topics covered in this chapter are:

- ❑ [Connecting Hitachi Storage Navigator Modular 2 to the Host](#)
- ❑ [Types of installations](#)
- ❑ [Installing Navigator 2](#)
- ❑ [Preinstallation information for Storage Features](#)
- ❑ [Installing storage features](#)
- ❑ [Uninstalling storage features](#)
- ❑ [Starting Navigator 2 host and client configuration](#)
- ❑ [Operations](#)
- ❑ [Understanding the Navigator 2 interface](#)
- ❑ [Performing Navigator 2 activities](#)
- ❑ [Description of Navigator 2 activities](#)

Connecting Hitachi Storage Navigator Modular 2 to the Host

You can connect Hitachi Storage Navigator Modular 2 to a host through a LAN with or without a switch.

When two or more LAN cards are installed in a host and a segment set in each LAN card is different from the others, Hitachi Storage Navigator Modular 2 can only access from the LAN card side specified by the installer. When accessing the array unit from the other segment, make the configuration that a router is used. Install one LAN card in the host to be installed.



NOTE: If an array unit is already connected with a LAN, a host is connected to the same network as the array unit.

Installing Hitachi Storage Navigator Modular 2

When Storage Navigator Modular is installed, Hitachi Storage Navigator Modular 2 cannot perform an update from Storage Navigator Modular.

Preparation

Make sure of the following on the host in which Hitachi Storage Navigator Modular 2 is to be installed before starting installation:

When the preparation items are not done correctly, installation may not be completed. It is usually completed in about 30 minutes. If it is not completed even one hour or more passes, terminate the installer forcibly and check that the preparation items are correctly done.

- For Windows, when you install Hitachi Storage Navigator Modular 2 to the C: partition. A filename "program" is required to be placed directly under the C: partition.
- For Windows, you are logged on to Windows as an Administrator or a member of the Administrators group.
For Linux and Solaris, you are logged on to as a root user.
- To install Hitachi Storage Navigator Modular 2, the following free disk capacity is required.

Table 3-1 details free disk capacity values.

Table 3-1: Free disk capacity

OS	Directory	Free Disk Capacity
Windows	Installed directory	1.5 GB
Linux	/opt/HiCommand	1.5 GB
Solaris	/opt/HiCommand	1.5 GB
	/var/opt/HiCommand	1.0 GB
	/var/tmp	1.0 GB

When install Hitachi Storage Navigator Modular 2, it is not required exists the above directory. If directories do not exist, above directories are required to have enough free space.

- For Linux and Solaris, when the /opt exists, the normal directory is required (not the symbolic link). However, the file system may be mounted as a mount point.
- For Linux and Solaris, the kernel parameters must be set correctly. For more details, see section 0 or 805518148.
- The following patch must be applied to Solaris 10 (SPARC).

The patch 120664-xx (xx: 01 or later)

The patch 127127-xx (xx: 11 or later)

Do not apply the patch 127111-02 and 127111-03

- The following patch must be applied to Solaris 10 (x64).

The patch 120665-xx (xx: 01 or later)

Do not apply the patch 127112-02 and 127112-03

- Products other than Hitachi Storage Command Suite Common Component are not using port numbers 1099, 23015 to 23018, 23032, and 45001 to 49000.

If other products are using these ports, you cannot start Hitachi Storage Navigator Modular 2, even if the installation of Hitachi Storage Navigator Modular 2 has finished normally. Make sure that no other products are using these ports, and then begin the installation. You can change the port numbers 1099 and 23015 after the installation. Refer to section 805518148 more details. If these port numbers have already been changed and used in an environment where Hitachi Storage Command Suite Common Component is installed, you can use the changed port numbers to install Hitachi Storage Navigator Modular 2. You do not have to change the port numbers back to the default.

- No other Hitachi Storage Command product is running.

When applications are running, stop the services (daemon process) according to the operation manual of each application.

- The installed Hitachi Storage Command Suite Common Components must not be operated in a cluster configuration.

When the host is in the cluster configuration, you cannot install Hitachi Storage Navigator Modular 2. In case of a cluster configuration, change it to the stand-alone configuration according to the manual.

- Dialog boxes used for operating Windows services, such as Computer Management or Services, are not displayed.

When you display a window, you may not able to install Hitachi Storage Navigator Modular 2. If the installation is not completed after one hour elapsed, terminate the installation forcibly and check if the window is displayed.

- Services (daemon process) such as process monitoring and virus monitoring must not be operating.

When the service (daemon process) is operating, you may not be able to install Hitachi Storage Navigator Modular 2. If the installation is not completed after one hour elapsed, terminate the installation forcibly and check what service (daemon process) is operating.

- When third-party-made firewall software other than Windows firewall is used, it must be invalidated during the installation or un-installation.

When you are using the third party- made firewall software, if the installation of Hitachi Storage Navigator Modular 2 is not completed after one hour elapsed, terminate the installation forcibly and check if the third party-made firewall software is invalidated.

- For Linux and Solaris environment, the firewall must be invalidated.

To invalidate the firewall, see the each firewall manual.

- Some of the firewall functions provided by the OS might terminate socket connections in the local host. You cannot install and operate Hitachi Storage Navigator Modular 2 in an environment in which socket connections are terminated in the local host. When setting up the firewall provided by the OS, configure the settings so that socket connections cannot be terminated in the local host.
- Windows must be set to produce the 8+3 form file name that is compatible with MS-DOS.

There is no problem because Windows creates the 8+3 form file name in the standard setting. When the tuning tool of Windows is used, the standard setting may have been changed. In that case, return the setting to the standard one.

- Hitachi Storage Navigator Modular 2 for Windows supports the Windows Remote Desktop functionality. Note that the Microsoft terms used for this functionality differ depending on the Windows OS. The following terms can refer to the same functionality:
 - Terminal Services in the Remote Administration mode
 - Remote Desktop for Administration
 - Remote Desktop connection

When using the Remote Desktop functionality to perform Hitachi Storage Navigator Modular 2 operation (including installation or un-installation), you need to connect to the console session of the target server in advance. However, even if you have successfully connected to the console session, the product might not work properly if another user connects to the console session.

- Windows must be used in the application server mode of the terminal service and must not be installed in the execution mode.

When installing Hitachi Storage Navigator Modular 2, do not use the application server mode of the terminal service. If the installer is executed in such an environment, the installation may fail or the installer may become unable to respond.



NOTE: Before installing Hitachi Storage Navigator Modular 2 on a host in which another Hitachi Storage Command product has already been installed, back up the database. However, you install Hitachi Storage Navigator Modular 2 only, it is not necessary to back up.



NOTE: When installing Hitachi Storage Navigator Modular 2 in Windows Server 2003 SP1 or Windows XP SP2 or later, you need to specify the following settings if Data Execution Prevention is being used:

Settings When Data Execution Prevention Is Enabled

If Data Execution Prevention (DEP) is enabled in Windows, sometimes installation cannot start. In this case, use the following procedure to disable DEP and then re-execute the installation operation.

To disable DEP

1. Choose **Start, Settings, Control Panel**, and then **System**.

The **System Properties** dialog box appears.

2. Select the **Advanced** tab, and under **Performance** click **Settings**.

The **Performance Options** dialog box appears.

3. Select the **Data Execution Prevention** tab, and select the **Turn on DEP for all programs and services except those I select** radio button.

4. Click **Add** and specify Hitachi Storage Navigator Modular 2 installer (HSNM2- xxxx-W-GUI.exe). (The portion “xxxx” of file names varies with the version of Hitachi Storage Navigator Modular 2, etc.)

Hitachi Storage Navigator Modular 2 installer (HSNM2-xxxx-W-GUI.exe or Setup.exe) is added to the list.

5. Select the checkbox next to Hitachi Storage Navigator Modular 2 installer (HSNM2-xxxx-W-GUI.exe) and click **OK**.

Automatic exception registration of Windows firewall:

When Windows firewall is used, the installer for Hitachi Storage Navigator Modular 2 automatically registers the file of Hitachi Storage Navigator Modular 2 and that included in Hitachi Storage Command Suite Common Components as exceptions to the firewall. Check that no problems of security exist before executing the installer.

Setting Linux kernel parameters

When you install Hitachi Storage Navigator Modular 2 to Linux, set the Linux kernel parameters. Otherwise, the installer ends without installing the software. The only exception is if Navigator 2 has already been installed and used in a Hitachi Storage Command Suite Common Component environment. In this case, you do not need to set the Linux kernel parameters.

To set the Linux kernel parameters

1. Back up the kernel parameters setting file (**/etc/sysctl.conf** and **/etc/security/limits.conf**).
2. Ascertain the IP address of the management console (for example, using **ipconfig** in a DOS environment). Then change its IP address to 192.168.0.x where x is a number from 1 to 254, excluding 16 and 17. Write this IP address on a piece of paper. You will be prompted for it during the Storage Navigator Modular 2 installation procedure.
3. Disable popup blockers in your Web browser. We also recommend that you disable anti-virus and proxy settings on the management console when installing the Storage Navigator Modular 2 software.
4. To log in to Storage Navigator Modular 2 with a Red Hat Enterprise Linux (RHEL) operating system, modify the kernel settings as follows:
 - **SHMMAX** parameter. This parameter defines the maximum size, in bytes, of a single shared memory segment that a Linux process can allocate in its virtual address space. If the RHEL default parameter is larger than both the SNM2 and Database values, you do not need to change it.
 - **SHMALL** parameter. This parameter sets the total amount of shared memory pages that can be used system wide. For SNM2, this value must equal the sum of the default value, SNM2, and Database values.
 - **Other** parameters. The following parameters follow the same rule as **SHMALL** and must be the higher of (**RHEL current value + value in Navigator 2 column**) or the value from the Database value.
 - **kernel.shmmni**
 - **kernel.threads-max**
 - **kernel.msgmni**
 - **kernel.sem** (second parameter)
 - **kernel.sem** (fourth parameter)
 - **fs.file-max** **nofile** **nproc**

Table 3-2 details recommended values for Linux kernel parameters.

Table 3-2: Linux kernel parameters

Parameter Name	Standard RHEL 5.x	Sample Customer Values	Storage Navigator Modular 2	SNM2 Database	Required New Value
kernel.shmmax	4294967295	4294967295	11542528	200000000	4294967295
kernel.shmall	268435456	268435456	22418432	22418432	22418432
kernel.shmmni	4096	4096	0	2000	2000
kernel.threads-max	65536	122876	184	574	123060

Table 3-2: Linux kernel parameters

kernel.msgmni	32	32	32	32	64
kernel.sem (Second parameters)	32000	32000	80	7200	32080
kernel.sem (Fourth parameters)	128	128	9	1024	1024
fs.file-max	205701	387230	53898	53898	441128
nofile	0	0	572	1344	1344
nproc	0	0	165	512	512

- Open the kernel parameters setting file (**/etc/sysctl.conf**) with a standard text editor and change referring to the following.

The parameters are specified using the form, which is **[name of parameter]=[value]**. Four values separated by space are specified in **kernel.sem**.

Then, the parameter must not exceed the maximum value that OS specifies.

The value can be checked by the following command.

cat /proc/sys/kernel/shmmax (Case: Check value of **kernel.shmmax**)

The default physical management port IP addresses are set to:

- Controller 0: 192.168.0.16
- Controller 1: 192.168.0.17

- Reboot host.

Setting Solaris 8 or Solaris 9 kernel parameters

When you install Hitachi Storage Navigator Modular 2 to Solaris 8 or Solaris 9, you must set the Solaris kernel parameters. If you not set the Solaris kernel parameters, Hitachi Storage Navigator Modular 2 installer terminates abnormally. Besides, when the application has already been installed and used in an environment that contains Hitachi Storage Command Suite Common Component, you do not need to set the Solaris kernel parameters.

To set the Solaris kernel parameters

- Back up the kernel parameters setting file (**/etc/system**).
- Open the kernel parameters setting file (**/etc/system**) with edit editor and add the following text line to bottom.

When a certain value has been set in the file, revise the existing value by adding the following value within the limit that the value does not exceed the maximum value which each OS specifies. For the maximum value, refer to the manual of each OS.

```
prctl -i project 'user.root'
prctl -i project 'system'
```



NOTE: The shmsys:shminfo_shmseg is not used in Solaris 9. But there is no influence even if sets it.

3. Reboot the Solaris host and then install Hitachi Storage Navigator Modular 2.

Setting Solaris 10 kernel parameters

When you install Hitachi Storage Navigator Modular 2 using Solaris 10, you must set the Solaris kernel parameters. If you do not set the Solaris kernel parameters, Hitachi Storage Navigator Modular 2 installer terminates abnormally. When the application has already been installed and used in an environment where Hitachi Storage Command Suite Common Component is present, you do not need to set the Solaris kernel parameters.

To set the Solaris kernel parameters

1. Back up the kernel parameters setting file (/etc/project).
2. From the console, execute the following command and then check the current parameter value.

```
prctl -i project 'user.root'  
prctl -i project 'system'
```

3. From the console, execute the following command and then set the parameters.

When a certain value has been set, revise the existing value by adding the following value within the limit that the value does not exceed the maximum value which each OS specifies. For the maximum value, refer to the manual of each OS.

The parameter must be set for the both projects, user.root and system.

```
projmod -a -K 'process.max-msg-messages=(priv,480,deny)' 'user.root'  
projmod -a -K 'process.max-sem-nsems=(priv,128,deny)' 'user.root'  
projmod -a -K 'process.max-sem-ops=(priv,128,deny)' 'user.root'  
projmod -a -K 'project.max-msg-ids=(priv,32,deny)' 'user.root'  
projmod -a -K 'project.max-sem-ids=(priv,1024,deny)' 'user.root'  
projmod -a -K 'project.max-shm-ids=(priv,2000,deny)' 'user.root'  
projmod -a -K 'project.max-shm-memory=(priv,26214400,deny)' 'user.root'  
projmod -a -K 'process.max-msg-messages=(priv,480,deny)' 'system'  
projmod -a -K 'process.max-sem-nsems=(priv,128,deny)' 'system'  
projmod -a -K 'process.max-sem-ops=(priv,128,deny)' 'system'  
projmod -a -K 'project.max-msg-ids=(priv,32,deny)' 'system'  
projmod -a -K 'project.max-sem-ids=(priv,1024,deny)' 'system'  
projmod -a -K 'project.max-shm-ids=(priv,2000,deny)' 'system'  
projmod -a -K 'project.max-shm-memory=(priv,26214400,deny)' 'system'
```

4. Reboot the Solaris host and then install Hitachi Storage Navigator Modular 2.



NOTE: In case of the setting of the kernel parameters is not enabled in Solaris 10, open the file (/etc/system) with text editor and change referring to the following before reboot host.

```
set msgsys:msginfo_msgmni=128
set msgsys:msginfo_msgtql=8192
set semsys:seminfo_semmni=1024
set semsys:seminfo_semmsl=512
set semsys:seminfo_semopm=512
set shmsys:shminfo_shmmax=262144000
set shmsys:shminfo_shmmni=2000
```

Operating or stopping Hitachi Command Suite (Device Manager, Tiered Storage Manager and Replication Manager) version 8.0.0 or later and navigator in the common server environment

Perform the following procedure to install or uninstall either of Hitachi Command Suite version 8.0.0 or later and the navigator by Windows or Linux 64-bit OS (excluding Solaris). Otherwise, collaboration with Hitachi Command Suite version 8.0.0 or later or uninstall might fail.

The collaboration described here refers to integrating Hitachi Command Suite and the user authentication when starting navigator into the user account management server of Hitachi Command Suite. By doing so, you can start the navigator using the user that was created by Hitachi Command Suite.

For this collaboration, registering the user account management server of the navigator in the user account management server of Hitachi Command Suite is required.



CAUTION! To collaborate with Hitachi Command Suite, do not change the port number and connection address. Otherwise, the collaboration setting is canceled.

Perform the operation by reference to the following cases.

Table 3-3: Cases

Case	Already installed		Operation to be performed
	Hitachi Command Suite version 8.0.0 or later	Navigator	
A	Yes	No	Install navigator
B	No	Yes	Install Hitachi Command Suite version 8.0.0 or later
C	Yes	Yes	Uninstall navigator

Table 3-3: Cases

Case	Already installed		Operation to be performed
	Hitachi Command Suite version 8.0.0 or later	Navigator	
D	Yes	Yes	Uninstall Hitachi Command Suite version 8.0.0 or later

Case A: When installing the navigator in the environment where Hitachi Command Suite version 8.0.0 or later is already installed

Set the following after installing the navigator.

1. Change the user account management server of the navigator to Hitachi Command Suite.

In case of Windows:

```
[Navigator install path]\bin>hcmdsprmset /host [IP address or host name] /port [Non-SSL port number]
```

In case of Linux:

```
[Navigator install path]/bin/hcmdsprmset -host [IP address or host name] -port [Non-SSL port number]
```

[Navigator install path]:

If it was installed by default, the following path is applied.

- Windows: C:\Program Files (x86)\HiCommand\Base
- Linux: /opt/HiCommand/Base

[IP address or host name]:

Specify one of the following:

Host name of the installed machine

IP address of the installed machine (127.0.0.1 cannot be specified)

Host name (excluding localhost) that is resolved in name for 127.0.0.1 by the hosts file



NOTE: We recommend a host name (excluding localhost) that is resolved in name for 127.0.0.1 by the hosts file to avoid that the communication between Hitachi Command Suite and the navigator is not done through a network.

[Non-SSL port number]:

Specify a non-SSL port for connecting to Hitachi Command Suite.

The default value is 22015.

2. Start Hitachi Command Suite again.

- Check that Hitachi Command Suite is running.
- When Hitachi Command Suite is running, stop Hitachi Command Suite and monitor Hitachi Command Suite until it stops.
- Start Hitachi Command Suite.
- Check that you can log into Hitachi Command Suite.



CAUTION! Refer to the Hitachi Command Suite manual version 8.0.0 or later to start, stop, and check the status of Hitachi Command Suite version 8.0.0 or later.

3. Start the common component of the navigator again.
 - Stop the service (daemon process) of the Hitachi Storage Command Suite common component.
 - Start the service (daemon process) of the Hitachi Storage Command Suite common component.
 - Stop or start the service (daemon process).

Case B: When installing Hitachi Command Suite version 8.0.0 or later in the environment where the navigator is already installed

Install according to the of Hitachi Command Suite manual version 8.0.0 or later.

Case C: When uninstalling the navigator while operating it in Hitachi Command Suite version 8.0.0 or later

Perform the following procedure to uninstall navigator.



CAUTION! To continue the application that uses the other Hitachi Storage Command Suite common component after uninstalling the navigator, see the manual of each product.

1. Stop the service (daemon process) of the Hitachi Storage Command Suite common component.
2. Stop the service of the common component of Hitachi Command Suite version 8.0.0 or later.



CAUTION! Follow the manual of Hitachi Command Suite version 8.0.0 or later to start, stop and check the status of Hitachi Command Suite version 8.0.0 or later.

3. Change the user account management server for navigator from Hitachi Command Suite to the default.



CAUTION! When the application that uses the other Hitachi Storage Command Suite common component is installed, setting the port number to the one that was set as the start port of Hitachi Command Suite version 8.0.0 or later is required after completing the uninstallation.

For Windows:

```
[Navigator install path]\bin>hcmdsprmset /host 127.0.0.1 /port 23015
```

For Linux:

```
[Navigator install path]/bin/hcmdsprmset -host 127.0.0.1 -port 23015
```

[Navigator install path]:

If it was installed by default, the following path is applied.

- Windows: C:\Program Files (x86)\HiCommand\Base
 - Linux: /opt/HiCommand/Base
4. Uninstall the navigator.
 5. Delete the navigator information from Hitachi Command Suite.
 6. Delete the navigator information in accordance with the Hitachi Command Suite manual version 8.0.0 or later.

Case D: When uninstalling Hitachi Command Suite version 8.0.0 or later while operating the navigator in Hitachi Command Suite version 8.0.0 or later

Change the navigator start port after uninstalling Hitachi Command Suite version 8.0.0 or later. The user information set in Hitachi Command Suite version 8.0.0 or later is deleted.

1. Change the navigator start port number. The default value of the non-SSL port is 23015.

For Windows:

```
[Navigator install path]\bin>hcmdsprmset /host 127.0.0.1 /port [Non-SSL port number]
```

For Linux:

```
[Navigator install path]/bin/hcmdsprmset -host 127.0.0.1 -port [Non-SSL port number]
```

[Navigator install path]:

If it was installed by default, the following path is applied.

- Windows: C:\Program Files (x86)\HiCommand\Base
- Linux: /opt/HiCommand/Base

Types of installations

Navigator 2 supports two types of installations:

- Interactive installations — attended installation that displays graphical windows and requires user input.
- Silent installations — unattended installation using command-line parameters that do not require any user input.

This chapter describes the interactive installation procedure. For information about performing silent installations using CLI commands, refer to the Hitachi Storage Navigator Modular 2 Command Line Interface (CLI) Reference Guide or the Navigator 2 online help.

Before proceeding, be sure you reviewed and completed all pre-installation requirements described earlier in this chapter in [Preinstallation information for Storage Features on page 3-20](#).

Installing Navigator 2

The following sections describe how to install Navigator 2 on a management console running one of the Windows, Solaris, or Linux operating systems that Navigator 2 supports.

During the Navigator installation procedure, the installer creates the directories `_HDBInstallerTemp` and `StorageNavigatorModular`. You can delete these directories if necessary.

To perform this procedure, you need the IP address (or host name) and port number that will be used to access Navigator 2. Avoid port number 1099 if this port number is available and use a port number such as 2500 instead.



NOTE: Installing Navigator 2 also installs the Hitachi Storage Command Suite Common Component. If the management console has other Hitachi Storage Command products installed, the Hitachi Storage Command Suite Common Component overwrites the current Hitachi Storage Command Suite Common Component.

Getting started (all users)

For all users, to start the Navigator 2 installation

1. Find out the IP address of the management console (e.g., using **ipconfig** on Windows or **ifconfig** on Solaris and Linux). This is the IP address you use to log in to Navigator 2, so long as it is a static IP address. Record this IP address. You will be prompted for it during the Navigator 2 installation procedure.



NOTE: On Hitachi storage systems, the default IP addressed for the management ports are 192.168.0.16 for Controller 0 and 192.168.0.17 for Controller 1.

2. Disable pop-up blockers in your Web browser. We also recommend that you disable anti-virus software and proxy settings on the management console when installing the Navigator 2 software.
3. Proceed to the appropriate section for the operating system running on your management console:
 - Microsoft Windows. See [Installing Navigator 2 on a Windows operating system](#), below.
 - Solaris. See [Installing Navigator 2 on a Sun Solaris operating system on page 3-15](#).

- Red Hat Enterprise Linux 4 and Red Hat Enterprise Linux 4 Linux. See [Installing Navigator 2 on a Red Hat Linux operating system on page 3-16](#).

The installation process takes about 15 minutes to complete. During the installation, the progress bar can pause for several seconds. This is normal and does not mean the installation has stopped.

Installing Navigator 2 on a Windows operating system

To install Navigator 2

1. Windows Vista users: Open the command prompt and then close it.
2. Insert the Navigator 2 CD in the management console CD drive and follow the installation wizard. If the CD does not auto-run, double-click the following file, where *nnnn* is the Navigator 2 version number:
`\program\hsnm2_win\HSNM2-nnnn-W-GUI.exe`
3. When prompted for an IP address, enter the IP address for your management console, which you obtained in step 1 and recorded in [Appendix B](#).
4. After you insert the Hitachi Storage Navigator Modular 2 installation CD-ROM into the management console's CD/DVD-ROM drive, the installation starts automatically and the Welcome window appears.
5. Click **Next** two times until the Choose Destination Location window appears.
6. Install Navigator 2 in the default destination folder shown or click the **Browse** button to select a different destination folder.
7. Click **Next**. The Input the IP address and port number of the PC window appears.
8. Enter the following information:
 - **IP Addr.** Enter the IP address or host name used to access Navigator 2 from your browser. Do not specify 127.0.0.1 and localhost.
 - **Port No.** Enter the port number used to access Navigator 2 from your browser. The default port number is 1099.



TIP: For environments using Dynamic Host Configuration Protocol (DHCP), enter the host name (computer name) for the IP address. If you are configuring Navigator 2 for one IP address, you can omit the **IP Addr.**

9. Click **Next**. The Start Copying Files window shows the installation settings you selected.
10. Review the settings to make sure they are correct. To change any, click **Back** until you return to the appropriate window, make the change, and click **Next** until you return to the Start Copying Files window.

11. In the Start Copying Files window, click **Next** to start the installation. During the installation, windows show the progress of the installation. When installation is complete, the InstallShield Wizard Complete window appears. You cannot stop the installation after it starts.
12. In the InstallShield Wizard Complete window, click **Finish** to complete the installation. Then proceed to for a description of the Navigator 2 interface. Navigator 2 is registered in All Programs in the Start menu.
13. Proceed to [Starting Navigator 2 host and client configuration on page 3-27](#) for instructions about logging in to Navigator 2.

If your Navigator 2 installation fails, see [If the installation fails on a Windows operating system on page 3-15](#).

If the installation fails on a Windows operating system

Data Execution Prevention (DEP) is a Windows security feature intended to prevent an application or service from executing code from a non-executable memory region. DEP perform checks on memory to prevent malicious code or exploits from running on the system by shut down the process once detected. However, DEP can accidentally shut down legitimate processes, like your Navigator 2 installation.

If your management console runs Windows Server 2003 SP1 or Windows XP SP2 or later, and your Navigator 2 installation fails, disable DEP.

To disable DEP

1. Click **Start**, and then click **Control Panel**.
2. Click **System**.
3. In the System Properties window, click the **Advanced** tab.
4. In the **Performance** area, click **Settings** and then click the **Data Execution Prevention** tab.
5. Click **Turn on DEP for all programs and services except those I select**.
6. Click **Add** and specify the Navigator 2 installer HSNM2-xxxx-W-GUI.exe, where xxxx varies with the version of Navigator 2. The Navigator 2 installer HSNM2-xxxx-W-GUI.exe is added to the list.
7. Click the checkbox next to the Navigator 2 installer HSNM2-xxxx-W-GUI.exe and click **OK**.

Installing Navigator 2 on a Sun Solaris operating system

The following procedure describes how to install Navigator 2 on a Navigator 2-supported version of Sun Solaris. Before you perform the following procedure, be sure that the following directories have at least the minimum of amount of available disk space shown in [Table 3-4](#).

Table 3-4: Solaris directories and disk space

Directory	Minimum Available Disk Space Required
/opt/HiCommand	1.5 GB
/var/opt/HiCommand	1.0 GB
/var/tmp	1.0 GB

To perform a new installation for Sun Solaris

1. Insert the Hitachi Storage Navigator Modular 2 installation CD-ROM into the management console's CD/DVD-ROM drive.



NOTE: If the CD-ROM cannot be read, copy the files `install-hsnm2.sh` and `HSNM2-XXXX-S-GUI.tar.gz` to a file system that the host can recognize.

2. Mount the CD-ROM on the file system. The mount destination is `/cdrom`.
3. Create a temporary directory with sufficient free space (more than 600 MB) on the file system and expand the compressed files. The temporary directory is `/temporary` here.
4. In the console, issue the following command lines. In the last command, `XXXX` varies with the version of Navigator 2.

```
mkdir /temporary
cd /temporary
gunzip < /cdrom/HSNM2-XXXX-S-GUI.tar.gz | tar xf -
```

5. In the console, issue the following command line:

```
/temporary/install-hsnm2.sh -a [IP address] -p [port number]
```

In this command line:

- `[IP address]` is the IP address used to access Navigator 2 from your browser. When entering an IP address, do not specify `127.0.0.1` and `localhost`. For DHCP environments, specify the host name (computer name).
- `[port number]` is the port number used to access Navigator 2 from your browser. The default, port number is 1099. If you use it, you can omit the `-p` option from the command line.



TIP: For environments using DHCP, enter the host name (computer name) for the IP address.

6. Proceed to [Starting Navigator 2 on page 3-33](#) for instructions about logging in to Navigator 2.

Installing Navigator 2 on a Red Hat Linux operating system

To install Navigator 2 on a Navigator 2-supported version of Red Hat Linux

1. Insert the Hitachi Storage Navigator Modular 2 installation CD-ROM into the management console's CD/DVD-ROM drive.



NOTE: If the CD-ROM cannot be read, copy the files `install-hsnm2.sh` and `HSNM2-XXXX-L-GUI.rpm` to a file system that the host can recognize.

2. Mount the CD-ROM on the file system. The mount destination is `/cdrom`.
3. In the console, issue the following command line:

```
sh /cdrom/install-hsnm2.sh -a [IP address] -p [port number]
```

In this command line:

- `[IP address]` is the IP address used to access Navigator 2 from your browser. When entering an IP address, do not specify `127.0.0.1` and `localhost`. For DHCP environments, specify the host name (computer name).
 - `[port number]` is the port number used to access Navigator 2 from your browser. The default port number is `1099`. If you use it, you can omit the `-p` option from the command line.
4. Proceed to [Chapter 4, Provisioning](#) for instructions about logging in to Navigator 2.

Updating Navigator 2

When you update the installed Navigator 2 to a newer version, you can perform the update installation using the installer. When you install the Navigator 2 of the same version as an installed instance of the software, the uninstaller starts, uninstalls the existing version and then installs the new version.

If you are using the installer by connecting with `https`, you need to set the server certificate and private key again after completing the update. When using the JRE by switching, it becomes the JRE in Navigator 2 by the update. Change the JRE again after completing the update.

Note the following restrictions for updating Navigator 2.

You cannot update Navigator 2 of the former version with any instance other than the installed version. When you need to return to the former version, uninstall Navigator 2 once. If you fail to reinstall it after uninstalling Navigator 2, restart the host, check that the installation of Navigator 2 is ready, and then install it again.

If you update Navigator 2 to version 5.00 or more, the login screen displays. If it does not display, perform the following tasks:

- Delete all files in the Temporary Internet Files folder.
- Clear entries in the History folder.
- Close all screens in the open browser.

To upgrade Navigator 2 for Windows:

1. Insert the Navigator 2 installation DVD-ROM.

The installation starts automatically. If you perform the installation later, terminate it. After terminating, display the contents of the DVD-ROM with Explorer and run HSNM2-xxxx-W-GUI.exe. The portion "xxxx" of filenames varies with the version of Navigator 2.

The Welcome to the InstallShield Wizard for Hitachi Storage Navigator Modular 2 dialog box displays.

2. Click Next.

Clicking Next displays the Confirm stopping the services dialog box.

3. Click Next.

Clicking Next displays the Confirm before updating dialog box.

4. Confirm the displayed information and then click Next.

Installation starts. During the installation, dialog boxes indicate the processing status appears. When installation completes, the Update Complete dialog box displays. The upgrade installation cannot be stopped after it starts.

5. Click Finish to finish the installation.

If you install Navigator 2 whose version is 23.50 or later by replacing an instance that is earlier than 23.50, Navigator 2 registers in All Programs in the Start menu.

When the installation completes normally, you can operate via a Web browser.

Setting the server certificate and private key

To use the storage system safely in your environment, use the server certificate and private key for SSL communication that you have created.

Since the `hcmdssslc` command is not provided, you cannot set it by using the `hcmdssslc` command. If it has been set by using the `hcmdssslc` command, set it in the following procedure:

To create and set the server certificate and private key:

1. Stop the services for Navigator 2. To completely halt services, stop the SNM2 Server service first, then stop the HiCommand Suite Common Components service.
2. Create a private key, self-signed certificate, and CSR. Create a private key, self-signed certificate, and CSR using the `hcmdsssltool` command
3. Open the command prompt (terminal console for Unix) and move to the following directory:
 - For Windows: *<Navigator 2 installation directory>\Base\bin*
 - For Unix: *<Navigator 2 installation directory>/Base/bin*

4. Using the `hcmdsssltool` command, create the directory where output for the private key will be directed. Issuing the following command line for the appropriate operating system:

For Windows and Unix:

```
hcmdsssltool /key <file name of private key> /car <file name of CSR>
/cert <file name of self-signed certificate file>
/certtest <file name of self-signed certificate
displaying contents of file>
/dname "CN=<server name>, OU=<organization unit>,
O=<organization name>, L=<city or locality>,
S=<state or province>, C=<country-code>
```

The following example displays a session issuing this command line.

```
>hcmdsssltool /key c:\ca\httpdkey.pem /car c:\ca\httpsd.csr /cert
c:\ca\httpsd.pem /certtext c:\ca\httpsd.pem.txt /dname
"CN=Hitachi,OU=hsnm2, O=Hitachi,L=SanJose,S=California,C=us"
KAPM0674-I The hcmdsssltool command ended successfully.
```

5. Issue a certificate in an external Certificate Authority (CA). Send the CSR file you created in step 2 to the CA which supports SHA256, and obtain the certificate of the SHA256 algorithm in PEM format. Note that when using the self-signing certificate, you do not need to send it to the CA, and use the self-signing certificate file created in step 2.
6. Edit the `httpsd.conf` file by following these steps:
7. Open the `httpsd.conf` file stored in *<installation directory>\Base\httpsd\conf* with the text editor and edit it.
8. Remove the pound (#) character of the following which are commented out by default and change the values of `SSLCertificateFile` and `SSLCertificateKeyFile`.
9. When connecting IPv6, remove the pound (#) character of `#Listen[::]:23016` on the `httpsd.conf` file and change the values of `SSLCertificateFile` and `SSLCertificateKeyFile`. When IPv6 is disabled, do not remove the pound (#) character.
10. Specify the signed certificate file obtained from the CA or self-signed certificate file for `SSLCertificateFile` and the full path of the private key file created in step 2 for `SSLCertificateKeyFile`.

The contents of the file are shown here:

```
SSLSessionCacheSize 0
#Listen 23016
#Listen [::]:23016
#<VirtualHost *:23016>
# ServerName pc-name
# SSLEnable
# SSLProtocol SSLv3 TLSv1
# SSLRequiredCiphers AES256-SHA:AES128-SHA:DES-CBC3-SHA
# SSLRequireSSL
# SSLCertificateFile "c:/ca/httpsd.pem"
# SSLCertificateKeyFile "C:/ca/httpsdkey.pem"
# SSLCACertificateFile "C:/Program
```

```
#Files/HiCommand/Base/httpsd/conf/ssl/cacert/anycert.pem"
# SSLSessionCacheTimeout 3600
#</VirtualHost>
```



NOTE: Since the SSLProtocol row specifies a protocol used for communication, the security vulnerability CVE-2014-3566 (<https://web.nvd.nist.gov/view/vuln/detail?vulnId=CVE-2014-3566>) exists in SSLv3 (SSL3.0). You can avoid the vulnerability by specifying the SSL protocol of the httpsd.conf file, as follows: **SSLProtocol TLSv1**

In this case, you need to set up the certificate so that TLSv2 can be used in a browser for Internet Explorer by enabling the option in the following menu path: **Tools -> Internet Options -> Advanced -> Security -> Use TLS1.0**). If you do not set up the browser, you cannot access Navigator 2.

```
# SSLProtocol TLSv1
```

11. Start the services for Navigator 2 by starting the SNM2 Server service first and then start the Hitachi Storage Command Suite Common Components service.

12. Confirm SSL communication by activating the browser and specifying the URL. An example of syntax for URL specification is:

```
http://#####:23016/StorageNavigatorModular/
```

↑ ↑
Host IP address or name Host port number

Preinstallation information for Storage Features

Before installing storage features, review the preinstallation information in the following sections.

Environments

Your system should be updated to the most recent firmware version and Navigator 2 software version to expose all the features currently available.

The current firmware, Navigator 2, and CCI versions applicable for this guide are as follows:

- Firmware version 0916/A (1.6A) or higher for the HUS storage system. (Version 0981/D is required for CCI and Copy-on-Write Snapshot.)
- Navigator 2 version 21.60 or higher for your computer. (Version 28.12 is required for CCI and Copy-on-Write Snapshot.)
- When using the command control interface (CCI), version 01-34-03/04 or higher is required for your computer.

Storage feature requirements

Before installing storage features, be sure you meet the following requirements.

- Storage feature license key.
- Controllers cannot be detached.
- When changing settings, reboot the array.
- When connecting the network interface, 10BASE-T, 100BASE-T, or 1000BASE-T (RJ-45 connector, twisted pair cable) is supported. The frame type must conform to Ethernet II (DIX) specifications.
- Two (2) controllers (dual configuration),
- Maximum of 128 command devices. Command devices are only required when the CCI is used for Volume Migration. The command device volume size must be 33 MB or more.
- One Differential Management Logical Unit (DMLU). The DMLU size must be 10 GB or more. Only one DMLU can be set for different RAID groups while the AMS 2000 supports two.
- The primary volume (P-VOL) size must equal the secondary volume (S-VOL) size.

Requirements for installing and enabling features

Before you install or enable your features:

- Verify that the array is operating in a normal state. If a failure (for example a controller blockade) has occurred, installing cannot be performed.
- Obtain the required key code or key file to install your feature. If you do not have it, obtain it from the download page on the HDS Support Portal: <http://support.hds.com>.

Account Authentication

- Account Authentication cannot be used with Password Protection. If Account Authentication is installed or enabled, Password Protection must be uninstalled or disabled.
- Password Protection cannot be used with Account Authentication. If Password Protection is installed or enabled, Account Authentication must be uninstalled or disabled.

Audit Logging requirements

- This feature and the Syslog server to which logs are sent require compliance with the BSD syslog Protocol (RFC3164) standard.
- This feature supports a maximum of two (2) syslog servers
- You must have an Account Administrator role (View and Modify).

- When disabling this feature, every account, except yours, is logged out.
- Uninstalling this feature deletes all the account information except for the built-in account password. However, disabling this feature does not delete the account information.

Cache Partition Manager requirements

If you plan to install Copy-on-Write Snapshot, True Copy Extended Distance (TCE), or Dynamic Provisioning after enabling and configuring Cache Partition Manager, note the following:

- SnapShot, TCE, and Dynamic Provisioning use a part of the cache area to manage array internal resources. As a result, the cache capacity that Cache Partition Manager can use becomes smaller than it otherwise would be.
- Check that the cache partition information is initialized properly when SnapShot, TCE, or Dynamic Provisioning is installed when Cache Partition Manager is enabled.
- Move the VOLs to the master partitions on the side of the default owner controller.
- Delete all of the sub-partitions and reduce the size of each master partition to one half of the user data area, the user data capacity after installing the SS/TCE/HDP.
- If you uninstall or disable this storage feature, sub-partitions, except for the master partition, must be deleted and the capacity of the master partition must be the default partition size (see [Table 7-1 on page 7-3](#)).

Data Retention requirements

- If you uninstall or disable this storage feature, you must return the volume attributes the Read/Write setting.

LUN Manager requirements

- If you uninstall or disable this storage feature, you must disable the host group and target security on every port.

Password Protection

- Password Protection cannot be used with Account Authentication. If Password Protection is installed or enabled, Account Authentication must be uninstalled or disabled.

SNMP Agent requirements

- We recommend that the SNMP Agent Support acquires Message Information Block (MIB) information periodically, because the User

Datagram Protocol (UDP) used for the SNMP Agent Support, does not guarantee correct error trap reporting to the SNMP manager.

- The array command processing performance is negatively affected if the interval for collecting MIB information is too short.
- If the SNMP manager is started after array failures, the failures are not reported with a trap. Acquire the MIB objects `dfRegressionStatus` after starting the SNMP manager, and verify whether failures occur.
- The SNMP Agent Support stops if the controller is blocked and the SNMP managers do not receive responses.
- When an array is configured from a dual system, hardware component failures (fan, battery, power supply, cache failure) during power-on before the array is Ready, or from the last power-off, are reported with a trap from both controllers. Failures in the array or while it is Ready, are reported with a trap from the controller that detects the failures.
- When an array is configured from a dual system, both controllers must be monitored by the SNMP manager. When only one of the controllers is monitored using the SNMP manager, monitor controller 0 and note the following:
 - Drive blockades detected by controller 1 are not reported with a trap.
 - Controller 1 is not reported as TRAP. The controller down is reported as `systemDown TRAP` by the controller that went down.
- After controller 0 is blocked, the SNMP Agent Support cannot be used.

Modular Volume Migration requirements

- To install and enable the Modular Volume Migration license, follow the procedure provided in [Installing storage features on page 3-23](#), and select the license **LU-MIGRATION**.
- If you uninstall or disable this storage feature, all the volume migration pairs must be released, including those with a Completed or Error status. You cannot have volumes registered as reserved.

Installing storage features

To install your features for each storage system

1. In Navigator 2, select the check box for the array where you want to install your feature, and then click **Show & Configure Array**.
2. On the Array screen under Common Array Tasks, click the **Licenses** in the Settings tree view.
3. In the Licenses list, click the feature name, for example, Data Retention.
4. In the Licenses list, click the **Key File** or **Key Code** button, then enter the file name or key code for the feature you want to install. You can browse for the Key File.

5. Click **OK**.
6. Follow the on-screen instructions. A message displays confirming the optional feature installed successfully. Mark the checkbox and click **Reboot Array**.
7. To complete the installation, restart the storage system. The feature will close upon restarting the storage system. The storage system cannot access the host until the reboot completes and the system restarts. Restarting usually takes from six to 25 minutes.



NOTE: The storage system may require more time to respond, depending on its condition. If it does not respond after 25 minutes, check the condition of the system.

Enabling storage features

To enable your features for each storage system

1. In Navigator 2, select the check box for the storage system where you are enabling or disabling your feature.
2. Click **Show & Configure Array**.
3. If Password Protection is installed and enabled, log in with the registered user ID and password for the array.
4. In the tree view, click **Settings**, and select **Licenses**.
5. Select the appropriate feature in the Licenses list.
6. Click **Change Status**. The Change License window appears.
7. Select the **Enable** check box.
8. Click **OK**.
9. Follow the on-screen instructions.

Disabling storage features

Before you disable storage features

- Verify that the array is operating in a normal state. If a failure (for example a controller blockade) has occurred, uninstalling cannot be performed.
- A key code is required to uninstall your feature. This is the same key code you used when you installed your feature.

To disable your features for each array

1. In Navigator 2, select the check box for the array where you are enabling or disabling your feature.
2. Click **Show & Configure Array**.
3. If Password Protection is installed and enabled, log in with the registered user ID and password for the array.
4. In the tree view, click **Settings**, and select **Licenses**.
5. Select the appropriate feature in the Licenses list.
6. Click **Change Status**. The Change License window appears.

7. Clear the **Enable** check box.
8. Click **OK**.
9. Follow the on-screen instructions.

Uninstalling storage features

Before you uninstall storage features

- Verify that the array is operating in a normal state. If a failure (for example a controller blockade) has occurred, uninstalling cannot be performed.
- A key code is required to uninstall your feature. This is the same key code you used when you installed your feature.

To uninstall your features for each array

1. In Navigator 2, select the check box for the array where you want to uninstall your feature, then click **Show & Configure Array**.
2. In the tree view, click **Settings**, then click **Licenses**.
3. On the Licenses screen, select your feature in the Licenses list and click **De-install License**.
4. When you uninstall the option using the key code, click the **Key Code** radio button, and then set up the key code. When you uninstall the option using the key file, click the **Key File** radio button, and then set up a path for the key filename.
5. Click **OK**.
6. Follow the on-screen instructions.
7. To complete uninstalling the option, restart the storage system. The feature will close upon restarting the storage system. The system cannot access the host until the reboot completes and the system restarts. Restarting usually takes 6 to 25 minutes.

If you uninstall Navigator 2, Hitachi Storage Navigator Modular 2 is deleted from **All Programs** in the Start Menu.



NOTE: The storage system may require more time to respond, depending on its condition. If it does not respond after 25 minutes, check the condition of the system.

8. Log out from the disk array.

Uninstalling of the feature is now complete.

Uninstalling Navigator 2

Check if the preparation for the un-installation is completed referring to section.



NOTE: When other applications using Hitachi Storage Command Suite Common Components have been installed in the host to uninstall Navigator 2, back up the data base of the applications before uninstalling Navigator 2.



NOTE: When you uninstall Navigator 2 and the message “forcible uninstallation is necessary” displays, restart the PC. Ensure the preparation for the uninstallation is completed, and uninstall Navigator 2 again. If you cannot uninstall Navigator 2 even after restarting the PC, reinstall the OS.



NOTE: When the Hitachi Command Suite version 8.0.0 or greater is installed, the procedures before and after the uninstallation are required. For these versions, refer to [Uninstalling when HCS V8.0.0 or greater is installed on page 3-27](#).

To uninstall Hitachi Storage Navigator Modular 2 for Windows:

1. In the Windows **Start** menu, choose **Settings, Control Panel, Add or Remove Programs**.

The **Add or Remove Programs** panel appears.

2. From the Currently installed programs list, select **Hitachi Storage Navigator Modular 2**.
3. Click **Remove**.

The **Uninstall Hitachi Storage Navigator Modular 2** dialog box appears.

4. Click **Next**.

Clicking **Next** displays the **Confirm stopping the services** dialog box.

5. Click **Next**.

Clicking **Next** displays the **Confirm before uninstallation** dialog box appears.

6. Click **Next**.

When un-installation is complete, the **Uninstall Complete** dialog box appears.

7. Click **Finish**.

When other applications using Hitachi Storage Command Suite Common Components have been installed, restart the services referring to the manual of each application.

To uninstall Hitachi Storage Navigator Modular 2 for Linux and Solaris:

1. When the application that uses Hitachi Storage Command Suite Common Components other than Hitachi Storage Navigator Modular 2 is installed, stop the daemon process of the Hitachi Storage Command Suite Common Components.
2. From the console, execute the following command:
`/opt/HiCommand/StorageNavigatorModular/uninstall-hsnm2.sh`

Uninstalling when HCS V8.0.0 or greater is installed

When you uninstall Navigator 2 in an environment where the Hitachi Command Suite (HCS) V8.0.0 or greater is installed in a 64-bit operating system of Windows or Linux, execute the following procedure before uninstalling. The HCS environment includes the Device Manager, tiered Storage Manager, and Replication Manager software.

1. Log out of the HCS environment.
2. Halt the HCS services.
 - for windows:
`<HCS installation path>\Base64\bin\hcnds64srv /stop`
 - for Linux:
`<HCS installation path> /Base64/bin/hcnds64srv -stop`
3. Set a management server for the user account.
 - for windows:
`<SNM2 installation path>\Base\bin\hcndsprmset /host 127.0.0.1 /port <non-SSL port number> [/sslport <SSL port number>]`
 - for Linux:
`<SNM2 installation path>/Base/bin/hcndsprmset /host 127.0.0.1 -port <non-SSL port number> [-sslport <SSL port number>]`

For <Port number>, specify a port number connecting to Navigator 2. The default value of the non-SSL port is 23,015 and the SSL port is 23,016. The specification of <SSL port number> is unnecessary if not using the connection to Navigator 2 by specifying the SSL port number.

The services of HCS stopped in step 2. After completing the uninstallation of Navigator 2, start the services of HCS in the following procedure. Otherwise, HCS does not operate normally.

4. Start the services of HCS.
 - for windows:
`<HCS installation path>\Base64\bin\hcnds64srv /start`
 - for Linux:
`<HCS installation path> /Base64/bin/hcnds64srv -sto`

Starting Navigator 2 host and client configuration

Host side

Verify through Control Panel -> **Administrative Tools** -> **Services** whether **HBase Storage Mgmt Common Service**, **HBase Storage Mgmt Web Service**, and **SNM2 Server** have started.

Start the **HBase Storage Mgmt Common Service**, **HBase Storage Mgmt Web Service**, and **SNM2 Server** if they have not started.

Client side

For Windows

When you use the JRE 1.6.0_10 or newer, setting the Java Runtime Parameters are not necessary in a client to start Navigator 2. When you use the JRE less than 1.6.0_10, setting the Java Runtime Parameters are necessary in a client to start Navigator 2.

When you using IE 8.0, IE 9.0, or IE 10.0, the Hitachi Storage Navigator Modular 2 operation may slow or terminate with the message “DMEG800007: The process is taking additional time. Please refresh and confirm the array status” due to executing the performance of firmware replacement. Make the environment possible to connect to the Internet or enable the SLL communication with Host Name:urs.microsoft.com / Tcp Port: 443. If the environment cannot be set, uncheck Enable SmartScreen Filter in the Internet Options setting.

When the Applet screen starts, a dialog to request the proxy connection may be displayed. Enter the user name and password for connecting to the internet. If the user name and password do not exist, set the following in the Advanced tab on the Java control panel. Perform this task corresponding to the JRE in use.

JRE6: Turn off Enable online certificate validation.

JRE7: Check Do not check (not recommended) in the item of Perform certificates revocation checks on.

To set **Java Runtime Parameters**

1. In the Windows **Start** menu, choose **Settings, Control Panel**.
2. From the **Control Panel**, select the **Java**.
3. Click **View** of the upper position in the **Java** tab.
4. Enter **-Xmx192m** to the Java Runtime Parameters field.

It is necessary to set the Java Runtime Parameters to display the Applet screen.

5. Click **OK**.
6. Click **OK** in the **Java** tab.
7. Close the **Control Panel**.

Changing JRE

The following procedure details changing the JRE used by Navigator 2 to JRE 1.7.45, JRE 1.7.0_51, JRE 1.7.0_60, JRE 1.7.0_67, JRE 1.8.0_5, or JRE 1.8.0_20.

Before changing the JRE, log out of Navigator 2 and close the browser. During the change, the operation on the Applet screen does not run normally because the services stop.

1. Install the JRE to be changed.
2. Execute the change tool.
3. Open the command prompt (terminal console for Unix) and move to the following directory.
 - For Windows:
`<installation_directory_name>\StorageNavigatorModular\bin`
 - For Unix: `<installation_directory_name> / StorageNavigatorModular/bin`
4. Execute the following command that indicates the storage folder path of the JRE executable to be stored:
 - For Windows:
`snmchgjre.bat <folder_path>`
 - For Unix:
`<installation_directory>/StorageNavigatorModular/bin`
5. Specify the folder path installed in step 1 for the storage destination folder path of the JRE to be stored. If installed in Windows by default, use the following path:
`C:\Program Files\Java\jre7`
The execution process displays.



NOTE: The change tool stops the services. If it takes time to stop the services, running the tool may cause an error. When an error occurs, run the change tool again.



NOTE: When configuring Java Runtime parameters in the previous section, the setting is disabled because of the JRE change you made in this section. To correct this, configure the setting again. If you perform the updated installation, the changed JRE becomes the JRE of the installed Navigator 2. Change the JRE again.

6. When returning the JRE to the installed Navigator 2, perform steps 1 and 2, and specify the following in a folder path to be specified in step 2.
 - For Windows:
`<installation_directory>\StorageNavigatorModular\server\jre1.6.0_instbk`
 - For Unix:
`<installation_directory>/StorageNavigatorModular/server/jre1.6.0_instbk`

Changing JDK

To change the JDK used by Navigator 2 to JDK 1.7.45, perform the following tasks:

1. Before changing the JDK, log out of the Navigator 2 and close the browser. During the update process, the JDK does not run normally because the services halt.
2. Install the JDK to be changed.
3. Stop the services for Navigator 2. To stop Navigator 2, first halt the service of the Navigator 2 Server first and then stop the service of the HiCommand Suite Common Components.
4. Execute the change tool by performing the following steps.
5. Open the command prompt (terminal console for Unix) and move to the following directory:

For Windows: `<installation_directory>\Base\bin`

For Unix: `<installation_directory>/Base/bin`

Execute the following commands

For Windows: `hcmdschgjdk`

For Unix: `hcmdschgjdk`

Select the JDK to be used on the displayed screen.

For Linux and Solaris

To set Java Runtime Parameters

1. Run the Java Control Panel from XWindow terminal executing the `<JRE installed directory>/bin/jcontrol`.
2. Click **View** of the upper position in the Java tab.
3. Enter “-Xmx192m” to the Java Runtime Parameters field.
It is necessary to set the Java Runtime Parameters to display the Applet screen.
4. Click **OK**.
5. Click **OK** in the **Java** tab.

Changing the Port Number for Applet Screen of Navigator 2

The following procedure details how to change the port number used to display the Applet screen of Navigator 2. The number specified at the time of the installation is set for the default port number. The default port number is 1099.

1. Stop the SNM2 Server service.
2. If there are other products that use the Hitachi Storage Command Suite Common Components, stop the service (the daemon process).

3. Edit the setting files (snmserver.properties) and change the port number.

For Windows: The `jp.co.Hitachi.strdiskarray.rmi.port` in the `C:\Program Files\HiCommand\StorageNavigatorModular\server\snmserver.properties` file specifies the port number. Rewrite to a port number you want to change

`jp.co.Hitachi.strdiskarray.rmi.port=1099`

For Linux and Solaris: The `jp.co.Hitachi.strdiskarray.rmi.port` in the `\opt\HiCommand\StorageNavigatorModular\server\snmserver.properties` file specifies the port number. Rewrite to a port number you want to change.

`jp.co.Hitachi.strdiskarray.rmi.port=1099`

4. Start the service for the Hitachi Storage Command Suite Common Components.
5. Start the SNM2 Server service.
6. If there are products that use the Hitachi Storage Command Suite Common Components, start the service.

Changing the Port Number for Starting Navigator 2

To change the port number you use for starting Navigator 2, operate it in the following procedure. For the default port number, 23015 is set.

1. Stop the SNM2 Server service (daemon process).
2. If there are other products to use Hitachi Storage Command Suite Common Components, stop the service (daemon process).
3. Stop the service (daemon process) for the Hitachi Storage Command Suite Common Components.
4. Edit setting file (httpsd.conf) and change the port number.

For Windows:

The listen in the `C:\Program Files\HiCommand\Base\httpsd\conf\httpsd.conf` file specifies the port number. Rewrite to a port number you want to change.

For Linux and Solaris:

The listen in the `/opt/HiCommand/Base/httpsd/conf/httpsd.conf` file specifies the port number. Rewrite to a port number you want to change.

Listen 23015

5. Execute the `hcmdsprmset` command and change the port number.

For Windows:

1. In the Windows Start menu, choose Programs, Accessories, and then Command Prompt.

2. The command prompt moves the current directory to the directory in which the commands of Hitachi Storage Command Suite Common Components have been installed.

C:\> cd /D C:\Program Files\HiCommand\Base\bin

3. The directory in which the commands of Hitachi Storage Command Suite Common Components have been installed, specify as following command and then press the Enter.

C:\Program Files\HiCommand\Base\bin> hcmdsprmset /host 127.0.0.1 /
port 23015

The underlined part of the execution example, C:\Program Files\HiCommand\Base\, is a path of the folder in which Hitachi Storage Command Base Common Components is installed. It is required to replace it according to the environment in which Hitachi Storage Command Base Common Components is installed. In the environment not using the products to use Hitachi Storage Command Suite Common Components other than Navigator 2, this folder is the same as that in which Navigator 2 has been installed. If other products are also installed, specify a folder that was specified when the first product was installed.

The underlined part of the execution example, 127.0.0.1, as it is.

If a different value is input, it is likely to become an error by the screen operations.

The underlined part of the execution example, 23015, is a changed port number.

For Linux and Solaris:

Execute the following command.

/opt/HiCommand/Base/bin/hcmdsprmset -host 127.0.0.1 -port 23015

The root authority is required to execute the command.

Any directory can be the current directory at the time of the command execution.

The underlined part of the execution example, /opt/HiCommand/Base/, is a path of the directory in which Hitachi Storage Command Base Common Components is installed. For Linux, it is required to replace it according to the environment in which Hitachi Storage Command Base Common Components is installed. In the environment not using the products to use Hitachi Storage Command Suite Common Components other than Navigator 2, this directory is /opt/HiCommand/Base/. If other products are also installed, specify a directory that was specified when the first product was installed. Refer to the manual of each product for the details.

For Solaris, it is not required to replace /opt/HiCommand/Base/.

The underlined part of the execution example, 127.0.0.1, as it is.

If a different value is input, it is likely to become an error by the screen operations.

The underlined part of the execution example, 23015, is a changed port number.

4. Start the service (daemon process) for the Hitachi Storage Command Suite Common Components.
5. Start the SNM2 Server service (daemon process).
6. If there are other products to use Hitachi Storage Command Suite Common Components, start the service (daemon process).

If you change the port number used by Hitachi Storage Command Suite Common Components, all the installed products to use Hitachi Storage Command Suite Common Components are affected. When other products to use Hitachi Storage Command Suite Common Components are installed, check that the port number change does not cause problems for operating other products

Starting Navigator 2

To start Navigator 2 with version 23.50 or later installed on a Windows PC

For Navigator 2 version 23.50 or later installed on a Windows PC, perform the following steps:

1. On the Start menu in Windows, point to **All Programs**, point to Hitachi Storage Navigator 2 and then click Login.

The login window displays with the following URL:

`http://127.0.0.1:23015/StorageNavigatorModular/`

2. Perform the following tasks based on your local considerations:
 - If you use https or IPv6, change the contents of the Start menu according to the descriptions in “For Windows, Linux and Solaris:”.
 - If you specify an IP address at the installation, use the IP address in the Start menu.
 - The Applet screen may not appear because 127.0.0.1 is used. In the Start menu, use the Ip address specified at the installation or a valid IP address in the PC where it is installed.

To start Navigator 2 with versions less than version 23.50 installed on Windows, Linux, and Solaris platforms

For Navigator 2 with versions less than version 23.50 installed on Windows, Linux, and Solaris platforms, perform the following steps:

1. Activate the browser and specify the URL as follows.



NOTE: The https is invalid in the status immediately after the installation. When connecting it with the https, it is required to set the server certificate and private key in advance referring to the section 0.

For the URL, specify a host name or IP address of Navigator 2. Do not specify a loop back address such as localhost and 127.0.0.1. When you specify a loop back address such as localhost or 127.0.0.1, the Web screen displays, but the Applet screen may not display.

The user login screen displays.

Example: `http://#####:23015/StorageNavigatorModular/`
 ↑ ↑
 The port number of host
IPv4 IP address or host name of host

Example: `http://[#####]:23015/StorageNavigatorModular/`
 ↑ ↑
 The port number of host
IPv6 IP address or host name of host: Put the IPv6 IP address in [].

Example: `https://#####:23016/StorageNavigatorModular/`
 ↑ ↑
 The port number of host
IPv4 IP address or host name of host

Example: `https://[#####]:23016/StorageNavigatorModular/`
 ↑ ↑
 The port number of host
IPv6 IP address or host name of host: Put the IPv6 IP address in [].




Figure 3-1: Navigator 2 login screen

2. Enter your login information and click Login.

When logging in Navigator 2 for the first time after installing it newly, login with the built-in user's system account. The default password of the system account is manager. If another user is registered, login with the registered user. Enter the user ID and password, and click **Login**.

To prevent unauthorized access, we recommend changing the default password of the system account. You cannot delete the system account or change the authority of the system account. The system account is the built-in account common to the Hitachi Storage Command Suite products.

The system account can use all the functions of Hitachi Storage Command Suite Common Component including Navigator 2 and access all the resources that each application manages. When Hitachi Storage Command Suite Common Component are already installed in the PC, etc. in which Navigator 2 is installed and the password of the system account is changed, login with the changed password.

Although you can login with the user ID registered in Hitachi Storage Command Suite Common Component, you cannot operate Navigator 2. Add the operation authority of Navigator 2 after logging in Navigator 2, and login again.

Navigator 2 starts and the Arrays screen displays.

3. Since the Arrays screen is displays, use Navigator 2 after registering the array unit in it.

Operations

Navigator 2 screens consist of the Web and Applet screens. When you start Navigator 2, the login screen is displayed. When you login, the Web screen that shows the Arrays list is displayed. On the Web screen, operations provided by the screen and the dialog box is displayed. When you execute Advanced Settings on the Arrays screen and when you select the HUS on the Web screen of the Arrays list, the Applet screen is displayed.

One user operates the Applet screen to run the HUS, and two or more users cannot access it at the same time.




Figure 3-2: Array Screen and HUS array screens

The following figure displays settings that appear in the Applet dialog box.




Figure 3-3: Applet dialog box

Screens such as the Array screens displays when you login are Web screens. When you click the tree, details of the screen are updated on the same screen. When you click a button on the screen, a dialog box displays. Two types of dialog boxes exist; one is displayed on the Web screen and the other is displayed on the Applet screen.

A dialog box on the Web screen is displayed on the same screen by clicking any buttons. This is the same with the Applet screen. Use each function of the Web or Applet screen after completing the dialog function and closing the dialog box.

You can operate another button while the dialogue box is open. In that case, the display in the dialog box currently open is changed. However, the function of the dialog box that was open has worked.

Refer to Help for the procedure for operating the Web screen. The Help is not described in this manual. Since the Help for the procedure for operating the Applet screen is not provided, refer to this manual when you need help.



NOTE: The Applet screen is displayed connected to the **SNM2 Server**. If 20 minutes elapse while displaying the Applet screen, you will not be able to operate it due to the automatic logoff function. If the operation is completed, close the screen.

The following table shows the troubleshooting steps to take when the Applet screen does not display.

Setting an attribute

To set an attribute

1. Start Navigator 2.
2. Log in as a registered user to Navigator 2.
3. Select the storage system in which you will set up an attribute.
4. Click **Show & Configure Array**.
5. Select the feature icon in the Security tree view. SNM2 displays the home feature window.
6. Consider the following fields and settings in the Data Retention window.

Additional guidelines

- Navigator 2 is used by service personnel to maintain the arrays; therefore, be sure they have accounts. Assign the Storage Administrator (View and Modify) for service personnel accounts.
- The Syslog server log may have omissions because the log is not reset when a failure on the communication path occurs.
- The audit log is sent to the Syslog server and conforms to the Berkeley Software Distribution (BSD) syslog protocol (RFC3164) standard.

- If you are auditing multiple arrays, synchronize the Network Time Protocol (NTP) server clock. For more details on setting the time on the NTP server, see the Hitachi Storage Navigator Modular 2 online help.
- Reboot the array when changing the volume cache memory or partition.

Help

Navigator 2 describes the function of the Web screen with Help. Display Help by the following operation.

The following two ways exist for starting Help:

- Starting it from the Help menu in the Arrays screen.
- Starting it with the Help button in the individual screen.

When starting the Help menu in the Arrays screen, the beginning of Help is displayed.




Figure 3-4: Help - Welcome screen

When starting it with the **Help** button in the individual screen, the description according to the function is displayed.




Figure 3-5: Help - context help screen shown in background

Understanding the Navigator 2 interface

Now that you have installed Navigator 2, you may want to develop a general understanding of the interface design. Review the following sections for a quick primer on the Navigator 2 interface.

Figure 3-6 shows the Navigator 2 interface with the Arrays window displayed. This window appears when you log in to Navigator 2. It also appears when you click **Arrays** in the Explorer panel.



Figure 3-6: Navigator 2 interface

Menu Panel



The Menu Panel appears on the left side of the Navigator 2 user interface.

The Menu Panel always contains the following menus, regardless of the window displayed in the Page Panel:

- **File** — contains commands for closing the Navigator 2 application or logging out. These commands are functionally equivalent to the **Close** and **Logout** buttons in the Button Panel, described on the next page.
- **Go** — lets you start the ACE tool, a utility for configuring older AMS 1000 family systems.
- **Help** — displays the Navigator 2 online help and version information.

Explorer Panel



The Explorer Panel appears below the Menu Panel. The Explorer Panel displays the following commands, regardless of the window shown in the Page Panel.

- **Resource** — contains the **Arrays** command for displaying the Arrays window.

- **Administration** — contains commands for accessing users, permissions, and security settings. We recommend you use **Administration > Security > Password > Edit Settings** to change the default password after you log in for the first time.
- **Settings** — lets you access user profile settings.

Button panel



The Button Panel appears on the right side of the Navigator 2 interface and contains two rows of buttons:

- Buttons on the top row let you close or log out of Navigator 2. These buttons are functionally equivalent to the **Close** and **Logout** commands in the **File** menu, described on the previous page.
- Buttons on the second row change, according to the window displayed in the Page Panel. In the example above, the buttons on the second row appear when the Arrays window appears in the Page Panel.

Page panel

Error Monitoring									
All Arrays Status				Error Monitoring		Stopped			
Refresh Mode		All Properties			Record in Event Viewer		No		
Arrays									
Rows/ Page: 25 Page 1 of 1									
☐	Array Name	Status	Group	Model	Serial No.	Capacity of All VOL	Raw Capacity of All Drives	Monitor Error	
☐	AMS2300_85040079	Normal		AMS2300	85040079	36.8TB	64.2TB	Yes	
☐	HUS110_91110021	Normal		HUS110	91110021	20.0GB	21.3TB	Yes	
☐	HUS130_92110028	Normal		HUS130	92110028	0.0MB	21.3TB	Yes	
☐	HUS150_93010015	Normal		HUS150	93010015	14.9TB	25.1TB	Yes	

The Page Panel is the large area below the Button Panel. When you click an item in the Explorer Panel or the Arrays Panel (described later in this chapter), the window associated with the item you clicked appears in the Page Panel.

Information can appear at the top of the Page Panel and buttons can appear at the bottom for performing tasks associated with the window in the Page Panel. When the Arrays window in the example above is shown, for example:

- Error monitoring information appears at the top of the Page Panel.
- Buttons at the bottom of the Page Panel let you reboot, show and configure, add, edit, remove, and filter Hitachi storage systems.

Performing Navigator 2 activities

To start performing Navigator activities, you click a Hitachi storage system on the Arrays window. When you click a storage system, an Arrays Panel appears between the Explorer Panel and Page Panel (see [Figure 3-7 on page 3-42](#)). At the top of the Arrays Panel are the type and serial number

of the storage system you selected to be managed from the Arrays window. If you click the type and serial number, common storage system tasks appear in the Page Panel.




Figure 3-7: Arrays panel

If you click a command in the Arrays Panel, the Page Panel shows the corresponding page or the Arrays Panel reveals a list of sub-commands for you to click. In [Figure 3-7](#), for example, clicking **Groups** reveals two sub-commands, **Volumes** and **Host Groups**. If you select either sub-command, the appropriate information appears in the Page Panel. [Figure 3-8](#) shows an example of how the Arrays Panel and Page Panel look after clicking **Volumes** in the Arrays Panel.




Figure 3-8: Example of volume information

Description of Navigator 2 activities

You use the Arrays Panel and Page Panel to manage and configure Hitachi storage systems. [Table 3-5](#) summarizes the Navigator 2 activities you can perform, and the commands and sub-commands you click in the Arrays Pane to perform them.

This document describes how to perform key Navigator 2 activities. If an activity is not covered in this document, please refer to the Navigator online help. To access the help, click the **Help** button in the Navigator 2 Menu Panel (see [Menu Panel on page 3-40](#)).

Table 3-5: Description of Navigator 2 activities

Arrays Pane	Description
Components — displays a page for accessing controllers, caches, interface boards, host connector, batteries, and trays, as described below.	
Components > Controllers	Lists each controller in the Hitachi storage system and the controller's status.
Components > Caches	Shows the status, capacity, and controller associated with the cache in the Hitachi storage system.
Components > Interface Boards	Shows status information about each interface board in the Hitachi storage system and its corresponding controller.
Components > Host Connectors	Shows the host connector and port ID, status, controller number, and type of host connector (for example, Fibre Channel) for each host connector in the Hitachi storage system.
Components > Batteries	Shows the batteries in the Hitachi storage system and their status.

Table 3-5: Description of Navigator 2 activities (Continued)

Arrays Pane	Description
Components > Trays	Shows the status, type, and serial number of the tray. The serial number is the same as the serial number of the Hitachi storage system.
Groups — displays a page for accessing volumes and host groups, as described below.	
Groups > Volumes	Shows the volumes, RAID groups, and Dynamic Provisioning pools defined for the Hitachi storage system. For information about Dynamic Provisioning, refer to the <i>Hitachi Unified Dynamic Provisioning Configuration Guide</i> (MK-91DF8277).
Groups > Host Groups	Lets you: • Create or edit host groups. • Enable host group port-level security. • Change or delete the WWNs and WWN nicknames.
Replication — displays a page for accessing local replication, remote replication, and setup parameters, as described below.	
Replication > Local Replication	Lets you create a copy of a volume in the storage system using: • ShadowImage to create a duplicate copy of a volume. • Copy on Write Snapshot to create a virtual point-in-time copy of a volume.
Replication > Remote Replication	Lets you back up information using TrueCopy remote replication and TrueCopy Extended Distance to create a copy of a volume (volume) in the Hitachi storage system.
Replication > Setup	Assists you in setting up components of both local and remote replication.
Settings — displays a page for accessing FC settings, spare drives, licenses, command devices, DMLU, volume migration, LAN settings, firmware version, email alerts, date and time, and advanced settings.	
Settings > FC Settings	Shows the Fibre Channel ports available on the Hitachi storage system and provides updated Transfer Rate, Topology, and Link Status information.
Settings > Spare Drives	Lets you select a spare drive from a list of assignable drives.
Settings > Licenses	Lets you enable licenses for Storage Features that require them.
Settings > Command Devices	Lets you add, change, and remove command devices (and their volumes and RAID manager protection setting) for Hitachi storage systems.

Table 3-5: Description of Navigator 2 activities (Continued)

Arrays Pane	Description
Settings > DMLU	Lets maintenance technicians and qualified users add and remove differential management volumes (DMLUs). DMLUs are volumes that consistently maintain the differences between two volumes: P-VOLS and S-VOLS.
Settings > Volume Migration	Lets you migrate data to other disk areas.
Settings > LAN	Shows user management port, maintenance port, port number and security (secure LAN) information about the Hitachi storage system being managed.
Settings > Firmware	Shows the firmware version installed on the Hitachi storage system and lets you upgrade the firmware.
Settings > Email Alert	Lets you configure the Hitachi storage system to send email alerts if a failure occurs.
Settings > Date & Time	Lets you set the Hitachi storage system date, time, timezone, and up to two NTP server settings.
Settings > Advanced Settings	Lets you access features available in Storage Navigator Modular.
Power Savings — displays a page for accessing RAID group power saving settings.	
Power Savings > RG Power Saving	Lets you control which RAID groups are in spin-up or spin-down mode to conserve power.
Security — displays a page for accessing Secure LAN and Account Authentication settings, as described below.	
Security > Secure LAN	Lets you view and refresh SSL certificates.
Security > Audit Logging	Lets you enable audit to collect Hitachi storage system event information and output the information to a configured Syslog server.
Performance — displays a page for monitoring the Hitachi storage system, configuring tuning parameters, and viewing DP pool trend and optimization information, as described below.	
Performance > Monitoring	Lets you monitor a Hitachi storage system's performance (for example, utilization rates of resources in a disk array and loads on the disks and ports) and output the information to a text file.
Performance > Tuning Parameters	Lets you set parameters to tune the Hitachi storage system for optimum performance.

Table 3-5: Description of Navigator 2 activities (Continued)

Arrays Pane	Description
Performance > DP Pool Trend	Lets you view the Dynamic Provisioning pool trend for the Hitachi storage system (for example, utilization rates of DP pools) and output the information to a CSV file. For information about Dynamic Provisioning, refer to the <i>Hitachi Unified Storage Dynamic Provisioning Configuration Guide</i> , MK-91DF8277.
Performance > DP Optimization	Lets you optimize DP optimization priority for the Hitachi storage system by resolving unbalanced conditions, optimizing DP, and reclaiming zero pages.
Alerts & Events — shows Hitachi storage system status, serial number and type, and firmware revision and build date. Also, displays events related to the storage system, including firmware downloads and installations, errors, alert parts, and event log messages.	

Stopping the Hitachi Storage Command Suite Common Components

For Windows:

To stop the Hitachi Storage Command Suite Common Components service, used the Command Prompt of Windows.

1. In the Windows Start menu, choose Programs, Accessories, and then Command Prompt.
2. The command prompt moves the current directory to the directory in which the commands of Hitachi Storage Command Suite Common Components have been installed.

`C:\> cd /D C:\Program Files\HiCommand\Base\bin`

3. The directory in which the commands of Hitachi Storage Command Suite Common Components have been installed, specify as following command and then press the Enter.

`C:\Program Files\HiCommand\Base\bin>hcmdssrv /stop`

4. It may take several seconds to several minutes to stop the service. Check that the service is stopped by executing the following command.

`C:\Program Files\HiCommand\Base\bin>hcmdssrv /status`

The underlined part of the execution example, C:\Program Files\HiCommand\Base\, is a path of the folder in which Hitachi Storage Command Base Common Components is installed. It is required to replace it according to the environment in which Hitachi Storage Command Base Common Components is installed. In the environment not using the products to use Hitachi Storage Command Suite Common Components

other than Navigator 2, this folder is the same as that in which Navigator 2 has been installed. If other products are also installed, specify a folder that was specified when the first product was installed.

Do not directly stop the service of Hitachi Storage Command Suite Common Components using the Services screen of the Control Panel or the net stop command. If you stop it directly, Hitachi Storage Command Suite Common Components may become unable to be used.

For Linux and Solaris:

To stop the Hitachi Storage Command Suite Common Components daemon process, execute the following command.

```
/opt/HiCommand/Base/bin/hcmdssrv -stop
```

The root authority is required to execute the command.

Any directory can be the current directory at the time of the command execution.

It may take several seconds to several minutes to stop the daemon process. Check that the daemon process is stopped by executing the following command.

```
/opt/HiCommand/Base/bin/hcmdssrv -status
```

The underlined part of the execution example, /opt/HiCommand/Base/, is a path of the directory in which Hitachi Storage Command Base Common Components is installed. For Linux, it is required to replace it according to the environment in which Hitachi Storage Command Base Common Components is installed. In the environment not using the products to use Hitachi Storage Command Suite Common Components other than Navigator 2, this directory is /opt/HiCommand/Base/. If other products are also installed, specify a directory that was specified when the first product was installed. Refer to the manual of each product for the details.

For Solaris, it is not required to replace /opt/HiCommand/Base/.

Starting the Hitachi Storage Command Suite Common Components

For Windows:

To start the Hitachi Storage Command Suite Common Components service, used the Command Prompt of Windows.

1. In the Windows Start menu, choose Programs, Accessories, and then Command Prompt.
2. The command prompt moves the current directory to the directory in which the commands of Hitachi Storage Command Suite Common Components have been installed.

C:\> cd /D C:\Program Files\HiCommand\Base\bin

3. The directory in which the commands of Hitachi Storage Command Suite Common Components have been installed, specify as following command and then press the Enter.

C:\Program Files\HiCommand\Base\bin>hcmdssrv /start

4. It may take several seconds to several minutes to start the service.
Check that the service is started by executing the following command.

C:\Program Files\HiCommand\Base\bin>hcmdssrv /status

The underlined part of the execution example, C:\Program Files\HiCommand\Base\, is a path of the folder in which Hitachi Storage Command Base Common Components is installed. It is required to replace it according to the environment in which Hitachi Storage Command Base Common Components is installed.

Do not directly start the service of Hitachi Storage Command Suite Common Components using the Services screen of the Control Panel or the net start command. If you start it directly, Hitachi Storage Command Suite Common Components may become unable to be used.

For Linux and Solaris:

To start the Hitachi Storage Command Suite Common Components daemon process, execute the following command.

/opt/HiCommand/Base/bin/hcmdssrv -start

The root authority is required to execute the command.

Any directory can be the current directory at the time of the command execution.

It may take several seconds to several minutes to start the daemon process. Check that the daemon process is stopped by executing the following command.

/opt/HiCommand/Base/bin/hcmdssrv -status

The underlined part of the execution command, /opt/HiCommand/Base/, is a path of the directory in which Hitachi Storage Command Base Common Components is installed. For Linux, it is required to replace it according to the environment in which Hitachi Storage Command Base Common Components is installed. In the environment not using the products to use Hitachi Storage Command Suite Common Components other than Navigator 2, this directory is /opt/HiCommand/Base. If other products are also installed, specify a directory that was specified when the first product was installed. Refer to the manual of each product for the details.

For Solaris, it is not required to replace /opt/HiCommand/Base/.

In the process of the Hitachi Storage Command Suite Common Components, do not start it by using the script below /etc/init.d or by sending a signal with the kill command, etc. If you stop it in such method, you may not use the Hitachi Storage Command Suite Common Components.

Provisioning

This chapter provides information on setting up, or provisioning, your storage systems so they are ready for use by storage administrators.

The topics covered in this chapter are:

- ❑ [Provisioning overview](#)
- ❑ [Provisioning wizards](#)
- ❑ [Hardware considerations](#)
- ❑ [Logging in to Navigator 2](#)
- ❑ [Selecting a storage system for the first time](#)
- ❑ [Provisioning concepts and environments](#)

Provisioning overview

To successfully establish a storage system that is running properly, you first must provision it. Provisioning refers to the pre-active state preparation of a storage system required to carry out desired storage tasks and functions and to make it available to administrators. Provisioning of HUS storage systems is easy and convenient because of the availability of provisioning wizards which automatically step you through stages of preparing the storage system for rollout. The following section details the main HUS SNM2 wizards.

Provisioning wizards

The following are features for provisioning Navigator 2.

- **Add Array Wizard**

Whenever Navigator 2. is launched, it searches the database for listings of existing arrays. If there are arrays listed in the database, the platform displays them in the Subsystems dialog box. If there are no arrays, Navigator 2. automatically launches the Add Array wizard.

This wizard works with only one array at a time. It guides users through the steps to set up e-mail alerts, management ports, iSCSI ports and setting the date and time.

- **Create & Map Volume Wizard**

This wizard helps you create a volume and map it to an iSCSI target. It includes the following steps: 1) Create a new volume or select an existing one. 2) Create a new iSCSI target or select an existing one. 3) Connect to Host 4) Confirm 5) Back up a volume to another volume in the same array.

- **LUN Wizard**

Enables you to configure volumes and corresponding unit numbers, and to assign segments of stored data to the volumes.

- **Create Local Backup Wizard**

This wizard helps you create a local backup of a volume. The wizard includes the following steps: 1). Select the volume to be backed up. 2) Select a volume to contain the copied data. You will have the option to allocate this volume to a host. 3) Name the pair (original volume and its backup), and set copy parameters.

- **User Registration Wizard**

The User Registration Wizard is available when using the Account Authentication feature, which secures selected arrays with roles-based authentication.

- **Simple DR Wizard**

This wizard helps you create a remote backup of a volume. The purpose is to duplicate the data and prevent data loss in case of a disaster such as the complete failure of the array on which the source volume is mounted. The wizard includes the following steps: 1) Introduction 2) Set up a Remote Path 3) Set Up Volumes 4) Confirm

Provisioning task flow

The following details the task flow of the provisioning process:

1. A storage administrator determines a new storage system needs to be added to the storage network for which he is responsible.
2. The administrator launches the wizard to discover arrays on the storage network to add them to the Navigator 2 database.
3. If this is the first time you are configuring the array, the Add Array Wizard launches automatically. If you are modifying an existing array configuration, then manually launch the array.



NOTE: If the wizard does not launch, disable the browser's popup blockers, then click the Add Array button at the bottom of the Array List dialog box to launch the wizard.

4. If you know the IP address of a specific array that you want to add, click either Specific IP Address or Array Name to Search: and enter the IP address of the array. The default IP addresses for each controller are as follows:
 - 192.168.0.16 - Controller 0
 - 192.168.0.17 - Controller 1
5. If you know the range of IP addresses that includes one or more arrays that you want to add, click Range of IP Addresses to Search and enter the low and high IP addresses of that range. The range of addresses must be located on a connected local area network (LAN).
6. This screen displays the results of the search that was specified in the Search Array screen. Use this screen to select the arrays you want to add to Navigator 2.
7. If you entered a specific IP address in the Search Array screen, that array is automatically registered in Navigator 2.
8. If you entered a range of IP addresses in the Search Array screen, all of the arrays within that range are displayed in this screen. To add an array whose name is displayed, click on the area to the left of the array name.

Hardware considerations

Before you log in to Navigator 2, observe the following considerations.

Verifying your hardware installation

Install your Hitachi Data Systems storage system according to the instructions in the system's hardware guide. Then verify that your Hitachi Data Systems storage system is operating properly.

Connecting the management console

After verifying that your Hitachi Data Systems storage system is operational, connect the management console on which you installed Navigator 2 to the storage system's management port(s).

Every controller on a Hitachi storage system has a 10/100BaseT Ethernet management port labeled **LAN**. Hitachi storage systems equipped with two controllers have two management ports, one for each controller. The management ports let you configure the controllers using an attached management console and the Navigator 2 software.

Your management console can connect to the management ports directly using an Ethernet cable or through an Ethernet switch or hub. The management ports support Auto-Medium Dependent Interface/Medium Dependent Interface Crossover (Auto-MDI/MDIX) technology, allowing you to use either standard (straight-through) or crossover Ethernet cables.



TIP: You can attach a portable (“pocket”) hub between the management console and storage system to configure both controllers in one procedure, similar to using a switch.

Use one of these methods to connect the management console to controller, then power up the storage system.

Logging in to Navigator 2

The following procedure describes how to log in to Navigator 2. When logging in, you can specify an IPv4 address or IPv6 address using a nonsecure (http) or secure (https) connection to the Hitachi storage system.

To log in to Navigator 2

1. Launch a Web browser on the management console.
2. In the browser's address bar, enter the IP address of the storage system's management port using IPv4 or IPv6 notation. You recorded this IP address in [Appendix B, Recording Navigator 2 Settings](#):
 - IPv4 http example:
<http://IP address:23015/StorageNavigatorModular/Login>
 - IPv4 https example:
<https://IP address:23016/StorageNavigatorModular/Login>
 - IPv6 https example (IP address must be entered in brackets):
[https://\[IP address\]:23015/StorageNavigatorModular/Login](https://[IP address]:23015/StorageNavigatorModular/Login)

You cannot make a secure connection immediately after installing Navigator 2. To connect using https, set the server certificate and private key (see [Setting the server certificate and private key on page 3-18](#)).
3. At the login page (see [Figure 4-1](#)), type **system** as the default User ID and **manager** as the default case-sensitive password.



NOTE: Do not type a loopback address such as **localhost** or 127.0.0.1; otherwise, the Web dialog box appears, but the dialog box following it does not.




Figure 4-1: Login page

4. Click **Login**. Navigator 2 starts and the Arrays dialog box appears, with a list of Hitachi storage systems (see [Figure 4-2 on page 4-5](#)).




Figure 4-2: Example of Storage Systems in the Arrays dialog box

5. Under **Array Name**, click the name of the storage system you want to manage. One of the following actions occurs:
 - If the storage system has not been configured using Navigator 2, a series of first-time setup wizards launch, starting with the Add Array wizard. See [Selecting a storage system for the first time](#), below.
 - Otherwise, the storage system uses the configuration settings previously defined.



NOTE: If no activity occurs during a Navigator 2 session for 20 minutes, the session ends automatically.

Selecting a storage system for the first time

With primary goals of simplicity and ease-of-use, Navigator 2 has been designed to make things obvious for new users from the get-go. To that end, Navigator 2 runs a series of first-time setup wizards that let you define the initial configuration settings for Hitachi storage systems. Configuration is as easy as pointing and clicking your mouse.

The following first-time setup wizards run automatically when you select a storage system from the Arrays dialog box. Use these wizards to define the basic configuration for a Hitachi storage system.

- Add Array wizard - lets you add Hitachi storage systems to the Navigator 2 database. See [Running the Add Array wizard on page 4-6](#).
- Initial (Array) Setup wizard - lets you configure e-mail alerts, management ports, Internet Small Computer Systems Interface (iSCSI) ports and setting the date and time. See [Running the Initial \(Array\) Setup wizard on page 4-8](#).
- Create & Map Volume wizard - lets you create a volume and map it to a Fibre Channel or iSCSI target. See [Using the Create & Map Volume Wizard to create a RAID group on page 4-17](#).

After you use these wizards to define the initial settings for your Hitachi storage system, you can use Navigator 2 to change the settings in the future if necessary.

Navigator 2 also provides the following wizard, which you can run manually to further configure your Hitachi storage system:

Running the Add Array wizard

When Navigator 2 launches, it searches its database for registered Hitachi storage systems. At initial login, there are no storage systems in the database, so Navigator 2 searches your storage network for Hitachi storage systems and lets you choose the ones you want to manage.

You can have Navigator 2 discover a storage system by specifying the system's IP address or host name if you know it. Otherwise, you can specify a range of IP addresses. Options let you expand the search to include IPv4 and IPv6 addresses. When Navigator discovers storage systems, it displays them under **Search Results**. To manage one, click to the left of its name and click **Next** to add it and display the Add Array dialog box. Click **Finish** to complete the procedure.

You can also run the Add Array wizard manually to add storage systems after initial log in by clicking **Add Array** at the bottom of the Arrays dialog box.

Initially, an introduction page lists the tasks you complete using this wizard. Click **Next >** to continue to the Search Array dialog box (see [Figure 4-3 on page 4-7](#)) to begin the configuration. [Table 4-1 on page 4-7](#) describes the fields in the Search Array dialog box. As you specify your settings, record them in [Appendix B](#) for future reference. Use the navigation buttons at the bottom of each dialog box to move forward or backward, cancel the wizard, and obtain online help.




Figure 4-3: Add Array Wizard - Search Array dialog box

Table 4-1: Add Array Wizard - Search Array dialog box

Field	Description
IP Address or Array Name	Discovers storage systems using a specific IP address or storage system name in the Controller 0 and 1 fields. The default IP addresses are: - Controller 0: 192.168.0.16 - Controller 1: 192.168.0.17 For directly connected consoles, enter the default IP address just for the port to which you are connected; you will configure the other controller later.
Range of IP Addresses	Discovers storage systems using a starting (From) and ending (To) range of IP addresses. Check Range of IPv4 Address and/or Search for IPv6 Addressees automatically to widen the search if desired.
Using Ports	Select whether communications between the console and management ports will be secure, nonsecure, or both.

Running the Initial (Array) Setup wizard

After you complete the Add Array wizard at initial log in, the Initial (Array) Setup wizard starts automatically.

Using this wizard, you can configure:

- E-mail alerts — see [page 4-9](#)
- Management ports — see [page 4-11](#)
- Host ports — see [page 4-12](#)
- Spare drives — see [page 4-14](#)
- System date and time — see [page 4-14](#)

Initially, an introduction page lists the tasks you complete using this wizard. Click **Next >** to continue to the Search Array dialog box (see [Figure 4-5 on page 4-10](#) and [Table 4-2 on page 4-10](#)) and begin the configuration. Use the navigation buttons at the bottom of each dialog box to move forward or backward, cancel the wizard, and obtain online help.

The following sections describe the Initial (Array) Setup wizard dialog boxes.



NOTE: To change these settings in the future, run the wizard manually by clicking the name of a storage system under the **Array Name** column in the Arrays dialog box and then clicking **Initial Setup** in the **Common Array Tasks** menu.

Registering the Array in the Hitachi Storage Navigator Modular 2

The Add Array wizard registers the storage system in the following steps:

1. Searches the storage system.
2. Registers the storage system.

- Displays the name of the storage system. Note the name of the storage system.




Figure 4-4: Recording the storage system

Initial Array (Setup) wizard — configuring email alerts

The Set up E-mail Alert dialog box is the first dialog box in the Initial (Setup) Array wizard. Using this dialog box, you can configure the storage system to send email notifications if an error occurs. By default, email notifications are disabled. To accept this setting, click **Next** and skip to [Initial Array \(Setup\) wizard — configuring management ports on page 4-11](#).

To enable email alerts

- Complete the fields in [Figure 4-5](#) (see [Table 4-2](#)).
- Click **Next** and go to [Initial Array \(Setup\) wizard — configuring management ports on page 4-11](#).



NOTE: This procedure assumes your Simple Mail Transfer Protocol (SMTP) server is set up correctly to handle email. If desired, you can send a test message to confirm that email notifications will work.

HSNM2 - Setup Array Wizard **HITACHI**

1. Introduction ▶ **2. Set up E-mail Alert** ▶ 3. Set up Management Ports ▶ 4. Set up Host Ports ▶ 5. Set up Spare Drive ▶ 6. Set up Date & Time ▶ 7. Confirm ▶ 8. Finish

Select Disable or Enable for E-mail Error Report. When enabling, enter the required information in the text boxes.

If an error occurs in the array, the array sends an E-mail to report the error to Send To Address(es) automatically.

* E-mail Error Report:
☒ Disable
☐ Enable

Domain Name:

Mail Server Address:

From Address:

Send To Address: (1) To

(2) To

(3) To

Reply To Address:

* Required field

< Back Next > Cancel Help

Figure 4-5: Set up E-mail Alert page

Table 4-2: Enabling email notifications

Field	Description
E-mail Error Report Disable / Enable	To enable email notifications, click Enable and complete the remaining fields.
Domain Name	Domain appended to addresses that do not contain one.
Mail Server Address	Email address or IP address that identifies the storage system as the source of the email.
From Address	Each email sent by the storage system will be identified as being sent from this address.
Send to Address	Up to 3 individual email addresses or distribution lists where notifications will be sent.
Reply To Address	Email address where replies can be sent.

Initial Array (Setup) wizard — configuring management ports

The Set up Management Ports dialog box lets you configure the management ports on the Hitachi storage system. These are the ports you use to manage the system using Navigator 2.

To configure the management ports

1. Complete the fields in [Figure 4-6](#) (see [Table 4-3](#)).
2. Click **Next** and go to [Initial Array \(Setup\) wizard — configuring host ports on page 4-12](#).



NOTE: If your management console is directly connected to a management port on one controller, enter settings for that controller only (you will configure the management port on the other controller later). If your console is connected via a switch or hub, enter settings for both controllers now.

HSNM2 - Setup Array Wizard

1. Introduction > 2. Set up E-mail Alert > 3. Set up Management Ports > 4. Set up Host Ports > 5. Set up Spare Drive > 6. Set up Date & Time > 7. Confirm > 8. Finish

Enter the information for the network settings of the management ports.

Protocol: ☒ IPv4 ☐ IPv6

Management Ports

☐ Use DHCP ☒ Set Manually

Controller 0

* IPv4 Address: 51.52.53.54

* IPv4 Subnet Mask: 61.62.63.64

* IPv4 Default Gateway: 71.72.73.74

* Negotiation: 100Mbps/Full

☒ Use DHCP ☐ Set Manually

Controller 1

* IPv4 Address: 151.152.153.154

* IPv4 Subnet Mask: 161.162.163.164

* IPv4 Default Gateway: 171.172.173.174

* Negotiation: 10Mbps/Half

* Required field

< Back Next > Cancel Help

Figure 4-6: Management Ports dialog box

Table 4-3: Management Ports dialog box

Field	Description
IPv4/IPv6	Select the IP addressing method you want to use. For more information about IPv6, see Using IPv6 Addresses on page 4-26 .
Use DHCP	Configures the management port automatically, but requires a Dynamic Host Control Protocol (DHCP) server. IPv6 users: note that IPv6 addresses are based on Ethernet addresses. If you replace the storage system, the IP address changes. Therefore, you can want to assign static IP addresses to the storage system using the Set Manually option instead of having them auto-assigned by a DHCP server.

Table 4-3: Management Ports dialog box (Continued)

Field	Description
Set Manually	Lets you complete the remaining fields to configure the management port manually.
IPv4 Address	Static Internet Protocol address that matches the subnet where the storage system will be used.
IPv4 Subnet Mask	Subnet mask that matches the subnet where the storage system will be used.
IPv4 Default Gateway	Default gateway that matches the gateway where the storage system will be used.
Negotiation	Use the default setting (Auto) to auto-negotiate speed and duplex mode, or select a fixed speed/duplex combination.

Initial Array (Setup) wizard — configuring host ports

The Set up Host Ports dialog box lets you configure the host data ports on the Hitachi storage system. The fields in this dialog box vary, depending on whether the host ports on the Hitachi storage system are Fibre Channel or iSCSI.

To configure the host data ports using the Initial Array wizard

1. Perform one of the following steps:
 - To configure the Fibre Channel host ports, complete the fields in [Figure 4-7 on page 4-13](#) (see [Table 4-4 on page 4-13](#)).
 - To configure iSCSI host ports, complete the fields in [Figure 4-6 on page 4-11](#) (see [Table 4-5 on page 4-13](#)).
2. Click **Next** and go to [Initial Array \(Setup\) wizard — configuring spare drives on page 4-14](#).

HSNM2 - Setup Array Wizard **HITACHI**

1. Introduction ▶ 2. Set up E-mail Alert ▶ 3. Set up Management Ports ▶ **4. Set up Host Ports** ▶ 5. Set up Spare Drive ▶ 6. Set up Date & Time ▶ 7. Confirm ▶ 8. Finish

Enter the information for the host ports.

FC Ports	
Port0A	* Port Address: 0000EF * Transfer Rate: Auto * Topology: Loop
Port0B	* Port Address: 0000EF * Transfer Rate: Auto * Topology: Loop
Port0C	* Port Address: 0000EF * Transfer Rate: Auto * Topology: Loop
Port0D	* Port Address: 0000EF * Transfer Rate: Auto * Topology: Loop

FC Ports	
Port1A	* Port Address: 0000EF * Transfer Rate: Auto * Topology: Loop
Port1B	* Port Address: 0000EF * Transfer Rate: Auto * Topology: Loop
Port1C	* Port Address: 0000EF * Transfer Rate: Auto * Topology: Loop
Port1D	* Port Address: 0000EF * Transfer Rate: Auto * Topology: Loop

* Required field

< Back Next > Cancel Help

Figure 4-7: Set up Host Ports dialog box for Fibre Channel host ports

Table 4-4: Set up Host Ports dialog box for Fibre Channel host ports

Field	Description
Port Address	Enter the address for the Fibre Channel ports.
Transfer Rate	Select a fixed data transfer rate from the drop-down list that corresponds to the maximum transfer rate supported by the device connected to the storage system, such as the server or switch.
Topology	Select the topology in which the port will participate: - Point-to-Point = port will be used with a Fibre Channel switch. - Loop = port is directly connected to the Fibre Channel port of an HBA installed in a server.

Table 4-5: Set up Host Ports dialog box for iSCSI host ports

Field	Description
IP Address	Enter the IP address for the storage system iSCSI host ports. The default IP addresses are: Controller 0, Port A: 192.168.0.200 Controller 0, Port B: 192.168.0.201 Controller 1, Port A: 192.168.0.208 Controller 1, Port B: 192.168.0.209
Subnet Mask	Enter the subnet mask for the storage system iSCSI host port.
Default Gateway	If a router is required for the storage system host port to reach the initiator(s), the default gateway must have the IP address of that router. In a network that requires a router between the storage system and the initiator, enter the router's IP address. In a network that uses only direct connection, or a switch between the storage system and the initiator(s), no entry is required.

Initial Array (Setup) wizard — configuring spare drives

Using the Set up Spare Drive dialog box, you can select a spare drive from the available drives. If a drive in a RAID group fails, the Hitachi storage system automatically uses the spare drive you select here. The spare drive must be the same type, for example, Serial Attached SCSI (SAS), or Solid State Disk (SSD), as the failed drive and have the same capacity as or higher capacity than the failed drive. When you finish, click **Next** and go to [Initial Array \(Setup\) wizard — configuring the system date and time on page 4-14](#).




Figure 4-8: Initial Array (Setup) wizard: Set up Spare Drive dialog box

Initial Array (Setup) wizard — configuring the system date and time

Using the Set up Date & Time dialog box, you can select whether the Hitachi storage system date and time are to be set automatically, manually, or not at all. If you select **Set Manually**, enter the date and time (in 24-hour format) in the fields provided. When you finish, click **Next**.

Initial Array (Setup) wizard — confirming your settings

Use the remaining dialog boxes to confirm your settings. As you confirm your settings, record them in [Appendix B, Recording Navigator 2 Settings](#) for future reference. To change a setting, click **Back** until you reach the desired dialog box, change the setting, and click **Next** until you return to the appropriate confirmation dialog box. At the final Confirm dialog box, click **Confirm** to commit your settings. At the Finish dialog box, click **Finish** and go to [Running the Create & Map Volume wizard on page 4-15](#).

Figure 4-9: Set up Date & Time dialog box

Running the Create & Map Volume wizard

After you complete the Initial (Array) Setup wizard, the Create & Map Volume wizard starts automatically. Using this wizard, you can create or select RAID groups, volumes, and host groups.

Initially, an introduction page lists the tasks that can be completed by this wizard. Click **Next >** to continue to the Search RAID Group dialog box (see [Figure 4-12 on page 4-17](#)) and begin the configuration. Use the navigation buttons at the bottom of each dialog box to move forward or backward, cancel the wizard, and obtain online help.



NOTE: To change these settings in the future, run the wizard manually by clicking the storage system in the Arrays dialog box, and then clicking **Create Volume and Mapping** in the **Common Array Tasks** menu.

Manually creating a RAID group

Use this function when you create, expand, delete, and refer to the RAID group. This function can be used in the device Ready state. The unit does not be rebooted.

To create a RAID group:

1. From the Arrays list in the Arrays dialog box, click the desired storage system name to display the information window for the specific storage system.
2. Confirm the storage system is in a ready state by checking the Status field.
3. From the left navigation pane, click **Groups**, then click **Volumes** to display the Volumes dialog box.

4. Click the **RAID Groups** tab to display the RAID Groups list as shown in [Figure 4-10](#). RAID groups and volumes defined for the storage system display.




Figure 4-10: Volumes dialog box - RAID Groups tab

5. Click the **Create RG**. The Create RAID Group dialog box displays as shown in [Figure 4-11](#).

RAID Group Property

Enter the information for the RAID group to be created.

* RAID Group: 002
From 0 to max (array model type dependent)

RAID Level: RAID5

Combination: 4D+1P

* Number of Parity Groups: 1
From 1 to max (based on drive count)

Drives: ☐ Automatic Selection: Drive Type: SAS, Drive Capacity: 600GB

☒ Manual Selection:

Tray	HDU	Drive Type	Status
01	00	SAS (600GB)	Mounted
01	01	SAS (600GB)	Mounted
01	02	SAS (600GB)	Mounted
01	03	SAS (600GB)	Mounted

Figure 4-11: Create RAID Group dialog box

6. Select or enter values for the following fields, listboxes, or text boxes:
 - RAID Group
 - RAID Level

- Combination
 - Number of Parity Groups
- In the Drives region, select one of the following radio buttons:
 - **Automatic Selection** to direct the system to automatically select a drive. Select a drive type and a drive capacity in the two list boxes in this region.
 - **Manual Selection** to manually select a desired drive in the Assignable Drives list. Select an assignable drive in the list.
 - Click **OK**.

Using the Create & Map Volume Wizard to create a RAID group

Using the Search RAID Group dialog box, create a new RAID group for the Hitachi storage system or make it part of an existing RAID group.

HSNM2 - Create & Map Volume Wizard **HITACHI**

1. Introduction > **2. Create or select RAID group/DP pool** > 3. Create or select volumes > 4. Create or select host group/iSCSI target > 5. Connect to hosts > 6. Confirm > 7. Finish

Select the field that create and select the volumes. Click Next to continue.

The field that create and select the volumes :

☒ **RAID Group :**

☒ **Create a new RAID group :**

Drive Type : SAS

RAID Level : RAID5

Combination : 4D+1P

☐ **Use an existing RAID group :**

N/A

Drive Type : N/A

RAID Level : N/A

Combination : N/A

☐ **DP Pool :**

☒ **Create a new DP pool :**

Drive Type : SAS

RAID Level : RAID6

Combination : 8D+2P

☐ **Use an existing DP pool :**

N/A

Drive Type : N/A

RAID Level : N/A

Combination : N/A

Figure 4-12: Create or select RAID group/DP pool dialog box

To create a new RAID group

- Click **Create a New RAID Group**.
- Use the drop-down lists to select a drive type, RAID level, and data + parity (D+P) combination for the RAID group
- Click **Next** to continue to the Create or Select volumes dialog box.

To select an existing RAID group

- Click **Use an Existing RAID Group**.
- Select the desired RAID Group from the **RAID Group** drop-down list.
- Click **Next** and go to [Create & Map Volume wizard — defining volumes](#).

Create & Map Volume wizard — defining volumes

Using the next dialog box in the Create & Map Volume wizard, you can create new volumes or use existing volumes for the Hitachi storage system.

HSNM2 - Create & Map Volume Wizard

1. Introduction > 2. Create or select RAID group/DP pool > 3. Create or select volumes > 4. Create or select host group/iSCSI target >
5. Connect to hosts > 6. Confirm > 7. Finish

Configure volumes to mapping. Click Next to continue.

* Volumes : ☒ Create a new volumes :

☒ Create many volumes : Volume Capacity : 10 GB Number of Volumes : 1

☐ Create one volume to assign the largest region of available free space

☐ Create one volume to assign all of the available free space in the selected RAID group

Existing volumes

VOL	Capacity	RAID Group	DP Pool	RAID Level	Drive Type	Number of Paths
No Object						

* Required field

< Back Next > Cancel Help

Figure 4-13: Create or Select Volumes dialog box

If you select a RAID group with a capacity less than 10 GB, select from the existing RAID group capacity or create RAID group capacity.

To create new volumes

1. Click the **Create a new volumes** check box.
2. Perform one of the following steps:
 - Enter the desired **Volume Capacity** and **Number of Volumes**. Each volume that will be created will be the same size that you specify in this field.
 - Click **Create one volume to assign one of the maximum free space in the selected RAID group** to create a single logical unit consisting of the maximum available free space in the selected RAID group.
3. Click **Next** and go to [Create & Map Volume wizard — defining host groups or iSCSI targets](#).

To select an existing volume

1. Select one or more volumes under **Existing volumes**.
2. Click **Next** and go to [Create & Map Volume wizard — defining host groups or iSCSI targets](#).

Create & Map Volume wizard — defining host groups or iSCSI targets

Using the next dialog box in the Create & Map Volume wizard, you can select:

- A physical port for a Fibre Channel host group or iSCSI target.
- Host groups for storage systems with Fibre Channel ports.
- iSCSI targets for storage systems with iSCSI ports.

[illegible]

Figure 4-14: Create or select host group/iSCSI target dialog box

To create or select a host group for Fibre Channel storage systems

1. Next to **Port**, select a physical port.
2. Create a new host group or select an existing one:
To create a new host group:
 - a. Click **Create a new host group**.
 - b. Enter a **Host Group No** (from 1 to 127).
 - c. Enter a host group **Name** (up to 32 characters).
 - d. Select **Platform** and **Middleware** settings from the drop-down lists (refer to the Navigator 2 online help).To select an existing host group:
 - a. Click **Use an existing host group**.
 - b. Select a host group from the **Host Group** drop-down list.
3. Click **Next** and go to [Create & Map Volume wizard — connecting to a host on page 4-20](#).

To create a new iSCSI target or select an existing one for iSCSI storage systems

1. Next to **Port**, select a port to map to from the available ports options.
2. Create a new iSCSI target or select an existing one:

To create a new iSCSI target:

- a. Click **Create a new iSCSI target**.
- b. Enter an **iSCSI Target No** (from 1 to 127).
- c. Enter an iSCSI target **Name** (up to 32 characters).
- d. Select **Platform** and **Middleware** settings from the drop-down lists (refer to the Navigator 2 online help).

To select an existing iSCSI target:

- a. Click **Use an existing iSCSI target**.
 - b. Select an iSCSI target from the **iSCSI Target** drop-down list.
3. Click **Next** and go to [Create & Map Volume wizard — connecting to a host](#), below.

Create & Map Volume wizard — connecting to a host

If LUN Manager is enabled, the Connect to Hosts dialog box lets you select the hosts to which the Hitachi storage system will be connected. If LUN Manager is not enabled, the wizard skips this dialog box and goes to the first confirm dialog box (see step 4 in the following procedure). The iSCSI target on the storage system will communicate with the iSCSI initiator on the host.




Figure 4-15: Connect to hosts dialog box

To map multiple hosts to volumes if the Connect to Hosts dialog box appears

1. To allow multiple hosts to map to the selected volumes, click **Allow multiple hosts**.
2. Check all of the hosts you want to connect to the Hitachi storage system.

3. When you finish, click **Next**.

Create & Map Volume wizard — confirming your settings

Use the remaining dialog boxes to confirm your settings. As you confirm your settings, record them in [Appendix B, Recording Navigator 2 Settings](#) for future reference. To change a setting, click **Back** until you reach the desired dialog box, change the setting, and click **Next** until you return to the appropriate confirmation dialog box. At the final Confirm dialog box, click **Confirm** to commit your settings.

To create additional RAID groups, volumes, and host groups, click **Create & Map More VOL** and repeat the wizard starting from the Search RAID Group dialog box. Otherwise, click **Finish** to close the wizard and return to the Array Properties dialog box.

This completes the first-time configuration wizards. If desired, you can run the remaining wizards described in this chapter to further configure your Hitachi storage system.

Provisioning concepts and environments

The following section detail important concepts and utilities involved with a standard provisioning of SNM2. Also, several key environments you will need to become acquainted with are discussed.

About DP-Vols

The DP-VOL is a virtual volume that consumes and maps physical storage space only for areas of the volume that have had data written. In Dynamic Provisioning, it is required to associated the DP-VOL with a DP pool.

The DP-VOL needs to specify a DP pool number, DP-VOL logical capacity and DP-VOL number. Many DP-VOLs can be defined for one pool. A given DP-VOL cannot be defined to multiple DP pools. The HUS can register up to 4,095 DP-VOLs. The maximum number of DP-VOLs is reduced by the number of RAID groups.

Changing DP-Vol Capacity

You can dynamically increase or decrease the defined logical capacity of a DP-VOL within certain limits. When decreasing a DP-VOL's logical capacity, any DP pool capacity mapped to the trimmed-away logical capacity is lost. Subsequent DP pool optimization processing may increase the free capacity of the DP pool.

The Dynamic Provisioning application, operating system, and file system must be able to recognize the increase or decrease in logical capacity to make it totally dynamic. Navigator 2 enables you to increase or decrease the capacity of the DP-Vol.

About volume numbers

A **volume number** is a number used to identify a **volume**, which is a device addressed by the protocol either Fibre Channel or iSCSI. A volume may be used with any device which supports read/write operations, such as a tape drive, but is most often used to refer to a logical disk as created on a SAN. Though not technically correct, the term "volume" is often also used to refer to the drive itself.

To provide a practical example, a typical disk array has multiple physical iSCSI ports, each with one SCSI target address assigned. Then the disk array is formatted as a RAID and then this RAID is partitioned into several separated storage volumes. To represent each volume, a SCSI target is configured to provide a volume. Each SCSI target may provide multiple volumes and thus represent multiple volumes, but this does *not* mean that those volumes are concatenated. The computer that accesses a volume on the disk array identifies which volume to read or write with the volume of the associated volume.

Another example is a single disk drive with one physical SCSI port. It usually provides just a single target, which in turn usually provides just a single volume whose volume is zero. This volume represents the entire storage of the disk drive.

In the current SCSI, a volume is a 64-bit identifier. It is divided into four 16-bit pieces that reflect a multilevel addressing scheme, and it is unusual to see any but the first of these used.

People usually represent a 16-bit single-level volume as a decimal number. In earlier versions of SCSI, and with some transport protocols, volumes can be restricted to 16, 6 or 3 bits.

How to select a volume: In the earliest versions of SCSI, an initiator delivers a Command Data Block (CDB) to a target (physical unit) and within the CDB is a 3-bit volume field to identify the volume within the target. In current SCSI, the initiator delivers the CDB to a particular volume, so the volume appears in the transport layer data structures and not in the CDB.

Volume vs. SCSI Device ID: The volume is not the only way to identify a volume. There is also the SCSI Device ID, which identifies a volume uniquely in the world. Labels or serial numbers stored in a volume's storage volume often serve to identify the volume. However, the volume is the only way for an initiator to address a command to a particular volume, so initiators often create, via a discovery process, a mapping table of volume to other identifiers.

Context sensitive: The volume identifies a volume only within the context of a particular initiator. So two computers that access the same disk volume may know it by different volumes.

Volume 0: There is one volume which is required to exist in every target: zero. The volume with volume zero is special in that it must implement a few specific commands which is how an initiator can find out all the other volumes in the target. But Volume zero need not provide any other services, such as a storage volume.

Many SCSI targets contain only one volume (so its volume is necessarily zero). Others have a small number of volumes that correspond to separate physical devices and have fixed volumes. A large storage system may have up to thousands of volumes, defined logically, by administrative command, and the administrator may choose the volume or the system may choose it.

About Host Groups

Host Groups are a class of object known as host storage domains. They are a feature of Hitachi LUN Manager and allow your array to be more easily managed. Hosts (WWNs) can be assigned to a Host Group and then the desired volumes can be associated with each host group. For more information on setting up a Host Group for Fibre Channel, go to the section that details setting up Host Groups.

An iSCSI target is a logical entity that associates a group of hosts communicated with iSCSI and volumes in the array. The iSCSI Targets menu item is displayed only if the array has an iSCSI interface to communicate with hosts and according to the model of the array.

Using the LUN Manager storage feature, you can add, modify, or delete iSCSI targets during system operation. For example, if an additional disk is installed or an additional host is connected in your iSCSI network, an additional iSCSI target can be created for them with LUN Manager.

Creating Host Groups

To add host groups, you must enable the host group security, and create a host group for each port.

To understand the host group configuration environment, you need to become familiar with the Host Groups Setting dialog box as shown in [Figure 4-7 on page 4-13](#).

The Host Groups Setting dialog box consists of the Host Groups, Host Group Security, and WWNs tabbed pages.

- **Host Groups** - Enables you to create and edit groups, initialize the Host Group 000, and delete groups.
- **Host Group Security** - Enable you to validate the host group security for each port. When the host group security is invalidated, only the Host Group 000 (default target) can be used. When it is validated, host groups following the host group 001 can be created, and the WWN of hosts to be permitted to access each host group can be specified.
- **WWNS** - Displays WWNs of hosts detected when the hosts are connected and those entered when the host groups are created. In this tabbed page, you can supply a nickname to each port name.

Displaying Host Group Properties

To display the properties of the host groups assigned to an array, perform the following steps:

In the Array List dialog box, select an array of interest and click Show and Configure Array.

In the Arrays tree, expand the Groups menu and select Host Groups. The Host Groups dialog box displays. It contains a table that lists the host groups that exist for the array.

The table includes the following data for each host group:

- Host group number and name, for example, 000-G000.
- Port number to which the host group belongs.
- Platform configured in the host group.
- Middleware configured in the host group.

To display detailed data for a single host group:

In the Host Groups dialog box, click the name of the host group you want to view. The Properties dialog box for the selected host group is displayed.

If the wizard does not launch, disable your browser's pop-up blockers, then click the Add Array button at the bottom of the Array List dialog box to launch the wizard.

About array management and provisioning

About array discovery

The Add Array wizard is used to discover arrays on a storage network and add them to the Navigator 2 database. The first time you configure the array, the Add Array wizard launches automatically.

Understanding the Arrays screen

Each time Navigator 2 starts after the initial startup, it searches its database for existing storage systems and displays them in the Arrays dialog box. If another Navigator 2 dialog box is displayed, you can redisplay the Arrays dialog box by clicking Resource in the Explorer pane.

The Arrays dialog box provides a central location for you to view the settings and status of the HUS Family storage systems that Navigator 2 is managing. Buttons at the top left side of the dialog box let you run, stop, and edit error monitoring.

There is also a Refresh Information button you can click to update the contents in the window. Below the buttons are fields that show the storage system array status and error monitoring status.

Below the status indications are a drop-down list for viewing the number of rows and pages (25, 50, or 100), and buttons for moving to the next, previous, first, last, and a specific page in the Arrays dialog box. Buttons at the bottom of the Arrays dialog box let you perform various tasks involving the storage systems shown in the dialog box. [Table 7-1](#) describes the tasks you can perform with these

Add Array screen

This screen displays the results of the search that was specified in the Search Array screen. Use this screen to select the arrays you want to add to Navigator 2.

If you entered a specific IP address in the Search Array screen, that array is automatically registered in Navigator 2. Click Next to continue to the Finish screen. A message box confirming that the array has been added is displayed.

If you entered a range of IP addresses in the Search Array screen, all of the arrays within that range are displayed in this screen. To add an array whose name is displayed:

1. Click the to the left of the array name.
2. Click Next to add the arrays and continue to the Finish screen.

Adding a Specific Array

To add a specific array

1. If you know the IP address of a specific array that you want to add, click either Specific IP Address or Array Name to Search: and enter the IP address of the array. The default IP addresses for each controller are as follows:
 2. 192.168.0.16 - Controller 0
 3. 192.168.0.17 - Controller 1
4. Click Next to continue and open the Add Array screen.

If your management console is directly connected to a management port, enter the default IP address just for that port. Configure the other controller after the current controller. Omit the IP address for Controller 1 if your array has only one controller.

Adding Arrays Within a Range of IP Addresses

To add arrays within a range of IP addresses

1. If you know the range of IP addresses that includes one or more arrays that you want to add, click Range of IP Addresses to Search and enter the low and high IP addresses of that range. The range of addresses must be located on a connected local area network (LAN).
2. Click Next to continue and open the Add Array screen.

3. If any of the IP addresses entered are incorrect, when you click Next, Navigator 2 displays the following message:

Failed to connect with the subsystem. Confirm the subsystem status and the LAN environment, and then try again.
4. When configuring the management port settings, be sure the subnet you specify matches the subnet of the management server or allows the server to communicate with the port via a gateway. Otherwise, the management server will not be able to communicate with the management port.

Using IPv6 Addresses

Observe the following guidelines when using IPv6 addresses:

- Servers that process the IPv6 protocol may contain many temporary IPv6 addresses and may require additional time to communicate with the array. We recommend that you do not use temporary IPv6 address for this system.
- IPv6 multicast is used when Navigator 2 searches for an array in an IPv6 environment, but is usable only within the same subnet.

Notes on Using Hitachi Storage Navigator Modular 2

When using Hitachi Storage Navigator Modular 2, consider the following:

- All of the output files from Applet screens of Hitachi Storage Navigator Modular 2 are sent to the host side. Specify all of the files that Applet screens of Hitachi Storage Navigator Modular 2 inputs Hitachi Storage Navigator Modular 2 activates as files on the host side.
- Hitachi Storage Navigator Modular 2 functions cannot be used unless a TCP/IP communication is made between the array unit and the host. Verify that the TCP/IP is set correctly.
- Some Hitachi Storage Navigator Modular 2 functions are not available while the array unit is online with a host. For more information, see Chapter 2.
- When a high I/O load exists, the functions that are available while online might cause a command time-out in the host or a recovering fault in Hitachi Storage Navigator Modular 2. It is recommended that these functions be executed while offline.
- When Hitachi Storage Navigator Modular 2 is installed in the host connected to an array unit, I/O loading from a host might cause a command time-out on the host side or an abnormal termination on Hitachi Storage Navigator Modular 2 side.
- Installing Hitachi Storage Navigator Modular 2 in a host that does not connect an array unit is recommended.
- When the host enters the suspension status (low power mode) while Hitachi Storage Navigator Modular 2 is running, Hitachi Storage

Navigator Modular 2 might not operate correctly after the host is released from the suspension status.

- When you operate Hitachi Storage Navigator Modular 2, disable power management through Windows® so that the host will not enter the suspension status.
- Hitachi Storage Navigator Modular 2 allows Array unit screens to open for one array unit. When performing operations by opening Array unit screens for one array unit, operations may terminate abnormally. Open only one Array unit screen for one array unit to operate an array unit.
- If Hitachi Storage Navigator Modular 2 does not succeed in connecting to the array unit, the following message may display:

An invalid response was received from the subsystem

This message indicates that Hitachi Storage Navigator Modular 2 may have been connected to the array unit while the array unit automatically rebooted. Reconnect to the array unit after 3 minutes.

- Hitachi Storage Navigator Modular 2 may hang up in the following cases:
- The communication with the connected array unit fails due to controller blockage, array unit failure, disconnected LAN connection, or the array unit receives a Reset/LIP from the host.
- Other applications are working concurrently and memory utilization or a CPU use rate is high.
- Disconnecting LAN connection of the host and client sides.

If Hitachi Storage Navigator Modular 2 hangs up, check the array unit status and the connection status of the LAN, terminate Hitachi Storage Navigator Modular 2, and restart Hitachi Storage Navigator Modular 2. When not recovered, stop the SNM2 Server service (daemon process), check the array unit status and the connection status of the LAN, restart the SNM2 Server service (daemon process).

- In the following case, display time of Hitachi Storage Navigator Modular 2 may be longer while operating Hitachi Storage Navigator Modular 2.
- The communication with the connected array unit fails
- Failures of the array unit (controller blockade, etc.)

If display time gets longer in Hitachi Storage Navigator Modular 2, check the status of the array unit and the connection status of the LAN, terminate Hitachi Storage Navigator Modular 2, and restart Hitachi Storage Navigator Modular 2. If the array unit has a failure, operate by registering the array unit as the controller of one side. If the controller is blocked, you cannot connect to Hitachi Storage Navigator Modular 2. Operate it by registering the normal controller side until the controller blockade is recovered.

- The Java™ Applet Window may be displayed at the bottom of a screen item (such as a dialog box, popup menu, combo box) on a Hitachi Storage Navigator Modular 2 panel. This is not an error. When Hitachi Storage Navigator Modular 2 version 6.00 or higher, The Java™ Applet Window will not be displayed at the bottom of a screen.

- When Hitachi Storage Navigator Modular 2 connects the array unit over a LAN, it uses the TCP/IP port number of 2000. The TCP/IP port number of the array unit is set as 2000 by default.
- The restrictions listed in Table 1.1 exist if Hitachi Storage Navigator Modular 2 is used in combination with other programs, for one array unit:

Table 4-6: Restrictions for Multiple Programs used Concurrently for One Array Unit

Program Name	Concurrent Use Supported	Concurrent Use Not Supported
Hitachi Storage Navigator Modular 2	• SNMP Function • Disk Array-built in Web Server Function	• Hitachi Storage Navigator Modular 2 • Storage Navigator Modular for Web
Storage Navigator Modular for Web	• SNMP Function • Disk Array-built-in Web Server Function	• Hitachi Storage Navigator Modular 2 • Storage Navigator Modular for Web
SNMP Function	• Hitachi Storage Navigator Modular 2 • Storage Navigator Modular for Web • SNMP Function • Disk Array-built-in Web Server Function	N/A
Disk Array-built-in Web Server Function	• Hitachi Storage Navigator Modular 2 • Storage Navigator Modular for Web • SNMP Function • disk Array-built-in Web Server Function	N/A



NOTE: You can have two concurrent instances of Navigator 2 provided they are not connected to the same storage system. The termination of one instance may affect the other instance.

Also, if one Navigator 2 instance terminates while another Navigator 2 instance is current, the other instance may terminate. If this occurs, relaunch Navigator 2.

If you run a combination of programs in which concurrent use is not allowed, when a program usage restriction placed on it has started, launch another program with that combination after terminating the running program.

Security

This chapter will cover Account Authentication and Audit Logging.
The topics covered in this chapter are:

- ❑ [Security overview](#)
- ❑ [Account Authentication overview](#)
- ❑ [Audit Logging overview](#)
- ❑ [Data Retention Utility overview](#)

Security overview

Storage security is the group of parameters and settings that make storage resources available to authorized users and trusted networks - and unavailable to other entities. These parameters can apply to hardware, programming, communications protocols, and organizational policy.

Several issues are important when considering a security method for a storage area network (SAN). The network must be easily accessible to authorized people, corporations, and agencies. It must be difficult for a potential hacker to compromise the system.

The network must be reliable and stable under a wide variety of environmental conditions and volumes of usage. Protection must be provided against online threats such as viruses, worms, Trojans, and other malicious code. Sensitive data should be encrypted. Unnecessary services should be disabled to minimize the number of potential security holes.

Updates to the operating system, supplied by the platform vendor, should be installed on a regular basis. Redundancy, in the form of identical (or mirrored) storage media, can help prevent catastrophic data loss if there is an unexpected malfunction. All users should be informed of the principles and policies that have been put in place governing the use of the network.

Two criteria can help determine the effectiveness of a storage security methodology. First, the cost of implementing the system should be a small fraction of the value of the protected data. Second, it should cost a potential hacker more, in terms of money and/or time, to compromise the system than the protected data is worth.

Security features

Navigator 2 uses four features to create a security solution:

- Account Authentication
- Audit Logging
- Data Retention Utility

Account Authentication

The Account Authentication feature enables your storage system to verify the authenticity of users attempting to access the system. You can use this feature to provide secure access to your site and leverage the database of many accounts.

Hitachi provides you with the information needed to track the user on the system. If the user does not have an account on the array, the information provided will be sufficient to identify and interact with the user.

Audit Logging

When an event occurs, it creates a piece of information that indicates the user, operation, location of the event, and the results produced. This information is known as an Audit Log entry. When a user accesses the storage system from a computer in which HSNM2 operates and creates a RAID group at the time of a setting operation outside the system, the disk creates a log entry. The log indicates the exact time in hours, minutes, and day of the month, that the operation occurred. It also indicates whether the operation succeeded or failed.

Data Retention Utility

The Data Retention Utility feature protects data in your disk array from I/O operations performed at open-systems hosts. Data Retention Utility enables you to assign an access attribute to each logical volume. If you use the Data Retention Utility, you will can use a logical volume as a read-only volume. You will also be able to protect a logical volume against both read and write operations.

Security benefits

Security on your storage system provides the following benefits:

- **User access control** - Only authorized parties can communicate with each other. Consequently, a management station can interact with a device only if the administrator configured the device to allow the interaction.
- **Fast transmission and receipt** - Messages are received promptly; users cannot save messages and replay them to alter content. This prevents users from sabotaging SNMP configurations and operations. For example, users can change configurations of network devices only if authorized to do so.

Account Authentication overview

The Account Authentication feature enables your storage system to verify the authenticity of users attempting to access the system. You can use this feature to provide secure access to your site and leverage the database of many accounts.

Hitachi provides you with the information needed to track the user on the system. If the user does not have an account on the array, the information provided will be sufficient to identify and interact with the user.

Account Authentication is the process of determining who the user is, then determining whether to grant that user access to the network. The primary purpose is to bar intruders from networks. RADIUS authentication uses a database of users and passwords.

A user who uses the storage system registers an account (user ID, password, etc.) before beginning to configure account authentication. When a user accesses the storage system, the Account Authentication feature verifies whether the user is registered. From this information, users who use the storage system can be discriminated and restricted.

A user who registered an account is given authority (role information) to view and modify the storage system resources according to each purpose of system management and the user can access each resource of the storage system within the range of the authority (Access control).

Account Authentication features

Account Authentication is a licensed, role-based storage security feature that allows you to manage what storage systems users can access who have a valid Navigator 2 account. From the Account Authentication dialog box, which is accessed from an array enabled with Account Authentication, you can configure users who may access and control the array.

The Account Authentication module supports the following features:

- **Quick view of registered storage systems** - All active Navigator 2 users can view all of the registered arrays from the Navigator 2 Array List dialog box, including arrays that are enabled with Account Authentication
- **Quick status retrieval** - Storage system status can be retrieved quickly. Arrays enabled with Account Authentication are identified by the  symbol in the Status column in the Array List dialog box
- **Secure account information** - User account information (user name and password) for the Account Authentication feature is separate from the account information for Navigator 2 and is configured and stored on a secured array itself.

Account Authentication benefits

Account Authentication provides the following benefits:

- **Authorized communication** - Only authorized parties can communicate with each other. Consequently, a management station can interact with a device only if the administrator configured the device to allow the interaction.
- **High performance of message transmission** - Messages are received promptly; users cannot save messages and replay them to alter content. This prevents users from sabotaging SNMP configurations and operations. For example, users can change configurations of network devices only if authorized to do so.
- **Role customization convenience** - You can tailor access to the role of the user. Typical roles are storage administrator, account administrator and audit log administrator. This protects and secures data from unauthorized access internally and externally. It also provides focus for the user.

Account Authentication caveats

Navigator 2 users do not have automatic access to a secured array until an account is created for them by an administrator from a secured array (see procedure below).

A user will not have to provide login information if you use the same user name and password for both Navigator 2 and the Account Authentication secured array.

The built-in or default root user account should only be used to create user names and passwords. We recommend that you change it immediately after enabling Account Authentication. Store your administration passwords according to your organization's security policies. There is no "back door" key to access an array in the event of a misplaced, lost, or forgotten password.

Account Authentication task flow

Since Account Authentication does not permit users who have not registered the accounts to access the storage system, it can prevent illegal break-in. Besides, since it can assign the authority to view and modify the resources according to each purpose of system management by the role information, it can place restrictions on illegal operation for another purpose other than the management of the storage system, even in the case of users who have registered their accounts.

The following steps detail the task flow of the Account Authentication configuration process:

1. You determine that selected users need to have access to your storage system and that all other users should be blocked from access to it.
2. You identify all users for access and all for denial, creating separate lists.
3. Configure the license for Account Authentication.
4. Log into HSNM2.

5. Go to the Access Control area in HSNM2 that controls the Authentication database.
6. Set a role-based permission to the administrator to whom you are granting access to the storage system. The three administrator roles supported on HSNM2 are:
 - **Account administrator.** This figure manages and provisions secure settings for individual accounts set for the storage system.
 - **Audit Log administrator.** This figure manages, retrieves, and provisions the Audit Log environment that is a record of all actions involved with the storage system.
 - **Storage administrator.** This figure manages and provisions storage configurations on the storage system.
7. The newly configured administrator sends a request (a security query packet) to a storage switch.
8. The storage switch forwards the packet to a location on the storage system that contains one of the following types of information.

In the instance of the account administrator

 - User account information
 - User role information

In the instance of the Audit Log administrator

 - Audit Log information
 - Array configuration information

In the instance of the storage administrator

 - General data
 - Storage configuration information
9. The packet travels to either a storage area network or directly to the storage system where the packets transmit header is evaluated for its source.
10. If the source is allowed to obtain the data the packet is attempting to locate, then it is granted permission to reach and retrieve the data.

Figure 5-1 provides an outline of the Account Authentication process.




Figure 5-1: Account Authentication task flow

The Account Authentication feature is preinstalled and enabled from the factory. Be sure to review carefully the information on the built-in default account in this section before you log in to the array for the first time. The following table details the settings in the built-in default account.

Hitachi recommends that you also create a service personnel account and assign the Storage Administrator (View and Modify) role.

We recommend that you create a public account and assign the necessary role to it when operating the disk array. Create a monitoring account to monitor possible failures by Navigator 2 for disk array operation. Assign the Storage Administrator (View and Modify) role.

For more information on Sessions and Resources, see [Session on page 5-12](#).

Account Authentication specifications

Table 5-1 details account authentication specifications.

Table 5-1: Account Authentication specifications

Item	Description
Account creation	The account information includes a user ID, password, role, and whether the account is enabled or disabled. The password must have at least six (6) characters.
Number of accounts	You can register 200 accounts.
Number of users	256 users can log in. This includes duplicate log ins by the same user.
Number of roles per account	6 roles can be assigned to an account. • Storage Administrator (View and Modify) • Storage Administrator (View) • Account Administrator (View and Modify) • Account Administrator (View) • Audit Log Administrator (View and Modify) • Audit Log Administrator (View)
Time before you are logged out	A log in can be set for 20-60 minutes in units of five minutes, 70-120 minutes in units of ten minutes, one day, or indefinitely (OFF).
Security mode	The Advanced Security Mode. Refer to Advanced Security Mode on page 5-14 for more details.

Accounts

The account is the information (user ID, password, role, and validity/invalidity of the account) that is registered in the array. An account is required to access arrays where Account Authentication is enabled. The array authenticates a user at the time of the log in, and can allow the user to refer to, or update, the resources after the log in. [Table 5-2](#) details registered account specifications.

Table 5-2: Registered account specifications

Item	Description	Specification
User ID	An identifier for the account.	Number of characters: 1 to 256. Usable characters: ASCII code (0 to 9, A to Z, a to z, ! # \$ % & ' * + - . / = ? @ ^ _ ` { } ~).
Password	Information for authenticating the account.	Number of characters: 6 to 256. Usable characters: ASCII code (0 to 9, A to Z, a to z, ! # \$ % & ' * + - . / = ? @ ^ _ ` { } ~).
Role	A role that is assigned to the account.	Assignable role number: 1 to 6. For more information, see Roles on page 5-9 .

Table 5-2: Registered account specifications

Item	Description	Specification
Information of Account (enable or disable)	Information on enabling or disabling authentication for the account.	Account: enable or disable.

Account types

There are two types of accounts:

- Built-in
- Public

The built-in default account is a root account that has been originally registered with the array. The user ID, password, and role are preset. Administrators may create “public” accounts and define roles for them. When operating the disk array, create a public account as the normally used account, and assign the necessary role to it. See [Table 5-3](#) for account types and permissions that may be created.

The built-in default account may only have one active session and should be used only to create accounts/users. Any current session is terminated if attempting to log in again under this account.



CAUTION! To maintain security, change the built-in default password after you first log in to the array. Be sure to manage your root account information properly and keep it in a safe place. Without a valid username and password, you cannot access the array without reinstalling the firmware. Hitachi Data Systems Technical Support cannot retrieve the username or password.

Table 5-3: Account types

Type	Initial User ID	Initial Password	Initial Assigned Role	Description
Built-In	root (cannot change)	storage (may change)	Account Administrator (View and Modify)	An account that has been registered with Account Authentication beforehand.
Public	Defined by administrator (cannot change)	Defined by administrator	Defined by administrator	An account that can be created after Account Authentication is enabled.

Roles

A role defines the permissions level to operate array resources (View and Modify or View Only). You can place restrictions by assigning a role to an account. [Table 5-4](#) details role types and permissions.

Table 5-4: Role types and permissions

Type	Permissions	Role Description
Storage Administrator (View and Modify)	You can view and modify the storage.	Assigned to a user who manages the storage.
Storage Administrator (View Only)	You can only view the storage.	Assigned to a user who views the storage information and a user who cannot log in with the Storage Administrator (View and Modify) in the modify mode.
Account Administrator (View and Modify)	You can view and modify the account.	Assigned to a user who authenticates the account information.
Account Administrator (View Only)	You can only view the account.	Assigned to a user who views the account information. and a user who cannot log in with the Account Administrator (View and Modify) in the modify mode.
Audit Log Administrator (View and Modify)	You can view and modify the audit log settings.	Assigned to a user who manages the audit log.
Audit Log Administrator (View Only)	You can only view the audit log.	Assigned to a user who views the audit log and a user who cannot log in with the Audit Log Administrator (View and Modify) in the modify mode.

Resources

The resource stores information (repository) that is defined by a role (for example, the function to create a volume and to delete an account).

[Table 5-5](#) details authentication resources.

Table 5-5: Resources

Resource Group	Repository	Description
Storage management	Role definition	Stores role information. What access a role has for a resource (role type, resource, whether or not you can operate).
Storage management	Key	Stores device authentication information (an authentication name for the CHAP authentication of the iSCSI and the secret (a password)).
Storage management	Storage resource	Stores storage management information such as that on the hosts, switches, volumes, and ports and settings.
Account management	Account	Stores user ID, password, etc. account information.
Account management	Role mapping	Stores information on the correspondence between an account and a role.
Account management	Account setting	Stores information on account functions For example, the time limit until the session times out, the minimum number of characters in a password, etc.

Table 5-5: Resources

Resource Group	Repository	Description
Audit log management	Audit log setting	A repository for setting Audit Logging. (IP address of the transfer destination log server, etc.)
Audit log management	Audit log	A file that stores the audit log in the array.

The relationship between the roles and resource groups are shown in the following table. For example, an account which is assigned the Storage Administrator role (View and Modify) can perform the operations to view and modify the key repository and the storage resource. [Table 5-6](#) details role and resource group relationships.

Table 5-6: Role and resource group relationships

Resource Group Name (Repository)	Role Definition	Key	Storage Resource	Account	Role Mapping	Account Setting	Audit Log Setting	Audit Log
Role Name								
Storage Administrator (View and Modify)	-	V/M	V/M	X	X	X	X	X
Storage Administrator (View Only)	-	V	V	X	X	X	X	X
Account Administrator (View and Modify)	-	X	X	V/M	V/M	V/M	X	X
Account Administrator (View Only)	-	X	X	V	V	V	X	X
Audit Log Administrator (View and Modify)	-	X	X	X	X	X	V/M	V
Audit Log Administrator (View Only)	-	X	X	X	X	X	V	V

Table Key:

- **V** = "View"
- **M** = "Modify"
- **V/M** = "View and Modify"

- **x** = “Cannot view or modify”
- **-** = “Not available”

Session

A session is the period that you logged in and out from an array. Every log in starts a session, so the same user can have more than one session.

When the user logs in, the array issues a session ID to the program they are operating. 256 users can log in a single array at the same time (including multiple log ins by the same user).

The session ID is deleted when the following occurs (note that after the session ID is deleted, the array is not operational):

- A user logs out
- A user is forced to log out
- The status without an operation exceeds the log in validity
- The planned shutdown is executed



NOTE: Pressing the **Logout** button does not immediately terminate an active session. The status for the array(s) remains “logged in” until the session timeout period is reached for either the array itself or by Navigator 2 reaching its timeout period.

One of two session timeout periods may be enforced from Navigator 2:

- Up to 17 minutes when a Navigator 2 session is terminated by pressing **Logout** from the main screen.
- Up to 34 minutes when a Navigator 2 session is terminated by closing the Web browser dialog box.

Session types for operating resources

A session type is used to avoid simultaneous resource updates by multiple users.

When multiple public accounts with the View and Modify role log in the array, the Modify role is given to the account that logs in first. The account that logs in after, only has the View role. However, if a user with the Storage Administrator (View and Modify) role logs in first, another user with the Account Administrator (View and Modify) role can still log in and have the Modify role because the roles are not duplicate. [Table 5-7](#) details authentication session types.

Table 5-7: Session Types

Type	Operation	Maximum Number of Session IDs
Modify mode	View and modify (setting) array operations.	3 (Only one log in for each role)
View mode	Only view the array setting information.	256

The built-in account for the Account Administrator role always logs in with the Modify mode. Therefore, after the built-in account logs in, a public account that has the same View and Modify role, is forced into the View mode.

Advanced Security Mode

The Advanced Security Mode is a feature that improves the strength of the password encryption registered in the array. By enabling the Advanced Security Mode, the password is encrypted in the next generation method which has the 128-bit strength.

Table 5-8: Advanced Security Mode description, specifications

Feature	Description	Specifications
Advanced Security Mode	You can select the strength of the encryption when you register the password in the array.	• Selection scope: Enable or disable (default). • Authority to operate. Built-in account only. • The encryption is executed using SHA256 when it is enabled and MD5 when it is disabled.

Advanced Security Mode can only be operated with a built-in account. Also, it can be set only when the firmware of version 0890/A or later is installed in the storage system and Navigator 2 of version 9.00 or later is installed in the management PC.

By changing the Advanced Security Mode, the following information is deleted or initialized. As necessary, check the set following information in advance, and set it again after changing the mode:

- All sessions during login (accounts during login are logged out)
- All public accounts registered in the storage system
- Role and password of the built-in account

Changing Advanced Security Mode

When you change the Advanced Security Mode, the following information will be deleted or initialized:

- All logged-in sessions. The logged-in account will log out.
- All public accounts registered to the storage system.
- The roles and password of the built-in account.

You can only change Advanced Security Mode using a built-in account.

To change Advanced Security Mode

1. From the command prompt, connect to the storage system to which you will change the Advanced Security Mode.
2. Execute the `aaccountopt` command to change the Advanced Security Mode.

Account Authentication procedures

The following sections describe Account Authentication procedures.

Initial settings

To configure initial settings

1. Verify that you have the environments and requirements for Account Authentication (see [Preinstallation information for Storage Features on page 3-20](#)).
2. Install the license.
3. Log in to Navigator 2.
4. Change the default password for the “built-in” account (see [Account types on page 5-9](#)).
5. Register an account (see [Adding accounts on page 5-17](#)).
6. Registering an account for the service personnel (see [Adding accounts on page 5-17](#)).

Managing accounts

The following sections describe how to:

- Display accounts — see [Displaying accounts](#), below.
- Add accounts — see [Adding accounts](#), below.
- Modify accounts — see [Modifying accounts on page 5-19](#).
- Delete accounts — see [Deleting accounts on page 5-21](#).

Displaying accounts

To display accounts, you must have an Account Administrator (View and Modify or View Only) role. See [Table 5-3 on page 5-9](#) for accounts types and permissions that may be created.

To display accounts

1. Start Navigator 2 and log in. The Arrays dialog box appears
2. Select the appropriate array and click **Show & Configure Array**.
3. Log in as an Account Administrator (View and Modify) or an Account Administrator (View Only)
4. Select the **Account Authentication** icon in the Security tree view.
5. The account information appears, as shown in [Figure 5-2 on page 5-16](#).




Figure 5-2: Account Information window

Review the areas of this dialog box as shown in [Table 5-9](#).

Table 5-9: Contents of the Account Information screen

Item	Description
User ID	Displays a standard ASCII string that identifies the user.
Account Type	Displays the account type.
Account Enable/Disable	Displays the administrative state of the account, either Enabled or Disabled.
Session Count	Displays the number of active sessions associated with the account. To obtain more information, refer to the session list.
Update Permission.	Displays the state of whether permissions can be updated or not. Allowed: The session ID is Modify mode. The session ID is View mode.

- When the Session Count value is one or more, you can refer to the session list. Click the numeric characters for the Session Count. The logged sessions count list appears as shown in




Figure 5-3: Sessions dialog box displaying root

Adding accounts

To add accounts, you must have an Account Administrator (View and Modify) role. After installing Account Authentication, log in with the built-in account and then add the account. When adding accounts, register an optional user ID and a password, and avoid the following strings:

Built_in_user, Admin, Administrator, Administrators, root, Authentication, Authentications, Guest, Guests, Anyone, Everyone, System, Maintenance, Developer, Supervisor.

To add accounts

1. Start Navigator 2 and log in. The Arrays dialog box appears
2. Select the appropriate array and click **Show & Configure Array**.
3. Log in as an Account Administrator (View and Modify) or an Account Administrator (View Only)
4. Select the **Account Authentication** icon in the Security tree view. Expand the Account Authentication list, and click **Account**. The Account Screen is displayed.
5. Click **Add Account**. as shown in




Figure 5-4: Account Authentication - Account Information tab, adding account

The Add Account screen is displayed. See [Figure 5-5 on page 5-18](#).

Add Account

[Help](#)

Add Account Settings

Enter the information to add the account.

* User ID:
From 1 to 256 characters(alphanumeric characters, special symbols
 "!", "#", "\$", "%", "&", "*", "+", "-", "/", "=", "?", "@", "^", "_", "`", "{", "|", "}", "~")

Account:
☒ Enable
☐ Disable

* Password:
From 6 to 256 characters(alphanumeric characters, special symbols
 "!", "#", "\$", "%", "&", "*", "+", "-", "/", "=", "<", ">", "?", "@", "[", "\", "]", "^",
 "_", "`", "{", "|", "}", "~")

* Retype Password:
 Please retype the password.

* Roles:

Available Roles	
☐	Roles
☒	Account Administrator (View and Modify)
☐	Account Administrator (View Only)
☐	Audit Log Administrator (View and Modify)
☐	Audit Log Administrator (View Only)
☒	Storage Administrator (View and Modify)
☐	Storage Administrator (View Only)

* Required field

Figure 5-5: Add Account dialog box

6. Type a new username in the **User ID** field.
7. Select **Enable** in Account to enable the account.
8. Type the old password in the Old password field. Then type the new password in the New password field. Then retype the new password in the Retype password field.

When skipping the password change, uncheck the Change Password Checkbox.

9. Click **Next**. The Confirm wizard appears.

Changing the Advanced Security Mode

To change security mode

1. Start Navigator 2 and log in. The Arrays dialog box appears
2. Select the appropriate array and click **Show & Configure Array**.
3. Log in as an Account Administrator (View and Modify) or an Account Administrator (View Only)
4. Select the **Account Authentication** icon in the Security tree view. Expand the Account Authentication list, and click **Account**. The Account Screen is displayed.

- Click **Change Security Mode**. The Change Security Mode screen displays as shown in [Figure 5-6](#).




Figure 5-6: Change Security Mode dialog box

- Change the Enable checkbox setting to enable or disable the Advanced Security Mode status.
 - To enable Advanced Security Mode, make sure the checkbox is checked.
 - To disable the Advanced Security Mode, make sure the checkbox is unchecked.
- Click **OK**.
- Observe any messages that display and click **Confirm** to continue. An example of a system message displays in [Figure 5-7](#)




Figure 5-7: Change Security Mode System Message

- Click **Close**.

Modifying accounts

If you are an Account Administrator (View and Modify), you can modify the account password, role, and whether the account is enabled or disabled.

Note the following:

- You cannot modify your account unless you are using the built-in account.
- A public account cannot modify a built-in account.
- The user ID of the public account and built-in account cannot be changed.

To modify accounts

- Start Navigator 2 and log in. The Arrays dialog box appears
- Select the appropriate array and click **Show & Configure Array**.
- Log in as an Account Administrator (View and Modify).

- [Change Security Mode](#)
[Change Password](#)
[Refresh Information](#)
[Help](#)

Account Authentication

HUS110_91100016 > [Security](#) > **Account Authentication**

Summary

Advanced Security Mode

Enable

Account Information

Option

Rows/Page: 25

14 34

Page 1 of 1

☐	User ID	Account Type	Account Enable/Disable	Session Count	Update Permission
☒	 root	Built-in	Enable	1	Allowed
☐	 User001	Public	Enable	0	---

[Add Account](#)
[Edit Account](#)
[Delete Account](#)
[Forced Logout](#)
[Filter](#)
[Filter Off](#)

The **Edit Account** dialog box appears, as shown in [Figure 5-9](#).

Edit Account Settings

Enter the information to edit the account.

User ID :
Account :

User001

☒ Enable

☐ Disable

* Password :

☐ Change the password

Password :

From 6 to 256 characters(alphanumeric characters, special symbols
!"#\$%&'()*+,-./:;<=>?@[\]^_`{|}~)

Retype Password :

Please retype the password.

* Roles :

Available Roles

☐ Roles

☒ Account Administrator (View and Modify)

☐ Account Administrator (View Only)

☐ Audit Log Administrator (View and Modify)

☐ Audit Log Administrator (View Only)

☒ Storage Administrator (View and Modify)

☐ Storage Administrator (View Only)

* Required field

OKCancel

OK

6. Select either **Account Enable / Disable** or **New Password** and **Retype Password**.
7. Select the Role to be modified, if any.

8. Click **OK**.
9. Review the information in the Confirmation screen and any additional messages, then click **Close**.
10. Follow the on-screen instructions.

Deleting accounts

If you are an Account Administrator (View and Modify), you can delete accounts. Note that you cannot delete the built-in, and your own, account.



NOTE: A user with active session is automatically logged out if you delete the account when they are logged in.

To delete accounts

1. Start Navigator 2 and log in. The Arrays dialog box appears
2. Select the appropriate array and click **Show & Configure Array**.
3. Log in as an Account Administrator (View and Modify) or an Account Administrator (View Only)
4. Select the **Account Authentication** icon in the Security tree view. Expand the Account Authentication list, and click **Account**. The Account Screen is displayed.
5. Select the account from the Account list to be deleted, then click **Delete Account** as shown in [Figure 5-10](#).




Figure 5-10: Account Authentication - Account Information tab, deleting account

6. Review the information in the Confirmation screen and any additional messages, then click **Close**.
7. Follow the on-screen instructions.

Changing session timeout length

If you are an Account Administrator (View and Modify or View Only), you can change how long a user can be logged in.

To change the session length

1. Start Navigator 2 and log in. The Arrays dialog box appears
2. Select the appropriate array and click **Show & Configure Array**.
3. Log in as an Account Administrator (View and Modify) or an Account Administrator (View Only)
4. Select the **Account Authentication** icon in the Security tree view. Expand the Account Authentication list, and click **Account**. The Account Screen is displayed.
5. Click the **Option** tab. The Account Authentication - Option tab displays as shown in [Figure 5-11](#)




Figure 5-11: Account Authentication Option tab

6. Click **Change session timeout time**. The Change session timeout time screen appears as shown in [Figure 5-12](#).




Figure 5-12: Change session timeout time dialog box

7. Under Session timeout, select **Enable** or **Disable**.
8. If you selected **Enable**, choose a session timeout value from the drop-down list.
9. Click **OK**.

Forcibly logging out

Log out forcibly when you want to log out other users except for the built-in account user.



NOTE: When a controller failure occurs in the array during a log in, a session ID can remain. Consequently, forcibly log out all accounts.

To forcibly log out of a specific account

1. Start Navigator 2 and log in. The Arrays dialog box appears
2. Select the appropriate array and click **Show & Configure Array**.
3. Log in as an Account Administrator (View and Modify) or an Account Administrator (View Only).
4. Select the **Account Authentication** icon in the Security tree view. Expand the Account Authentication list, and click **Account**. The Account Screen is displayed.
5. Select the account you want to forcibly log out from Account list, then click **Forced Logout** as shown in [Figure 5-13](#).




Figure 5-13: Account Authentication - Account Information tab, forcing logout

6. Observe any messages that appear and click Confirm to continue.
7. Review the information in the Confirmation screen and any additional messages, then click **Close**.

Setting and deleting a warning banner

A warning banner is a mechanism that lists recent messages that have been generated by the Account Authentication mechanism.

To set a warning banner

1. Start Navigator 2 and log in. The Arrays dialog box appears
2. SLog in as an Account Administrator (View and Modify) or an Account Administrator (View Only).
3. Select the **Security** icon in the **Administration** menu in the Explorer.

4. Select the Warning Banner option in the Security menu. The Warning Banner screen displays. Then click Edit Message in the Warning Banner screen displays as shown in [Figure 5-14](#).




Figure 5-14: Warning Banner window, editing messages

The Edit Message screen displays as shown in [Figure 5-15](#).

Edit Message




Figure 5-15: Edit Message window

5. Enter a text to the Message frame, and click **Preview**.

6. Review the preview contents and click **Ok**. A set message displays in the Warning Banner view as shown in Figure 5-16.




Figure 5-16: Set Message text in the Warning Banner view

7. Click **Logout** and restart Navigator 2.

To delete a warning banner

1. Start Navigator 2 and log in. The Arrays dialog box appears
2. SLog in as an Account Administrator (View and Modify) or an Account Administrator (View Only).
3. Select the **Security** icon in the **Administration** menu in the Explorer.
4. Click **Edit Message**. The Edit Message screen displays.
5. Click **Delete**, then click **Ok**, and then click **Logout**.

Troubleshooting

Problem: The permission to modify (View and Modify) cannot be obtained for a user who has the proper privileges.

Description and Solution: Log out of the account and then log back in. The account may become View Only.

If this problem occurs, the login status of the array is retained until the time-out of the array session occurs or while the login to Navigator 2 is valid (up to 17 minutes when Navigator 2 is terminated by pressing the **Logout** button or up to 34 minutes when Navigator 2 is terminated by clicking the **Close** or **X** button.

When a change of the settings of the array is required immediately after the logout, return to the Arrays screen by clicking the Resources button on the left side of the screen, and then terminate Navigator 2 by clicking the button.

See the section [Displaying accounts on page 5-15](#) and confirm the account has update permissions. When the number of sessions is more than one, you can confirm update permissions and IP addresses per session. Also, issue a forced logout operation to log out of this account forcibly because a user using the account requiring updated permissions cannot be specified.

Problem: Error message DMED1F0029 is received. You have no permission to modify.

Description and Solution: Please contact the Account Administrator and confirm your permission.

If your Modify permissions are confirmed and you are unable to modify:

- Failure monitoring is being performed using the built-in account.
- Another user/PC has logged in to the array under the built-in account.

When logging in by the built-in account, the permission to modify shifts to the built-in account, and the permission to modify of the public account under login is removed. Since the target of the built-in account is to be used as the host administrator (super user), create a public account having the necessary operation permission and use it for everyday use.

When monitoring failures, we recommend creating a failure monitoring account having only the Storage Administrator permission.

Problem: Session time-outs occur frequently.

Description and Solution: When logging in Navigator 2 by the built-in account, and session time-out occurs frequently during the operation, the following causes are possible:

- Failure monitoring is being performed using the built-in account.
- Another user/PC has logged in to the array under the built-in account.

When logging in by the built-in account, any current session of the built-in account is terminated. Since the target of the built-in account is to be used as the host administrator (super user), create a public account having the necessary operation permission and use it for everyday use.

When monitoring failures, we recommend creating a failure monitoring account having only the Storage Administrator permission.

Audit Logging overview

When an event occurs, it creates a piece of information that indicates the user, operation, location of the event, and the results produced. When a user accesses the storage system from a computer in which HSNM2 operates and creates a RAID group at the time of a setting operation outside the system, the disk creates a log entry. The log indicates the exact time in hours, minutes, and day of the month, that the operation occurred. It also indicates whether the operation succeeded or failed.

If the storage system enters the Ready status at the time of a status change (system event) inside the disk, the storage system creates a log indicating the exact time and success state of the Array Ready operation. It then sends a log to the Syslog server.

Audit Logging features

Audit Logging provides the following features:

- **History** - Provides a history of all operations performed on your storage system.
- **Timestamping** - Provides a series of timestamps that give you markers identifying when certain events occurred.

Audit Logging benefits

Audit Logging provides the following benefits.

- **Compliance** - More and more companies are required to show historical data for compliance of being able to identify a moment when a hacking event occurred on a system. The guidelines also indicate that a company has to prove they can trace irregular actions. Audit Log perform both functions.
- **Accountability** – Log data can identify what accounts are associated with certain events. This information then can be used to highlight where training and/or disciplinary actions are needed.
- **Reconstruction** – Log data can be reviewed chronologically to determine what was happening both before and during an event. For this to happen, the accuracy and coordination of system clocks are critical. To accurately trace activity, clocks need to be regularly synchronized to a central source to ensure that the date/time stamps are in synch.
- **Intrusion detection** – Unusual or unauthorized events can be detected through the review of log data, assuming that the correct data is being logged and reviewed. The definition of what constitutes unusual activity varies, but can include failed login attempts, login attempts outside of designated schedules, locked accounts, port sweeps, network activity levels, memory utilization, key file/data access, etc.
- **Problem detection** – In the same way that log data can be used to identify security events, it can be used to identify problems that need

to be addressed. For example, investigating causal factors of failed jobs, resource utilization, trending and so on.

- **Creates an audit trail** - Enables you to problem solve and trace back to where a potential mistake has been made

Audit Logging task flow

The following steps detail the task flow of the Audit Logging configuration process:

1. You determine that security using Audit Logging would be helpful in tracking intrusions and potential hacking.
2. Log in to HSNM2.
3. Install the license key for Audit Logging.
4. Identify a syslog server to which you want Audit Log entries to be forwarded.
1. An host on a storage area network performs an action and sends a packet recording that action.
2. A PC that has Storage Navigator Modular 2 installed on it sends a packet to a domain that executes the setting of the Audit Log operation.
3. The Startingterminating process begins.
4. The output of the logged data is stored on an internal Audit Log database.
5. The logged data then is forwarded over to an External Syslog server.
6. The Audit Log record of the action is now ready for an Audit Log administrator to retrieve.
7. In HSNM2, go to the Audit Log area and indicate the IP address of the syslog server.
8. Events captured on the storage system are tracked and sent to the syslog.
9. Obtain these events off the box in real time so you have an external records of the actions taken. A typical instance of this is someone breaks into an array and has many failed attempts to log in. This generates a series of Audit Log entries that are forwarded on from the syslog server to the Event Management server.

Figure 5-17 figure details the sequence of events that occur when an audit log is created.




Figure 5-17: Audit Logging outline

Audit Logging specifications

Table 5-10 describes specifications for Audit Logging.

Table 5-10: Audit Logging specifications

Item	Description
Number of external Syslog server	Two IPv4 or IPv6 IP addresses can be registered.
External Syslog server transmission method	UDP port number 514 is to be used. The log conforms to the BSD syslog Protocol (RFC3164).
Audit log length	Less than 1,024 bytes per log. If the log (output) is more, the message may be incomplete. For the log of 1,024 bytes or more, only the first 1,024 bytes is output.
Audit log format	The end of a log is expressed with the LF (Line Feed) code. For more information, see the <i>Hitachi Unified Storage Command Line Interface Reference Guide (MK-91DF8276)</i> .
Audit log occurrence	The audit log is sent when any of the following occurs in the array. - Starting and stopping the array. - Logging in and out using an account created with Account Authentication. - Changing an array setting (for example, creating or deleting a volume). - Initializing the log.
Sending the log to the external Syslog server	The log is sent when an audit event occurs. However, depending on the network traffic, there can be a delay of some seconds.

Table 5-10: Audit Logging specifications (Continued)

Item	Description
Number of events that can be stored	2,048 events (fixed). When the number of events exceeds 2,048, they are wrapped around. The audit log is stored inside the system disk.

What to log?

Essentially, for each system monitored and likely event condition there must be enough data logged for determinations to be made. At a minimum, you need to be able to answer the standard who, what and when questions.

The data logged must be retained long enough to answer questions, but not indefinitely. Storage space costs money and at a certain point, depending on the data, the cost of storage is greater than the probable value of the log data.

Security of logs

For the log data to be useful, it must be secured from unauthorized access and integrity problems. This means there should be proper segregation of duties between those who administer system/network accounts and those who can access the log data.

The idea is to not have someone who can do both or else the risk, real or perceived, is that an account can be created for malicious purposes, activity performed, the account deleted and then the logs altered to not show what happened. Bottom-line, access to the logs must be restricted to ensure their integrity. This necessitates access controls as well as the use of hardened systems.

Consideration must be given to the location of the logs as well – moving logs to a central spot or at least off the sample platform can give added security in the event that a given platform fails or is compromised. In other words, if system X has catastrophic failure and the log data is on X, then the most recent log data may be lost. However, if X's data is stored on Y, then if X fails, the log data is not lost and can be immediately available for analysis. This can apply to hosts within a data center as well as across data centers when geographic redundancy is viewed as important.

Pulling it all together

The trick is to understand what will be logged for each system. Log review is a control put in place to mitigate risks to an acceptable level. The intent is to only log what is necessary and to be able to ensure that management agrees, which means talking to each system's stakeholders. Be sure to involve IT operations, security, end-user support, the business and the legal department.

Work with the stakeholders and populate a matrix wherein each system is listed and then details are spelled out in terms of: what data must be logged for security and operational considerations, how long it will be retained, how it will be destroyed, who should have access, who will be responsible to review it, how often it will be reviewed and how the review will be evidenced. The latter is from a compliance perspective – if log reviews are a required control, how can they be evidenced to auditors?

Finally, be sure to get senior management to formally approve the matrix, associated policies and procedures. The idea is to be able to attest both that reviews are happening and that senior management agrees with the activity being performed.

Summary

Audit logs are beneficial to have for a number of reasons. To be effective, IT must understand log requirements for each system, then document what will be logged for each system and get management’s approval. This will reduce ambiguity over the details of logging and facilitate proper management.

The audit log for an event has the format shown in [Figure 5-18](#).




Figure 5-18: Audit Log format

The output of an audit log is shown in [Figure 5-19](#). Items are separated by commas. When there is no item to be output, nothing is output.




Figure 5-19: Log example

For more details about Audit log format, see the *Hitachi Unified Storage Command Line Interface Reference Guide (MK-91DF8276)*.

Audit Logging procedures

The following sections describe the Audit Logging procedures.

Initial settings

To configure initial settings

1. Verify that you have the environments and requirements for Audit Logging (see [Preinstallation information for Storage Features on page 3-20](#)).
2. Set the Syslog Server (see [Table 5-10 on page 5-29](#)).

Optional operations

To configure optional operations

1. Export the internal logged data.
2. Initialize the internal logged data (see [Initializing logs on page 5-35](#)).

Enabling Audit Log data transfers

To transfer data to the Syslog server

1. Start Navigator 2 and log in. The Arrays dialog box appears
2. Select the appropriate array and click **Show & Configure Array**.
3. Log in as an **Audit Log Administrator** (View and Modify).
4. Select the **Audit Logging** icon in the Security tree view. The Audit Logging dialog box is displayed.
5. Click **Configure Audit Log**. The Configure Audit Log dialog box is displayed. See [Figure 5-20](#).




Figure 5-20: Configure Audit Log dialog box

6. Select the **Enable transfer to syslog server** check box.
7. Select the **Server 1** checkbox and enter the IP address for server 1. To add a second Syslog server, select the **Server 2** checkbox and enter the IP address for server 2.

8. To save a copy of the log on the array itself, select **Yes** under **Enable Internal Log**.



NOTE: This is recommended, because the log is sent to the Syslog server uses UDP, may not record all events if there is a failure along the communication path. See *Hitachi Unified Storage Command Line Interface Reference Guide (MK-91DF8276)* for information on exporting the internal log.

9. Click **OK**.

If the Syslog server is successfully configured, a confirmation message is sent to the Syslog server. If that confirmation message is not received at the server, verify the following:

- The IP address of the destination Syslog server
- The management port IP address
- The subnet mask
- The default gateway

Viewing Audit Log data

This section describes how to view audit log data.



NOTE: You must be logged on to the array as an **Audit Log Administrator** (View or View and Modify) to perform this task if the array is secured using Account Authentication.

To display the audit log

1. Start Navigator 2 and log in. The Arrays dialog box appears
2. Select the appropriate array and click **Show & Configure Array**.
3. Log in as an **Account Administrator** (View and Modify) or an Account Administrator (View Only).
4. Select the **Audit Logging** icon in the Security tree view. The Audit Logging dialog box is displayed.
5. Click **Show Internal Log**. The Show Internal Log confirmation screen appears as shown in Figure 5-21.

Show Internal Log




Figure 5-21: Show Internal Log confirmation

6. Select the **Yes, I have read the above warning and wish to continue** check box and press **Confirm**. The Internal Log screen opens (see Figure 5-22).

Show Internal Log




Figure 5-22: Internal Log window

7. Click **Close** when you are finished viewing the internal log.



NOTE: The output can only be executed by one user at a time. If the output fails due to a LAN or controller failure, wait 3 minutes and then execute the output again.

Initializing logs

When logs are initialized, the stored logs are deleted and cannot be restored. Be sure you export logs before initializing them. For more information, see *Hitachi Unified Storage Command Line Interface Reference Guide (MK-91DF8276)*.

To initialize logs

1. Start Navigator 2 and log in. The Arrays dialog box appears
2. Select the appropriate array and click **Show & Configure Array**.
3. Log in to Navigator 2. If the array is secured with Account Authentication, you must log on as an **Account Administrator** (View and Modify) or an Account Administrator (View Only).
4. Select the **Audit Logging** icon in the Security tree view. The Audit Logging dialog box is displayed (see [Figure 5-23](#)).

Initialize Internal Log




Figure 5-23: Initialize Internal Log message window

5. Select the **Yes, I have read the above warning and wish to continue** check box and click **Confirm**.
6. Review the confirmation message and click **Close**.



NOTE: All stored internal log information is deleted when you initialize the log. This information cannot be restored.

Configuring Audit Logging to an external Syslog server

If you are configuring Audit Logging to send log information from the array to an external syslog server, observe the following key points:

- Edit the syslog configuration file for the OS under which the syslog server runs to specify an output log file you name.
For example, under Linux syslogd, edit **syslog.conf** and add a proper path to the target log file, such as “/var/log/Audit_Logging.log”.
- Configure the syslog server to accept external log data
- Restart the syslog services for the OS under which the syslog server runs

We recommend that you refer to the user documentation for the OS that you use for your syslog data for more information on managing external log data transfers.

Data Retention Utility overview

The Data Retention Utility feature protects data in your disk array from I/O operations performed at open-systems hosts. Data Retention Utility enables you to assign an access attribute to each logical volume. If you use the Data Retention Utility, you will can use a logical volume as a read-only volume. You will also be able to protect a volume against both read and write operations.

Once data has been written, it can be retrieved and read only by authorized applications or users.

Data Retention Utility features

The following are Data Retention Utility features:

- **Data lock-down for authorized access** - Lock disk volumes as read-only for a prescribed period of time and ensure authorized-only access.
- **Data protection from standard I/O**- The Data Retention Utility protects data in your disk array from I/O operations performed at open-systems hosts.
- **Logical volume access** - Data Retention Utility enables you to assign an access attribute to each logical volume.
- **Read-only volumes** - If you use the Data Retention Utility, you will can use a logical volume as a read-only volume. You will also be able to protect a logical volume against both read and write operations.
- **Data tamper blocking** - Makes data tamper proof by making it non-erasable and non-rewritable.
- **Data retention period manageability** - Provides flexible retention periods where data cannot be altered or deleted during the specified interval.
- **WORM support** - Supports Write Once Read Many protocol for security of high number of records.

Data Retention Utility benefits

The following are Data Retention Utility benefits:

- **Sensitive data safety** - Protects sensitive information for compliance and legal purposes.
- **Casual data removal prevention** - Protects data from being accidentally removed.
- **Compliance** - facilitate compliance with government and industry regulations.

Data Retention Utility specifications

Table 5-11 shows the specifications of the Data Retention Utility.

Table 5-11: Specifications of the Data Retention Utility

Parameter	Specifications
Unit of setting	The setting is made for each unit. (However the expiration Lock is set for each disk array.)
Number of settable volumes	HUS 110: 2,048 volumes HUS 130/150: 4,096 volumes
Kinds of access attributes	Defines the following types of attributes: • Read/Write (default setting) • S-VOL Disable • Read Only • Protect • Read Capacity 0(can be set or reset by CCI only) • Invisible from Inquiry Command Can be set or reset by CCI only)
Guard against a change of an access attribute	A change from Read Only, Protect, Read Capacity 0, or invisible from Inquiry Command to Read/Write is rejected when the Retention Term does not expire or the Expiration Lock is set to ON.
Volumes not supported.	The following volumes are not supported: • Command device • DMLU • Sub-volume of a unified volume • Unformatted volume • Volume set as a data pool of SnapShot or TCE
Relation with ShadowImage/ SnapShot/TrueCopy/ TCE	If the S-VOL Disable is set for an volume, a volume pair using the volume as an S-VOL (data pool) is suppressed. A setting of the S-VOL Disable of a volume that has already become an S-VOL (V-VOL or data pool) is not suppressed only when the pair status is Split. Besides, when the S-VOL Disable is set for a P-VOL, restoration of SnapShot, restoration of ShadowImage is suppressed but a swapping of TrueCopy is not suppressed.
Powering off/on	An access attribute that has been set is retained even when the power is turned off/on.
Controller detachment	An access attribute that has been set is retained even following a controller detachment.
Relation with drive restoration	A correction copy, dynamic sparing, and copy back are performed like a usual volume.
Volume detachment	An access attribute that has been set for an volume is retained even when the volume is detached.
Restriction of firmware replacement	When the Data Retention Utility is enabled, initial setup and initialization of the feature's settings (Configuration Clear) are suppressed.

Restriction of access attribute setting	The following operations for a volume whose access attribute is other than Read/Write and for a RAID group that includes the volume are suppressed: • Volume deletion • Volume formatting • RAID group deletion
Setting by Navigator 2	Navigator 2 can set an access attribute, one volume at a time.
Unified VOL	A unified volume whose access level is a value other than Read/Write can neither be composed nor dissolved.
Deleting, growing, or shrinking of VOL	A volume for which an access attribute has been set cannot be deleting, growing, or shrinking. An access attribute can be set for a volume being grown or shrunken volume.
Expansion of RAID group	You can expand the RAID group to which the volumes that the access attribute is set belong.
Cache Residency Manager	An volume for which an access attribute has been set can be used for the Cache Residency Manager. On the other hand, an access attribute can be set for an volume being used for the Cache Residency Manager.
Concurrent use of LUN Manager	Available.
Concurrent use of Volume Migration	Available. The volume which executed the migration carries over the access attribute and the retention term set by the Data Retention Utility to the volume of the migration destination of the data and releases the access attribute and the retention term of migration resource (see Note below). When the access attribute is other than Read/Write, the volume cannot be specified as an S-VOL of Volume Migration.
Concurrent use of Password Protection	Available.
Concurrent use of SNMP Agent	Available.
Concurrent use of Cache Partition Manager	Available.
Concurrent use of Dynamic Provisioning	Available. The DP-VOLs that creating by Dynamic Provisioning cannot be used. The Data Retention Utility can be executed to the normal volume.
Setting range of Retention Term	From the 0th to 21,900 days (60 years) or unlimited.



NOTE: Figure 5-24 shows the status where the migration is performed for a volume which set the Read Only attribute. When the migration of the VOL0 which set the attribute of Read Only to the VOL1 in the RAID group 1 is executed, the Read Only attribute carries over to the volume of the migration destination of the data. Therefore, the VOL0 is in the status that the Read Only attribute is set irrespective of the execution of the migration. The Read Only attributes not copied to the VOL1. When the migration pair is released and the VOL1 is deleted from the reserved volume, a host can Read/Write to the VOL1.




Figure 5-24: Volume Migration of Read Only attribute

Data Retention Utility task flow

The following steps detail the task flow of the Data Retention Utility configuration process:

1. You find that some of your data is vulnerable to accidental loss or removal.
2. You determine that you want to deploy the Data Retention Utility to protect your volatile data.

3. You define time intervals, or *retention periods* for which you want data protected.
4. You configure the Data Retention Utility to apply to volumes that contain volatile data.
5. You enable the Data Retention Utility.

Assigning access attribute to volumes

By default, all the open-systems volumes are subject to read and write operations by open-systems hosts. For this reason, data on open-systems volumes might be damaged or lost if an open-systems host performs erroneous write operations. Also, confidential data on open-systems volumes might be stolen if an operator without approved access performs read operations on open-systems hosts.

By using the Data Retention Utility, you can use volumes as read-only volumes to protect the volumes against write operations. You can also protect logical volumes against both read and write operations. The Data Retention Utility enables you to restrict read operations and write operations on logical volumes and prevents data from being damaged, lost, and stolen.

To restrict read and write operations, you must assign an access attribute to each logical volume. Set the access attribute by using Command Control Interface (CCI) and/or Hitachi Storage Navigator Modular 2 (Navigator 2). A system administrator can set or reset one of the following access attributes for the each volume.

When the Read Only or Protect attribute is set using Navigator 2, the S-VOL Disable attribute for prohibiting a copy operation is set automatically. However, the S-VOL Disable attribute is not set automatically when CCI is used. When setting the Read Only, Protect, Report Zero Read Cap. mode, or Invisible mode using the CCI, specify the S-VOL Disable attribute for prohibiting a copy operation at the same time.

Read/Write

If a logical volume has the Read/Write attribute, open-systems hosts can perform both read and write operations on the logical volume.

ShadowImage, SnapShot, TrueCopy, and TCE can copy data to logical volumes that have Read/Write attribute. However, if necessary, you can prevent copying data to logical volumes that have the Read/Write attribute.

The Read/Write attribute is set by default for every volume.

Read Only

If a logical volume has the Read Only attribute, open-systems hosts can perform read operations but cannot perform write operations on the volume.

ShadowImage, SnapShot, TrueCopy, and TCE cannot copy data to volumes that have Read Only attribute.

Protect

If a logical volume has the Protect attribute, open-systems hosts cannot access the logical volume. Open-systems hosts cannot perform either read nor write operations on the volume.

ShadowImage, SnapShot, TrueCopy, and TCE cannot copy data to logical volumes that have Protect attribute.

Report Zero Read Cap. (Mode)

Report Zero Read Cap. mode can be set or reset by CCI only. When the Report Zero Read Cap. mode is set for the volume, the Read Capacity of the volume becomes zero. The host becomes unable to access the volume; it can neither read nor write data from/to it.

ShadowImage, SnapShot, TrueCopy, and TCE cannot copy data to an volume with an attribute that is Read Capacity 0.

Invisible (Mode)

The Invisible mode can be set or reset by CCI only. When the Invisible mode is set for the volume, the Read Capacity of the volume becomes zero and the volume is invisible from the Inquiry command. The host becomes unable to access the volume; it can neither read nor write data from/to it. The Read Capacity of the volume becomes zero and the volume is hidden from the Inquiry command.

ShadowImage, SnapShot, TrueCopy, and TCE cannot copy data to an volume with an attribute that is in Invisible mode.

Retention terms

When the access attribute is changed to Read Only, Protect, Read Capacity 0, or Invisible from Inquiry Command, another change to Read/Write is prohibited for a certain period. In the Data Retention Utility, the prohibited change period is called Retention Term. When the Retention Term of an volume is "2,190 days," the access attribute of the volume cannot be changed for 2,190 days ahead.

The Retention Term is specified when the access attribute changes to Read Only, Protect, Read Capacity 0, or Invisible from Inquiry Command from Read/Write. The Retention Term that has been specified once can be extended, but cannot be shortened.

When the Retention Term expires, the Retention Term of the volume, with an attribute is Read Only, Protect, Red Capacity 0, or Invisible from Inquiry Command, can be changed to Read/Write.



NOTE: The Retention Term interval is updated only when the disk array is in the Ready status. Therefore, the Retention Term may become longer than the specified term when the disk array power is turned on/off by a user. Also, the Retention Term interval may generate errors depending on the environment.

However, when the Expiration Lock is set to ON by Navigator 2, all the volume attributes, which are Read Only, Protect, Read Capacity 0, and Invisible from Inquiry Command, are unable to be changed to Read/Write.

When a host tries to write data to a Read Only volume, the write operation fails. The write failure is reported to the host. This occurs even when the Retention Term expires.

Also, when the Data Retention Utility is started for the first time, the Expiration Lock is set to OFF. When a host tries to read data from or write data to a logical volume that has the Protect attribute, the attempted access fails. The access failure is reported to the host.

Protecting volumes from copy operations

When ShadowImage, SnapShot, TrueCopy, or TCE copies data, the data on the copy destination volume (also known as the secondary volume) is overwritten. If a volume containing important data is specified as a secondary volume by mistake, ShadowImage, SnapShot, TrueCopy, or TCE can overwrite important data on the volume and you could suffer loss of important data. The Data Retention Utility lets you avoid potential data losses.

If you assign Read Only attribute or Protect attribute to a volume, ShadowImage, SnapShot, TrueCopy, and TCE cannot copy data to that logical volume. Any other write operations are prohibited on that logical volume. For example, business application software will be unable to write data to such a volume.

To block ShadowImage, SnapShot, TrueCopy, and TCE from assigning the volume as a secondary volume and permit the volume to be used by other data writing, set the access attribute of the volume as Read/Write. Additionally, when "Inhibition of S-VOL Making with Simplex volume (S-VOL Disable)" is set for the primary volume of ShadowImage, SnapShot, TrueCopy, or TCE, the following copy procedures in the primary volume can be prevented.

- Restoration by ShadowImage or SnapShot
- Takeover by TrueCopy



NOTE: In the ShadowImage, TrueCopy, and TCE manuals, the term "S-VOL" is used in place of the term "secondary volume".



NOTE: SnapShot has two types of secondary volumes: a virtual volume (V-VOL) and an area where differential data is stored (DP pool).

Usage

This section provides notes on using Data Retention.

Volume access attributes

Do not modify volume access attributes while operations are performed on the data residing on the volume, or the operation may terminate abnormally.

You cannot change access attributes for the following logical volumes:

- A volume assigned to command device
- A volume assigned to a DMLU
- An uninstalled volume
- A unformatted volume

Unified volumes

You cannot combine logical volumes that do not have a Read/Write attribute. Unification of a unified volume, whose access attribute is not Read/Write, cannot be dissolved.

SnapShot and TCE

A volume whose access attribute is not Read/Write, cannot be assigned to a DP pool. Additionally, an access attribute that is not Read/Write cannot be set for a volume that has been assigned to a DP pool.

SYNCHRONIZE CACHE command

When a SYNCHRONIZE CACHE command is received from a host, it usually writes the write pending data stored in the cache memory to drives. However, with Data Retention, the write pending data is not written to drives on the SYNCHRONIZE CACHE command.

When you need to write the write pending data stored in the cache memory, turn on the Synchronize Cache Execution Mode through Navigator 2. When you are done, turn it off, or the host application may fail.

Host Side application example

Uses IXOS-eCONserver.

Operating System (OS) restrictions

This section describes the restrictions of each operating system.

Volume attributes set from the operating system

If you set access attributes from the OS, you must do so before mounting the volume. If the access attributes are set after the volume is mounted, the system may not operate properly.

When a command (create partition, format, etc.) is issued to a volume with access attributes, it appears as if the command ended normally. However, although the information is written to the host cache memory, the new information is not reflected on the volume.

A OS may not recognize a volume when the volume number (volume) is larger than the one on which Invisible mode was set.

Windows 2000

A volume with a Read Only access attribute cannot be mounted.

Windows Server 2003/Windows Server 2008

When mounting a volume with a Read Only attribute, do not use the **diskpart** command to mount and unmount a volume. Use the **-x mount** and **-x umount** CCI commands.

Windows 2000/Windows Server 2003/Windows Server 2008

When setting a volume, Data Retention can only be used for basic disks. When Data Retention is applied to dynamic disks, volumes are not correctly recognized.

Unix

When mounting a volume with a Read Only attribute, mount it as Read Only (using the mount **-r** command).

Hewlett Packard Unix (HP-UX)

If there is a volume with a Read Only attribute, host shutdown may not be possible. When shutting down the host, change the volume attribute from Read Only to Protect.

If there is a volume with Protect attribute, host startup time may be lengthy. When starting the host, change the volume attribute to Read Only, or make the volume unrecognizable from the host by using mapping functions.

If a write is completed on the volume with a Read Only attribute, this may result in no response; therefore, do not perform write commands (e.g., dd command).

If Read/Write is done on a volume with a Protect attribute, this may result in no response; therefore, do not perform read or write commands (e.g. dd command).

Logical Volume Manager (LVM)

When changing the LVM configuration, the specified volume must be temporarily suspended using the **raidvchkset -vg** command. Place the volume again in the status in which it is checked when the LVM configuration change is completed.

HA Cluster Software

At times, a volume cannot be used as a resource for the HA cluster software (such as the MSCS), because the HA cluster software periodically writes management information in the management area to check resource propriety.

Notes on usage

The access attribute for a volume should not be modified while an operation is performed on the data residing on the volume. The operation may terminate abnormally.

Logical volume for which the access attribute cannot be changed:

The Data Retention Utility does not enable you to change the access attributes of the following logical volumes:

- A volume assigned to command device
- A volume assigned to DMLU
- An uninstalled volume
- A un-formatted volume

Notes about unified LU

You cannot combine logical volumes that do not have a Read/Write attribute. A unified volume whose access attribute is not Read/Write cannot be dissolved.

Notes About SnapShot and TCE

An volume, whose access attribute is not Read/Write, cannot be assigned to a data pool. Additionally, an access attribute other than Read/Write cannot be set for an volume that has been assigned to a data pool.

Notes and restrictions for each operating system

- Use a volume whose access attributes have been set from the OS:
 - If access attributes are set from the OS, they must be set before mounting the volume. If the access attributes are set to the volume after it is mounted, the system may not operate properly.
 - If a command (create partition, format, etc.) is issued to an volume with access attributes, from the operating system, it appears as if the command ended normally. The information is written to the host cache memory, the new information is not reflected in the volume.
 - An OS may not recognize a volume when the volume is larger than the one on which the Invisible mode was set.
- Microsoft Windows® 2000:
 - An volume with a Read Only access attribute cannot be mounted.
- Microsoft Windows Server 2003/Windows Server 2008
 - When mounting an volume with a Read Only attribute, do not use the **diskpart** command to mount and un-mount a volume. Use the -x mount and -x umount commands of CCI.
- Using Windows® 2000/Windows Server 2003/Windows Server 2008:
 - When setting a volume used by Windows® 2000/Windows Server 2003/Windows Server 2008 as the Data Retention Utility Volume, the Data Retention Utility can be applied to a basic disk only. When the Data Retention Utility is applied to a dynamic disk, an volume is not correctly recognized.
- Unix® OS
 - When mounting an volume with a is Read Only attribute, mount it as Read Only (using the mount -r command).
- HP-UX®
 - If there is an volume with a Read Only attribute, host shutdown might not be possible. When shutting down the host, change the attribute of volume from Read Only to Protect in advance.
 - A volume with a Protect attribute, host startup time may be lengthy. When starting the host, either change the attribute of the volume from Protect to Read Only, or use mapping functions to make the volume unrecognizable from the host.
 - If a write is completed on the volume with a Read Only attribute, it can results in no response; therefore, do not perform write commands (e.g. dd command).
 - If a Read/Write operation is performed on an volume with a Protect attribute, this may result in no response; therefore, do not perform read or write commands (for example, dd command).
- Using LVM

- If you change the LVM configuration, including Data Retention Volume, the specified volume must be temporarily blocked by the **raidvchkset -vg** command. Place the volume again in the status in which it is checked when the LVM configuration change is completed.
- Using HA cluster software
 - There may be times when an volume to which the Data Retention Utility is applied might not be used as a resource of the HA cluster software (such as the MSCS). This is because the HA cluster software (such as the MSCS) writes management information in the management area periodically to check propriety of the resource.

Operations example

The operations procedure to use of the Data Retention Utility are shown in the following sections.

Initial settings

[Table 5-12](#) indicates what chapters contain topics on initial settings.

Configuring and modifying key settings

Configuring and modifying key settings in the DRU software can help customize the data retention process so it fits your needs. Attributes that set access privileges and the secondary volume (S-VOL) object, which acts as a active standby storage system, both enable you to tune your storage system to perform in a desired manner.

Also both the retention term and expiration lock objects enable you to define how long the storage system holds specific data, enabling you to create the appropriate amount of space on the system and to optimize its performance.

Data Retention Utility procedures

To configure initial settings for the Data Retention Utility

1. Verify that you have the environments and requirements for Data Retention (see [Preinstallation information for Storage Features on page 3-20](#)).
2. Set the command device using the CCI. Refer to documentation for more information on the CCI.
3. Set the configuration definition file using the CCI. Refer to the appropriate CCI end-user document (see list above).
4. Set the environment variable using the CCI. Refer to the appropriate CCI end-user document (see list above).

Optional procedures

To configure optional operations

1. Set an attribute (see [Setting S-VOLs on page 5-50](#)).
2. Changing the retention term (see [Setting S-VOLs on page 5-50](#)).
3. Set an S-VOL (see [Setting S-VOLs on page 5-50](#)).
4. Set the expiration lock (see [Setting expiration locks on page 5-50](#)).

Opening the Data Retention dialog box

To open the Data Retention dialog box

1. Start Navigator 2 and log in. The Arrays dialog box appears
2. Click the appropriate array.
3. Click **Data Retention**. [Figure 5-25](#) appears.




Figure 5-25: Data Retention dialog box

4. The following options are available:
 - **VOL** - Volume number

- **Attribute** - Read/Write, Read Only, Protect, or Can't Guard
- **Capacity** - Volume size
- **S-VOL** - Whether the volume can be set to S-VOL (Enable) or not (Disable)
- **Mode** - The retention mode
- **Retention Term** - How long the data is retained



NOTE: When the attribute Read Only or Protect is set, the S-VOL is disabled.

5. Select the volume and click **Edit Retention**. The Edit Retention screen displays as shown in [Figure 5-26](#).

Edit Retention - 0000 Help

Edit Retention Property

Enter the information to edit the retention attribute. When the retention attribute changes from Read/Write to Read Only or Protect, the available secondary volume automatically changes from Enable to Disable. However, the volume does not automatically change when the attribute changes from Read Only or Protect to Read/Write.

Retention Attribute : ☒ Read/Write
☐ Read Only
☐ Protect

Secondary Volume Available : ☒ Enable

Retention Term : ☐ * Term
☒ Unlimited

From 0 day to 21900 days (60 years)

* Required field

OK Cancel

Figure 5-26: Edit Retention Property dialog box

6. Select the **Read Only** or **Protect** option from the Retention Attribute list.
7. Select **Term** or **Unlimited** from the Retention Term list. If you select Term, set a retention term in years (0 to 60) and/or days (0 to 21,900) and click **OK**.
8. Continue with the following sections to configure the desired Data Retention attributes.

Setting S-VOLs

To set S-VOLs

1. Select a volume, and click **Edit Retention**. The Edit Retention screen displays as shown in [Figure 5-27](#).




Figure 5-27: Edit Retention dialog box

2. Uncheck the **Enable** checkbox from the Secondary Volume Available area, and click **OK**.
3. Follow the on-screen instructions.

Setting expiration locks

To set expiration locks

1. Select the **Data Retention** icon in the Security tree view.
2. Click **Change Lock**. The Change Expiration Lock screen displays as shown in [Figure 5-28](#).




Figure 5-28: Change Expiration Lock window

3. Follow the on-screen instructions.

Setting an attribute

To set an attribute

1. Start Navigator 2.
2. Log in as a registered user to Navigator 2.
3. Select the storage system in which you will set up an attribute.
4. Click **Show & Configure Array**.
5. Select the **Data Retention** icon in the Security tree view.
6. Consider the fields and settings in the Data Retention dialog box as shown in [Table 5-12](#).

Table 5-12: Fields in the Data Retention dialog box

Item	Description
VOL	Displays the volume number.
Retention Attribute	Displays the attribute associated with managing the data. Values: Read/Write, Read Only, Protect, Can't Guard
Capacity	Displays the volume capacity.
Secondary Volume Available	Displays whether the volume can be set to S-VOL (Enable) or is prevented from being set to S-VOL (Disable).
Retention Term	Displays the length of time associated with the retention. Values: Unlimited or N/A.
Retention Mode	Displays the mode associated with retaining data. This field is for reference only. Values: Read Capacity 0 (Zero), Hiding from Inquiry Command Mode (Zero/Inv), or unspecifying (N/A).



NOTE: When Read only or Protect is set as the attribute, S-VOL will be disabled.

7. Select the volume and click **Edit Retention**.

The Edit Retention dialog box displays.




Figure 5-29: Edit Retention dialog box

8. Select **Read Only** or **Protect** from the Retention Attribute region.
9. Select **Term** or **Unlimited** from the Retention Term region.
If you select **Term**, set a retention term in years (0 to 60) and/or days (0 to 21,900).
10. Click **Ok** to display a confirmation message. Click **Confirm** and follow the screen instructions.

Changing the retention term



NOTE: The Data Retention Utility cannot shorten the Retention Term.

The retention term is the length of time that the storage system keeps the desired content. It can be either Unlimited or an integer value. If no retention time is specified, the notation for three dotted lines (---) displays as output.

To change the retention term

1. Select the volume, and then click **Edit Retention**.
The Edit Retention dialog box appears as shown in [Figure 5-29](#).
2. Select **Term** or **Unlimited** from Retention Term. If you select **Term**, set a Retention Term in years (0 to 60) and days (0 to 21,900).
A term of six years has been entered in default.
3. Click **OK** to display a confirmation message. Click **Confirm** and follow the screen instructions.

Setting the expiration lock

The expiration lock sets the time limit on when the data in your storage system is no longer needed.

To set the expiration lock:

1. Select the **Data Retention** icon in the Security tree view.
2. Click **Change Lock**.

The Change Expiration Lock dialog box displays.




Figure 5-30: Change Expiration Lock dialog box

3. Select **Enable**.
4. Click **OK** to display a confirmation message. Click **Confirm** and follow the screen instructions.

Setting S-VOL Disable

To set S-VOL Disable

1. Select the VOL, click Edit Retention.

The Edit Retention dialog box displays.




Figure 5-31: Edit Retention Property dialog box

2. Uncheck the **Enable** checkbox from the Secondary Volume Available area and click **OK**.
3. Click **Confirm** on the confirmation messages that display.

Provisioning volumes

This chapter will cover provisioning volumes.

- ❑ [LUN Manager overview](#)
- ❑ [Design configurations and best practices](#)
- ❑ [LUN Manager procedures](#)
- ❑ [Fibre Channel operations using LUN Manager](#)
- ❑ [iSCSI operations using LUN Manager](#)

LUN Manager overview

Volumes are user-designated partitions of the free storage space in a storage system and are used by a host to manage the data in the storage space they define. A volume can include all of the free storage space on a storage system or only part of it.

For example, you can create a volume for the free space on each drive, or divide the free space on a drive into parts and create a volume for each part. The parts can be any size you want. You could also create a volume that includes part of the free space on each of the drives.

The number of volumes you can create depend on your system. Refer to the user's guides for your system's specifications.

LUN Manager manages access paths between hosts and volumes for each port. With LUN Manager, two or more systems or operating systems (also called host groups) may be connected to one port of a Hitachi disk array, and volumes may be freely assigned to each host system.

With LUN Manager, illegal access to volumes from any host system may be prevented, and each host system may safely use a disk array as if it were connected to several storage systems.



NOTE: The term volume previously was referred to as a logical unit (volume). Most of the references to the term “logical unit” or “volume” have been changed to the term “volume,” although, in some instances, the term volume persists, especially in many of the figures in this chapter. These references will be changed progressively over the next several releases of HSNM2.

LUN Manager features

LUN Manager for Fibre Channel provides the following features.

- **Prevents illegal access.** LUN Manager for Fibre Channel prevents illegal access from other hosts. volumes are grouped and each group is registered in a port. LUN Manager specifies which host may access which volume by assigning hosts and volumes to each host group.
- **Host Connection Mode set for each host.** The Host Connection Mode can be set for each connected host. Also, the host connection mode can be set for each group.
- **Volume mapping set for each host.** The volume mapping feature can be set for each connected host. The volume numbers (H-LUN) recognized by a host can be assigned to each host group. By virtue of this, two or more hosts that require VOL0 can be connected to the same port.

You can connect additional hosts to one port, although more connections increases traffic on the port. When you use LUN Manager, design the system configuration appropriately to evenly distribute traffic at the port, controller, and drive.

Navigator 2 supports the following LUN Manager types:

- Standard volumes are just designated partitions of storage space.
- Differential Management Logical Units (DMLUs). DMLUs are volumes that consistently maintain the differences between them.
- SnapShot volumes are virtual volumes and are specified as the secondary volume of a SnapShot pair when you create a pair. See [Create SnapShot volume](#) for more information.

LUN Manager benefits

- **Ease of provisioning** - Enables you to divide up content on your storage system into units of a manageable size, enabling you to provisioning and manage your system with ease.
- **Ease of content identification** - Enables you to create a scheme that helps you easily identify where specific content resides in your storage system.

LUN Manager task flow

LUN Manager manages access paths between hosts and volumes for each port. With LUN Manager, two or more host systems or operating systems (also called host groups) may be connected to one port of a Hitachi disk array, and volumes may be freely assigned to each host system.

With LUN Manager, illegal access to volumes from any host system may be prevented, and each host system may safely use a disk array as if it were connected to several storage systems.

The following steps detail the task flow of the LUN Manager configuration process:

1. A system administrator determines that volumes are required for operating on a currently configured storage system in the data center.
2. Determine which protocol is being used in the storage system: either Fibre Channel or iSCSI.
3. Configure the license for LUN Manager.
4. Log into HSNM2.

For Fibre Channel

1. Assign volumes to the host.
2. Group hosts into a host group. Assign properties to the host group.
3. Assign volumes to RAID groups.
4. Determine how to prevent unauthorized access to the storage system, using Account Authentication.
5. Determine input/output paths for data passing through host and into storage system.
6. Determine queue depths for storage system.

For iSCSI

1. Use Storage Navigator Modular 2 to set up volumes on the array.
2. Use LUN Manager to set up the following on the array:
 - For each array port that will connect to the network, add one or more targets and set up target options.
 - Map the volumes to targets.
 - Register CHAP users that are authorized to access the volumes.
 - Keep a record of the iSCSI names and related settings to simplify making any changes later.
3. Physically connect the array to the network.
4. Connect hosts to their targets on the array by using the Initiator function in LUN Manager to select the host's initiator driver or the initiator iSCSI name of the HBA.
5. As a security measure, use LUN Manager in assignment mode to determine input/output paths between hosts and volumes. The input/output path is a route through which access from the host is permitted.
6. When connecting multiple hosts to an array port, verify and set the queue depth. If additional commands from the additional hosts exceed the port's limit, increase the queue depth setting.
7. Test host connections to the volumes on the array.
8. Perform maintenance as needed: host and HBA addition, volume addition, HBA replacement, and switch replacement. Refer to your HBA vendor's documentation and Web site.

Figure 6-1 illustrates a port being shared by multiple host systems with volumes created in the host:




Figure 6-1: Setting access paths between hosts and volumes for Fibre Channel

Understanding preconfigured volumes

The HUS storage systems are set up at the factory with one or more volumes, depending on the model. This helps users by making the storage systems easier and faster to configure. The factory configurations are described below, by model.

HUS storage systems are set up at the factory with one pre-configured volume. [Table 6-1](#) lists the parameters of that volume. If desired, you can create additional volumes.

Table 6-1: Preconfigured volume on HUS storage systems

No. Ctlrs	No. Volumes	Volume Nos.	Volume Type	Size	Port	Purpose/Notes
2	1	Volume 0	Volume	50 GB	0A	• Normal use. • Can be allocated to a host. May be spread across multiple drives.

LUN Manager specifications for Fibre Channel

[Table 6-2](#) details specifications for LUN Manager Fibre Channel.

Table 6-2: LUN Manager Fibre Channel specifications

Item	LUN Manager Fibre Channel Specifications
Host Group	128 host groups can be set for each port, and host group 0 (zero) is required.
Setting and Deleting Host Groups	• Host groups 1-127 can be set or deleted. • Host group 0 cannot be deleted. To delete the World Wide Name (WWN) and volume mapping of Host group 0, initialize Host group 0.
Host Group Name	A name is assigned to a host group when it is created, and this name can be changed.
WWN (Port Name)	• Up to 128 WWNs can be set for each port. • 128 WWNs for host bus adaptors (HBAs) and be set for a host group or port. • The WWN cannot be assigned to another host group on the same port. • A WWN may also be set to the host group by selecting it from an HBA WWN connected to the port.
Nickname	• An optional name may be assigned to a WWN allocated to a host group. • A name assigned to a WWN is valid until the WWN is deleted.
Host Connection Mode	The host connection mode of a host group can be changed.
Volume Mapping	• Volume mapping can be set to the host group. • 2,048 volume mappings can be set for a host group • 16,384 volume mappings can be set for a port.

Table 6-2: LUN Manager Fibre Channel specifications

Item	LUN Manager Fibre Channel Specifications
Enable and Disable Port Settings	- LUN Manager can be enabled or disabled for each port. - When LUN Manager is disabled, the information is available when it is enabled again.
Online Setting	When adding, modifying, or deleting settings, restarting the array is not required. To modify settings, Navigator 2 is required.
Maximum Queue Depth	32 commands per volume, and 1,024 commands per port.

About iSCSI

iSCSI makes it possible to construct an IP Storage Area Network (SAN), connecting many hosts and storage systems at low cost. However, iSCSI greatly increases the I/O workload of the network and the array. When using iSCSI, to obtain the advantages of using iSCSI, you must configure the network in a way where the workload evenly distributes the network, port, controller and drive.

While LAN switches and Network Interface Cards (NICs) are viewed in networks as equivalent nodes, some important differences exist between them with the LAN connection when you use iSCSI. Pay attention to the following:

iSCSI consumes almost all of the available Ethernet bandwidth, unlike a conventional LAN connection. The high consumption significantly degrades the performance of both the iSCSI traffic and the LAN. Make sure to separate the iSCSI IP SAN and the office LAN to ensure the network your group performs tasks on continues to enjoy good network performance.

The Host I/O load affects the iSCSI response time. Expect that when the Host I/O load increases, your iSCSI environment performance will degrade.

Create a backup path between the host and iSCSI where the active connection can switch to another path so that you can update the firmware without stopping the system. [Table 6-3](#) details LUN Manager iSCSI specifications.

Table 6-3: LUN Manager iSCSI specifications

Item	LUN Manager Fibre Channel Specifications
Target	255 targets can be set for each port, and target 0 (zero) is required.
Setting/Deleting a Target	Targets 1 through 254 can be set or deleted. Target 0 (zero) cannot be deleted. To delete the initiator iSCSI Name, options, and volume mapping of target 0 (zero), initialize target 0.
Target alias	A name is assigned to a target upon creation. This alias can be changed.

Table 6-3: LUN Manager iSCSI specifications (Continued)

Item	LUN Manager Fibre Channel Specifications
iSCSI Name	- Used for identifying initiators and targets. - iSCSI Name needs to have a World Wide Name (World Wide Unique), and iqu and eui are supported. - The iSCSI Name of a target is set as a World Wide Unique name when initializing the target.
Initiator iSCSI Name	- Up to 255 Initiator iSCSI Names can be set for each port. 256 initiator drivers or HBA iSCSI names can be per target per port. - The same Initiator iSCSI Name can be used by both targets on the same port. - The Initiator iSCSI Name to be set to the target can also be selected from the initiator drivers connected to the port, and the detected Initiators of the HBA.
Target iSCSI Name	- The Target iSCSI Name can be set for each target. - The same Target iSCSI Name cannot be set to another target on the same port.
Initiator Name	- An Initiator Name can be assigned to an initiator iSCSI - Name allocated to the target. An Initiator Name can be deleted. - An Initiator Name assigned to an initiator iSCSI Name is valid until the initiator iSCSI Name is deleted.
Discovery	SendTargets and iSNS are supported.
Authentication of login	None and CHAP are supported.
User Authentication Information	- User authentication may can be set for 512 ports. - The user authentication information can be set to the target that has been set by the LUN Manager. - The same user authentication information can also be set to other targets on the same port.
Host Connection Mode	The Host Connection Mode of the target can be changed.
Volume Mapping	- A volume can be set to the target. 2,048 volume mappings can be set for a target. - Up to 16,384 volume mappings can be set for a port.
Enable/Disable Settings for Each Port	When LUN Manager is disabled, the LUN Manager information is saved.
Online Setting	When adding, modifying, or deleting settings, you do not have to restart the array.
Other Settings	Navigator 2 is required.
Using LUN Manager with Other Features	The maximum number of configurable hosts is 239 if TrueCopy is installed on the array.
iSCSI target settings copy function	iSCSI target settings can be copied to the other ports to configure an alternate path.

Table 6-4 detail the acceptable combinations of operating systems and Host Bus Adapter (HBA) iSCSI entities.

Table 6-4: Operating System (OS) and host bus adapter (HBA) iSCSI combinations

Operating System	Software Initiator/Host Bus Adapter
Windows XP®	Microsoft iSCSI Software initiator + NIC
Windows® Server™ 2003	Microsoft iSCSI Software initiator + NIC Qlogic® HBA
Linux®	SourceForge iSCSI Software initiator + NIC Qlogic HBA

For additional OS support information, please review the following document located at the Hitachi Data Systems support site. Alternatively, go to <http://www.hds.com/products/interoperability/>. Or go to: <http://www.hds.com/assets/pdf/simple-modular-storage-100-sms100.pdf>

Design configurations and best practices

The following sections provide some basic design configurations and best practices information on setting up arrays under the Fibre Channel and iSCSI protocols.

When connecting multiple hosts to one port of the storage system, the storage system must be designed to accommodate the following:

System design. For proper system design, ensure the following tasks have been performed:

- Assign volumes to hosts
- Assign volumes to RAID groups
- Determine the system configuration
- Determine the method of illegal access prevention
- Determine queue depth

System configuration. For proper system configuration, ensure the following tasks have been performed:

- Set LUN Manager
- Set switch zoning

Component addition and replacement. For proper addition or replacement of components, ensure the following tasks have been performed:

- Host and HBA addition
- Volume addition
- HBA replacement
- Switch replacement

Fibre Channel configuration

The array is connected to the host with an optical fibre cable. The end of the cable on the host side is connected to a host bus adapter (HBA) and the end of the cable on the array side is connected to the array port.

Volumes can be grouped and assigned to a port as a host group. You can specify which HBA can access that group by assigning the WWNs of the HBAs to each host group. [Table 6-5](#) details combinations of OS and HBA for Fibre Channel.

Table 6-5: Combinations of OS and HBA for Fibre Channel

Operating System	HBA	Remarks
HP-UX [®]	HP HBA	When you are in HP-UX mode, you have selected Enable.
IRIX [®]	SGI [®] HBA	--
Windows [®]	Emulex [®] HBA (with Miniport Driver) Qlogic [®] HBA	--
Linux [®]	Emulex [®] HBA Qlogic [®] HBA	--

Conditions for Using LUN Manager for Fibre Channel

[Table 6-6](#) displays the conditions for using LUN Manager for Fibre Channel.

Table 6-6: Conditions for Using LUN Manager for Fibre Channel

Item	Conditions
Making Settings	- Firmware: Version 0915/B or more is required. - Navigator 2 Version 21.50 or more is required for management.
Using LUN Manager with other Optional Functions	Optional functions can be used together with LUN Manager.
Queue Depth	- Maximum 32 commands per volume. - Maximum 512 commands per port. Note in firmware version 0935/A or more, it is extensible to the maximum of 1,024 commands per port by the port option setting

Identify which volumes you want to use with a host, and then define a host group on that port for them (see [Figure 6-2 on page 6-10](#)).




Figure 6-2: Fibre Channel system configuration

Examples of configurations for creating host groups in multipathed and clustered environments appear in [Figure 6-3](#) and [Figure 6-4](#).




Figure 6-3: One host group Fibre Channel configuration




Figure 6-4: Two host groups Fibre Channel configuration

Fibre Channel design considerations

When connecting multiple hosts to an array port, make sure you do the following.

Fibre system configuration

To specify the input/output paths between hosts and volumes, set the following for each array. Keep a record of the array settings. For example, if an HBA is replaced, change the WWN name accordingly.

- Host group
- WWN of HBA
- Volume mapping
- Host connection mode

Connect the hosts and the array to a switch, and set a zone for the switch. Create a diagram and keep a record of the connections between the switch and hosts, and between the switch and the array. For example, when the switch is replaced, replace the connections.

iSCSI system design considerations

This section provides information on what you should consider when setting up your iSCSI network using LUN Manager.



CAUTION! To prevent unauthorized access to the array during setup, perform the first two bullets with the array not connected to the network.

iSCSI network port and switch considerations

This section provides information on when to use switches and what type of network ports you should use for your application.

- Design the connections of the hosts and the arrays for constructing the iSCSI environment. When connecting the array to more hosts than its ports, design the Network Switch connection and the Virtual LAN (VLAN).
- Choose a network interface for each host, either an iSCSI HBA (host bus adapter) or a NIC (network interface card) with a software initiator driver. The NIC and software initiator combination costs less. However, the HBA, with its own processor, minimizes the demand on the host from protocol processing.
- If the number of hosts to connect is greater than the number of iSCSI ports, network switches are needed to connect them.
- Array iSCSI cannot connect directly to a switch that does not support 1000BASE-T (full-duplex). However, a switch that supports both 1000BASE-T (full-duplex) and 1000BASE-SX or 100BASE-TX, will allow communication with 1000BASE-SX or 100BASE-TX.
- All connections direct to iSCSI in the IP-SAN should be 1000BASE-T (full-duplex).
- 100BASE-T decreases IP-SAN performance. Instead, use 1000BASE-T (full-duplex) for all connections.
- Array iSCSI does not support direct or indirect connections to a network peripheral that only supports 10BASE.
- The network switch is available as long as it is transparent to the arrays (port base VLAN, etc.).
- Array iSCSI does not support tagged VLAN or link aggregation. The packets to transfer such protocols should be filtered out in switches.
- When IP-SAN is designed, it is similar to construct the traditional network. Overlapping of addresses or a loop made in a subnet will cause serious degrade of communication performance and even cause disconnections.
- Network switches with management functions such as SNMP can facilitate network troubleshooting.
- To achieve the performance or security of iSCSI communication, you need to separate an IP-SAN (i.e., the network on which iSCSI communication is done) from the other network (management LAN, office LAN, other IP-SAN, etc.). The switch port VLAN function will be able to separate the networks logically.
- When multiple NICs are installed in a host, they should have addresses that belong to different network segments.

For iSCSI port network settings, note the following:

- Make sure to set the IP address (IPv4) to each iSCSI port so that it does not overlap the other ports (including other network equipment ports). Then set the appropriate subnet mask and default gateway address to each port.

- Targets are set to the subordinate of iSCSI ports. Target 0 is made in default for each iSCSI ports.
- Each iSCSI target is assigned its iSCSI name automatically.
- When connecting hosts and one port of the array using the network switch, a control to distinguish accessible host is required for each volume.

Additional system design considerations

Consider the following before configuring the array for your iSCSI network.

- Network boot disk is not supported. You cannot use an array as a “netboot” device as it does not support operation as a network boot disk
- Array reboot is not required for LUN Manager changes.

With LUN Manager, you can add, modify, or delete a target during system operation. For example, if an additional disk is installed or an additional host is connected, an additional target may still be created. If removing an existing host, the target that is connected to the host is deleted first and then the host is removed.

- Ensure that the host demand on an array does not exceed bandwidth.
- Use redundant paths to help ensure array availability if hardware components fail.
- Multiple host connects can affect performance.

Up to 255 hosts can be connected to an iSCSI port. It is possible to connect up to 255 hosts to an iSCSI port. Too many hosts, however, can increase network traffic beyond the processing capacity of the port. When using LUN Manager, you should design a system configuration to evenly distribute traffic concentrated at the port, controller, and disk drive.

- Use iSNS where possible to facility target discovery and management. Doing so eliminates the need to know IP addresses. Hosts must be connected to the IP-SAN to implement iSNS.
- iSCSI digests and performance.

For arrays that support both an iSCSI Header digest and an iSCSI Data digest, you can enable the digests to verify the integrity of network data. However, the verification has a modest cost in processing power at the hosts and arrays, in order to generate and check the data digest code. Typically data transfer decreases to about 90%. (This rate will be affected by network configuration, host performance, host application, and so forth).



NOTE: Enable digests when using an L3 switch (including router) to connect the host to the array iSCSI port.

To enable header and data digests, refer to your iSCSI initiator documentation, which may describe it as Cyclical Redundancy Checking (CRC), CRC32, or a checksum parameter Host Competition for Disk Access within a RAID Group Lowers Performance.

- Providing iSCSI network security. To provide network security, consider implementing one or more of the following:
 - Closed IP-SAN

It is best to design IP-SANs completely isolated from the other external networks.
 - CHAP authentication

You must register the CHAP user who is authorized for the connection and the secret in the array. The user can be authenticated for each target by using LUN Manager.

The user name and the secret for the user authentication on the host side are first set to the port, and then assigned to the target. The same user name and secret may be assigned to multiple targets within the same port.

You can import CHAP authentication information in a CSV format file. For security, you can only import, and not export CHAP authentication files with LUN Manager. Always keep CSV files secure in order to prevent others from using the information to gain unauthorized access.

When registering for CHAP authentication you must use the iSCSI name, acquiring the iSCSI Name for each platform and each HBA. Set the port-based VLAN of the network switch if necessary.
- Verify host/volume paths with LUN Manager

Determine input/output paths between hosts and volumes according to the assignment mode using LUN Manager. The input/output path is a route through which access from the host is permitted.

System topology examples

The array is connected to a host with an Ethernet cable (category 6). The end of the cable on the host side is connected to an iSCSI HBA or Network Interface Card (NIC). The end of the cable on the array side is connected to a port of the array.

Direct Attached and the Network Switch (Network Attached) are supported connection methods, and an IP-SAN connection using a Layer 2 or Layer 3 switch is also supported.

The following illustrations show possible topologies for direct attached connections.




Figure 6-5: Direct attached type 1 for iSCSI




Figure 6-6: Direct attached type 2 for iSCSI




Figure 6-7: Direct attached type 3 for iSCSI




Figure 6-8: Direct attached type 4 for iSCSI




Figure 6-9: Direct attached type 5 for iSCSI

The following figures show possible topologies for switch-attached connections.




Figure 6-10: Switch attached type 1 for iSCSI




Figure 6-11: Switch attached type 2 for iSCSI




Figure 6-12: Switch attached type 3 for iSCSI

Assigning iSCSI targets and volumes to hosts

The host recognizes volume between H-LUN0 and H-LUN255. When you assign volumes of more than 256 volumes to the host, you must set the target volume mapping to be between H-LUN0 and H-LUN255.

- Up to 2,048 volume mappings can be set for a target.
- Up to 16,384 volume mappings can be set for a port.




Figure 6-13: Mapping volumes between LU256-511 to the host

When assigning VOL3 to Host 1 and VOL4 to Host 2, both hosts can access the same volume if the volume mapping is set alone as shown in [Figure 6-14 on page 6-19](#). When LUN Manager or CHAP is used in this case, the host (iSCSI Name) access to each volume can be distinguished even in the same port as shown in [Figure 6-15 on page 6-19](#).




Figure 6-14: LUN mapping—different hosts can access volumes




Figure 6-15: Volume target assignment—separate host access to volumes

Preventing unauthorized SAN access

When connecting hosts to one port of an array using a switch, you must assign an accessible host for each volume.

When assigning VOL3 to Host 1 and VOL4 to Host 2 as in [Figure 6-16 on page 6-20](#), both hosts can access the same volume if the mapping is set separately.




Figure 6-16: Volume mapping—no host access restrictions

When LUN Manager or CHAP is used, the host (iSCSI Name) access to each volume can be distinguished even within in the same port as shown in [Figure 6-17 on page 6-20](#).




Figure 6-17: LUN Manager/CHAP—restricted host access

To prevent ports of the array from being affected by other hosts even when LUN Manager is used, it is recommended that zoning be set, as shown in [Figure 6-18 on page 6-21](#).




Figure 6-18: Switch zoning

Avoiding RAID Group Conflicts

When multiple hosts are connected to an array and the volumes assigned to each host belong to the same RAID group, concurrent access to the same disk can occur and performance can decrease. To avoid conflicts, only have one host access multiple volumes in one RAID group.

The number of RAID groups that can be created is determined by the number of mounted drives and the RAID level of the RAID groups you are creating. If you cannot create as many RAID groups as hosts to be connected, organize the RAID groups according to the operational states of the hosts (see [Figure 6-19 on page 6-21](#) and [Figure 6-20 on page 6-22](#)).




Figure 6-19: Hosts connected to the same RAID group




Figure 6-20: Hosts connected to different RAID groups

SAN queue depth setting

A host can queue array commands, and queue depth is the number of times commands are issued. When more than one host is connected to an array port, the number of queue commands increases because the host issues commands to each array separately.

Multiple hosts can be connected to a single port. However, the queue depth that can be handled by one port is limited, and performance drops if that limit is exceeded. To avoid performance drops, specify the queue depth so that the sum for all hosts does not exceed the port's limit.



NOTES: If the queue depth is increased, array traffic also increases, and host and switch traffic can increase. The formula for defining host queue depth depends on the operating system or HBA. When determining the host queue depth, consider the port limit. The formula for defining queue depth on the host side varies depending on the type of operating system or HBA. When determining the overall queue depth settings for hosts, consideration should be given to the port limit.

For iSCSI configurations, each operating and HBA configuration has an individual queue depth value unit and setting unit, as shown in [Table 6-7 on page 6-22](#).

Table 6-7: iSCSI queue depth configuration

Platform	Product	Queue Depth (Unit)	Queue Depth (Default)	Unit of Setting
Windows	Microsoft Initiator			
	Qlogic	Port	16	HBA
Linux	Software initiator			
	Qlogic	Port	16	HBA



NOTE: If the host operating system is either Microsoft Windows NT or Microsoft Windows 2000/2003 and is connected to a single array port, you must set the Queue Depth to a maximum of 16 commands per port for the QLogic HBA.

Note that the maximum queue depth for the SAS LU is 32.

Increasing queue depth and port sharing

Figure 6-21 on page 6-23 shows how to determine the queue depth when a port is shared. In this example, Host 1, 2, 3, and 4, are connected to a port with a 512 command limit. Specify the queue depth so that the queue depth for Hosts A, B, C, and D, does not exceed X.




Figure 6-21: Queue depth does not exceed port limit

Increasing queue depth through path switching

Figure 6-22 on page 6-24 shows how to determine queue depth when an alternative path is configured. Host 1 and 2 are assigned to the primary and secondary paths, respectively.

Commands are issued to a volume via the primary path on Host 1. In this configuration, commands to be issued via the primary path are moved to the secondary path because of path switching, and the queue depth for a port connected to a host on the secondary path is increased. You must specify the appropriate queue depth for each host so that the number does not exceed its limit after the path switching.




Figure 6-22: Queue depth increase from path switching

Queue depth allocation according to host job priority

Figure 6-23 on page 6-24 shows how to determine the queue depth when priority is given connected hosts. To increase the priority of the host job individually, increase the host queue depth. When the host queue depth is increased, the port cannot exceed its limit. If the array does not have a prioritized order, allocate the host queue depth.




Figure 6-23: Host job priority



NOTE: We recommend that you execute any ping command tests when there is no I/O between hosts and controllers.

LUN Manager procedures

This section describes LUN Manager operations for Fibre Channel and iSCSI.

Using Fibre Channel

To use Fibre Channel

1. Verify that you have the environments and requirements for LUN Manager (see [Preinstallation information for Storage Features on page 3-20](#)).

For the array:

2. Set up a fibre channel port (see [Fibre Channel operations using LUN Manager on page 6-29](#)).
3. Create a host group (see [Adding host groups on page 6-30](#)).
4. Set the World Wide Name (WWN).
5. Set the host connection mode.
6. Create a volume.
7. Set the volume mapping.
8. Set the fibre channel switch zoning.

For the host:

9. Set the host bus adapter (HBA).
10. Set the HBA driver parameters.
11. Set the queue depth (repeat if necessary).
12. Create the disk partitions (repeat if necessary).

Figure 6-24 details the flow of tasks involved with configuring LUN Manager using Fibre Channel.




Figure 6-24: Operations flow (Fibre Channel)

Using iSCSI

The procedure flow for iSCSI below. For more information, see the *Hitachi Unified Storage Hardware Installation and Configuration Guide MK-91DF8273*.

To configure iSCSI

1. Verify that you have the environments and requirements for LUN Manager (see [Preinstallation information for Storage Features on page 3-20](#)).

For the array:

2. Set up the iSCSI port (see [iSCSI operations using LUN Manager on page 6-39](#)).
3. Create a target (see [Adding and deleting targets on page 6-44](#)).

4. Set the iSCSI host name (see [Setting the iSCSI target security on page 6-43](#)).
5. Set the host connection mode. For more information, see the *Hitachi Unified Storage Hardware Installation and Configuration Guide MK-91DF8273*).
6. Set the CHAP security (see [CHAP users on page 6-53](#)).
7. Create a volume.
8. Set the volume mapping.
9. Set the network switch parameters. For more information, see the *Hitachi Unified Storage Hardware Installation and Configuration Guide MK-91DF8273*).

For the host:

10. Set the host bus adapter (HBA). For more information, see the *Hitachi Unified Storage Hardware Installation and Configuration Guide MK-91DF8273*).
11. Set the HBA driver parameters. For more information, see the *Hitachi Unified Storage Hardware Installation and Configuration Guide MK-91DF8273*).
12. Set the queue depth. For more information, see the *Hitachi Unified Storage Hardware Installation and Configuration Guide MK-91DF8273*).
13. Set the CHAP security for the host (see [CHAP users on page 6-53](#)).

14. Create the disk partitions. For more information, see the *Hitachi Unified Storage Hardware Installation and Configuration Guide MK-91DF8273*).




Figure 6-25: Operations flow (iSCSI)

Fibre Channel operations using LUN Manager

LUN Manager allows you to perform fibre channel operations. With LUN Manager enabled, you can

- Add, edit, and delete host groups
- Initialize host group 000
- Change nicknames
- Delete World Wide Names
- Copy settings to other ports

About Host Groups

A storage administrator uses LUN Manager to connect a port of a disk array to a host using a storage switch, and then sets a data input/output path between the host and the volume. This setting indicates which host may access a specific volume.

To set a data input/output path, the authorized hosts for the volume are required to be classified as a host group. Then the classified host group is set to the port. For example, if a Windows host and a Linux host are connected to port A, you must create host groups of volumes that can be accessed by other operating systems.

A host group option (host connection mode) may be set for each host group you create. Hosts connected to different ports cannot share the same host group. Even if the volume to be accessed is the same, separated host groups should be created for each port to which the hosts are connected.




Figure 6-26: Setting access paths between hosts and volumes for Fibre Channel

Adding host groups

To add host groups, you must enable the host group security, and create a host group for each port.

To understand the host group configuration environment, you need to become familiar with the Host Groups Setting Window as shown in [Figure 6-27](#).

The Host Groups Setting window consists of the Host Groups, Host Group Security, and WWNs tabbed pages.

- **Host Groups**
Enables you to create and edit groups, initialize the Host Group 000, and delete groups.
- **Host Group Security**
Enable you to validate the host group security for each port. When the host group security is invalidated, only the Host Group 000 (default target) can be used. When it is validated, host groups following the host group 001 can be created, and the WWN of hosts to be permitted to access each host group can be specified.
- **WWNS**
Displays WWNs of hosts detected when the hosts are connected and those entered when the host groups are created. In this tabbed page, you can supply a nickname to each port name.

Enabling and disabling host group security

By default, the host group security is disabled for each port.



NOTE: When changing the host group security to a port with online host groups, please stop all host access to the port and restart hosts after making the change.

To enable or disable host group security

1. Start Navigator 2 and log in. The Arrays window appears
2. Click the appropriate array.
3. Expand the **Groups** list, and click **Host Groups**. The **Host Groups** window appears (see [Figure 6-27](#)).




Figure 6-27: Host Groups window



NOTE: The number of ports displayed in the Host Groups and Host Group Security windows can vary. SMS systems may display only four ports.

4. Click the **Host Group Security** tab.
5. Select the port you want to configure and click **Change Host Group Security**.
6. Select the port whose security you are changing, and click **Change Host Group Security**.
7. In the **Enable Host Group Security** field, select the **Yes** checkbox to enable security, or clear the checkbox to disable security.
8. Follow the on-screen instructions.
 - After enabling host group security, **Detected Hosts** is displayed.
 - The WWN of the HBA connected to the selected port is displayed in the **Detected Hosts** field.



NOTE: When changing the drop-down list of the platform or middleware, the items in the detail setting list count automatically and the correct setting redisplay. While performing the automatic calculation and redisplay, the drop-down menu changes to inactive.

If you execute other operations in the status where the drop-down list is inactive (lighter color), the automatic calculation and redisplay of the detail setting list are not completed and the list may not be set correctly.

When you run other operations, ensure the pull-down menu changes are active.

Creating and editing host groups

If you click **Create Host Group** without selecting a port, you can apply the same setting for multiple ports.

To create and edit host groups

1. In the Host Groups tab, click **Create Host Group** or **Edit Host Group**.
[Figure 6-28](#) appears

WWNs Volumes Options

Check WWN or enter new port name form and click Add when assign WWN to host group.
Check selected WWN and push Remove when disassign WWN from host group.

☒ Select From List ☐ Enter WWNs Manually

Detected WWNs

Rows/Page: 25 Page 0 of 0

	Nickname	Port Name	Port
No Object			

Filter Filter Off

Selected WWNs

Rows/Page: 25 Page 0 of 0

	Nickname	Port Name	Port
No Object			

Filter Filter Off

Add > < Remove

Figure 6-28: Host Group Property window — WWNs tab

With the **WWNs** tab, you specify the WWNs of hosts permitted to access the host group for each host group. You can specify the WWNs of hosts in two ways:

- Select the WWNs from the **Detected WWNs** list.
- Enter the WWNs manually.

The WWN is not a copy target in the case of selecting two or more ports for the **Create to** (or **Edit to**) field used for setting the alternate path. The **WWNs** list assigned to the host group of the **Host Group No.** field associated with each port selected in the **Available Ports** list is displayed in the **Selected WWNs** list.

2. Specify the appropriate information.
 - **Host Group No.** — This number can be 1 through 127.
 - **Name:** — One name for each port, and the name cannot be more than 32 alphanumeric characters (excluding \, /, : , , , ; , * , ? , " , < , > , | and ').
3. Click the **WWN** tab and specify the appropriate host information.

- Click the **Options** tab. The Create Host Group options dialog box appears.

Create Host Group Help

Host Group Property

Enter the information for the host group to be created.

* Host Group No.: * Create to: **Available Ports**
 From 1 to 127 .

* Name:
 32 characters or less (alphanumeric characters, '!', '@', '#', '\$', '%', '&', '(', ')', '+', ',', '-', '.', ':', ';', '=', '<', '>', '^', '_', '{', '}', '~', ' ', '(', ')', '[', ']', ']' or ' ').

Options:

Platform:
 Middleware:

Available Ports

Port
☒ 0A
☐ 0B
☐ 0C
☐ 0D

WWNs Volumes **Options**

Select or check options for host group.
 When Platform or Middleware is set, Mode Setting follow automatically.

Platform: Middleware:

Mode Settings:

Common Setting:

Additional Setting:

HP-UX Mode:	☐ Yes
PSUE Read Reject Mode:	☐ Yes
Mode Parameters Changed Notification Mode:	☐ Yes
NACA Mode:	☐ Yes
Task Management Isolation Mode:	☐ Yes
Unique Reserve Mode 1:	☐ Yes
Port-ID Conversion Mode:	☐ Yes
Tru Cluster Mode:	☐ Yes
Product Serial Response Mode:	☐ Yes
Same Node Name Mode:	☐ Yes
CCHS Mode:	☐ Yes
NOP-In Suppress Mode:	☐ Yes
S-VOL Disable Advanced Mode:	☐ Yes
Discovery CHAP Mode:	☐ Yes
Unique Extended COPY Mode:	☐ Yes
Unique Write Same Mode:	☐ Yes
DP Depletion Detail Reply Mode:	☐ Yes
Unit Attention Change Mode:	☐ Yes
Unique Compare and Write Mode:	☐ Yes

* Required field

OK Cancel

Figure 6-30: Host Group Property window - Options tab

- From the **Platform** and **Middleware** pull-down lists, select the appropriate platform and middleware, and click **OK**. When you want to apply the changed contents to other ports, select the desired port in the Available Ports list. Two or more ports can be selected. The following items display when selecting or not selecting the checkbox of the Forced set to all selected ports:
 - Selecting the checkbox.** The current settings are replaced by the edited contents.
 - Not selecting the checkbox.** The current settings of the selected ports cannot be changed. An error occurs.

10. Click **OK**.
11. When two or more ports are selected and the host group already exists in the ports, at the time you select the Forced set to all selected ports checkbox, the following message appears.
12. Follow the on-screen instructions.



NOTE: When changing the drop-down list of the platform or middleware, the items in the detail setting list count automatically and the correct setting redisplay. While performing the automatic calculation and redisplay, the drop-down menu changes to inactive.

If you execute other operations in the status where the drop-down list is inactive (lighter color), the automatic calculation and redisplay of the detail setting list are not completed and the list may not be set correctly.

When you run other operations, ensure the pull-down menu changes are active.

Observe the following caveats regarding the drop-down list of the platform or middleware:

- When changing the drop-down list of the platform or middleware, the items on the detail setting list are counted automatically and the correct setting redisplay after the calculation process.
- If you run other windows when the status of the drop-down list is inactive, the automatic calculation and redisplay of the detail setting list are not completed and the list may not be set correctly.
- When you run other windows, ensure the pull-down menu status becomes active before executing it.

Initializing Host Group 000

When you reset Host Group 000 to its default, its WWNs and volume settings are deleted and the host group name is reset to G000.

To initialize Host Group 000

1. In the Hosts Groups window ([Figure 6-27 on page 6-31](#)), select the appropriate host group, and click **Initialize Host Group 000**.
2. Follow the on-screen instructions.
3. Specify the copy destination of the edited host group setting.
4. Select the port of the copy destination in Available Ports for editing and click **OK**.

Deleting host groups

Host group 000 cannot be deleted. When deleting all the WWNs and volumes in Host Group 000, initialize it (see [Initializing Host Group 000 on page 6-35](#)).

To delete host groups

1. In the Host Groups window ([Figure 6-27 on page 6-31](#)), select the appropriate host group and click **Delete Host Group**.
2. Follow the on-screen instructions.

Changing nicknames

To change nicknames

1. In the Host Groups window (Figure 6-27 on page 6-31), click the WWNs tab. The WWNS tab appears (see Figure 6-31).




Figure 6-31: Edit Host Group - WWNs tab, changing nickname

2. Select the appropriate WWN, and click **Change Nickname**.




Figure 6-32: Change Nickname Dialog Box

3. Specify the nickname (up to 32 alphanumeric characters) and click **OK**.
4. Follow the on-screen instructions.

Deleting World Wide Names

To delete World Wide Names

1. In the Host Groups window (Figure 6-27 on page 6-31), click the WWNs tab. Figure 6-31 on page 6-37 appears.
2. Select the appropriate WWN, and click **Delete WWN**.

3. Follow the on-screen instructions.

Copy settings to other ports

The host group setting can be copied to the other port for the alternate path setting, and so forth. To specify the copy destination, select **Available Ports** when creating host groups.

Settings required for copying

The settings for copying is as follows:

- Setting the created/edited host group
- Setting the assignment of the volume of the created/edited host group
- Setting the options of the volume of the created/edited host group

The setting created in the Create Host Group screen and the setting corrected in the Edit Host Group screen can be copied.

Copying during host group creation

To copy to the other port at the time of the host group creation

1. In the Host Groups tab, click **Create Host Group**. The Create Host Group screen appears.
2. Set the host group according to the procedure under [Adding host groups on page 6-30](#).
3. Specify the copy destination of the created host group setting.
4. Select the port of the copy destination in the Available Ports for creation.
5. The port concerned that created the host group is already selected for the Available Ports for creation. Add the port of the copy destination and select it.
6. To copy to all the ports, select the Port.
7. Click **OK**.

If the host group of the same host group number as the host group concerned is created in the copy destination port, this operation will end.

Copying when editing a host group

To copy to the other port at the time of the host group editing

1. In the Host Groups tab, click **Edit Host Group**. The Edit Host Group screen appears.
2. Set the host group according to the procedure for the section for Editing a Host Group on page [6-31](#).
3. Specify the copy destination of the edited host group setting.
4. Select the port of the copy destination in the Available Ports for editing.

5. The port concerned that edited the host group is already selected for the available ports for editing. Add the port of the copy destination and select it.
6. To copy to all the ports, select the port.
7. When you select the Forced set to all selected ports checkbox, the current settings are replaced by the edited contents.
8. Click **OK**.
9. Confirm the appeared message.
10. When executing it as is, click **Confirm**.

You will receive a warning message to verify your actions when:

- The host group of the same host group number as the host group concerned is not created in the copy destination port.
- The host group of the same host group number as the host group concerned is created in the copy destination port.

iSCSI operations using LUN Manager

LUN Manager allows you to perform various iSCSI operations from the iSCSI Targets setting window (see [Figure 6-33 on page 6-40](#)), which consists of the following tabs:

- iSCSI Targets

With this tab, you can create and edit targets, edit the authentication, initialize target 000, and delete targets.

- iSCSI Target Security

With this tab, you specify the validation of the iSCSI target security for each port. When the iSCSI target security is invalidated, only the Target 000 (default target) can be used. When it is validated, targets following the Target 001 can be created, and the iSCSI Names of hosts to be permitted to access each target can be specified.

- Hosts

This tab displays the iSCSI Names of hosts detected when the hosts are connected and those entered when the targets are created. In this tabbed page, you can give a nickname to each iSCSI Name.

- CHAP Users

With this tab, you register user names and secrets for the CHAP authentication to be used for authentication of initiators and assign the user names to targets.



Figure 6-33: iSCSI Targets window - iSCSI Targets tab

The following sections provide details on using LUN Manager to configure your iSCSI settings.

Creating an iSCSI target

To create a target for each port, you must create a target.

Using LUN Manager, you must connect a port of the disk array to a host using the switching-hub or connecting the host directly to the port, and then set a data input/output path between the host and the volume. This setting specifies which host can access which volume.

For example, when a Windows Host (initiator iSCSI Name A) and a Linux Host (initiator iSCSI Name B) are connected to Port A, you must create targets of volumes to be accessed from the Windows Host (initiator iSCSI Name A) and by the Linux Host (initiator iSCSI Name B) as shown in [Figure 6-5 on page 6-15](#).

Set a **Target** option (Host Connection Mode) to the newly created target to confirm the setting.

With the **Hosts** tab, you specify the iSCSI names of hosts to be permitted to access the target. For each target, you can specify the iSCSI names in two ways:

- Select the names from the **Detected Hosts** list.
- Enter the names manually.

The iSCSI name of the host is not a copy target in case you have selected two or more ports for either the Create to or Edit to field used for setting the alternate path. The iSCSI name assigned to the iSCSI target of the iSCSI Target No. field concerned with each port selected by the Available Ports field is displayed in the Selected Hosts list.

Using the iSCSI Target Tabs

In addition to the Hosts tab, the iSCSI Target Property window contains several tabs that enable you to customize the configuration of the iSCSI target to a finer degree.

The Volumes tab enables you to assign volumes to volume numbers (H-LUNs) that are recognized by hosts. [Figure 6-34](#) displays the iSCSI Target Properties - Volumes tab.

Create iSCSI Target Help

iSCSI Target Property

Enter the information for the iSCSI target to be created.

* iSCSI Target No.: From 1 to 254 .

* Alias: 32 characters or less (alphanumeric characters, '!', '@', '#', '\$', '%', '&', '*', '+', '=', '~', '^', '_, '(', ')', '[', ']', ' ' or space).

* iSCSI Name: ☒ Use default value for iSCSI Name ☐ Enter iSCSI Name Manually
 223 characters or less (alphanumeric characters, '!', ' ', '-' or '_').

Options:

Platform: not specify
 Middleware: not specify

Available Ports

Port
☒ 0E
☐ 0F
☐ 1E
☐ 1F

Hosts **Volumes** **Options**

To assign volumes to an iSCSI target, choose an H-LUN, then select one or more Available Volumes and click Add. An assigned volume will be seen from hosts as H-LUN.
 To remove volumes from Assigned Volumes, make selections from the list and click Remove.

LU Mappings:

H-LUNs

Rows/Page: 25 1 of 82

H-LUN
☒ 0000
☐ 0001
☐ 0002
☐ 0003
☐ 0004

Available Volumes

Rows/Page: 25 1 of 1

VOL	Capacity	RAID Group	DP Pool	RAID
☐ 0000	100.0GB	000	N/A	RAID5
☐ 1000	100.0GB	000	N/A	RAID5

Assigned Volumes

Rows/Page: 25 0 of 0

H-LUN	VOL	Capacity	RAID Group	DP Pool	RAID
No Object					

Add Remove

* Required field

OK Cancel

Figure 6-34: iSCSI Target Property window - Volumes tab

The iSCSI Target Property - Options tab enables you to select a platform and middleware that suit the environment of each host to be connected. You do not need to set the mode individually. [Figure 6-35](#) displays the iSCSI Target Property - Volumes tab.

Observe the following caveats regarding the drop-down list of the platform or middleware:

- When changing the drop-down list of the platform or middleware, the items on the detail setting list are counted automatically and the correct setting redisplay after the calculation process.
- If you run other windows when the status of the drop-down list is inactive, the automatic calculation and redisplay of the detail setting list are not completed and the list may not be set correctly.
- When you run other windows, ensure the pull-down menu status becomes active before executing it.

Create iSCSI Target Help

iSCSI Target Property

Enter the information for the iSCSI target to be created.

* iSCSI Target No.: From 1 to 254.

* Alias: 32 characters or less (alphanumeric characters, '!', '@', '#', '\$', '%', '&', '*', '^', '_', '{', '}', '~', ' ', '(', ')', '[', ']' or space).

* iSCSI Name: ☒ Use default value for iSCSI Name ☐ Enter iSCSI Name Manually
 223 characters or less (alphanumeric characters, '!', ' ', ' ' or ' ').

Options:
Platform:
Middleware:

* Create to:

Available Ports	
	Port
☒	0E
☐	0F
☐	1E
☐	1F

Hosts Volumes **Options**

Select or check options for iSCSI target.
When Platform or Middleware is set, Mode Setting appears automatically.

Platform: Middleware:

Mode Settings:

Common Setting:

Additional Setting:

HP-UX Mode:	☐ Yes
PSUE Read Reject Mode:	☐ Yes
Mode Parameters Changed Notification Mode:	☐ Yes
NACA Mode:	☐ Yes
Task Management Isolation Mode:	☐ Yes
Unique Reserve Mode 1:	☐ Yes
Port-ID Conversion Mode:	☐ Yes
Tru Cluster Mode:	☐ Yes
Product Serial Response Mode:	☐ Yes
Same Node Name Mode:	☐ Yes
CCHS Mode:	☐ Yes
NOP-In Suppress Mode:	☐ Yes
S-VOL Disable Advanced Mode:	☐ Yes
Discovery CHAP Mode:	☐ Yes
Unique Extended COPY Mode:	☐ Yes
Unique Write Same Mode:	☐ Yes
Report iSCSI Full Portal List Mode:	☐ Yes
DP Depletion Detail Reply Mode:	☐ Yes
Unit Attention Change Mode:	☐ Yes
Unique Compare and Write Mode:	☐ Yes

* Required field

OK Cancel

Figure 6-35: iSCSI Target Property window - Options tab



NOTE: When changing the drop-down list of the platform or middleware, the items in the detail setting list count automatically and the correct setting redisplay. While performing the automatic calculation and redisplay, the drop-down menu changes to inactive.

If you execute other operations in the status where the drop-down list is inactive (lighter color), the automatic calculation and redisplay of the detail setting list are not completed and the list may not be set correctly.

When you run other operations, ensure the pull-down menu changes are active.

Setting the iSCSI target security

The target security default setting is **disabled** for each port.

To enable or disable the target security for each port

1. Start Navigator 2 and log in. The Arrays window appears.
2. Click the appropriate array.
3. Expand the **Groups** list, and click **iSCSI Targets** to display the iSCSI Targets window as shown in [Figure 6-34](#).



Figure 6-36: iSCSI Targets Setting window - iSCSI Targets tab

4. Click the **iSCSI Target Security** tab, which displays the security settings for the data ports on your Hitachi Unified Storage system.

Yes = security is enabled for the data port.

No = security is disabled for the data port.




Figure 6-37: iSCSI Target Security tab

5. Click the port whose security setting you want to change.
6. Click **Change iSCSI Target Security**
7. Select (or deselect) the **Enable iSCSI Target Security** check box to enable (or disable) security, the click **OK**.
8. Read the confirmation message and click **Close**.



NOTE: If iSCSI target security is enabled, the iSCSI host name specified in your iSCSI initiator software must be added to the **Hosts** tab in Storage Navigator Modular 2.

1. From the iSCSI Targets screen, check the name of an iSCSI target and click **Edit Target**.
2. When the Edit iSCSI Target screen appears, go to the **Hosts** tab and select **Enter iSCSI Name Manually**.
3. When the next Edit iSCSI Target window appears, enter the iSCSI host name in the **iSCSI Host Name** field of the **Hosts** tab.
4. Click the **Add** button followed by the **OK** button.

Editing iSCSI target nicknames.

You can assign a nickname to each iSCSI target.

To edit a nickname to an iSCSI target

1. Start Navigator 2 and log in. The Arrays window appears.
2. Click the appropriate array.
3. Expand the **Groups** list, and click **iSCSI Targets** to display the iSCSI Targets window.
4. Click the **Hosts** tab, which displays an iSCSI target nickname, an indication of whether it has been assigned to any iSCSI targets, an associated port number and an associated iSCSI name.
5. Figure 6-38 displays the **Hosts** tab.

Nickname	Assigned to iSCSI Targets	Port	iSCSI Name
	No	0E	

Figure 6-38: Hosts tab

6. To edit a nickname, click on the nickname you want to change and click the **Change Nickname** button.
7. Type in a new nickname and click **OK**. Note the new nickname displayed in the **Hosts** tab.
8. Read the confirmation message and click **Close**.

Adding and deleting targets

The following section provides information for adding and deleting targets.

Adding targets

When you add targets and click **Create Target** without selecting a port, multiple ports are listed in the **Available Ports** list. Doing so allows you to use the same setting for multiple ports. By editing the targets after making the setting, you can omit the procedure for creating the target for each port.

To create targets for each port

1. In the **iSCSI Targets** tab, click **Create Target**. The iSCSI Target Property screen is displayed.

Figure 6-39: iSCSI Target Property window

2. Enter the iSCSI Target No., Alias, or iSCSI Name. [Table 6-8](#) describes these value types.

Table 6-8: iSCSI Target Number, Alias, and iSCSI Name

Value Type	Description	Value
iSCSI Target No.	The iSCSI bus address of the target, the system component that receives an iSCSI I/O command.	Range: An integer from 1 through 254.
Alias	An alternate, friendly, name for the iSCSI target. Notes • Spaces at the top or end are ignored. • The same name cannot be used in the same port.	Length: Less than or equal to 32 ASCII characters. Type: !, #, \$, %, &, ', +, -, ., =, @, ^, _ {, }, -, (,), [,], (space) Spaces at the top or end are ignored.

Value Type	Description	Value
iSCSI Name	The name of the iSCSI initiator or iSCSI target. iSCSI names are long and can be created with the following two naming types: - an iSCSI qualified name (iqn). - an extended unique identifier (eui). Notes When many iSCSI targets are created, entering the iqn type iSCSI Name with a maximum (223) characters, the host may be unable to recognize any iSCSI targets. In this case, type the iqn type iSCSI Name using the default iSCSI name that contains 47 characters.	Length: 223 or less characters. Type: Alphanumeric, these special characters allowed: - a period (.) - a hyphen (-) - a colon (:) Naming Type: - iqn: Consists of the following data components: - type identifier - domain acquisition date - domain name - character string assigned by person who acquired the domain Example: iqn.1994-04.jp.co.hitachi:rsd.d9b.t.00026.1e000 - eui: (64-bit identifier) Consists of the following data components: - type identifier - eui - ASCII coded hexadecimal eui-64 identifier. Example: eui.0123456789abcdef

Note that the **Hosts** tab displays only when iSCSI Target Security is enabled.

- If the iSCSI Target Security is enabled, set the host information in the Hosts tab. [Figure 6-40](#) displays an example of creating targets by selecting the **Enter iSCSI Name Manually** button.

The screenshot shows the 'Hosts' tab of a storage management application. At the top, there are tabs for 'Hosts', 'Volumes', and 'Options'. Below them, a message states: 'Check Select from Detected Hosts and select hosts, or check Enter iSCSI Name Manually and enter iSCSI Name. Then click Add to assign the hosts to the iSCSI target. When de-assign hosts, select from Selected Hosts and click Remove.' Two radio buttons are present: 'Select from Detected Hosts' (unselected) and 'Enter iSCSI Name Manually' (selected). Below the radio buttons, the 'iSCSI Name' field contains the text 'iqn.1994-04.jp.co.hitachi:rsd.d9b.t.00026.1e000'. A small tooltip below the field indicates '223 characters or less (alphanumeric characters, '.', '-' or ':')'. To the right of the main form is a 'Selected Hosts' panel. It has a table with columns 'Nickname', 'Port', and 'iSCSI Name'. The table is currently empty, showing 'No Object'. Below the table are 'Filter' and 'Filter Off' buttons. At the bottom of the main form, there is an 'Add' button with a plus icon. A legend at the bottom left indicates '* Required field'. At the bottom right, there are 'OK' and 'Cancel' buttons.

Figure 6-40: Setting Host Information in the Hosts tab

Using the Hosts tab, you can specify for each target the iSCSI Names of the hosts to be permitted to access the target. There are two ways to specify the iSCSI Names:

- You can select the names from the list of Detected Hosts as shown in [Figure 6-41](#), or
- You can enter the names manually.

For the initial configuration, write down the name and enter the name manually.

4. Click **Add**. The added host information is displayed in the **Selected Hosts** list.


The image shows the 'iSCSI Target Property' dialog box. At the top, it says 'Enter the information for the iSCSI target to be created.' Below this, there are several fields: 'iSCSI Target No.: 001', '* Alias: T001' (with a note about character restrictions), and '* iSCSI Name:'. The iSCSI Name field has two radio buttons: 'Use default value for iSCSI Name' and 'Enter iSCSI Name Manually'. The 'Enter iSCSI Name Manually' option is selected, and the text 'iqn.1994-04.jp.co.hitachiirsd.d9a.t00026.0e001' is entered. Below the iSCSI Name field is a note about character restrictions. To the right of the main form is an 'Available Ports' section with a table of ports: Port, 0E, 0F, 1E, 1F. The 'Port' column has checkboxes, and the '0E' row is selected. At the bottom right of the ports section is a checkbox for 'Forced set to all selected ports'. At the bottom of the main form is an 'Options' section with two fields: 'Platform: not specify' and 'Middleware: not specify'.

Figure 6-41: iSCSI Target Properties dialog box



NOTES: Up to 256 Hosts can be assigned for a port. The total of the number of Hosts that have been already assigned (Selected Hosts) and the number of Hosts that can be assigned (Selected Hosts) further is 256 for a Port. If the number of Hosts assigned to a port exceeds 256 and further input is impossible, delete a Host that is not assigned to a target.

In some cases, the Host is not listed in the Detected Hosts list, even though the port is connected to a host. When the Host to be assigned to a target is not listed in the Detected Hosts list, input and add it.

Not all targets may display when executing Discovery on the host and may depend on the HBA in use due to the restriction of the number of characters set for the iSCSI Name.

5. Click the **Volumes** tab.

6. Select an available Host Volume Number from the H-LUN list. The host uses this number to identify the volume it can connect to and click **Add**. The added volumes are displayed in the Selected Volumes list as shown in Figure 6-42.

The screenshot shows the 'LU Mapping' window with three tabs: 'Hosts', 'Volumes', and 'Options'. The 'Volumes' tab is active. It contains two main sections: 'H-LUNs' and 'Available Volumes'. The 'H-LUNs' section has a table with columns 'H-LUN' and 'Capacity'. The 'Available Volumes' section has a table with columns 'VOL', 'Capacity', 'RAID Group', 'DP Pool', and 'RAID'. The 'Assigned Volumes' section on the right is currently empty, showing 'No Object'. An 'Add' button is visible between the 'Available Volumes' and 'Assigned Volumes' sections. The 'Add' button is highlighted, and a tooltip shows 'Add' and 'Remove' options. The 'OK' and 'Cancel' buttons are at the bottom right.

LU Mapping:

To assign volumes to an iSCSI target, choose an H-LUN, then select one or more Available Volumes and click Add. An assigned volume will be seen from hosts as H-LUN.
To remove volumes from Assigned Volumes, make selections from the list and click Remove.

H-LUNs

H-LUN	Capacity
0000	
0001	
0002	
0003	
0004	
0005	

Available Volumes

VOL	Capacity	RAID Group	DP Pool	RAID
☒ 0000	100.0GB	000	N/A	RAID5
☐ 1000	100.0GB	000	N/A	RAID5

Assigned Volumes

H-LUN	VOL	Capacity	RAID Group	DP Pool
No Object				

* Required field

OK Cancel

Figure 6-42: Added contents to assigned volumes list

To remove an item from the list, select it and click **Remove**.

7. Click the **Options** tab.
8. From the Options tab, select **Platform** and **Middleware** from the pull-down lists.
 - **Platform** Options
Select either **HP-UX**, **Solaris**, **AIX**, **Linux**, **Windows**, **VMware** or **not specified** from the pull-down list.
 - **Middleware** Options
Select either **VCS**, **True Cluster** or **not specified** from the pull-down list.
9. Click **OK**. The confirmation message is displayed.
10. Click **Close**.

The new settings are displayed in the **iSCSI Targets** window.



NOTE: When changing the drop-down list of the platform or middleware, the items in the detail setting list count automatically and the correct setting redisplay. While performing the automatic calculation and redisplay, the drop-down menu changes to inactive.

If you execute other operations in the status where the drop-down list is inactive (lighter color), the automatic calculation and redisplay of the detail setting list are not completed and the list may not be set correctly.

When you run other operations, ensure the pull-down menu changes are active.

About iSCSI target numbers, aliases, and names

Consult [Table 6-9](#) when entering target numbers, aliases, or names.

Table 6-9: iSCSI target numbers, aliases, and names

Item	Requirements
iSCSI Target No.	Enter a numeral from 1 through 254.
Alias	Enter the alias of the target with less than or equal to 32 ASCII characters (alphabetic characters, numerals, and the following symbols) can be used: (!, #, \$, %, &, ', +, -, ., =, @, ^, _, {, }, -, (,), [,], (space). Spaces at the top are ignored. The same name cannot be used in the same port.
iSCSI Name	When entering an iSCSI Name manually, enter the name of the iSCSI Name with 223 or less alphanumeric characters. A period (.), hyphen (-), and colon (:), can be used. For the iSCSI name, both the iqn and eui types are supported. iqn (iSCSI qualified name): The iqn consists of a type identifier, "iqn", a date of domain acquisition, a domain name, and a character string given by a person who acquired the domain. Example: iqn.1994-04.jp.co.hitachi:rsd.d9b.t.00026.1a000 eui (64-bit extended unique identifier): The eui consists of a type identifier "eui" and an ASCII coded hexadecimal eui-64 identifier. Example: eui.0123456789abcdef

Deleting Targets



NOTE: Target 000 cannot be deleted. When deleting all the hosts and all the Volumes in Target 000, initialize Target 000 (see section Initializing Target 000).

To delete a target

1. Select the Target to be deleted and click **Delete Target**.
2. Click **OK**. The confirmation message appears.
3. Click **Confirm**. A deletion complete message appears.
4. Click **Close**.

The new settings are displayed in the **iSCSI Targets** window.

Editing target information

When editing targets, if you select multiple targets and click **Edit Target** multiple ports are listed in the Available Ports list. You can apply the same setting to the all of the selected targets at the same time.

To edit the target information

1. Select the Target requiring the target information and click **Edit Target**. The Edit iSCSI Target screen appears as shown in [Figure 6-43](#).

Figure 6-43: iSCSI Target Property window - Hosts tab

2. Type the Alias or iSCSI Name, as required.
3. Set the host information from the **Hosts** tab.
4. Select the **Volumes** tab.
5. Set the volumes information if necessary.
6. Select the **Options** tab.
7. Set the **Platform** and **Middleware** as required.

8. From the **Platform** and **Middleware** pull-down lists, select the appropriate platform and middleware, and click **OK**. When you want to apply the changed contents to other ports, select the desired port in the Available Ports list. Two or more ports can be selected. The following items display when selecting or not selecting the checkbox of the Forced set to all selected ports:
 - **Selecting the checkbox.** The current settings are replaced by the edited contents.
 - **Not selecting the checkbox.** The current settings of the selected ports cannot be changed. An error occurs.
9. Click **OK**.
10. When two or more ports are selected and the host group already exists in the ports, at the time you select the Forced set to all selected ports checkbox, a confirmation message appears.
11. When you select the Forced set to all selected ports checkbox, the current settings are replaced by the edited contents.
12. Click **OK**. The confirmation message is displayed.
13. Click **Close**.

The new settings are displayed in the **iSCSI Targets** window.

Editing authentication properties

To edit authentication properties

1. Select the Target requiring the target information and click **Edit Authentication**. The Edit Authentication screen is displayed as shown in Figure 6-44 on page 6-51.

Figure 6-44: Edit Authentication window

2. Select or enter the **Authentication Method**, **Enable Mutual Authentication**, or **For Mutual Authentication**.
 - **Authentication Method** options

- Select the **CHAP, None**, or **CHAP, None**.
 - CHAP Algorithm option
MD5 is always displayed.
 - **Enable Mutual Authentication** settings
Select (or deselect) the check box. If you select the check box, complete the parameters for **User Name** and **Secret**.
3. Click **OK**. The confirmation message appears.
 4. Click **Close**.
- The new settings appear in the **iSCSI Targets** window.

Initializing Target 000

You can reset target 000 to the default state by initializing it. If Target 000 is reset to the default state, hosts that belong to Target 000 and the settings of the volumes that belong to Target 000 are deleted. The Target options of Target 000 are reset to the default state and the target name is reset to T000.

To initialize Target 000

1. Select Target 000 to be initialized and click **Initialize Target 000**.
2. Click **OK**. The confirmation message appears.
3. Click **Confirm**. The initialization confirmation screen appears.
4. Click **Close**.

Changing a nickname

To change a nickname

1. From the iSCSI Targets window, click the **Hosts** tab as shown in [Figure 6-45 on page 6-52](#).




Figure 6-45: iSCSI Target window — Hosts tab

2. Select the Hosts information and click **Change Nickname**.
3. Type the new Nickname and click **OK**. The changed nickname confirmation screen appears.
4. Click **Close**.

CHAP users

CHAP is a security mechanism that one entity uses to verify the identity of another entity, without revealing a secret password that is shared by the two entities. In this way, CHAP prevents an unauthorized system from using an authorized system's iSCSI name to access storage.

User authentication information can be set to the target to authorize access for the target and to increase security.

The **User Name** and the **Secret** for the user authentication on the host side are first set to the port, and then assigned to the Target. The same **User Name** and **Secret** may be assigned to multiple targets within the same port.

The **User Name** and the **Secret** for the user authentication are set to each target.

Adding a CHAP user

To add a CHAP User

1. Select the **CHAP User** tab. The CHAP Users screen appears as shown in [Figure 6-43 on page 6-50](#).
2. Click **Create CHAP User**. The Create CHAP User window appears as shown in [Figure 6-46 on page 6-53](#).

Create CHAP User [Help]

CHAP User Property

Enter the information for the CHAP user to be created.

* User Name: * Create to:

256 characters or less (alphanumeric characters and symbols including: " ", "+", "@", "_", "=", "!", "/", "[", "]", "\\", or space).

* Secret:

12 to 32 characters (alphanumeric characters and symbols including: " ", "+", "@", "_", "=", "!", "/", "[", "]", "\\", or space).

* Retype Secret:

Same to Secret

* Required field

Port	Port
☒	0E
☒	0F
☒	1E
☒	1F

[OK] [Cancel]

Figure 6-46: Create CHAP User window

3. In the **Create CHAP User** screen, type the **User Name** and **Secret**, then re-type the Secret.
4. Select the port to be created from the **Available Ports** list.
5. Click **OK**. The created CHAP user message appears.
6. Click **Close**.

Changing the CHAP user

To change the CHAP User

1. Select the **CHAP User** tab.
2. Select a CHAP User to be changed from the CHAP User list and click **Edit CHAP User**. The Edit CHAP User window appears. [Figure 6-47 on page 6-54](#) shows the Edit CHAP User Window.

Edit CHAP User - OE : User001 [Help]

CHAP User Property

Enter the information for the CHAP user to be edited.

iSCSI Port: OE

User Name: User001
256 characters or less (alphanumeric characters and symbols including: ".,_-+=@%&*/'[]~", or space).

Change Secret: ☐ Yes

Secret: [Masked]
12 to 32 characters (alphanumeric characters and symbols including: ".,_-+=@%&*/'[]~", or space).

Retype Secret: [Masked]
Same to Secret

Assign Targets: ☐ Yes

Available Targets

iSCSI Target	iSCSI Name
☐ 000:T000	iqn.1994-04.jp.co.hitachi:rsd.d9a.t

[Add] [Remove]

Assigned Targets

iSCSI Target	iSCSI Name
No Object	

[Filter] [Filter Off]

[OK] [Cancel]

Figure 6-47: Edit CHAP User window

3. Type the **User Name** and **Secret**, then re-type the Secret as required.
4. Select the iSCSI Target from the **Available Targets** list and click **Add** as required. The selected target is displayed in the Assigned Targets list.
5. Click **OK**. The changed CHAP user message appears.
6. Click **Close**.

Deleting the CHAP user

To delete the CHAP User

1. Click the **CHAP User** tab.
2. Select the CHAP User to be deleted from the **CHAP User** list and click **Delete CHAP User**.
3. A screen appears, requesting a confirmation to delete the **CHAP User**, select the check box and click **Confirm**.
4. Click **OK**. The deleted CHAP user message appears.
5. Click **Close**.

Setting Copy to the Other Ports

The iSCSI target setting can be copied to the other port for the alternate path setting, etc. To specify the copy destination, select the **Available Ports** for creation at the time of operating the iSCSI target creation and iSCSI target edit.

Setting Information for Copying

The setting information for copying is shown below.

- Setting the created/edited iSCSI target
- Setting the assignment of the volume of the created/edited iSCSI target
- Setting the options of the volume of the created/edited iSCSI target

The setting created in the **Create iSCSI Target** screen and the setting corrected in the **Edit iSCSI Target** screen can be copied.

Copying when iSCSI Target Creation

To copy to the other port at the time of the iSCSI target creation

1. In the **iSCSI Targets** tab, click **Create Target**.

The **Create iSCSI Target** screen appears.

2. Set the iSCSI target according to the procedure for the section [Adding and deleting targets on page 6-44](#).
3. Specify the copy destination of the created iSCSI target setting.

Select the port of the copy destination in the **Available Ports** for creation.

The port concerned that created the iSCSI target is already selected for the **Available Ports** for creation. Therefore, add the port of the copy destination and select it.

To copy to all the ports, select the **Port**.

4. Click **OK**.

When the iSCSI target of the same target group number as the iSCSI target concerned is created in the copy destination port, this operation will be terminate abnormally.

Copying when iSCSI Target Editing

To copy to the other port at the time of the iSCSI target editing

1. In the **iSCSI Targets** tab, click **Edit Target**.

The **Edit iSCSI Target** screen appears.

2. Set the iSCSI target according to the procedure for the section [Editing target information on page 6-50](#).

3. Specify the copy destination of the edited iSCSI target setting.
Select the port of the copy destination in the **Available Ports** for creation.
The port concerned that created the iSCSI target is already selected for the **Available Ports** for creation. Therefore, add the port of the copy destination and select it.
4. To copy to all the ports, select the **Port**.
5. Click **OK**.
6. Confirm the appeared message.
 - When executing it as is, click **Confirm**.
 - When the iSCSI target of the same iSCSI target number as the iSCSI target concerned is not created in the copy destination port, the following message displays.

Edit Host Group - Port 0A:001



Figure 6-48: Instance: target not created in copy destination port

- When the iSCSI target of the same iSCSI target number as the iSCSI target concerned is created in the copy destination port, the following message displays.

Edit Host Group - Port 0A:000




Figure 6-49: Instance: target created in copy destination port

Capacity

This chapter provides detail on managing, provisioning, and sectioning capacity on your storage system into partitions in the storage system cache, using both Cache Partition Manager and Cache Residency Manager.

The topics covered in this chapter are:

- ❑ [Capacity overview](#)
- ❑ [Cache Partition Manager overview](#)
- ❑ [Partition capacity](#)
- ❑ [Supported partition capacities](#)
- ❑ [Cache Partition Manager procedures](#)
- ❑ [Cache Residency Manager overview](#)
- ❑ [Supported Cache Residency capacities](#)
- ❑ [Cache Residency Manager procedures](#)

Capacity overview

The cache memory on a disk array is a gateway for receiving/sending data from/to a host. In the disk array, the cache memory is used being divided into a system control area and a user data area. For sending and receiving data, the user data area is used.

Cache Partition Manager overview

Cache Partition Manager is a priced optional feature of the disk array that enables the user data area in the disk array to be used being divided more finely. Each of the divided portions of the cache memory is called a partition. A volume defined in the disk array is used being assigned to the partition.

A user can specify a size of the partition and a segment size (size of a unit of data management) of a partition can be changed also. Therefore you can optimize the data reception/sending from/to a host by assigning the most suitable partition to a volume according to a kind of data to be received from a host.



NOTE: Before using Cache Partition Manager, be sure to refer to the section.

Cache Partition Manager features

Cache Partition Manager has the following features:

- **Cache division function** - This function divides the cache into two or more partitions. The cache capacity to be assigned to the partition to be created can be specified. Besides, a partition to be used for each volume can be selected.
- **Segment size change function** - This function can change a segment size for each partition. This function can optimize the segment sizes to be used according to the application and the use and can enhance the effective use and performance of the cache.
- **Specifying a pair cache partition** - When using Cache Partition Manager, you can specify a partition to be changed. (When the Load Balancing is **Disable**, it is not necessary to specify.)

Cache Partition Manager benefits

- **Increased manageability of storage content** - Enables you to divide up storage units into multiple partitions, enabling ease of use, manageability and addressability.
- **Convenient application mapping** - Allows you to partition the storage cache to map to various features.
- **I/O interruption protection** - Makes the volume less affected by the condition of I/O loads on the other volumes.
- **Increased performance** - Optimizes the segment sizes to be used according to the application and the use and can enhance the effective

use and performance of the cache. Enables applications to have access to applications and data in the cache. By doing this, your retrieval time of content is less, improving performance. Ordinarily applications are swapped in and out of cache. As soon as we need the information

- **Volume independence** - Cache division enhances the independence between the volumes that use each cache partition and can make the volume less affected by the condition of I/O loads on the other volumes.

Cache Partition Manager feature specifications

Table 7-2 details Cache Partition Manager feature specifications.

Table 7-1: Cache Partition Manager feature specifications

Item	Description
Supported cache memory	HUS110: 4 GB/controller HUS130: 8 GB/controller HUS150: 8, 16 GB/controller
Number of partitions	HUS110 (4 GB/controller): 2 to 7 HUS130 (8 GB/controller): 2 to 11 HUS150 (8 GB/controller): 2 to 151 HUS150 (16 GB/controller): 2 to 27
Partition capacity	The partition capacity depends on the array model and the capacity of the cache memory installed in the controller. For more information, see Cache Partition Manager procedures on page 7-14 .
Memory segment size	• Master partition: Fixed 16 KB • Sub partition: 4, 8, 16, 64, 256, or 512 KB When changing the segment size, make sure you refer to Specifying partition capacity on page 7-10 .
Pair cache partition	The default setting is "Auto" and you can specify the partition. It is recommended that you use Load Balancing in the "Auto" mode. For more information, see Segment and stripe size restrictions on page 7-9 .
Partition mirroring	Always On (it is always mirrored).

Confirming Environments

Before using

Cache Partition Manager task flow

The following steps detail the task flow of the Cache Partition Manager configuration process:

1. You determine you need to create partitions in your storage cache to map to your applications for fasted access.
2. Map out a system of partitions on paper that you will apply to configuration in HSNM2.

3. Install a license key for Cache Partition Manager.
4. Launch HSNM2.
5. Launch Cache Partition Manager.
6. Create a series of partitions that you will map to applications.
7. Create a system of pairing that you apply to the partitions.

Operation task flow

The following steps detail the flow of tasks for the operating procedure to use Cache Partition Manager.

1. Install Cache Partition Manager.
2. Change the partition size of the master partition. the change of the settings of a partition (addition, deletion, partition size change, and segment size change) and the change of a partition to which the volume belongs are validated after the storage system starts.
3. Add a sub-partition.
4. Change the partition to the volume to which it belongs. You perform this task when the existing volume is used by the new partition.
5. Restart the disk array. Validate the newly created partition and change the partition to map to the volume to which it belong after the restart.
6. Create a volume when newly adding a volume belonging to the new partition.
7. Begin operating with the cache partition active.

To create a volume using the additionally created partition, you need to determine the partition beforehand. Add the volume after the disk array restarts and the partition is validated. For a change in a partition or a volume after the operation starts, see the related section.

Stopping Cache Partition Manager

The storage system must restart before stopping the use of Cache Partition Manager. The same cautions indicated in the previous section about starting the system apply to restarting.

Note that performing any setting changes of the partition (addition, deletion, partition size change, and segment size change) and the change of a partition to which the volume belongs are validated after the storage system restarts.

The following steps detail the flow of tasks for stopping Cache Partition Manager from running.

1. Change the partitions to the ones which all the volumes belong to the master partitions.
2. Delete the sub-partitions.
3. Return the partition sizes of the master partitions to the default size.

4. Restart the disk array. This event has the result of deleting and validating the change of the partitions sizes after the restart.
5. Uninstall the Cache Partition Manager.

Pair cache partition

The pair cache partition is a partition to be changed in the Load Balancing mode. By configuring controllers in the way detailed in [Figure 7-1](#), partitions can be used continuously in the way that partition numbers 0 and 1 are for the SAS drives and partition numbers 2 and 3 are for the SAS7.2K drives even if Load Balancing occurs.

Also, a case exists where an operation of I/O to and from a volume consisting of the SAS drives is performed in partition numbers 0 and 1 and an operation of I/O to and from a volume consisting of the SAS7.2K drives is performed in the partition numbers 2 and 3 as shown in [Table 7-2](#). The settings shown in [Table 7-2](#) makes it possible to specify the partition to be used by each volume expressly when a controller that controls the volume is changed due to Load Balancing.

Table 7-2: Pair Cache Partition Policy Example

Vol Number	Drive Type	Belonging to Partition	Pair Cache Partition
0	SAS	0 (Ownership is controller 0)	1 (Ownership is controller 1)
1	SAS	1 (Ownership is controller 1)	0 (Ownership is controller 0)
2	SAS7.2K	2 (Ownership is controller 0)	3 (Ownership is controller 1)
3	SAS7.2K	3 (Ownership is controller1)	2 (Ownership is controller 0)

By creating the settings shown in [Table 7-3](#), partitions can be used continuously in the way that partition numbers 0 and 1 are for the SAS drives and partition numbers 2 and 3 for the SAS7.2K drives even if Load Balancing occurs.




Figure 7-1: Cache Partition Manager Task Flow

Partition capacity

The partition capacity depends on the following entities.

- **User data area** - The user data area depends on the array model, controller configuration (dual or single), and the controller cache memory. You cannot create a partition that is larger than the user data area.
- **Default partition size** - The tables in the partitioning sections show partition sizes in MB for Cache Partition Manager. When you stop using Cache Partition Manager, you must set the partition size to the default size. The default partition size is equal to one half of the user data area for dual controller configurations, and the whole user data area for single controller configurations.
- **Partitions size for small segments** - This applies to partitions using 4 KB or 8 KB segments, and the value depends on the array model. Sizes of partitions using all 4 KB or 8 KB segments must meet specific criteria for maximum partitions size of small segments.

The following formulas should be observed.

$[(\text{The size of partitions using all 4 KB segments in MB}) + (\text{The size of partitions using all 8 KB segments is shown in MB}/3)]$ has to be less or equal to maximum partition size of small segments (in MB) from the table.

If you are using Copy-on-Write SnapShot, True Copy Extended Distance (TCE), or Dynamic Provisioning, the supported capacity of the partition that can be created is changed because a portion of the user data area is needed to manage the internal resources.

Supported partition capacities

The supported partition capacity is determined depending on the user data area of the cache memory and a specified segment size and the supported partition capacity (when the hardware revision is 0100). All units are in Megabytes (MB). [Table 7-3](#) describes the supported partition capacity tables for instances of a Dual Controller Configuration and Dynamic Provisioning being disabled.

Table 7-3: Supported partition capacity (dual controller configuration and Dynamic Provisioning and Dynamic Tiering are disabled)

Array Model	Cache	User Data Area	Default Partition Size	Default Minimum Size	Default Maximum Size	Partition Capacity for Small Segment
HUS 110	4 GB/CTL	1,420	710	200	1,220	1,020
HUS 130	8 GB/CTL	4,660	2,330	400	4,260	3,860
	16 GB/CTL	11,280	5,640		10,880	4,990
HUS 150	8 GB/CTL	4,540	2,270		4,140	3,740
	16 GB/CTL	11,160	5,580		10,760	4,870

When Dynamic Provisioning or Dynamic Tiering is used, the supported capacity of the partition that can be created is changed because a part of the user data area is used to manage the internal resources. T shows the supported capacity in the case where Dynamic Provisioning is used. [Table 7-4](#) displays the supported capacity in a case where Dynamic Tiering is used.

Table 7-4: Supported partition capacity dual controller configuration and Dynamic Provisioning (Regular Capacity) are enabled

Array Model	Cache	User Data Area	Default Partition Size	Default Minimum Size	Default Maximum Size	Partition Capacity for Small Segment
HUS 110	4 GB/CTL	1,000	500	200	800	600
HUS 130	8 GB/CTL	4,020	2,010	400	3,620	3,220
	16 GB/CTL	10,640	5,320		10,240	4,990
HUS 150	8 GB/CTL	2,900	1,450		2,500	2,100
	16 GB/CTL	9,520	4,760		9,120	4,870

Table 7-5: Supported Partition Capacity Dual Controller Configuration and Dynamic Provisioning (Maximum Capacity) is

Enabled						
Array Model	Cache	User Data Area	Default Partition Size	Default Minimum Size	Default Maximum Size	Partition Capacity for Small Segment
HUS 110	4 GB/CTL	-	-	-	-	-
HUS 130	8 GB/CTL	3,000	1,500	400	2,600	2,200
	16 GB/CTL	9,620	4,810		9,220	4,990
HUS 150	8 GB/CTL	-	-		-	-
	16 GB/CTL	7,860	3,930		7,460	4,870

Table 7-6 details capacity values for an instance of a Dual Controller configuration where Dynamic Provisioning is enabled.

Table 7-6: Supported Partition Capacity Dual Controller Configuration and Dynamic Provisioning (Maximum Capacity) and Dynamic Tiering are enabled

Array Model	Cache	User Data Area	Default Partition Size	Default Minimum Size	Default Maximum Size	Partition Capacity for Small Segment
HUS 110	4 GB/CTL	960	480	200	760	560
HUS 130	8 GB/CTL	3,820	1,910	400	3,420	3,020
	16 GB/CTL	10,440	5,220		10,040	4,990
HUS 150	8 GB/CTL	2,700	1,350		2,300	1,900
	16 GB/CTL	9,320	4,660		8,920	4,870

Table 7-7 displays the supported capacity in a case where Dynamic Tiering is used.

Table 7-7: Supported Partition Capacity Dual Controller Configuration and Dynamic Tiering (Maximum Capacity) is enabled

Array Model	Cache	User Data Area	Default Partition Size	Default Minimum Size	Default Maximum Size	Partition Capacity for Small Segment
HUS 110	4 GB/CTL	-	-	-	-	-
HUS 130	8 GB/CTL	2,800	1,400	400	2,400	2,000
	16 GB/CTL	9,420	4,710		9,020	4,990
HUS 150	8 GB/CTL	-	-		-	-
	16 GB/CTL	7,660	3,830		7,260	4,870

Table 7-8 details supported partition capacity for a single controller configuration.

Table 7-8: Supported partition capacity (single controller configuration)

Array Model	Cache	User Data Area	Default Partition Size	Default Minimum Size	Default Maximum Size	Partition Capacity for Small Segment
HUS 110	4 GB/CTL	1,430	1,430	400	1,430	1,020

[Table 7-9](#) details segment and stripe size combinations.

Table 7-9: Segment And stripe size combinations

Segment	64 KB Stripe	256 KB Stripe	512 KB Stripe
4 KB	Yes	No	No
8 KB	Yes	Yes	No
16 KB	Yes	Yes (Default)	Yes
64 KB	Yes	Yes	Yes
256 KB	No	Yes	Yes
512 KB	No	No	Yes

The sum capacities of all the partitions cannot exceed the capacity of the user data area. The maximum partition capacity above is a value that can be calculated when the capacity of the other partition is established as the minimum in the case of a configuration with only the master partitions. You can calculate the residual capacity by using Navigator 2. Also, sizes of partitions using all 4 Kbyte and 8 Kbyte segments must be within the limits of the relational values shown in the next section.

Segment and stripe size restrictions

A volume stripe size depends on the segment size of the partition, as shown in [Table 7-10 on page 7-9](#). The default stripe size is 256 KB. [Table 7-10](#) details Cache Partition Manager restrictions.

Table 7-10: Segment and stripe size restrictions

Item	Description
Modifying settings	If you delete or add a partition, or change a partition or segment size, you must restart the array.
Pair cache partition	The segment size of a volume partition must be the same as the specified partition. When a cache partition is changed to a pair cache partition, the other partition cannot be specified as a change destination.
Changing single or dual configurations	The configuration cannot be changed when Cache Partition Manager is enabled.
Concurrent use of ShadowImage	When using ShadowImage, see Using ShadowImage, Dynamic Provisioning, or TCE on page 7-11 .

Table 7-10: Segment and stripe size restrictions

Item	Description
Concurrent use of Dynamic Provisioning/Dynamic Tiering	When Dynamic Provisioning or Dynamic Tiering is enabled, the partition status is initialized. When using Dynamic Provisioning or Dynamic Tiering, see Using ShadowImage, Dynamic Provisioning, or TCE on page 7-11 .
Concurrent use of a unified volume	All the default partitions of the volume must be the same partition.
Volume Expansion	You cannot expand volumes while making changes with the Cache Partition Manager.
Concurrent use of RAID group Expansion	- You cannot change the Cache Partition Manager configuration for volumes belonging to a RAID group that is being expanded. - You cannot expand RAID groups while making changes with the Cache Partition Manager.
Concurrent use of Cache Residency Manager	Only the master partition can be used together. A segment size of the partition to which a Cache Residency volume belongs to, cannot be changed.
Concurrent use of Volume Migration	A volume that belongs to a partition cannot carry over. When the migration is completed, the volume belonging to a partition is changed to destination partition.
Copy of partition information by Navigator 2	Not available. Cache partition information cannot be copied.
Load Balancing	Load balancing is not available for volumes where there is no cache partition with the same segment size available on the destination controller.
DP-VOLs	The DP-VOLs can be set as a partition the same as the normal volume. The DP pool cannot be set as a partition.



NOTE: You can only make changes when the cache is empty. Restart the array after the cache is empty.

Specifying partition capacity

When the number of RAID group drives (to which volumes belong to) increases, the use capacity of the Cache also increases. When a volume exceeds 17 (15D+2P or more) of the number of disk drives that configure the RAID group, using a partition with the capacity of the minimum partition capacity +100 MB or more is recommended.

Using a large segment

When a large segment is used, performance can deteriorate if you do not have enough partition capacity. The recommended partition capacity when changing the segment size appears in [Table 7-11](#).

Table 7-11: Partition capacity when changing segment size

Segment Size	Partition Capacity	
	HUS 110 / 130	HUS 150
64 KB	More than 300 MB	More than 600 MB
256 KB	More than 500 MB	More than 1,000 MB
512 KB	More than 1,000 MB	More than 2,000 MB

Using load balancing

The volume partition can be automatically moved to a pair partition according to the array CPU load condition of the CPU. If you do not want to move the volume partition, invalidate the load balance.

Using ShadowImage, Dynamic Provisioning, or TCE

The recommended segment size of the ShadowImage S-VOL, Dynamic Provisioning, TCE, or Volume Migration is 16 KB. When a different segment size is used, the performance and copy pace of the P-VOL may deteriorate.

You must satisfy one of the following conditions when using these features with Cache Partition Manager to pair the volumes:

- The P-VOL and S-VOL (V-VOL in the case of Dynamic Provisioning) belong to the master partition (partition 0 or 1).
- The volume partitions that are used as the P-VOL and S-VOL are controlled by the same controller.

You can check the information on the partitions, to which each volume belongs, and the controllers that control the partitions in the setup window of Cache Partition Manager. The detail is explained in the Chapter 4. For the pair creation procedures, and so forth, please refer to the Hitachi ShadowImage In-system Replication User's Guide or Hitachi Dynamic Provisioning User's Guide.

The P-VOL and S-VOL/V-VOL partitions that you want to specify as volumes must be controlled by the same controller. See page 4 17 for more information.

After creating the pair, monitor the partitions for each volume to ensure they are controlled by the same controller.

Installing Dynamic Provisioning/Dynamic Tiering when Cache

Partition Manager is Used

Dynamic Provisioning or Dynamic Tiering uses a part of the cache area to manage internal resources. Because of this, the cache capacity that Cache Partition Manager can use becomes smaller than the usual one.

Make sure that the cache partition information is initialized as shown below when Dynamic Provisioning or Dynamic Tiering is installed in the status where Cache Partition Manager is already in use.

- All the volumes are moved to the master partitions on the side of the default owner controller.
- All the sub-partitions are deleted and the size of each master partition is reduced to a half of the user data area after installing Dynamic Provisioning or Dynamic Tiering.

An example of the case where Cache Partition Manager is used is shown in [Figure 7-2](#).




Figure 7-2: Standard case where Cache Partition Manager is used

An example of the case where Dynamic Provisioning or Dynamic Tiering is installed in the context of obtaining the status that Cache Partition Manager is used is shown in [Figure 7-3](#).




Figure 7-3: Case where Dynamic Provisioning or Dynamic Tiering is installed for use with Cache Partition Manager

Adding or reducing cache memory

You can add or reduce the cache memory used by Cache Partition Manager, unless the following conditions apply.

- A sub-partition exists or is reserved.
- For dual controllers, the master partitions 0 and 1 sizes are different, or the partition size reserved for the change is different.

Cache Partition Manager procedures

The following sections describe Cache Partition Manager settings.

If a cache partition is added, deleted, or modified during power down, power down can fail. If this happens, power down again and verify that no RAID group in the Power Saving Status of Normal (Command Monitoring) exists. Then, you can add, delete, or modify the Cache Partition.

When you set, delete or change Cache Partition Manager settings when the storage system is used on other remote side of TrueCopy or TCE, the following activity results when you restart the system:

- Both paths of TrueCopy or TCE are blocked. In an instance of a blocked path, the system generates a trap to the SNMP Agent Support function. The path of TrueCopy or TCE is automatically recovered from the block after the system restarts.
- When the pair status of TrueCopy or TCE is in either a Paired or Synchronizing state, it changes to the Failure state.

Confirming Environments

Confirm the following before configuring your storage system using Cache Partition Manager.

When the Power Saving instruction of the non I/O link is executed with the priced option, Power Saving or Power Saving Plus are used together. If a Cache Partition is added, deleted, or changed while the Power Saving status is Normal (Command Monitoring), the status is changed to "Normal (Spin down Failure: PS OFF/ON)" by the array reboot which works at the time of the setting change and then the spin-down may fail.

When the spin-down fails, run a spin-down session again. Before adding, deleting, or changing the Cache Partition, check that the spin-down instruction has not been issued or there is no RAID group where the Power Saving status is Normal (Command Monitoring) by the Power Saving instruction of the non I/O link.

Initial settings

To configure initial settings

1. Verify that you have the environments and requirements for Cache Partition Manager (see [Preinstallation information for Storage Features on page 3-20](#)).
2. Change the partition size of the master partition (*Note 1*).
3. Add a sub partition (*Note 1*).
4. Change the partition the volume belongs to (*Note 1*).
5. Restart the array (*Note 1*).
6. Create a volume (*Note 3*).

7. Operate the cache partition.



NOTE: 1. When you modify partition settings, the change is validated after the array is restarted.



NOTE: 2. You only have to restart the array once to validate multiple partition setting modifications.



NOTE: 3. To create a volume with the partition you created, determine the partition beforehand. Then, add the volume after the array is restarted and the partition is validated.

Stopping Cache Partition Manager

The array must be restarted before you stop using Cache Partition Manager.

To stop Cache Partition Manager

1. In the master partition, change volume partitions.
2. Delete sub partitions.
3. Return the master partition size (#0 and #1) to their default size.
4. Restart the array.
5. Disable or remove Cache Partition Manager.

Working with cache partitions

Cache Partition Manager helps you segregate the workloads within an array. Using Cache Partition Manager allows you to configure the following parameters in the system memory cache:

- Selectable segment size — Allows the customization of the cache segment size for a user application
- Partitioning of cache memory — Allows the separation of workloads by dividing cache into individually managed, multiple partitions. A partition can then be customized to best match the I/O characteristics of the assigned volumes.
- Selectable stripe size — Helps increase performance by customizing the disk access size.



NOTE: If you are using the Power Savings feature and make any changes to the cache partition during a spin-down of the disks, the spin-down process may fail. In this case, re-execute the spin-down.

We recommend that you verify that the array is not in spin-down mode and that no RAID group is in Power Savings **Normal** status before making any changes to a cache partition.

After making changes to cache partitions, you must restart the array.

Adding cache partitions

To add cache partitions:

1. Start Navigator 2 and log in. The Arrays window appears
2. Click the appropriate array.
3. In the Navigation bar, click **Performance**, then click **Cache Partition**. [Figure 7-4](#) appears.




Figure 7-4: Cache Partition dialog box

4. Click **Set**. The Cache Partition dialog box appears.
5. Select cache partition **00** and click **Add Partition**. The Add Cache Partition Property window displays as shown in [Figure 7-5](#).




Figure 7-5: Add Cache Partition Property dialog box

6. Specify the following for partition 02:
 - Select **0** or **1** from the **CTL** drop-down menu.
 - Double-click the **Size** field and specify the size. The actual size is 10 times the specified number.
 - Select the segment size from the **Segment Size** drop-down menu.See [Cache Partition Manager procedures on page 7-14](#) for more information about supported partition sizes.
7. Click **OK** and follow the on-screen instructions.

Deleting cache partitions

Before deleting a cache partition, move the volume that has been assigned to it, to another partition.

To delete cache partitions

1. Start Navigator 2 and log in. The Arrays window appears
2. Click the appropriate array.
3. In the Navigation bar, click **Performance**, then click **Cache Partition**. [Figure 7-4 on page 7-16](#) appears.
4. Click **Set**. The Cache Partition dialog box appears, as shown in [Figure 5 on page 7-16](#).
5. Select the cache partition number that you are deleting, and click **Delete** as shown in [Figure 7-6](#).




Figure 7-6: Cache Partitions window - deleting a cache partition

6. Click **OK** and follow the on-screen instructions. Restarting the storage system takes approximately seven to 25 minutes.

Assigning cache partitions

If you do not assign a volume to a cache partition, it is assigned to the master partition. Also, note that the controllers for the volume and pair cache partitions must be different.

To assign cache partitions

1. Start Navigator 2 and log in.
2. Select the appropriate array.

- Click **Show & Configure Array**. The Show and Set Reservation window displays as shown in Figure 7-7.

Show and Set Reservation




Figure 7-7: Show and Set Reservation window

-



NOTE: The maximum number of persistent reservations does not vary based on model. However, two reservation maximum values exist, one for an individual volume, the other for an entire array. One volume can have up to 128 reservations while an array can have up to 8,194. One volume can receive up to 128 paths while one array can receive up to 8,194 paths.

- Under Arrays, click **Groups**.
- Click the **Volumes** tab. Figure 7-8 appears.




Figure 7-8: Volumes tab

7. Select a volume from the volume list, and click **Edit Cache Partition**. The Edit Cache Partition window displays as shown in [Figure 7-9](#).

Edit Cache Partition -00

Figure 7-9: Edit Cache Partition Window

8. Select a partition number from the **Cache Partition** drop-down menu, and click **OK**.
9. Follow the on-screen instructions. Restarting the storage system takes approximately seven to 25 minutes.



NOTE: The rebooting process will execute after you change the settings.

Setting a pair cache partition

This section describes how to configure a pair cache partition.

We recommend you observe the following when setting a pair cache partition:

- Use the default “Auto” mode.
- Set **Load Balancing** to **Disable** (use **Enable** if you want the partition to change with Load Balancing)



NOTE: The owner controller must be different for the partition where the volume is located and the partition pair cache is located.

To set a pair cache partition

1. Start Navigator 2 and log in.
2. Select the appropriate array.
3. Click **Show & Configure Array**.
4. Under Arrays, click **Groups**.
5. Click the **Volumes** tab. (See [Figure 7-8 on page 7-18](#))
6. Select a volume from the volume list and click **Edit Cache Partition**.

7. Select a partition number from the **Pair Cache Partition** drop-down list and click **OK**.
8. Click **Close** after successfully creating the pair cache partition.

Changing cache partitions

Before you change a cache partition, please note the following:

- You can only change the size of a cache sub-partition
- You must reboot the array for the changes to take affect

To change cache partitions

1. Start Navigator 2 and log in. The Arrays window appears
2. Click the appropriate array.
3. In the Navigation bar, click **Performance**, then click **Cache Partition**. [Figure 7-4 on page 7-16](#) appears.
4. Click **Set**. The Cache Partition dialog box appears, as shown in [Figure 5 on page 7-16](#).
5. Select a cache partition number that you want to edit and click Edit Partition as shown in [Figure 7-9](#).




Figure 7-10: Editing a cache partition

- To change capacity, double-click the **Size (x10MB)** field and make the desired change as shown in [Figure 7-9](#).

Edit Cache Partition -02

Figure 7-11: Edit Cache Partition Property window with segment size selection

- To change the segment size, select **segment size** from the drop-down menu to the left of **Segment Size**.
- Follow the on-screen instructions.

Changing cache partitions owner controller

The controller that processes the I/O of a volume is referred to as the owner controller.

To change cache partitions owner controllers:

- Start Navigator 2 and log in. The Arrays window appears
- Click the appropriate array.
- In the Navigation bar, click **Performance**, then click **Cache Partition**. [Figure 7-4 on page 7-16](#) appears.
- Click **Set**. The Cache Partition dialog box appears, as shown in [Figure 5 on page 7-16](#).
- Select a partition number for which you want to change the owner controller and click **Edit Partition**. The Edit Cache Partition screen displays.

6. Select the Cache Partition number and the controller (CTL) number (0 or 1) from the drop-down menu and click **OK** as shown in [Figure 7-12](#).

Edit Cache Partition -02

Help

Edit Cache Partition Property

Enter the information to edit the cache partition.

Cache partition: 02

Available Space:

Controller 0: 400MB

Controller 1: 400MB

Controller:

☐ Controller 0

☒ Controller 1

* Size: 200 MB

From 200 or 400 to max(depend on Available Space), 10MB range

Segment Size: 8KB

* Required field

OK Cancel

Figure 7-12: Edit Cache Partition Property window with new cache partition owner controller selected

7. Follow the on-screen instructions.
8. The **Automatic Pair Cache Partition Confirmation** message box displays.

Depending on the type of change you make, the setting of the pair cache partition may be switched to Auto. Verify this by checking the setting after restarting the storage system.

Click **OK** to continue. The **Restart Array** message is displayed. You must restart the storage system to validate the settings, however, you do not have to do it at this time. Restarting the storage system takes approximately seven to 25 minutes.
9. To restart now, click **OK**. Restarting the storage system takes approximately seven to 25 minutes. To restart later, click **Cancel**.

Your changes will be retained and implemented the next time you restart the array.

Installing SnapShot or TCE or Dynamic

SnapShot, TrueCopy Extended Distance (TCE), and Dynamic Provisioning use a portion of the cache to manage internal resources. This means that the cache capacity available to Cache Partition Manager becomes smaller (see [Table 7-17 on page 7-29](#) for additional details).

Note the following:

- Make sure that the cache partition information is initialized as shown below when SnapShot, TCE, or Dynamic Provisioning is installed under Cache Partition Manager.

- All the volumes are moved to the master partitions on the side of the default owner controller.
- All the sub-partitions are deleted and the size of the each master partition is reduced to a half of the user data area after the installation of either SnapShot, TCE, or Dynamic Provisioning.

VMWare and Cache Partition Manager

The VMWare ESX has a function that clones the virtual machine. If the source volume and the target volume cloning are different, and the volumes belong to subpartitions, the time required for a clone to occur may become too long when vStorage APIs for array integration (VAAI) function is enabled. If you need to clone between volumes which belong to subpartitions, please disable the VAAI function of ESX to achieve higher performance.

Cache Residency Manager overview

The Cache Residency Manager function ensures that all data in a volume is stored in cache memory. All read/write commands to the volume can be executed at a 100% cache hit rate without accessing the drive. The system throughput is improved when this function is applied to a volume that contains data accessed frequently because no latency period is needed to access the disk drive.

If a cache residency setting is added, deleted, or modified during power down, power down can fail. If this happens, power down again and verify that no RAID group in the Power Saving Status of Normal (Command Monitoring) exists. Then, you can add, delete, or modify the Cache Partition.

When you set, delete or change Cache Residency Manager settings when the storage system is used on other remote side of TrueCopy or TCE, the following activity results when you restart the system:

- Both paths of TrueCopy or TCE are blocked. In an instance of a blocked path, the system generates a trap to the SNMP Agent Support function. The path of TrueCopy or TCE is automatically recovered from the block after the system restarts.
- When the pair status of TrueCopy or TCE is in either a Paired or Synchronizing state, it changes to the Failure state.

When the Power Saving instruction of the non I/O link is executed with the priced option, Power Saving or Power Saving Plus are used together. If a Cache Partition is added, deleted, or changed while the Power Saving status is Normal (Command Monitoring), the status is changed to "Normal (Spin down Failure: PS OFF/ON)" by the array reboot which works at the time of the setting change and then the spin-down may fail.

Cache Residency Manager features

The following are Cache Residency Manager features:

- **Data and Applications Available in Cache** - Cache Residency loads a volume into the cache.

Cache Residency Benefits

The following are Cache Residency Manager benefits:

Improves read-write performance for a specific volume that has been loaded into cache.

Write data is mirrored and protected on disk asynchronously, or in its own time. If you have a power outage, cache only works when you have power to it because it contains dynamic memory. Saved permanently.

All Read/Write activity occurs in cache.

- **Application Management Ease** - Enables ease of management of applications as they are easily executable from DRAM cache rather than general memory.
- **Application Portability** - Enables portability of applications as they are easily retrievable from DRAM cache rather than general memory.
- **Enhanced Performance** - Enables higher performance of applications as they can launch more quickly from DRAM cache rather than general memory.

Cache Residency Manager task flow

The following steps detail the task flow of the Cache Partition Manager configuration process:

1. You determine you need to create partitions in your storage cache to map to your applications for fasted access.
2. Map out a system of partitions on paper that you will apply to configuration in HSNM2.
3. Install a license key for Cache Residency Manager.
4. Launch HSNM2.
5. Launch Cache Residency Manager and configure Cache Residency Manager. The controller executes read/write commands to the volume using the Cache Residency Manager as follows:
6. Read data accessed by the host is stored in the cache memory until the array is turned off. Subsequent host access to the previously accessed area is transferred from the cache memory without accessing the disk drives.
7. Write data from the host is stored in the cache memory, and not written to the disk drives until the array is turned off.
8. The cache memory utilizes a battery backup and the write data is duplicated (stored in the cache memory on both controllers).
9. Write data stored in the cache memory is written to disk drives when the array is turned off and when the Cache Residency Manager is stopped by failures.

The internal controller operation is the same as that of the commands issued to other volumes, except that the read/write command to the volume with the Cache Residency Manager can be transferred from/to the cache memory without accessing the disk drives.

A delay can occur in the following cases even if Cache Residency Manager is applied to the volumes.

1. The command execution may wait for the completion of commands issued to other volumes.
2. The command execution may wait for the completion of commands other than read/write commands (such as the Mode Select command) issued to the same volume.
3. The command execution may wait for the completion of processing for internal operation such as data reconstruction, etc.

Figure 7-13 shows how part of cache memory installed in the controller is used for the Cache Residency Manager function. Cache memory utilizes a battery backup on both controllers, and the data is duplicated on each controller for safety against power failure and cache package failure.




Figure 7-13: Cache Residency Manager task flow

Cache Residency Manager Specifications

Table 7-12 details the equipment required for Cache Residency Manager.

Table 7-12: Cache Residency Specifications

Item	Description
Controller configuration	Dual Controller configuration and controller is not blockaded.
RAID level	RAID 5, RAID 6, or RAID 1+0.
Cache partition	Only the volume belonging to a master partition.

Table 7-12: Cache Residency Specifications (Continued)

Item	Description
Number of volumes with the Cache Residency function	1/controller (2/arrays)

Termination Conditions

Cache Residency Manager restarts when the failures are corrected. [Table 7-13](#) details the conditions that terminate Cache Residency Manager.

Table 7-13: Cache Residency Manager Termination

Condition	Description
The array is turned off	Normal case.
The cache capacity is changed and the available capacity of the cache memory is less than volume size	Cache uninstallation.
A controller failure	Failure.
The battery alarm occurs	Failure.
A battery backup circuit failure	Failure.
The number of PIN data (data unable to be written to disk drives because of failures) exceeds the threshold value	Failure.

Cache Residency Manager operations are restarted after failures are corrected.

Disabling Conditions

[Table 7-14](#) details conditions that disable Cache Residency Manager.

Table 7-14: Cache Residency Manager Disabling

Condition	Description
The Cache Residency Manager setting is cleared	Caused by the user.
The Cache Residency Manager is disabled or uninstalled (locked)	Caused by the user.
The Cache Residency Manager volume or RAID group is deleted	Caused by the user.
The controller configuration is changed (Dual/Single)	Caused by the user.

Table 7-14: Cache Residency Manager Disabling

Condition	Description
The array restarted with a status exceeding the resident volume size that can be specified after changing the enabled status of Dynamic Provisioning and Dynamic Tiering.	
The array was restarted with a status exceeding the resident volume size that can be specified after changing the setting of the DP Capacity Mode of Dynamic Provisioning.	



NOTE: When the controller configuration is changed from single to dual after setting up the Cache Residency volume, the Cache Residency volume is cancelled. You can open the Cache Residency Manager in single configuration, but neither setup nor operation can be performed.

Equipment

Table 7-15 details equipment required for Cache Residency Manager.

Table 7-15: Cache Residency Manager Equipment

Item	Description
Controller configuration	Dual Controller configuration and controller is not blockaded.
RAID level	RAID 5, RAID 6, or RAID 1+0.
Cache partition	Only the volume belonging to a master partition.
Volume size	To be resident able volume size.
Number of volumes with the Cache Residency function	1/controller (2/arrays)

Volume Capacity

The maximum size of the Cache Residency Manager volume depends on the cache memory. Note that the Cache Residency volume is only assigned a master partition.

The capacity varies with Cache Partition Manager with the existence of a setting by Cache Partition Manager and Dynamic Provisioning/Dynamic Tiering. There are four scenarios:

- Cache Partition Manager and Dynamic Provisioning are disabled
- Cache Partition Manager is disabled and Dynamic Provisioning (Regular Capacity) is enabled
- Cache Partition Manager is enabled and Dynamic Provisioning (Regular Capacity) and Dynamic Tiering are valid.

- When Cache Partition Manager is invalid and Dynamic Provisioning (Maximum Capacity) and Dynamic Tiering are valid.
- Only when Dynamic Provisioning is valid.

The following restrictions are also detailed in the section below:

- When Cache Partition Manager, SnapShot/TCE/Dynamic Provisioning are disabled when the hardware revision is 0100.
- When Cache Partition Manager, SnapShot/TCE/Dynamic Provisioning are disabled when the hardware revision is 0200.

Supported Cache Residency capacities

This section details Cache Residency capacities.

[Table 7-16](#) details maximum supported capacity for Cache Residency Volume when Cache Partition Manager and Dynamic Provisioning/Dynamic Tiering are disabled.

Table 7-16: Supported Capacity of Cache Residency Volume (Cache Partition Manager and Dynamic Provisioning / Dynamic Tiering are disabled)

Array Model	Installed Cache Memory	Maximum Capacity of Cache Residency Volume
HUS 110	4 GB/CTL	1,028,160 blocks (approx. 502 MB)
HUS 130	8 GB/CTL	3,890,880 blocks (approx. 1,899 MB)
	16 GB/CTL	10,563,840 blocks (approx. 5,158 MB)
HUS 150	8 GB/CTL	3,769,920 blocks (approx. 1,840 MB)
	16 GB/CTL	10,442,880 blocks (approx. 5,099 MB)

[Table 7-17](#) details supported capacity for Cache Residency Manager where Cache Partition Manager is disabled and Dynamic Provisioning - Regular Capacity is enabled.

Table 7-17: Supported Capacity of Cache Residency Volume (Cache Partition Manager is disabled and Dynamic Provisioning - Regular Capacity is Enabled)

Array Model	Installed Cache Memory	Maximum Capacity of Cache Residency Volume
HUS 110	4 GB/CTL	604,800 blocks (approx. 295 MB)
HUS 130	8 GB/CTL	3,245,760 blocks (approx. 1,584 MB)
	16 GB/CTL	9,918,720 blocks (approx. 4,843 MB)
HUS 150	8 GB/CTL	2,116,800 blocks (approx. 1,033 MB)
	16 GB/CTL	8,789,760 blocks (approx. 4,291 MB)

[Table 7-18](#) details supported capacity of Cache Residency Volume where Cache Partition Manager is disabled and Dynamic Provisioning (Maximum Capacity) is enabled.

Table 7-18: Supported Capacity of Cache Residency Volume (Cache Partition Manager is Disabled and Dynamic Provisioning (Maximum Capacity) is Enabled)

Array Model	Installed Cache Memory	Maximum Capacity of Cache Residency Volume
HUS 110	4 GB/CTL	--

HUS 130	8 GB/CTL	2,217,600 blocks (approx. 1,082 MB)
	16 GB/CTL	8,890,560 blocks (approx. 4,341 MB)
HUS 150	8 GB/CTL	--
	16 GB/CTL	7,116,480 blocks (approx. 3,474 MB)

[Table 7-19](#) details supported capacity for Cache Residency Manager where Cache Partition Manager is disabled and Dynamic Provisioning and Dynamic Tiering are enabled.

Table 7-19: Supported Capacity of Cache Residency Volume Cache Partition Manager is disabled and Dynamic Provisioning (Regular Capacity) Dynamic Tiering (Regular Capacity) is enabled

Array Model	Installed Cache Memory	Maximum Capacity of Cache Residency Volume
HUS 110	4 GB/CTL	564, 480 blocks (approx. 275 MB)
HUS 130	8 GB/CTL	3,044,160 blocks (approx 1,486 MB)
	16 GB/CTL	9,717,120 blocks (approx. 4,744 MB)
HUS 150	8 GB/CTL	1,915,200 blocks (approx. 935 MB)
	16 GB/CTL	8,588,180 blocks (approx. 4,193 MB)

Table 7-20: Supported Capacity of Cache Residency Volume Cache Partition Manager is Disabled and Dynamic Provisioning (Maximum Capacity) and Dynamic Tiering (Maximum Capacity) is Enabled

Array Model	Installed Cache Memory	Maximum Capacity of Cache Residency Volume
HUS 110	4 GB/CTL	-
HUS 130	8 GB/CTL	2,016,000 blocks (approx. 984 MB)
	16 GB/CTL	8,668,960 blocks (approx. 4,242 MB)
HUS 150	8 GB/CTL	-
	16 GB/CTL	6,914,880 blocks (approx. 3,376 MB)

[Table 7-21](#) details supported capacity for Cache Residency Manager where Cache Partition Manager is disabled and Dynamic Provisioning is Enabled.

Table 7-21: Supported capacity of Cache Residency Volume (Cache Partition Manager is disabled and Dynamic Provisioning is enabled)

Array Model	Installed Cache Memory	Maximum Capacity of Cache Residency Volume
HUS 110	4 GB/CTL	806,400 blocks (approx. 393 MB)
HUS 130	8 GB/CTL	3,245,760 blocks (approx. 1,584 MB)
HUS 150	8 GB/CTL	2,116,800 blocks (approx. 1,033 MB)
	16 GB/CTL	8,789,760 blocks (approx. 4,291 MB)

Table 7-22 details supported capacity where Cache Partition Manager is disabled and Dynamic Tiering is enabled.

Table 7-22: Supported capacity of Cache Residency Volume (Cache Partition Manager disabled and Dynamic Tiering enabled)

Array Model	Cache	Volume Capacity
HUS 110	4 GB/CTL	564,480 blocks (approx. 275 MB)
HUS 130	8 GB/CTL	3,044,160 blocks (approx. 1,486 MB)
	16 GB/CTL	9,717,120 blocks (approx. 4,744 MB)
HUS 150	8 GB/CTL	1,915,200 blocks (approx. 935 MB)
	16 GB/CTL	8,588,160 blocks (approx. 4,2193 MB)

Table 7-23 details supported capacity where Cache Partition Manager is enabled and where Dynamic Provisioning/Dynamic Tiering is enabled or disabled.

Table 7-23: Supported capacity of Cache Residency Volume with Cache Partition Manager enabled

Array Model	Cache	Volume Capacity
HUS 110	4 GB/CTL	(The master partition size (MB) Note 1 - 200 MB) x 2,016 (Blocks)
HUS 130	8 GB/CTL	
	16 GB/CTL	
HUS 150	8 GB/CTL	(The master partition size (MB) Note 1 - 400 MB) x 2,016 (blocks)
	16 GB/CTL	



NOTE: 1. The size becomes effective next time you start and is the master partition size. Use the value of the smaller one in a formula.



NOTE: 2. One (1) block = 512 bytes, and a fraction less than 2,047 MB is omitted.

Restrictions

Table 7-24 details Cache Residency Manager restrictions.

Table 7-24: Cache Residency Manager restrictions

Item	Description	Remarks
Concurrent use of SnapShot	Cache Residency Manager and SnapShot can be used together at the same time, but the volume specified for Cache Residency Manager (volume cache residence) cannot be set to P-VOL, V-VOL.	
Concurrent use of Cache Partition Manager	You cannot change a partition affiliated with the Cache Residency volume. After you cancel the Cache Residency volume, you must set it up again.	After you cancel a Cache Residency volume, you must reconfigure the environment deploying concurrent use of Cache Residency Manager and Cache Partition Manager.
Concurrent use of Volume Migration	The Cache Residency Manager volume (volume cache residence) cannot be set to P-VOL or S-VOL. After you cancel the Cache Residency volume, you must set it up again.	After you cancel a Cache Residency volume, you must reconfigure the environment deploying concurrent use of Cache Residency Manager and Volume Migration.
Concurrent use of Power Saving/Power Saving Plus	A RAID group volume that has powered down can be specified as the Cache Residency volume. However, if a RAID group spins down by non I/O link power saving instructions, an error occurs in host access in the Cache Residency volume belonging to the RAID Group.	
Concurrent use of TCE	The volume specified for Cache Residency Manager (volume cache residence) cannot be set to P-VOL or S-VOL. When using TCE concurrently, volume capacity is limited.	
Concurrent use of Volume Expansion	The unified volume cannot be set to the Cache Residency volume. The Cache Residency volume cannot be used as a unified volume.	
Concurrent use of RAID group expansion	You cannot configure an volume as a Cache Residency volume while executing a RAID group expansion. You cannot execute a RAID group expansion for a RAID group that contains a Cache Residency volume.	

Table 7-24: Cache Residency Manager restrictions (Continued)

Item	Description	Remarks
Volume Expansion	You cannot configure an volume as a Cache Residency volume if that volume has been expanded. growing as a Cache Residency volume. You cannot expand volumes that have been configured as Cache Residency volumes.	
Volume Reduction (shrinking)	You can specify the volume after the volume reduction as a Cache Residency volume. However, you cannot execute an volume reduction for a Cache Residency volume.	
Load balancing	The volume specified for Cache Residency Manager is out of the range of load balancing.	
DP-VOLs	You cannot specify the DP-VOLs created by Dynamic Provisioning.	

Cache Residency Manager procedures

The procedure for Cache Residency Manager appears below.

Confirming environments

When the Power Saving instruction of the non I/O link is executed with the priced option, Power Saving or Power Saving Plus are used together. If a Cache Residency Manager instance is installed, uninstalled, or changed while the Power Saving status is Normal (Command Monitoring), the status is changed to “Normal (Spin down Failure: PS OFF/ON)” by the array reboot which works at the time of the setting change and then the spin-down may fail.

When the spin-down fails, run a spin-down session again. Before adding, deleting, or changing the Cache Partition, check that the spin-down instruction has not been issued or there is no RAID group where the Power Saving status is Normal (Command Monitoring) by the Power Saving instruction of the non I/O link before installing, uninstalling, or changing the Cache Residency Manager instance.

Initial settings

To configure initial settings

1. Verify that you have the environments and requirements for Cache Residency Manager (see [Preinstallation information for Storage Features on page 3-20](#)).
2. Set the Cache Residency Manager (see section below).

Stopping Cache Residency Manager

To stop Cache Residency Manager

- 1. Cancel the volume (see section below).
- 2. Disable Cache Residency Manager (see section below).

Before managing cache residency volumes, make sure that they have been defined.

Setting and canceling residency volumes

To set and cancel residency volumes

- 1. Start Navigator 2 and log in. The Arrays window appears
- 2. Click the appropriate array.
- 3. Click **Cache Residency** from the Performance option in the tree view. The Cache Residency dialog box displays as shown in [Figure 7-14](#).




Figure 7-14: Cache Residency dialog box

- 4. Click **Change Residency**. The Change Residency screen displays as shown in [Figure 7-15](#).




Figure 7-15: Change Residency dialog box

5. Click the **Enable** checkbox of the Controller 0 or Controller 1. To cancel Cache Residency, uncheck the **Enable** checkbox for the selected controller.
6. Select a volume and click **Ok**. A message box displays.
7. Follow the on-screen instructions. A message displays confirming the optional feature installed successfully. Mark the checkbox and click **Reboot Array**.
8. To complete the installation, restart the storage system. The feature will close upon restarting the storage system. The storage system cannot access the host until the reboot completes and the system restarts. Restarting usually takes from seven to 25 minutes.



NOTE: The storage system may require more time to respond, depending on its condition. If it does not respond after 25 minutes, check the condition of the system.

NAS Unit Considerations

The following items are considerations for using the NAS unit when it is connected to the storage system.

- Check the following items in advance:
- NAS unit is connected to the storage system. (*1).
- NAS unit is in operation (*2).
- A failure has not occurred on the NAS unit. (*3).

- Confirm with the storage system administrator to check whether the NAS unit is connected or not.
- Confirm with the NAS unit administrator to check whether the NAS service is operating or not.
- Ask the NAS unit administrator to check whether failure has occurred or not by checking with the NAS administration software, NAS Manager GUI, List of RAS Information, etc. In case of failure, execute the maintenance operation together with the NAS maintenance personal.
- Correspondence when connecting the NAS unit:
If the NAS unit is connected, ask the NAS unit administrator for termination of NAS OS and planned shutdown of the NAS unit.
- Points to be checked after completing this operation:
Ask the NAS unit administrator to reboot the NAS unit. After rebooting, ask the NAS unit administrator to refer to “Recovering from FC path errors” in “Hitachi NAS Manager User’s Guide” and check the status of the Fibre Channel path and to recover the FC path if it is in a failure status.

In addition, if there are any personnel for the NAS unit maintenance, ask the NAS unit maintenance personnel to reboot the NAS unit.

VMware and Cache Residency Manager

The VMware ESX has a function to clone the virtual machine. If the source volume or the target volume of cloning is set the Residency volume, the time required for the clone may become long when vStorage APIs for Array Integration (VAAI) function is enabled. If you need to clone the Residency volume, please disable the VAAI function of ESX.

Performance Monitor

This chapter provides details on monitoring your HUS storage system using Performance Monitor, an event tracking system provided in Navigator 2.

The topics covered in this chapter are:

- ❑ [Performance Monitor overview](#)
- ❑ [Launching Performance Monitor](#)
- ❑ [Performance Monitor procedures](#)
- ❑ [Optimizing system performance](#)
- ❑ [Dirty Data Flush](#)

Performance Monitor overview

Performance Monitor is a program that is used to monitor various activities on a storage system such as disk usage, transfer time, port administrative states, and memory usage.

When the disk array is monitored using Performance Monitor, utilization rates of resources in the disk array (such as loads on the disks and ports) can be measured. When a problem such as slow response occurs in a host, the system administrator can quickly determine the source of the difficulty by using Performance Monitor.

Performance Monitor can display information as a graph, a bar chart, or numeric values and can update information using a range of time intervals. The categories of information that you can monitor depend on which networking and storage services are installed on your system. Other possible categories include Microsoft Network Client, Microsoft Network Server, and protocol categories.

This application is usually used to determine the cause of problems on a local or remote computer by measuring the performance of hardware, software services, and applications. Performance Monitor is not installed automatically during setup, you must install it using a license key.

Three main areas of performance you measure are:

- CPU activity
- Memory activity
- I/O operations

Monitoring features

- **Graphing utility** - Performance Monitor provides a mechanism to create graphs that represent activity that occurs using a specific system trend or event as a criterion. An example of a trend that you can generate a graph from is CPU usage.
- **Flexible data collecting criteria** - Performance Monitor enables you to change data collecting criteria like interval time and using combinations of criteria objects.
- **Multiple output types** - Performance Monitor enables you to display monitored data in various forms in addition to a graph, including bar and pie charts.
- **Tree view** - Performance Monitor provides its own menuing system in the form of a navigation tree called a Tree View. The various items you can display in the Tree View include volumes, data pools, and ports.
- **Collection status utility** - Performance Monitor provides a mechanism where data generated by the monitor. displays according to the Change Measurement Items utility. It provides a status of the current snapshot of the trend or event.
- **Ability to save monitored data** - Performance Monitor enables you to save data generated through monitoring sessions by exporting it to various file types.

- **Dirty Data Flush** - A mode that improves the read response performance when the I/O load is light. If the write I/O load is heavy, a timeout may occur because not enough dirty jobs exist to process the conversion of dirty data as the number of jobs is limited to one.

Monitoring benefits

The following are benefits of the Performance Monitor system.

- **Adjustment elimination** - Eliminates ongoing adjustment of storage system and storage network.
- **Rapid diagnosis** - Enables users to more rapidly diagnose performance capabilities of host based systems and applications
- **Increased efficiency** - Enables increased efficiency by locating and recommending solutions to impasses in the storage system and SAN performance. Decreases problem determination time and diagnostic analysis

Monitoring task flow

The following is a typical task flow of monitoring trends and events on your storage system.

1. An event or trend occurs where retrieval of data from your storage system has either slowed or has yielded inaccurate or partial renderings of data.
2. Attempts at troubleshooting the problem are unsuccessful.
3. Enter Performance Monitor home screen.
4. Display a graph of recent performance.
5. Change trend or event criteria settings for monitoring performance.
6. Set an interval time for obtaining data on performance.
7. Display a new graph.
8. Export the data to a .CSV file.

The following figure details the flow of tasks involved with Performance Monitor:




Figure 8-1: Performance Monitor task flow

Monitoring feature specifications

Table 8-1 lists the Performance Monitor specifications.

Table 8-1: Performance Monitor specifications

Item	Description
Information	Acquires array performance and resource utilization.
Graphic display	Information is displayed with line graphs. Information displayed can be near-real time.
Information output	The information can be output to a CSV file.
Management PC disk capacity	Navigator 2 creates a temporary file to the directory where it is installed to store the monitor output data. The disk capacity of the maximum of 2.4 GB is required. For CSV file output, a free disk capacity of at least 750 MB is required.
Performance information acquisition	Performance Monitor acquires information on performance and resource utilization of the disk array.
Disk capacity of management PC	When outputting the monitoring data, Hitachi Storage Navigator Modular 2 creates a temporary file to the directory where Hitachi Storage Navigator Modular 2 is installed. The disk capacity of the maximum of 2.4 GB is required. When outputting the CSV files, the disk capacity of the maximum of 750 MB is required.
Concurrent use other price-cost optional feature	Concurrent use together with other all price-cost optional feature.

Analysis bottlenecks of performance

Rising processor and drive usage in the storage system may create a bottleneck for performance on the system. In addition, the performance bottleneck may occur when there is an imbalanced load. When conditions result in slowing performance, you may want to change the environment on your system.

The following table [Table 8-2](#) details criteria for judging a high load.

Table 8-2: High load performance limitations

No.	Type	Performance	Description
1	Processor	Usage (%)	When the operating rate of the processor exceeds 90 percent.
2	Drive Operation	Operating Rate (%)	When the operating rate of the drive exceeds 80 percent.
3		Tag Count	When the drive is the SAS/SSD/SAS7.2K drive and the multiplicity of commands is more than 20 tags.
4		Tag Average	

Note that these limitations are measured during normal operation when hardware failures have not occurred.

To launch Performance Monitor, click a storage system in the navigation tree, then click Performance, and click Monitoring to launch the Monitoring - Performance Measurement Items window as shown in Figure 8-2. Note that Dynamic Provisioning is valid in the following figure.




Figure 8-2: Performance Monitor - Monitoring window

By clicking the Show Graph button, Performance Monitor displays the non-graph Performance Monitor - Show Graph screen with Dynamic Provisioning enabled.




Figure 8-3: Performance Monitor window (non graph: Dynamic Provisioning enabled)

The following table provides summary information for each item in the Performance Monitor screen.

Figure 8-4: Performance Monitor window summary information

Item	Description
Graph item	The objects of the information acquisition and the graphic display occur with icons. When you click on a radio button, details of the icon display in the Detailed Graph Item.
Detailed Graph Item	Details of items selected in the Graph Item display. The most recent performance information of each item displays for the array configuration and the defined configuration.
Graph Item Information	Specify items to be graphically displayed by selecting them from the listed items. Items to be displayed are determined according to the selection that is made in the Graph Item.
Interval Time	Specify an interval for acquiring the information. Specify it in units of minutes within a range from one minute to 23 hours and 59 minutes. The default interval is one minute. In the above-mentioned interval time, the data for a maximum of 1,440 times can be stored. If it exceeds 1,440 times, it will be overwritten from the old data.

Performance Monitor procedures

The procedure for Performance Monitor appears below.

Initial settings

To configure initial settings

1. Verify that you have the environments and requirements for Performance Monitor (see [Preinstallation information for Storage Features on page 3-20](#)).
2. Collect the performance monitoring data (see [Obtaining information on page 8-8](#)).

Optional operations

1. Use the graphic displays (see [Using graphic displays on page 8-8](#)).
2. Output the performance monitor information to a file.
3. Optimize the performance (see [Troubleshooting performance on page 8-36](#)).

Optimizing system performance

This section describes how to use Performance Monitor to optimize your system.

Obtaining information

The information is obtained for each controller.

To obtain information for each controller

1. Start Navigator 2 and log in. The Arrays window opens
2. Click the appropriate array.
3. Click **Performance** and click **Monitoring**. The Monitor - Performance Measurement Items window displays.
4. Click **Show Graph**.
5. Specify the interval time.
6. Select the items (up to 8) that you want to appear in the graph.
7. Click **Start**. When the interval elapses, the graph appears.



NOTE: If the array is turned off or cannot acquire data, or a controller failure occurs, incorrect data can appear.

Using graphic displays

You must have the license key installed to display performance graphs. When installed, the **Show Graph** button is available from the Performance Monitor window.

To display graphs

1. Obtain the information. Note that if you close the Performance Monitor window, the information is lost.
2. Select the appropriate item, and click **Show Graph**. The Performance Monitor Graph window appears (see [Figure 8-3 on page 8-6](#)).
3. To change the item that is being displayed, select the appropriate values from the drop-down menus.



NOTE: The graphic display data cannot be saved. However, you can copy the information in a comma-separated values (CSV) file. For more information, see [Dirty Data Flush is a mode that improves the read response performance when the I/O load is light. If the write I/O load is heavy, a timeout may occur because not enough dirty jobs exist to process the conversion of dirty data as the number of jobs is limited to one. So the mode should be changed when the I/O load is light.](#) on page 8-37.

An example of a Performance Monitor graph (CPU usage) is shown in [Figure 8-5 on page 8-9](#).



Figure 8-5: Performance Monitor — sample graph (CPU usage)

Table 8-3 shows the summary of each item in the Performance Monitor.

Table 8-3: Summary of Performance Monitor window

Item	Description
Collection Status of Performance Statistics	Data in the Category and Status columns are displayed according to the selection that is made in the Change Measurement Items. Start is displayed in the Status column.
Interval Time	Specify an interval for acquiring information. Specify the interval in minute time units within a range from one minute to 23 hours and 59 minutes. The default interval is one minute. A maximum of 1,440 instances of interval time can be stored. If the number of instances exceeds 1,440 times, Performance Monitor, overwrites the old data.
Tree View	The objects associated with performance measurement display as a list in the navigation bar to the right of the main region of the Performance Monitor Window. The objects display as text strings accompanied by mnemonic icons to the left of the strings. The object types are associated with information acquisition and graphic display.
List	Details of items selected in the Tree View display as a list. The most recent performance information of each item displays for the storage system configuration and the defined configuration.
Displayed Items	Specify items to be graphically displayed by selecting them from the listed items. Items displayed in the drop-down list to be displayed are determined according to the selection that is made in the Tree View.

Working with the Performance Monitor Tree View

The Tree View is the list of objects Performance Monitor measures displayed in the navigation bar to the right of the main portion of the Performance Monitor Window. The objects display as text strings accompanied by icons to the left of the strings. The objects are associated with information acquisition and graphic display. Table 8-4 provides descriptions of Tree View icons.

Table 8-4: Tree View icons

Icon	Item Name	Description
	Registered array name.	Represents the array.

Table 8-4: Tree View icons

Icon	Item Name	Description
	Controller 0/Controller 1 Information	Represents the controller on the storage system. In the case of the single controller system, an icon of the Controller 1 is not displayed. When one of the controllers is registered with Navigator 2 in the case of the dual controller system, only an icon of the connected controllers display. Clicking this icon displays a Tree view of icons that belong to the controller. Information on this icon is not displayed in the list. In the case of a single controller system, an icon of CTL 1 is not displayed. When one of the controllers is registered with HSNM2 in the case of the dual controller system, only an icon of the connected controller displays.
	Port Information	Represents the selected port number on the current storage system. Information on the port displays in the list.
	RAID Groups Information	Represents RAID groups that have been defined for the current storage system. Information on the RAID groups display in the list.
	DP Pool Information	Represents the Dynamic Provisioning pools that have been defined for the current storage system. Information on the DP pool displays in the list.
	Volume Information	Represents the volumes defined for the current storage system. Information on the volumes displays in the list.
	Cache Information	Represents the cache resident in the current storage system. Information on the cache displays in the list.
	Processor Information	Represents the processor in the current storage system. Information on the processor displays in the list.
	Drive Information	Represents the disk drive in the current storage system. Information on the drive displays in the list.
	Drive Operation Information	Represents the drive operation in the current storage system. Information on the drive displays in the list.
	Back-End Information	Represents the back-end of the current storage system. Information on the back-end displays in the list.

Note that procedures in this guide frequently refer to the Tree View as a list, for example, the Volume Migration list.

More about Tree View items in Performance Monitor

The following tables detail items selected in the Tree View. The most recent performance information of each item displays for the storage system configuration and the defined configuration.

During the monitoring process, the display updates automatically at regular intervals. Even if the definition of the RAID group or volume changes during the monitoring, the change produces no effect on the list. Before the monitoring starts, the list is blank.

After the monitoring begins, the agent may not acquire the information to run the application. This may occur because of traffic problems on the LAN when the specified interval elapses. In cases of blocked information acquisition, a series of three dash symbols (---) displays. For a list of items that have blocked information acquisition, the N/A string displays.

Specify items to be graphically displayed by selecting them from the drop-down list launched from the top level list of objects in the Tree View. Items displayed in the drop-down list of objects to be displayed are determined according to the selection that is made in the Tree View.

The following tables display the relationship between the Tree View and the display in the list.

[Table 8-5](#) details items in the Port item.

Table 8-5: Expanded Tree View of port item

Displayed Items	Description
Port	Port number (The maximum numbers of resources that can be installed in the array are displayed).
IO Rate (IOPS)	Received number of Read/Write commands per second.
Read Rate (IOPS)	Received number of Read commands per second.
Write Rate (IOPS)	Received number of Write commands per second.
Read Hit (%)	Rate of cache-hitting within the received Read command.
Write Hit (%)	Rate of cache-hitting within the received Write command.
Trans. Rate (MB/s)	Transfer size of Read/Write commands per second.
Read Trans. Rate (MB/s)	Transfer size of Read commands per second.
Write Trans. Rate (MB/s)	Transfer size of Write commands per second.
CTL CMD IO Rate (IOPS)	Sent number of control commands of TrueCopy Initiator per second (acquired local side only).
Data CMD IO Rate (IOPS)	Sent number of data commands of TrueCopy initiator per second (acquired local side only).
CTL CMD Trans. Rate (KB/s)	Transfer size of control commands of TrueCopy Initiator per second (acquired local side only).
Data CMD Trans. Rate (MB/s)	Transfer size of data commands of TrueCopy Initiator per second (acquired local side only).
CTL CMD Time (microsec.)	Average response time of commands of TrueCopy Initiator (acquired local side only).

Displayed Items	Description
Data CMD Time (microsec.)	Average response time of data commands of TrueCopy Initiator (acquired local side only).
CTL CMD Max Time (microsec.)	Maximum response time of control commands of TrueCopy Initiator (acquired local side only)
Data CMD Max Time (microsec.)	Maximum response time of data commands of TrueCopy Initiator (acquired local side only)
XCOPY Rate (IOPS)	Received number of XCOPY commands per second
XCOPY Time (microsec.)	Average response time of XCOPY commands
XCOPY Max Time (microsec)	Maximum response time of XCOPY commands
XCOPY Read Trans Rate (MB/s)	Transfer size of XCOPY Read commands per second
XCOPY Write Rate (IOPS)	Received number of XCOPY Write commands per second
XCOPY Write Trans Rate (MB/s)	Transfer size of XCOPY Write commands per second

Table 8-6 details items in the RAID Groups DP Pools item.

Table 8-6: Expanded Tree View of RAID groups DP Pool items

Displayed Items	Description
RAID Group/DP Pool	The RAID group/DP Pool number that has been defined for the current storage system.
IO Rate (IOPS)	Received number of read/write commands per second.
Read Rate (IOPS)	Received number of read commands per second.
Write Rate (IOPS)	Received number of write commands per second.
Read Hit (%)	Rate of cache-hitting within the received Read command.
Write Hit (%)	Rate of cache-hitting within the received Write command.
Trans. Rate (MB/s)	Transfer size of read/write commands per second.
Read Trans. Rate (MB/s)	Transfer size of read commands per second.
Write Trans. Rate (MB/s)	Transfer size of write commands per second.
XCOPY Rate (IOPS)	Received number of XCOPY commands per second
XCOPY Time (microsec.)	Average response time of XCOPY commands
XCOPY Max Time (microsec.)	Maximum response time of XCOPY commands.
XCOPY Read Rate (IOPS)	Received number of XCOPY Read commands per second
XCOPY Read Trans Rate (MB/s)	Transfer size of XCOPY Read commands per second
XCOPY Write Trans Rate (MB/s)	Transfer size of XCOPY Write commands per second

Table 8-7 details items in the Volume, Cache, and Processor items.

Table 8-7: Expanded Tree View of volume, cache, and processor items

Item	Displayed Items	Description
Volume DP Pool	Volume	Volume number defined for the current storage system.
	IO Rate (IOPS)	Received number of read/write commands per second.
	Read Rate (IOPS)	Received number of read commands per second.
	Write Rate (IOPS)	Received number of write commands per second.
	Read Hit (%)	Rate of cache-hitting within the received read command.
	Write Hit (%)	Rate of cache hitting within the received write command.
	Trans. Rate (MB/s)	Transfer size of read/write commands.
	Read Trans. Rate (MB/s)	Transfer size of read commands per second.
	Write Trans. Rate (MB/s)	Transfer size of write commands per second.
	Tag Count (only volume)	Maximum multiplicity of commands between intervals.
	Tag Average (only volume)	Average multiplicity of commands between intervals.
	Data CMD IO Rate (IOPS)	Sent number of data commands of TrueCopy Initiator per second (acquired local side only).
	Data CMD Trans. Rate (MB/s)	Transfer size of data commands of TrueCopy Initiator per second (acquired local side only)
	XCOPY Max Time (microsec.)	Maximum response time of XCOPY commands
	XCOPY Read Rate (IOPS)	Received number of XCOPY Read commands per second.
	XCOPY Read Trans. Rate (MB/s)	Transfer size of XCOPY Read commands per second
	XCOPY Write Rate (IOPS)	Received number of XCOPY Write commands per second
	XCOPY Write Trans Rate (MB/s)	Transfer size of XCOPY Write commands per second
Cache	Write Pending Rate (%)	Rate of cache usage capacity within the cache capacity.
	Clean Queue Usage Rate (%)	Clean cache usage rate.
	Middle Queue Usage Rate (%)	Middle cache usage rate.
	Physical Queue Usage Rate (%)	Physical cache usage rate.
	Total Queue Usage Rate (%)	Total cache usage rate.
Processor	Usage (%)	Operation rate of the processor.



NOTE: Total cache usage rate and cache usage rate per partition display.

Table 8-8 details items in the Volume, Cache, and Processor items.

Table 8-8: Expanded Tree View of drive and back-end items

Item	Displayed Items	Description
Drive	Unit	Operation rate of the processor.
	HDU	Hard Drive Unit number, the maximum number of resources that can be installed in the array display.
	IO Rate (IOPS)	Received number of read/write commands per second.
	Read Rate (IOPS)	Received number of read commands per second.
	Write Rate (IOPS)	Received number of write commands per second.
	Trans. Rate (MB/s)	Transfer size of read/write commands per second.
	Read Trans. Rate (MB/s)	Transfer size of read commands per second.
	Write Trans. Rate (MB/s)	Transfer size of write commands per second.
	Online Verify Rate (IOPS)	Number of Online Verify commands per second.
Drive Operation	Unit	Unit number, the maximum number of resources that can be installed in the array display.
	HDU	Hard Drive Unit number, the maximum number of resources that can be installed in the storage system display.
	Operating Rate (%)	Operation rate of the drive.
	Tag Count	Maximum multiplicity of drive commands between intervals.
	Tag Average	Average multiplicity of drive commands between intervals.
Back-End	Path	Path number, the maximum number of resources that can be installed in the storage system display.
	IO Rate (IOPS)	Received number of read/write commands per second.
	Read Rate (IOPS)	Received number of read commands per second.
	Write Rate (IOPS)	Received number of write commands per second.
	Trans. Rate (MB/s)	Transfer size of read/write commands per second.
	Read Trans. Rate (MB/s)	Transfer size of read commands per second.

Table 8-8: Expanded Tree View of drive and back-end items

Item	Displayed Items	Description
	Write Trans. Rate (MB/s)	Transfer size of write commands per second.
	Online Verify Rate (IOPS)	Number of Online Verify commands per second.

For the cache hit of the write command, the command performs the operation (write after) to respond to a host with the status at the time of completing write to the cache memory. Because of this response type, two exception cases exist that are worth noting where a write to the cache memory is viewed by the application variously as a hit and a miss:

- A case where the write to the cache memory is immediately performed is defined as a hit.
- A case where the write to the cache memory is delayed because of heavy cache memory use is defined as a miss.

Using Performance Monitor with Dynamic Provisioning

When using Performance Monitor with Dynamic Provisioning enabled, the output displayed is slightly different. [Figure 8-6](#) displays a sample Performance Monitor Window when Dynamic Provisioning is valid.

[Show Graph](#)

Graph Property

Select the items that you want to show on the graph. After completing settings, click Start. Click Show Graph to show performance graph. If you also want to save the performance data to the CSV file, move to CSV Files tab after measuring data in the Graph tab.

Graph **CSV Files**

Array Name: HUS110_91100026
 Serial Number: 91100026
 Interval Time: 1 minute(s)
 Preservation Period: From 2011/12/15 13:35 to 2011/12/15 14:34

* Time range to save in the CSV file:
 From: 2011 / 12 / 15 13 : 35
 To: 2011 / 12 / 15 14 : 34

Items to be saved in the CSV file:

- ☒ HUS110_91100026
- ☒ Controller 0
- ☐ Port Information
- ☐ RAID Group Information
- ☐ DP Pool Information
- ☐ Volume Information
- ☒ Cache Information
- ☐ Processor Information

Measurement Item Information

Item
☒ Write Pending Rate
☒ Clean Queue Usage Rate
☒ Middle Queue Usage Rate
☒ Physical Queue Usage Rate
☒ Total Queue Usage Rate

[Get CSV File](#) [Get CSV File](#)

Figure 8-6: Performance Monitor: Dynamic Provisioning is valid

Working with Graphing and Dynamic Provisioning

The Performance Monitor graph application also behaves differently when Dynamic Provisioning is valid. [Figure 8-7 on page 8-17](#) displays a sample graph when Dynamic Provisioning is valid.

[Show Graph](#)



Figure 8-7: Performance Monitor graph: Dynamic Provisioning enabled

The time and date when the information was acquired is displayed on the axis of the abscissa. the axis of the ordinate is determined by selecting the maximum value on the Y-axis. Selectable values vary according to the item selected.

In the graph, five data points corresponding to particular intervals are plotted per on graduation. the name of the item being displayed is show below the graph. The example shown in the figure is CTL0-Processor-Usage(%).

Invalid data may display if any of the following events occur during monitoring:

- Storage system power is off or shuts down
- Controller failure
- Storage system could not acquire data by a network obstacle
- Firmware in the process of updating

Table 8-9 displays selectable Y axis values.

Table 8-9: Selectable Y axis values in SNM2 versions less than V22.50

Selected Item	Displayed Items	Selectable Y Axis Values
Port Information	IO Rate	10, 20, 50, 100, 200, 500, 1,000, 2,000, 5,000,
	Read Rate	10,000, 20,000 , 50,000, 100,000, 150,000,
	Write Rate	300,000
	Read Hit	20, 50, 100
	Write Hit	
	Trans. Rate	0, 20, 50, 100, 200 , 500, 1,000, 2,000
	Read Trans. Rate	
	Write Trans. Rate	
	CTL Command IO Rate	10, 50, 100, 200, 500, 1,000, 2,000 , 5,000, 10,000, 20,000, 50,000
	Data Command IO Rate	10, 50, 100, 500, 1,000, 2,000, 5,000, 10,000, 20,000 , 50,000
	CTL Command Trans. Rate	10, 20, 50, 100 , 200, 500, 1,000, 5,000, 10,000, 50,000, 100,000, 150,000
	Data Command Trans. Rate	10, 20, 50, 100 , 200, 400
	CTL Command Time	100, 500, 1,000, 5,000, 10,000, 20,000, 50,000, 100,000 , 200,000, 500,000, 1,000,000, 5,000,000, 10,000,000, 60,000,000
	Data Command Time	100, 500, 1,000, 2,000, 5,000, 10,000, 20,000, 50,000, 100,000 , 500,000, 1,000,000, 5,000,000, 10,000,000, 60,000,000
	CTL Command Max Time	100, 500, 1,000, 5,000, 10,000, 50,000, 100,000, 200,000, 500,000
	Data Command Max Time	1,000,000 , 2,000,000, 5,000,000, 10,000,000, 20,000,000, 60,000,000
	XCOPY Rate	10, 20, 50, 100 , 200, 500, 1,000, 5,000, 10,000, 20,000, 50,000, 10,000, 150,000
	XCOPY Time	100, 500, 1,000, 2,000, 5,000, 10,000 , 20,000, 50,000, 100,000, 500,000, 1,000,000, 5,000,000, 10,000,000, 60,000,000
	XCOPY Max Time	100, 500, 1,000, 5,000, 10,000, 50,000, 100,000 , 200,000, 500,000, 1,000,000, 2,000,000, 5,000,000, 10,000,000, 60,000,000
	XCOPY Read Rate	10, 20, 50, 100, 200, 500, 1,000, 2,000, 5,000, 10,000, 20,000, 50,000, 100,000, 150,000
	XCOPY Write Rate	
	XCOPY Read Trans. Rate	10, 20, 50, 100, 200, 500 , 1,000, 2,000
	XCOPY Write Trans. Rate	

Table 8-9: Selectable Y axis values in SNM2 versions less than V22.50

Selected Item	Displayed Items	Selectable Y Axis Values
RAID Group Information DP Pool Information	I/O Rate	10, 20, 50, 100, 200, 500, 1,000, 2,000, 5,000, 10,000, 20,000 , 50,000, 100,000, 150,000, 300,000
	Read Rate	
	Write Rate	
	Read Hit	20, 50, 100
	Write Hit	
	Trans. Rate	0, 20, 50, 100, 200 , 500, 1,000, 2,000
	Read Trans. Rate	
	Write Trans. Rate	
	XCOPY Rate	10, 20, 50, 100 , 200, 500, 1,000, 5,000, 10,000, 20,000, 50,000, 100,000, 150,000
	XCOPY Time	100, 500, 1,000, 2,000, 5,000, 10,000, 20,000
	XCOPY Max Time	100, 500, 1,000, 2,000, 5,000, 10,000 , 20,000, 50,000, 100,000, 500,000, 1,000,000, 2,000,000, 5,000,000, 10,000,000, 60,000,000
	XCOPY Read Rate	10, 20, 50, 100 , 200, 500, 1,000, 2,000, 5,000, 10,000, 20,000, 50,000, 100,000, 150,000
	XCOPY Write Rate	
	XCOPY Read Trans. Rate	10, 20, 50, 100, 200, 500 , 1,000, 2,000
	XCOPY Write Trans. Rate	
Volume Information	I/O Rate	10, 20, 50, 100, 200, 500, 1,000, 2,000, 5,000, 10,000, 20,000 , 50,000, 100,000, 150,000, 300,000
	Read Rate	
	Write Rate	
	Read Hit	20, 50, 100
	Write Hit	
	Trans. Rate	0, 20, 50, 100, 200 , 500, 1,000, 2,000
	Read Trans. Rate	
	Write Trans. Rate	
	Max Tag Count	500, 1,000, 2,000, 5,000, 10,000, 20,000, 50,000, 100,000
	Average Tag Count	
	Data Command I/O Rate	10, 50, 100, 500, 1,000, 2,000, 5,000, 10,000, 20,000 , 50,000
	Data Command Trans. Rate	10, 20, 50, 100 , 200, 400
	XCOPY Rate	10, 20, 50, 100 , 200, 500, 1,000, 5,000, 10,000, 20,000, 50,000, 100,000, 150,000
	XCOPY Time	100, 500, 1,000, 2,000, 5,000, 10,000 , 20,000, 50,000, 100,000, 500,000, 1,000,000, 5,000,000, 10,000,000, 60,000,000
	XCOPY Max Time	100, 500, 1,000, 5,000, 10,000, 50,000, 100,000 , 200,000, 500,000, 1,000,000, 2,000,000, 5,000,000, 10,000,000, 60,000,000

Table 8-9: Selectable Y axis values in SNM2 versions less than V22.50

Selected Item	Displayed Items	Selectable Y Axis Values
	XCOPY Read Rate	10, 20, 50, 100, 200, 500, 1,000, 2,000, 5,000, 10,000 , 20,000, 50,000, 100,000, 150,000
	XCOPY Write Rate	
	XCOPY Read Trans. Rate	10, 20, 50, 100, 200, 500 , 1,000, 2,000
	XCOPY Write Trans. Rate	
Cache Information	Write Pending Rate Note	20, 50, 100
	Clean Queue Usage Rate Note	
	Middle Queue Usage Rate Note	
	Physical Queue Usage Rate Note	
	Total Queue Usage Rate	
Processor Information	Usage	
Drive Information	I/O Rate	10, 20, 50, 100, 200 , 500, 1,000, 2,000, 5,000, 10,000, 20,000, 50,000
	Read Rate	
	Write Rate	
	Trans. Rate	
	Read Trans. Rate	
	Write Trans. Rate	10, 20, 50, 100, 200, 500, 1,000, 2,000, 5,000, 10,000, 20,000 , 50,000
	Online Verify Rate	
Drive Operation Information	Operating Rate	20, 50, 100
	Max Tag Count	
	Average Tag Count	
Back-end Information	I/O Rate	10, 20, 50, 100, 200, 500, 1,000, 2,000, 5,000, 10,000, 20,000 , 50,000, 100,000
	Read Rate	
	Write Rate	
	Trans. Rate	10, 20, 50, 100, 200 , 500, 1,000, 2,000
	Read Trans. Rate	
	Write Trans. Rate	
	Online Verify Rate	10, 20, 50, 100, 200, 500, 1,000, 2,000, 5,000, 10,000, 20,000 , 50,000

Total cache usage rate and cache usage rate per partition are displayed.

Select the maximum value on the Y-axis judging from the look of the line graph displayed. When the maximum value on the Y-axis is too small, data bigger than the maximum value cannot be displayed because it is beyond the limits of display. When the Show Graph button is clicked, the maximum

value on the Y-axis is set as the default value. However, when the item to be displayed is not changed, the graph is displayed based on the maximum value on the Y-axis used immediately before.

Table 8-10: Selectable Y axis values in SNM2 versions greater than V22.50

Selected Item	Displayed Items	Selectable Y Axis Values
Port Information	IO Rate	10, 20, 50, 100, 200, 500, 1,000, 2,000, 5,000,
	Read Rate	10,000, 20,000 , 50,000, 100,000, 150,000,
	Write Rate	300,000
	Read Hit	20, 50, 100
	Write Hit	
	Trans. Rate	0, 20, 50, 100, 200 , 500, 1,000, 2,000
	Read Trans. Rate	
	Write Trans. Rate	
	CTL Command I/O Rate	10, 50, 100, 200, 500, 1,000, 2,000 , 5,000, 10,000, 20,000, 50,000
	Data Command I/O Rate	10, 50, 100, 500, 1,000, 2,000, 5,000, 10,000, 20,000 , 50,000
	CTL Command Trans. Rate	10, 20, 50, 100 , 200, 500, 1,000, 5,000, 10,000, 50,000, 100,000, 150,000
	Data Command Trans. Rate	10, 20, 50, 100 , 200, 400
	CTL Command Time	100, 500, 1,000, 5,000, 10,000, 20,000, 50,000, 100,000 , 200,000, 500,000, 1,000,000, 5,000,000, 10,000,000, 60,000,000
	Data Command Time	100, 500, 1,000, 2,000, 5,000, 10,000, 20,000, 50,000, 100,000 , 500,000, 1,000,000, 5,000,000, 10,000,000, 60,000,000
	CTL Command Max Time	100, 500, 1,000, 5,000, 10,000, 50,000, 100,000, 200,000, 500,000
	Data Command Max Time	1,000,000 , 2,000,000, 5,000,000, 10,000,000, 20,000,000, 60,000,000
	XCOPY Rate	10, 20, 50, 100 , 200, 500, 1,000, 5,000, 10,000, 20,000, 50,000, 10,000, 150,000
	XCOPY Time	100, 500, 1,000, 2,000, 5,000, 10,000 , 20,000, 50,000, 100,000, 500,000, 1,000,000, 5,000,000, 10,000,000, 60,000,000
	XCOPY Max Time	100, 500, 1,000, 5,000, 10,000, 50,000, 100,000 , 200,000, 500,000, 1,000,000, 2,000,000, 5,000,000, 10,000,000, 60,000,000
	XCOPY Read Rate	10, 20, 50, 100, 200, 500, 1,000, 2,000, 5,000, 10,000, 20,000, 50,000, 100,000, 150,000
	XCOPY Write Rate	
	XCOPY Read Trans. Rate	10, 20, 50, 100, 200, 500 , 1,000, 2,000
	XCOPY Write Trans. Rate	

Table 8-10: Selectable Y axis values in SNM2 versions greater than V22.50

Selected Item	Displayed Items	Selectable Y Axis Values
RAID Group Information DP Pool Information	I/O Rate	10, 20, 50, 100, 200, 500, 1,000, 2,000, 5,000, 10,000, 20,000 , 50,000, 100,000, 150,000, 300,000
	Read Rate	
	Write Rate	
	Read Hit	20, 50, 100
	Write Hit	
	Trans. Rate	0, 20, 50, 100, 200 , 500, 1,000, 2,000
	Read Trans. Rate	
	Write Trans. Rate	
	XCOPY Rate	10, 20, 50, 100 , 200, 500, 1,000, 5,000, 10,000, 20,000, 50,000, 100,000, 150,000
	XCOPY Time	100, 500, 1,000, 2,000, 5,000, 10,000, 20,000
	XCOPY Max Time	100, 500, 1,000, 2,000, 5,000, 10,000 , 20,000, 50,000, 100,000, 500,000, 1,000,000, 2,000,000, 5,000,000, 10,000,000, 60,000,000
	XCOPY Read Rate	10, 20, 50, 100 , 200, 500, 1,000, 2,000, 5,000, 10,000, 20,000, 50,000, 100,000, 150,000
	XCOPY Write Rate	
	XCOPY Read Trans. Rate	10, 20, 50, 100, 200, 500 , 1,000, 2,000
	XCOPY Write Trans. Rate	
Volume Information	I/O Rate	10, 20, 50, 100, 200, 500, 1,000, 2,000, 5,000, 10,000, 20,000 , 50,000, 100,000, 150,000, 300,000
	Read Rate	
	Write Rate	
	Read Hit	20, 50, 100
	Write Hit	
	Trans. Rate	0, 20, 50, 100, 200 , 500, 1,000, 2,000
	Read Trans. Rate	
	Write Trans. Rate	
	Max Tag Count	500, 1,000, 2,000, 5,000, 10,000, 20,000, 50,000, 100,000
	Average Tag Count	
	Data Command I/O Rate	10, 50, 100, 500, 1,000, 2,000, 5,000, 10,000, 20,000 , 50,000
	Data Command Trans. Rate	10, 20, 50, 100 , 200, 400
	XCOPY Rate	10, 20, 50, 100 , 200, 500, 1,000, 5,000, 10,000, 20,000, 50,000, 100,000, 150,000
	XCOPY Time	100, 500, 1,000, 2,000, 5,000, 10,000 , 20,000, 50,000, 100,000, 500,000, 1,000,000, 5,000,000, 10,000,000, 60,000,000
	XCOPY Max Time	100, 500, 1,000, 5,000, 10,000, 50,000, 100,000 , 200,000, 500,000, 1,000,000, 2,000,000, 5,000,000, 10,000,000, 60,000,000

Table 8-10: Selectable Y axis values in SNM2 versions greater than V22.50

Selected Item	Displayed Items	Selectable Y Axis Values
	XCOPY Read Rate	10, 20, 50, 100, 200, 500, 1,000, 2,000, 5,000, 10,000 , 20,000, 50,000, 100,000, 150,000
	XCOPY Write Rate	
	XCOPY Read Trans. Rate	10, 20, 50, 100, 200, 500 , 1,000, 2,000
	XCOPY Write Trans. Rate	
Cache Information	Write Pending Rate Note	20, 50, 100
	Clean Queue Usage Rate Note	
	Middle Queue Usage Rate Note	
	Physical Queue Usage Rate Note	
	Total Queue Usage Rate	
Processor Information	Usage	
Drive Information	I/O Rate	10, 20, 50, 100, 200 , 500, 1,000, 2,000, 5,000, 10,000, 20,000, 50,000
	Read Rate	
	Write Rate	
	Trans. Rate	
	Read Trans. Rate	
	Write Trans. Rate	10, 20, 50, 100, 200, 500, 1,000, 2,000, 5,000, 10,000, 20,000 , 50,000
	Online Verify Rate	
Drive Operation Information	Operating Rate	20, 50, 100
	Max Tag Count	
	Average Tag Count	
Back-end Information	I/O Rate	10, 20, 50, 100, 200, 500, 1,000, 2,000, 5,000, 10,000, 20,000 , 50,000, 100,000
	Read Rate	
	Write Rate	
	Trans. Rate	10, 20, 50, 100, 200 , 500, 1,000, 2,000
	Read Trans. Rate	
	Write Trans. Rate	
	Online Verify Rate	10, 20, 50, 100, 200, 500, 1,000, 2,000, 5,000, 10,000, 20,000 , 50,000

Total cache usage rate and cache usage rate per partition are displayed.

Select the maximum value on the Y-axis judging from the look of the line graph displayed. When the maximum value on the Y-axis is too small, data bigger than the maximum value cannot be displayed because it is beyond the limits of display. When the Show Graph button is clicked, the maximum

value on the Y-axis is set as the default value. However, when the item to be displayed is not changed, the graph is displayed based on the maximum value on the Y-axis used immediately before.

Displayed Items

The following are displayed items in the Port tree view.

- IO Rate
- Read Rate
- Write Rate
- Read Hit
- Write Hit
- Trans. Rate
- Read Trans. Rate
- Write Trans. Rate
- CTL CMD IO Rate
- CTL CMD Trans. Rate
- Data CMD Trans. Rate
- CTL CMD Time
- Data CMD Time
- CTL CMD Max Time
- Data CMD Max Time
- XCOPY Rate
- XCOPY Time
- XCOPY Max Time
- XCOPY Read Rate
- XCOPY Read Trans.Rate
- XCOPY Write Rate
- XCOPY Write Trans.Rate

The following are displayed items in the RAID Groups DP Pool tree view.

- IO Rate
- Read Rate
- Write Rate
- Read Hit
- Write Hit
- Trans. Rate
- Read Trans. Rate
- Write Trans. Rate
- XCOPY Time

- XCOPY Max Time
- XCOPY Read Rate
- XCOPY Read Trans.Rate
- XCOPY Write Rate
- XCOPY Write Trans.Rate

The following are displayed items in the Volume tree view.

- IO Rate
- Read Rate
- Write Rate
- Read Hit
- Write Hit
- Trans. Rate
- Read Trans. Rate
- Write Trans. Rate
- Max Tag Count
- Average Tag Count
- Data CMD IO Rate
- Data CMD Trans. Rate
- XCOPY Rate
- XCOPY Time
- XCOPY Max Time
- XCOPY Read Rate
- XCOPY Read Trans.Rate
- XCOPY Write Rate
- XCOPY Write Trans.Rate
- CacheWrite Pending Rate Note
- Clean Queue Usage Rate Note
- Middle Queue Usage Rate Note
- Physical Queue Usage Rate Note
- Total Queue Usage Rate
- ProcessorUsage
- Drive
- Back-endIO Rate
- Read Rate
- Write Rate
- Trans. Rate
- Read Trans. Rate
- Write Trans. Rate

- Online Verify Rate
- Drive OperationOperating
- Rate
- Max Tag Count
- Average Tag Count

Determining the ordinate axis

The Y axis is a control object in the graphing feature in Performance Monitor because it determines value information conveyed in the graph. Most importantly, the axis of the ordinate is determined by selecting the maximum value on the Y-axis.

Table 8-11 shows the relationship between displayed items for selected objects and the maximum values on the Y axis. The three objects to which the displayed items belong are Port, RAID Groups DP Pools, and Volumes. The bolded values are default settings.

While the table is inclusive to the three object types, Note displayed items for volumes only extend between IO Rate and Write Hit in the table. Also, displayed items for RAID Groups DP Pools only extend between IO Rate and Write Trans. Rate in the table.

Table 8-11: Selectable Y axis values for RAID Group and DP Pool Information object

Displayed Items	Selectable Y Axis Values
IO Rate	10, 20, 50, 100, 200, 500, 1,000, 2,000, 5,000, 10,000, 20,000 , 50,000, 100,000, 150,000, 300,000
Read Rate	
Write Rate	
Read Hit	20, 50, 100
Write Hit	
Trans. Rate	0, 20, 50, 100, 200 , 500, 1,000, 2,000
Read Trans. Rate	
Write Trans. Rate	
CTL CMD IO Rate	10, 50, 100, 200, 500, 1,000, 2,000 , 5,000, 10,000, 20,000, 50,000
Data CMD IO Rate	10, 50, 100, 200, 500, 1,000, 2,000, 5,000, 10,000, 20,000 , 50,000
CTL CMD Trans. Rate	10, 50, 100 , 200, 500, 1,000, 2,000, 5,000, 10,000, 20,000, 50,000, 100,000, 150,000
Data CMD Trans. Rate	10, 20, 50, 100 , 200, 400
CTL CMD Time	10, 50, 100, 200, 500, 1,000, 2,000, 5,000, 10,000, 20,000, 50,000, 100,000 , 200,000, 500,000, 1,000,000, 5,000,000, 10,000,000, 60,000,000

Table 8-11: Selectable Y axis values for RAID Group and DP Pool Information object

Displayed Items	Selectable Y Axis Values
Data CMD Time	10, 50, 100, 200, 500, 1,000, 2,000, 5,000, 10,000, 20,000, 50,000, 100,000 , 200,000, 500,000, 1,000,000, 5,000,000, 10,000,000, 60,000,000
CTL CMD Max Time	10, 50, 100, 200, 500, 1,000, 2,000, 5,000, 10,000, 20,000, 50,000, 100,000 , 200,000, 500,000, 1,000,000, 5,000,000, 10,000,000, 60,000,000
Data CMD Max Time	
XCOPY Rate	10, 20, 50, 100, 200, 500, 1,000, 5,000, 10,000, 20,000, 50,000, 10,000, 150,000
XCOPY Time	100, 500, 1,000, 2,000, 5,000, 10,000, 20,000, 50,000, 100,000, 500,000, 1,000,000, 5,000,000, 10,000,000, 60,000,000
XCOPY Max Time	100, 500, 1,000, 5,000, 10,000, 20,000, 50,000, 100,000, 200,000, 500,000, 1,000,000, 2,000,000, 5,000,000, 10,000,000, 60,000,000
XCOPY Read Rate	10, 20, 50, 100, 200, 500, 1,000, 2,000, 5,000, 10,000, 20,000, 50,000, 100,000, 150,000
XCOPY Write Rate	
XCOPY Read Trans. Rate	10, 20, 50, 100, 200, 500, 1,000, 2,000
XCOPY Write Trans. Rate	

Table 8-12 details Y axis values for the RAID Groups DP Pools item.

Table 8-12: Selectable Y-axis values for objects, RAID groups DP Pools

Displayed Items	Selectable Y Axis Values
IO Rate	10, 20, 50, 100, 200, 500, 1,000, 2,000, 5,000, 10,000, 20,000, 50,000
Read Rate	
Write Rate	
Read Hit	20, 50, 100
Write Hit	
Trans. Rate	0, 20, 50, 100, 200, 500, 1,000, 2,000
Read Trans. Rate	
Write Trans. Rate	
XCOPY Rate	10, 20, 50, 100, 200, 500, 1,000, 5,000, 10,000, 20,000, 50,000, 100,000, 150,000
XCOPY Time	100, 500, 1,000, 2,000, 5,000, 10,000, 20,000, 50,000, 100,000, 500,000, 1,000,000, 5,000,000, 10,000,000, 60,000,000
XCOPY Max Time	100, 500, 1,000, 2,000, 5,000, 10,000, 20,000, 50,000, 100,000, 200,000, 500,000, 1,000,000, 2,000,000, 5,000,000, 10,000,000, 60,000,000
XCOPY Read Rate	10, 20, 50, 100, 200, 500, 1,000, 2,000, 5,000, 10,000, 20,000, 50,000, 100,000, 150,000
XCOPY Write Rate	

**Table 8-12: Selectable Y-axis values for objects, RAID groups
DP Pools**

Displayed Items	Selectable Y Axis Values
XCOPY Read Trans. Rate	10, 20, 50, 100, 200, 500, 1,000, 2,000
XCOPY Write Trans. Rate	

Table 8-13 details Y axis values for the volume item.

Table 8-13: Selectable Y-Axis values for Volume Information

Displayed Items	Selectable Y Axis Values
IO Rate	10, 20, 100, 200, 500, 1,000, 5,000, 10,000, 20,000, 50,000, 100,000, 150,000, 300,000
Read Rate	
Write Rate	
Read Hit	20, 50, 100
Write Hit	
Trans. Rate	0, 20, 50, 100, 200, 500, 1,000, 2,000
Read Trans. Rate	
Write Trans. Rate	
Max Tag Count	500, 1,000, 2,000, 5,000, 10,000, 20,000, 50,000, 100,000
Average Tag Count	
Data CMD IO Rate	10, 50, 100, 500, 1,000, 2,000, 5,000, 10,000, 20,000, 50,000
Data CMD Trans. Rate	10, 20, 50, 200, 200, 400
XCOPY Rate	10, 20, 50, 100, 200, 500, 1,000, 5,000, 10,000, 20,000, 50,000, 100,000, 150,000
XCOPY Time	100, 500, 1,000, 2,000, 5,000, 10,000, 20,000, 50,000, 100,000, 500,000, 1,000,000, 5,000,000
XCOPY Max Time	100, 500, 1,000, 5,000, 10,000, 50,000, 100,000, 200,000, 500,000, 1,000,000, 5,000,000, 10,000,000, 60,000,000
XCOPY Read Rate	10, 20, 50, 100, 200, 500, 1,000, 2,000, 5,000, 10,000, 20,000, 50,000, 100,000, 150,000
XCOPY Write Rate	
XCOPY Read Trans. Rate	10, 20, 50, 100, 200, 500, 1,000, 2,000
XCOPY Write Trans. Rate	

Table 8-14 details Y-axis values for Cache Information, Processor Information, Drive Information, Drive Operation Information, and Back-End Information.

Table 8-14: Y-Axis details for Cache, Drive, Drive Operation,

and Back-End information

Cache Information	Write Pending Rate Note	20, 50, 100
	Clean Queue Usage Rate Note	
	Middle Queue Usage Rate Note	
	Physical Queue Usage Rate Note	
	Total Queue Usage Rate	
Processor Information	Usage	
Drive Information Back-end Information	I/O Rate	10, 20, 50, 100, 200, 500, 1000, 2000, 5000, 10000, 20000, 50000
	Read Rate	
	Write Rate	
	Trans. Rate	10, 20, 50, 100, 200, 1000, 2000
	Read Trans. Rate	
	Write Trans. Rate	
	Online Verify Rate	10, 20, 50, 100, 200, 500, 1000, 2000, 5000, 10000, 20000, 50000
Drive Operation Information	Operating Rate	20, 50, 100
	Max Tag Count	
	Average Tag Count	
Back-end Information	I/O Rate	10, 20, 50, 100, 200, 500, 1000, 2000, 5000, 10000, 20000, 50000, 100,000
	Read Rate	
	Write Rate	
	Trans. Rate	10, 20, 50, 100, 200, 500, 1,000, 2,000
	Read Trans. Rate	
	Write Trans. Rate	
	Online Verify Rate	10, 20, 50, 100, 200, 500, 1,000, 2,000, 5,000, 10,000, 20,000 , 50,000

Saving monitored data

To save the settings you changed for Performance Monitor

1. Click the Options tab. Performance Monitor displays the Options Window that contains two sub tabs: Output Monitoring Data and Save Monitoring Data.
2. Click the Save Monitoring Data checkbox to place a check in the box.
3. Obtain your data and click **Stop**.
4. Click **Close** to exit the Options Window.

Exporting Performance Monitor information

To copy the monitored data to a CSV file

1. In the Performance Monitor window, click the **Option** tab.
2. Select the **Save Monitoring Data** checkbox.
3. Obtain your data, and click **Stop**.
4. Click the **Output CSV** tab and select the items you want to output.
5. Click **Output**. Performance Monitor displays the Output CSV Window as shown in [Figure 8-8](#).

Show Graph

Graph Property

Select the items that you want to show on the graph. After completing settings, click Start. Click Show Graph to show performance graph. If you also want to save the performance data to the CSV file, move to CSV Files tab after measuring data in the Graph tab.

Graph CSV Files

Array Name: HUS110_91100026
Serial Number: 91100026
Interval Time: 1 minute(s)
Preservation Period: From 2011/12/15 13:35 to 2011/12/15 14:34

* Time range to save in the CSV file:
From: 2011 / 12 / 15 13 : 35
To: 2011 / 12 / 15 14 : 34

Items to be saved in the CSV file:

- ☒ HUS110_91100026
- ☒ Controller 0
- ☐ Port Information
- ☐ RAID Group Information
- ☐ DP Pool Information
- ☐ Volume Information
- ☒ Cache Information
- ☐ Processor Information

Measurement Item Information

Item
☒ Write Pending Rate
☒ Clean Queue Usage Rate
☒ Middle Queue Usage Rate
☒ Physical Queue Usage Rate
☒ Total Queue Usage Rate

Get CSV File Get CSV File

* Required field

Close

Figure 8-8: Output CSV tab: Dynamic Provisioning valid

[Table 8-15](#) provides descriptions of objects displayed in the Output CSV Window.

Table 8-15: Descriptions of Output CSV tab objects

Displayed Items	Description
Array Unit	A name of the storage system from which the data was collected.
Serial Number	A serial number of the storage system from which the data was collected.
Output Time	Specifies the period when the data to be output is produced, using the From and To sliders.
Interval Time	The range of time between data collections.
Output Item	Checks the items you want to export.
Output Directory	Specifies a target directory to where the CSV file will be exported.

Once you have exported content to a CSV file, the files take default filenames each with a .CSV extension. The following tables detail filenames for each object type.

[Table 8-16](#) lists filenames for the Port object.

Table 8-16: CSV filenames: port object

List Items	CSV Filename
IO Rate	CTL0_Port_IORate.csv
Read Rate	CTL0_Port_ReadRate.csv
Write Rate	CTL0_Port_WriteRate.csv
Read Hit	CTL0_Port_ReadHit.csv
Write Hit	CTL0_Port_WriteHit.csv
Trans. Rate	CTL0_Port_TransRate.csv
Read Trans. Rate	CTL0_Port_ReadTransRate.csv
Write Trans. Rate	CTL0_Port_WriteTransRate.csv
CTL CMD IO Rate	CTL0_Port__CTL_CMD_IORate.csv
Data CMD IO Rate	CTL0_Port_Data_CMD_TransRate.csv
CTL CMD Trans. Rate	CTL0_Port_CTL_CMD_TransRate.csv
Data CMD Trans. Rate	CTL0_Port_data_CMD_Trans_Time.csv
CTL CMD Max Time	CTL0_Port_CTL_CMD_Max_Time.csv
Data CMD Max Time	CTL0_Port_Data_CMD_Max_Time.csv
XCOPY Rate	CTL0_Port_XcopyRate.csv
XCOPY Time	CTL0_Port_XcopyTime.csv
XCOPY Max Time	CTL0_Port_XcopyMaxTime.csv
XCOPY Read Rate	CTL0_Port_XcopyReadRate.csv
XCOPY Read Trans. Rate	CTL0_Port_XcopyReadTransRate.csv

[Table 8-17](#) details CSV filenames for list items for RAID Groups and DP Pool objects.

Table 8-17: CSV filenames: RAID groups and DP Pool objects

Object	List Items	CSV Filename
RAID Groups	IO Rate	CTL0_Rg_IORatenn.csv
	Read Rate	CTL0_Rg_ReadRatenn.csv
	Write Rate	CTL0_Rg_WriteRatenn.csv
	Read Hit	CTL0_Rg_ReadHitnn.csv
	Write Hit	CTL0_Rg_WriteHitnn.csv
	Trans. Rate	CTL0_Rg_TransRatenn.csv
	Read Trans. Rate	CTL0_Rg_ReadTransRatenn.csv
	Write Trans. Rate	CTL0_Rg_WriteTransRatenn.csv

Table 8-17: CSV filenames: RAID groups and DP Pool objects

Object	List Items	CSV Filename
DP Pools	IO Rate	CTL0_DPPool_IORatenn.csv
	Read Rate	CTL0_DPPool_ReadRatenn.csv
	Write Rate	CTL0_DPPool_WriteRatenn.csv
	Read Hit	CTL0_DPPool_ReadHitnn.csv
	Write Hit	CTL0_DPPool_WriteHitnn.csv
	Trans. Rate	CTL0_DPPool_TransRatenn.csv
	Read Trans. Rate	CTL0_DPPool_ReadTransRatenn.csv
	Write Trans. Rate	CTL0_DPPool_WriteTransRatenn.csv
	XCOPY Rate	CTL0_DPPool_XcopyRatenn.csv
	XCOPY Time	CTL0_DPPool_XcopyTimenn.csv
	XCOPY Max Time	CTL0_DPPool_XcopyMaxTimenn.csv
	XCOPY Read Rate	CTL0_DPPool_XcopyReadRatenn.csv
	XCOPY Read Trans. Rate	CTL0_DPPool_XcopyReadTransRatenn.csv
	XCOPY Write Rate	CTL0_DPPool_XcopyWriteRatenn.csv
	XCOPY Write Trans. Rate	CTL0_DPPool_XcopyWriteTransRatenn.csv

Table 8-18 details CSV filenames for list items associated with Volumes and Processor objects.

Table 8-18: CSV filenames: volumes and processor objects

Object	List Items	CSV Filename
Volume	IO Rate	CTL0_Lu_IORatenn.csv
	Read Rate	CTL0_Lu_ReadRatenn.csv
	Write Rate	CTL0_Lu_WriteRatenn.csv
	Read Hit	CTL0_Lu_ReadHitnn.csv
	Write Hit	CTL0_Lu_WriteHitnn.csv
	Trans. Rate	CTL0_Lu_TransRatenn.csv
	Read Trans. Rate	CTL0_Lu_ReadTransRatenn.csv
	Write Trans. Rate	CTL0_Lu_WriteTransRatenn.csv
	CTL CMD IO Rate	CTL0_Lu_CTL_CMD_IORatenn.csv
	Data CMD IO Rate	CTL0_Lu_CMD_TransRatenn.csv
	CTL CMD Trans. Rate	CTL0_Lu_CTL_CMD_TransRatenn.csv
	Data CMD Trans. Rate	CTL0_Lu_data_CMD_Trans_Timenn.csv
	XCOPY Rate	CTL0_Lu_XcopyRatenn.csv
	XCOPY Time	CTL0_Lu_XcopyTimenn.csv
	XCOPY Max Time	CTL0_Lu_XcopyMaxTimenn.csv
	XCOPY Read Rate	CTL0_Lu_XcopyReadRatenn.csv
	XCOPY Read Trans. Rate	CTL0_Lu_XcopyReadTransRatenn.csv
	XCOPY Write Rate	CTL0_LuXcopyWriteRatenn.csv

Table 8-18: CSV filenames: volumes and processor objects

Object	List Items	CSV Filename
	XCOPY Write Trans. Rate	CTL0_Lu_XcopyWriteTransRatenn.csv
Processor	Usage	CTL0_Processor_Usage.csv

Table 8-19 details CSV filenames for list items associated with Cache, Drive, and Drive Operation objects.

Table 8-19: CSV filenames: cache, drive, drive operation objects

Object	List Items	CSV Filename
Cache	Write Pending Rate (per partition)	CTL0_Cache_WritePendingRate.csv
		CTL0_CachePartition_WritePendingRate.csv
	Clean Usage Rate (per partition)	CTL0_Cache_CleanUsageRate.csv
		CTL0_CachePartition_CleanUsageRate.csv
	Middle Usage Rate (per partition)	CTL0_Cache_MiddleUsageRate.csv
		CTL0_CachePartition_MiddleUsageRate.csv
	Physical Usage Rate (per partition)	CTL0_Cache_PhysicalUsageRate.csv
		CTL0_CachePartition_PhysicalUsageRate.csv
	Total Usage Rate	CTL0_Cache_TotalUsageRate.csv
Drive	IO Rate	CTL0_Drive_IORatenn.csv
	Read Rate	CTL0_Drive_ReadRatenn.csv
	Write Rate	CTL0_Drive_WriteRatenn.csv
	Trans. Rate	CTL0_Drive_TransRatenn.csv
	Read Trans. Rate	CTL0_Drive_ReadTransRatenn.csv
	Write Trans. Rate	CTL0_Drive_WriteTransRatenn.csv
	Online Verify Rate	CTL0_Drive_OnlineVerifyRatenn.csv
Drive Operation	Operating Rate	CTL0_DriveOpe_OperatingRatenn.csv
	Max Tag Count	CTL0_DriveOpe_MaxtagCountnn.csv

Enabling performance measuring items

The Performance Measuring tool enables you to enable specific types of performance monitoring.

To access the Performance Measuring tool

1. Start Navigator 2 and log in. The Arrays window opens
2. Click the appropriate array.

- Click **Performance** and click **Monitoring**. The Monitoring - Performance Measurement Items window displays as shown in [Figure 8-9 on page 8-34](#).




Figure 8-9: Monitoring - Performance Measurement items

- Click on the Change Measurement Name Button. The Change Measurement Items dialog box displays with six performance statistics. [Table 8-20](#) describes each of the performance statistics.

Table 8-20: Performance statistics

Item	Description
Port Information	Displays information about the port.
RAID Group, DP VOL and Volume Information	Displays information about RAID groups, Dynamic provisioning pools and volumes.
Cache Information	Displays information about cache on the storage system.
Processor Information	Displays information about the storage system processor.
Drive Information	Displays information about the administrative state of the storage system disk drive.
Drive Operation Information	Displays information about the operation of the storage system disk drive.
Back-end Information	Displays information about the back-end of the storage system.
Management Area Information	Displays cache hit rates and access count of management data in stored drives acquired by the array. This information is used only for acquiring performance data. This information cannot be graphed.

The default setting for each of the performance statistics is Enabled (acquire). If one of the item settings is Disabled, the automatic load balance function does not work. The load balance function failure occurs

because the internal performance monitoring does not perform. To ensure that load balancing works, set all performance statistics to Enabled.

5. To disable one of the performance statistics, click in the checkbox to the right of the statistic to remove the checkmark.

Working with port information

The storage system acquires port I/O and data transfer rates for all Read and Write commands received from a host. It can also acquire the number of commands that made cache hits and cache-hit rates for all Read and Write commands.

Working with RAID Group, DP Pool and volume information

The storage system acquires all array RAID group/DP pool information of volumes. It also acquires the I/O and data transfer rates for all Read and Write commands received from a host. In addition, it also acquires the number of commands that made cache hits and cache-hit rates for all Read and Write commands.

Working with cache information

The storage system displays the ratio of data in a write queue to the entire cache and utilization rates of the clean, middle, and physical queues.

The clean queue consists of a number of segments of data that have been read from the drives and exist in cache.

The middle queue consists of a number of segments that retain write data, have been sent from a host, exist in cache, and have no parity data generated.

The physical queue consists of a number of segments that retain data, exist in cache, and have parity data generated, but not written to the drives.

For the Cache Hit parameter of the Write command, a *hit* is a response to the host that has completed a Write to the Cache (Write-After). A miss is a response to the host that has completed a Write to the Drive (Write-Through). When the cache use volume is large or the battery unit fails, Write-Through is more likely.

Working with processor information

The storage system can acquire and display the utilization rate for each processor.

Troubleshooting performance

If there are performance issues, refer to [Figure 8-10](#) for information on how to analyze the problem.




Figure 8-10: Performance Optimization Analysis

Performance imbalance and solutions

Performance imbalance can occur between controllers, ports, RAID groups, and back-ends.

Controller imbalance

The controller load information can be obtained from the processor operation rate and its cache use rate.

The volume load can be obtained from the I/O and transfer rate of each volume.

When the loads between controllers differ considerably, the array disperses the loads (load balancing). However, when this does not work, change the volume by using the tuning parameters.

Port imbalance

The port load in the array can be obtained from the I/O and transfer rate of each port.

If the loads between ports differ considerably, transfer the volume that belongs to the port with the largest load, to a port with a smaller load.

RAID group imbalance

The RAID group load in the array can be obtained from the I/O and transfer rate of the RAID group information.

If the load between RAID group varies considerably, transfer the volume that belongs to the RAID group with the largest load, to a Raid group with a smaller load.

Back-end imbalance

The back-end load in the array can be obtained from the I/O and transfer rate of the back-end information.

If the load between back-ends varies considerably, transfer the RAID group and volume with the largest load, to a back-end with a smaller load. For the back-end loop transfer, you can change the owner controller of each volume; however controller imbalance can occur.

Dirty Data Flush

You may require that your storage system has the best possible I/O performance at all times. When ShadowImage or SnapShot environments are introduced, the system's internal resource allocation to support the current task load may not meet the your performance objectives. The switch intends to support the best possible performance requirements while supporting ShadowImage and SnapShot.

HDS provides a system tool that reprioritizes the internal I/O in the system processor in favor of a production I/O. This feature is the Dirty Data Flush.

Dirty Data Flush is a mode that improves the read response performance when the I/O load is light. If the write I/O load is heavy, a timeout may occur because not enough dirty jobs exist to process the conversion of dirty data as the number of jobs is limited to one. So the mode should be changed when the I/O load is light.

The mode is effective when the following conditions are met:

- The new mode is enabled while one of the following features is enabled:
 - Modular Volume Migration
 - SnapShot
 - ShadowImage
- Only volumes from RAID0, RAID1, and RAID1+0 exist in the system.
- Only volumes from SAS drives exist in the system.
- Remote replications such as TrueCopy and TrueCopy Extended Distance are disabled.

To set the mode, perform the following steps:

1. Go to the Array Home screen.
2. In the Navigation Tree, click **Performance**. HSNM2 displays the Performance window as shown in [Figure 8-11](#).




Figure 8-11: Performance window

3. Click **Tuning Parameters**. HSNM2 displays the Tuning Parameters window as shown in Figure 8-12.




Figure 8-12: Tuning Parameters window

4. Click **System Tuning**. HSNM2 displays the System Tuning window as shown in Figure 8-13. Note that the Dirty Data Flush Number Limit field in the System Tuning list has a setting of Disabled, the default value.




Figure 8-13: System Tuning window

5. In the System Tuning list, click on the **Edit System Tuning Parameters** button to display the Edit System Tuning Parameters dialog box as shown in [Figure 8-14](#).




Figure 8-14: Edit System Tuning Parameters dialog box

6. In the Dirty Data Flush Number Limit radio button box, click **Enable** to change the setting from Disabled to Enabled. Note that the setting is a toggle between the Disabled and Enabled radio buttons.
7. Click **OK**. HSNM2 displays the System Tuning window with the Enabled setting in the Dirty Data Flush Number Limit field.

SNMP Agent Support

This chapter describes the Hitachi SNMP Agent Support function, a software process that interprets Simple Network Management Protocol (SNMP) requests, performs the actions required by that request, and produces an SMNP reply.

The key topics in this chapter are:

- ❑ [SNMP overview](#)
- ❑ [Supported configurations](#)
- ❑ [Supported configurations](#)
- ❑ [Hitachi SNMP Agent Support procedures](#)
- ❑ [Operational guidelines](#)
- ❑ [MIBs](#)
- ❑ [Additional resources](#)

SNMP overview

SNMP is an open Internet standard for managing networked devices. SNMP is based on the manager/agent model consisting of:

- A manager
- An agent
- A database of management information
- Managed objects, such as the Hitachi modular storage arrays
- The network protocol

The manager is the computer or workstation that lets the network administrator perform management requests. The agent acts as the interface between the manager and the physical devices being managed, and makes it possible to collect information on the different objects.

The SNMP agent provided for the HUS systems is designed to provide SAN information to MIB browsers that support SNMP v1.X Using Hitachi SNMP Agent Support, you can monitor inventory, configuration, service indicators, and environmental and fault reporting on Hitachi modular storage arrays using SNMP network management systems such as IBM Tivoli, CA Unicenter, and HP OpenView.

SNMP features

- **Availability of MIBs** - All SNMP-compliant devices include a specific text file called a Management Information Base (MIB). A MIB is a collection of hierarchically organized information that defines what specific data can be collected from that particular device.
- **Common language of network monitoring** - SNMP (Simple Network Management Protocol) is the common language of network monitoring—it is integrated into most network infrastructure devices today, and many network management tools include the ability to pull and receive SNMP information.
- **Data collection services** - SNMP extends network visibility into network-attached devices by providing data collection services useful to any administrator. These devices include switches and routers as well as servers and printers. The following information is designed to give the reader a general understanding of what SNMP is, the benefits of SNMP, and the proper usage of SNMP as part of a complete network monitoring and management solution.
- **Standard application layer protocol** - The Simple Network Management Protocol (SNMP) is a standard application layer protocol (defined by RFC 1157) that allows a management station (the software that collects SNMP information) to poll agents running on network devices for specific pieces of information. What the agents report is dependent on the device. For example, if the agent is running on a server, it might report the server's processor utilization and memory usage. If the agent is running on a router, it could report statistics such as interface utilization, priority queue levels, congestion notifications,

environmental factors (i.e. fans are running, heat is acceptable), and interface status.

- **Protocol for device information access** - SNMP is the protocol used to access the information on the device the MIB describes. MIB compilers convert these text-based MIB modules into a format usable by SNMP management stations. With this information, the SNMP management station queries the device using different commands to obtain device-specific information.
- **Small command set for information retrieval** - There are three principal commands that an SNMP management station uses to obtain information from an SNMP agent:
- **Reporting and analysis of device status** - The SNMP management console reviews and analyzes the different variables maintained by that device to report on device uptime, bandwidth utilization, and other network details. However, the switch maintains a count of the discarded error frames and this counter can be retrieved via an SNMP query.

SNMP benefits

The following are SNMP benefits:

- **Distributed model of management** - Enables a centralized, distributed way to manage nodes on a network across multiple domains. This provides an efficient way to manage devices where one administrator can have visibility to many locations.
- **System portability** - Enables portability to other vendors to develop applications to the main platform.
- **Industry-wide common compliance** - SNMP delivers management information in a common, non-proprietary manner, making it easy for an administrator to manage devices from different vendors using the same tools and interface. Its power is in the fact that it is a standard: one SNMP-compliant management station can communicate with agents from multiple vendors, and do so simultaneously. Illustration 1 shows a sample SNMP management station screen displaying key network statistics.
- **Data transparency** - The type of data that can be acquired is transparent. For example, when using a protocol analyzer to monitor network traffic from a switch's SPAN or mirror port, physical layer errors are invisible. This is because switches do not forward error packets to either the original destination port or to the analysis port.

Environments and requirements

Item	Description
Environments	Firmware: Version 0915/B or more. Hitachi Storage Navigator Modular 2: Version 21.50 or more is required for the management PC When using the accumulated operating time (sysUpTime) of the SNMP agent, the firmware of 0940/A or more is required.
Requirements	The license key for the SNMP Agent Support Function.

SNMP task flow

The following details the task flow of the SNMP process:

1. You determine that you want to establish an environment for network management of your storage system on the network. selected users need to have access to your storage system and that all other users should be blocked from access to it.
2. You identify all users for access as network managers.
3. Configure the license for SNMP.
4. Install and enable SNMP.

SNMP along with the associated Management Information Base (MIB), encourage trap-directed notification.

The idea behind trap-directed notification is that if a manager is responsible for a large number of devices, and each device has a large number of objects, it is impractical for the manager to poll or request information from every object on every device. The solution is for each agent on the managed device to notify the manager without solicitation. It does this by sending a message known as a trap of the event.

After the manager receives the event, the manager displays it and can choose to take an action based on the event. For instance, the manager can poll the agent directly, or poll other associated device agents to get a better understanding of the event.

Trap-directed notification can result in substantial savings of network and agent resources by eliminating the need for frivolous SNMP requests. However, it is not possible to totally eliminate SNMP polling. SNMP requests

are required for discovery and topology changes. In addition, a managed device agent can not send a trap, if the device has had a catastrophic outage.




Figure 9-1: SNMP request, response, and trap generation

SNMP versions

Like other Internet standards, SNMP is defined by a number of Requests for Comments (RFCs) published by the Internet Engineering Task Force (IETF).

There are three SNMP versions that define approved standards:

- SNMP version 1 (SNMP v1)
- SNMP version 2 (SNMP v2)
- SNMP version 3 (SNMP v3)

SNMP v1 was introduced in 1988. SNMPv2 followed in 1993 and included further protocol operations and data types for additional security. Limitations in the security model led to the SNMPv2c standard. Experimental versions, known as SNMPv2usec and SNMPv2*, followed, but have not been widely adopted. SNMPv3, defined in 1999, calls out the SNMP management framework supporting pluggable components, including security.

For more information about SNMP standards, see [Additional resources on page 9-62](#). The SNMP Agent Support Function complies with SNMP v1. Hitachi modular storage arrays support SNMP v2.

SNMP managers and agents

SNMP is a network protocol that allows networked devices to be managed remotely by a network management station (NMS), also called a *manager*. To be managed, a device must have an SNMP agent associated with it.

The purpose of the SNMP agent is to:

- Receive requests for data representing the state of the device from the manager and provide an appropriate response.
- Accept data from the manager to enable control of the device state.
- Generate SNMP traps, which are unsolicited messages sent to one or more selected managers to signal significant events relating to the device.

Management Information Base (MIB)

The SNMP agent itself does not define which information a managed device should offer. Rather, the agent uses an extensible design, where the available information is defined by a Management Information Base (MIB).

The MIB is a tree-like data dictionary used to assemble and interpret SNMP messages. The manager accesses the MIB content using Get and Set operations.

For example, if an SNMP manager wants to know the value of an object, such as the status of a Hitachi modular storage array controller and drive, it assembles a Get packet that includes the object identifier (OID) for each object of interest.

- In response to a Get operation, the agent provides data maintained either locally or directly from the managed device.
- In response to a Set operation, the agent typically performs an action that affect the state of either itself or the managed device.



NOTE: MIBs are defined using Abstract Syntax Notation number one (ASN.1), an international standard notation that describes data structures for representing, encoding, transmitting, and decoding data. Discussion of ASN.1 exceeds the scope of this chapter. For more information, refer to the IETF Web site at <http://www.ietf.org>.

Object identifiers (OIDs)

An OID consists of a hierarchically arranged sequence of numbers separated by decimal points that defines a unique name space. Each assigned number has an associated text name. The numeric form is used within SNMP protocol transactions, while the text form is used in user interfaces to enhance readability.

Figure 9-2 shows an example of the Hitachi SNMP Agent Support MIB-II hierarchy that defines all OIDs residing below the series of integers beginning with 1.3.6.1.2.1.




Figure 9-2: Example of an OID

SNMP command messages

SNMP is a packet-oriented protocol that uses the following basic messages, or *protocol data units* (PDUs), for communicating between the SNMP manager and SNMP agent.

- *Get*
- *GetNext*
- *GetResponse*
- *GetNextResponse*
- *Set*
- *Trap*

The SNMP manager sends a *Get* or *GetNext* message to request the status of a managed object. The agent's *GetResponse* message contains the requested information if managed or an error indication as to why the request cannot be processed.

The SNMP manager sends a *Set* to change a Managed object to a new value. The agent's *GetResponse* message confirms the change if allowed or an error indication as to why the change cannot be made.

The agent sends a *Trap* when a specific event occurs. The *Trap* message allows the agent to spontaneously inform the manager about an important event.

[Figure 9-3](#) shows the core PDUs that the SNMP Agent Support Function supports and [Table 9-1 on page 9-8](#) summarizes them.




Figure 9-3: Core PDUs supported by Hitachi SNMP Agent Support

Table 9-1: Supported core PDUs

PDU	Description
<i>GetRequest</i>	A manager-to-agent request to retrieve the value of a MIB object. A Response with current values is returned.
<i>GetResponse</i>	If an error in a request from the SNMP manager is detected, the storage array sends a <i>GetResponse</i> to the manager, together with the error status, as shown in Table 9-2 on page 9-8 .
<i>GetNextRequest</i>	A manager-to-agent request to discover available MIB objects continuously. The entire MIB of an agent can be walked by iterative application of <i>GetNextRequest</i> , starting at OID 0.
<i>GetNextResponse</i>	SNMP agent response to a <i>GetNextRequest</i> operation.
<i>Trap</i>	An asynchronous notification from the agent to the manager. If an event occurs, the agent sends a Trap to the manager, regardless of SNMP manager's request. A trap notifies the manager about status changes and error conditions that may not be able to wait until the next interrogation cycle. The SNMP Agent Support Function supports standard and extended traps (see SNMP traps on page 9-9).

[Table 9-2](#) details the status of SNMP errors.

Table 9-2: SNMP error status

Error Status Code	Description
noError (0)	Normal operation, no error detected. The requested MIB object value is placed in the SNMP message to be sent.
tooBig (1)	SNMP message is too large (exceeds 484 bytes) to contain the operation result. To avoid this problem, configure the SNMP manager to send messages that request a response less than 485 bytes.

Table 9-2: SNMP error status

Error Status Code	Description
noSuchName (2)	Requested MIB object could not be found. The GetNextRequest specified was received. However, the requested MIB object value is not in the SNMP message and the requested process (SetRequest) did not execute.
badValue (3)	N/A (does not occur)
readOnly (4)	N/A (does not occur)
genErr (5)	The requested operation cannot be performed for a reason other than one of the reasons above.

If the following errors are detected in the SNMP manager's request, the Hitachi modular storage array does not respond.

- The community name does not match the setting. The array does not respond and sends the standard trap Authentication Failure (incorrect community name) to the manager.
- The SNMP request message exceeds 484 bytes. The array cannot send or receive SNMP messages larger than 484 bytes, and does not respond to received SNMP messages that exceed this limit.

SNMP traps

Traps are the method an agent uses to report important, unsolicited information to a manager. Trap responses are not defined in SNMP v1, so each managed element must have one or more trap receivers defined for the trap to be effective.

In SNMP v2 and higher, the concept of a trap was extended using another SNMP message called Inform. Like a trap, an Inform message is unsolicited. However, Inform enables a manager running SNMP v2 or higher to send a trap to another manager. It can also be used by an SNMP v2 or higher managed node to send an SNMP v2 trap. The receiving node sends a response, telling the sending manager that the receiving manager received the Inform message. Both messages are sent on UDP Port 161.

The SNMP Agent Support Function reports SNMP v1 standard traps and SNMP v2 extended traps. The following list shows the standard traps that are supported.

- Start up SNMP Agent Support Function (when installing or enabling SNMP Agent Support Function)
- Changing SNMP Agent Support Function setting
- Incorrect community name when acquiring MIB information

[Figure 9-4](#) shows an example of an SNMP trap within the Hitachi modular storage array. For more information, see [SNMP traps on page 9-9](#).




Figure 9-4: Example of a trap in a Hitachi modular storage array

The following list shows the extended traps that are supported. The superscripted numbers correspond to the numbers in the legend following the table.

• Own controller failure ^{1 2}	• Path blockade ⁴	• Failure (TrueCopy Extended)
• Drive blockage (data drive)	• Host connector failure	• Failure (Modular Volume Migration)
• Fan failure	• Interface board failure	• Data pool threshold over
• Power supply failure	• Host I/O module failure	• Data pool no free
• Battery failure	• Drive I/O module failure	• Cycle time threshold over
• Cache memory failure	• Management module failure	• Volume data is not recoverable (multiple failures of drives) ⁵
• UPS failure	• Side Card failure	• Replace the air filter of DC power supply
• Cache backup circuit failure	• Controller failure by related parts	• DP pool consumed capacity early alert
• Slave controller failure ²	• Additional battery failure	• DP pool consumed capacity depletion alert
• Warning disk array ³	• Failure (ShadowImage)	• DP pool consumed capacity over
• Spare drive failure	• Failure (SnapShot)	• Over provisioning warning threshold
• Enclosure controller failure	• Failure (TrueCopy)	• Over provisioning limit threshold
• Over replication depletion alert threshold	• Over replication data released threshold	• Over SSD write count threshold
• SSD write count exceeds threshold	• The HDD mounting location error has occurred in the DBW	• The destage of the page management information failed in the DT pool relocation processing.

Legend:

- 1: Depending on the contents of the failure, this trap might not be reported.
- 2: If a controller blockage occurs, the storage array issues Traps that show the blockage. The controller blockage may recover automatically, depending on the cause of the failure.
- 3: The Trap that shows the warning status of the storage array may be issued via preventive maintenance, periodic part replacement, or field work conducted by Hitachi service personnel.

- 4: Path blockage is reported when TrueCopy or TrueCopy Extended is enabled.
- 5: If multiple failures occur in drives and the volume data in the RAID group is not recoverable, a Trap is reported. For example, if a failures occurs in three drives in a RAID 6 configuration (two drives in RAID 5), a Trap is issued.

Supported configurations

The SNMP Agent Support Function can be used in two configurations.

- Direct-connect — where a local computer or workstation acting as an SNMP manager is directly connected to the Hitachi modular storage array being managed within a private Local Area Network (LAN). [Figure 9-5](#) shows an example of this configuration.
- Public network — where gateways allow a remote computer or workstation acting as an SNMP manager to connect to the Hitachi modular storage array being managed. [Figure 9-6](#) shows an example of this configuration.

Both configurations support 10BaseT, 100BaseT, and 1000BaseT connections to Hitachi modular storage arrays over twisted-pair cable.




Figure 9-5: Example of a direct connect configuration




Figure 9-6: Example of a public network configuration

Frame types

The SNMP Agent Support Function supports Ethernet Version 2 frames (IEEE 802.3 frames, etc.) only. Other frames are not supported.

License key

The SNMP Agent Support Function requires a license key before it can be used. To obtain the required license key, please contact your Hitachi representative.


Installing Hitachi SNMP Agent Support

After obtaining a license key, use the following procedure to install the SNMP Agent Support Function.



NOTE: Hitachi SNMP Agent Support can also be installed from a command line. Refer to the *Hitachi Unified Storage Command Line Interface Reference Guide*.

1. Start Navigator 2 and log in as a registered user.
2. From the Arrays page, check the check box in the left column that corresponds to the Hitachi modular storage array on which you want to install the SNMP Agent Support Function.
3. At the bottom of the page, click **Show & Configure Array**.
4. Under **Common Array Task**, click the **Install License** icon:



The Install License page appears.

Figure 9-7: Install License - License Property dialog box

5. Perform one of the following steps at the **Install with** field:
 - To install the option using a key file, click **Key File**, and either enter the path where the key file resides or click the **Browse** button and select the path where the key file resides.
 - To install the option using a key code, click **Key Code** and enter the key code in the field provided.
6. Click **OK**.
7. When the confirmation page appears, click **Confirm**.
8. When the next page tells you that the license installation was complete, click **Close**.

This completes the procedure for installing Hitachi SNMP Agent Support. Proceed to [Hitachi SNMP Agent Support procedures](#), below, to confirm that Hitachi SNMP Agent Support is enabled.

Hitachi SNMP Agent Support procedures

The following sections describe how to:

- Prepare the SNMP manager for Hitachi SNMP Agent Support. See [Preparing the SNMP manager](#), below.
- Prepare the Hitachi modular storage array for Hitachi SNMP Agent Support. See [Preparing the Hitachi modular storage array](#), below.
- Confirm your setup. See [Confirming your setup on page 9-25](#).

Preparing the SNMP manager

To prepare the SNMP manager for use with Hitachi SNMP Agent Support

1. Provide the SNMP manager with the MIB definition file supplied with the Hitachi SNMP Agent Support function. For more information, refer to the documentation for your SNMP manager.
2. Register the Hitachi modular storage array with the SNMP manager. For more information, refer to the documentation for your SNMP manager.

Preparing the Hitachi modular storage array

To prepare the Hitachi modular storage array for use with Hitachi SNMP Agent Support

1. Use Navigator 2 to configure the array's LAN settings, such as the IP address, subnet mask, and default gateway. For more information, refer to the *AMS Installation, Upgrade, and Routine Operations Guide*.
2. Confirm that the SNMP Agent Support Function is enabled. See [Hitachi SNMP Agent Support procedures on page 9-15](#).
3. Create the following SNMP environment information files:
 - An operating environment file named `Config.txt`. This file contains the IP address and community information where the SNMP manager can send traps. See [Creating an operating environment file on page 9-16](#).
 - A storage array name file named `Name.txt`. This file contains the names of the Hitachi modular storage arrays to be managed. See [Creating a storage array name file on page 9-22](#).



NOTE: Hitachi modular storage arrays with dual controllers require only one operating environment file and one storage array name file. You cannot have separate environment information files for each controller.

4. Using Navigator 2, take the SNMP environment information file created in step 3 and register it with the storage array. See [Registering the SNMP environment information on page 9-22](#).

Creating an operating environment file

The operating environment file `Config.txt` is a text file you create using a text editor such as Notepad or WordPad. [Figure 9-8](#) and [Figure 9-9](#) show examples of this file using different IP addressing methods. Instructions for creating this file appear after the figures.

See step 1 →	INITIAL sysContact "Taro Hitachi"
See step 2 →	INITIAL sysLocation "Computer Room A on Hitachi STR HSP 10F north"
See step 3 →	COMMUNITY tagmastore ALLOW ALL OPERATIONS
See step 4 →	MANAGER 123.45.67.89 SEND ALL TRAPS TO PORT 162 WITH COMMUNITY "HITACHI DF800"
→	MANAGER 123.45.67.90 SEND ALL TRAPS TO PORT 162 WITH COMMUNITY "HITACHI DF800"

Figure 9-8: Sample file using IPv4 addressing

See step 1 →	INITIAL sysContact "Taro Hitachi"
See step 2 →	INITIAL sysLocation "Computer Room A on Hitachi STR HSP 10F north"
See step 3 →	COMMUNITY tagmastore ALLOW ALL OPERATIONS
See step 4 →	MANAGER 123.45.67.89 SEND ALL TRAPS TO PORT 162 WITH COMMUNITY "HITACHI DF800"
→	MANAGER 2001::1::20a:87ff:fec6:1928 SEND ALL TRAPS TO PORT 162 WITH COMMUNITY "HITACHI DF800"

Figure 9-9: Sample file using IPv6 addressing

To create the operating environment settings file Config.txt

1. Add a sysContact value by adding a line beginning with **INITIAL**. See the following example and the description in [Table 9-3 on page 9-20](#).

INITIAL sysContact user set information

2. Add a sysLocation value by adding a line beginning with **INITIAL**. See the following example and the description in [Table 9-3 on page 9-20](#).

INITIAL sysLocation user set information

When entering the information in steps 1 and 2:

- Do not exceed 255 alphanumeric characters.
 - To add special characters, such as a space, tab, hyphen, or quotation mark, enclose them in double quotation marks (for example "-").
 - Do not type line feeds when entering this information.
3. Below the sysContact value, add a line beginning with **COMMUNITY** to specify the community name with which the disk array allows receiving of requests. See the following example and the description in [Table 9-3 on page 9-20](#).

COMMUNITY community name ALLOW ALL OPERATIONS

When entering the community name:

- If these two lines are omitted, the Hitachi modular storage array accepts all community names.
 - Enter the community name using alphanumeric characters only.
 - To add special characters, such as a space, tab, hyphen, or quotation mark, enclose them in double quotation marks (for example "-").
 - Do not type line feeds when entering this information.
4. Port setting of the transmitting agency:

The disk array issues traps from port number 161. However, you need to issue traps from the dynamic port (optional port number 49152 to 65535), add the following line to the environment setting file.

SEND ALL TRANS FROM DYNAMIC PORT

```

INITIAL sysContact "Taro Hitachi"

INITIAL sysLocation "Computer Room A on Hitachi STR HSP 10F north"

COMMUNITY tagmastore
ALLOW ALL OPERATIONS

SEND ALL TRANS FROM DYNAMIC PORT

MANAGER 123.45.67.89
SEND ALL TRAPS TO PORT 162
WITH COMMUNITY "HITACHI DF800"

MANAGER 2001::1::20a:87ff:fec6:1928
SEND ALL TRAPS TO PORT 162
WITH COMMUNITY "HITACHI DF800"

```

5. Below the community name, specify up to three SNMP managers to whom the disk array will issue traps. Each line begins with **MANAGER**. If specifying more than one SNMP manager, use a line feed to separate the managers. See the following example and the description in [Table 9-3 on page 9-20](#).

```

MANAGER SNMP manager IP address
SEND ALL TRAPS TO PORT Port No.
WITH COMMUNITY Community name

MANAGER SNMP manager IP address
SEND ALL TRAPS TO PORT Port No.
WITH COMMUNITY Community name

```


When specifying SNMP managers:

- Enter the IP address for the object SNMP manager. Do not specify a host name. IP addresses can be entered in IPv4 or IPv6 format. Omit leading zeros in the IP address (to specify the IP address 111.022.003.055, for example, enter 111.22.3.55).
- Enter the UDP destination port number to be used when sending a trap to the SNMP manager. Typically, SNMP managers use the well-known port number 162 to receive traps.
- Enter a community name that will be contained in SNMP messages when traps are sent. Use alphanumeric characters only. To add special characters, such as a space, tab, hyphen, or quotation mark, enclose them in double quotation marks (for example "-").
- This information cannot contain line feeds.
- If the community name information does not contain a line beginning with **WITH COMMUNITY**, add **public** to the community name.

6. Setting sysUpTime

The accumulated time (sysUpTime) since the SNMP agent started is set to 0 by default. However, when setting the accumulated time for sysUpTime, add the following line to the environment setting file:

```
SET SYSUPTIME
```

The SNMP agent starts at the time of starting the array, rebooting the controller, and enabling the SNMP function. If you disable the SNMP function and then enable it, the time starts to be measured when the function is enabled.

```
INITIAL sysContact "Taro Hitachi"

INITIAL sysLocation "Computer Room A on Hitachi STR HSP 10F north"

COMMUNITY tagmastore
ALLOW ALL OPERATIONS

SET SYSUPTIME

MANAGER 123.45.67.89
SEND ALL TRAPS TO PORT 162
WITH COMMUNITY "HITACHI DF800"

MANAGER 2001::1::20a:87ff:fec6:1928
SEND ALL TRAPS TO PORT 162
WITH COMMUNITY "HITACHI DF800"
```

7. LAN port link-up determination when issuing a trap.

A trap is normally issued by the port of the controller which detected a failure. However, when issuing a trap, it is possible to determine whether the LAN port is linked up and issue the trap from the linked up port of the controller. To determine the linked-up LAN port, add the following line to the environment setting file:

LAN PORT CHECK

The following example shows and a setup of the LAN port link-up determination when issuing a trap.

```
INITIAL sysContact "Taro Hitachi"

INITIAL sysLocation "Computer Room A on Hitachi STR HSP 10F north"

COMMUNITY tagmastore
ALLOW ALL OPERATIONS

LAN PORT CHECK

MANAGER 123.45.67.89
SEND ALL TRAPS TO PORT 162
WITH COMMUNITY "HITACHI DF800"

MANAGER 2001::1::20a:87ff:fec6:1928
SEND ALL TRAPS TO PORT 162
WITH COMMUNITY "HITACHI DF800"
```



NOTE: The operating environment settings file cannot exceed 1,140 bytes. If the community name is less than 10 characters, the total length of the sysContact, sysLocation, and sysName values should not exceed 280 characters. Otherwise, all of the objects in the MIB-II system group cannot be obtained with one *GET* request. Keeping the total length of these values to less than 280 also prevents *tooBig* error messages from being generated.

Table 9-3 details SNMP operation environment file items.

Table 9-3: Operation environment file

Item	Description	Comments
sysContact (MIB information)	Manager information for contact (name, department, extension number., and so on)	Internal object value of MIB-II system group in ASCII form, not exceeding 255 characters (optional item)
sysLocation (MIB information)	Location where the device is installed.	
Community information setting (MIB information)	Name of the community permitted access.	A number of names of the community can be set. (optional item)
Transmitting agency port setting	Trap agency port setting	Default port is 161. Only one set can be set. (Item can be omitted).
sysUpTime	If using this function, add the SET SYSUPTIME line.	If this is not set, sysUpTime is fixed to 0.
LAN port link-up determination setting when issuing Trap	If using this function, add the LAN PORT CHECK line.	If this is not set, the trap is issued from the controller that detected the failure.

Table 9-3: Operation environment file

Item	Description	Comments
Trap sending (Trap report)	Setting of information for sending a trap: • Destination manager IP address • Destination port number • Community name given to a trap	Several combinations of information can be set. (Required item)

Creating a storage array name file

The storage array name file named `Name.txt` is a text file you create using a text editor such as Notepad or WordPad. [Table 9-4](#) lists the contents of this file.

Table 9-4: Storage array name file

Item	Description	Comments
sysName	Name of the Hitachi modular storage array to be managed	Internal object value of MIB-II system group in ASCII character string, not to exceed 255 characters. Example: AMS-01 Hitachi Disk Array

Observe the following guidelines:

- Use only alphanumeric characters:
- Do not use line feeds in this file. No line feed is necessary at the end of a sentence.
- To set the value of *sysName*, register the information continuously. Since the entire contents of this file are recognized as the *sysName* value, the file should not exceed 255 characters.

Registering the SNMP environment information

To register the SNMP environment information file

1. Start Navigator 2 and log in as a registered user.
2. From the Arrays page, check the check box in the left column that corresponds to the Hitachi modular storage array on which you will set up SNMP Agent Support Function.
3. At the bottom of the page, click **Show & Configure Array**.
4. In the center pane, under **Settings**, click **SNMP Agent**. The SNMP Agent page appears.




Figure 9-10: SNMP Agent window

5. Click **Edit SNMP Settings**. The Edit SNMP Settings screen appears.




Figure 9-11: Edit SNMP Settings window

6. Next to **Environment Settings**, click either **Enter SNMP settings manually** or **Load from file**.
 - If you clicked **Enter SNMP settings manually**, enter the SNMP registration information directly in the screen and see [Creating an operating environment file on page 9-16](#).

- If you select the **Load from file**, either enter the path to the SNMP environmental information file named config.txt or click the **Browse** button and select the path to this file.
7. Next to **Array Name**, click either **Enter array name manually** or **Load from file**.
 - If you clicked **Enter array name manually**, enter the name of the array and see [Creating a storage array name file on page 9-22](#).
 - If you select the **Load from file**, either enter the path to the SNMP environmental information file named config.txt or click the **Browse** button and select the path to this file.
 8. Click **OK**. The following confirmation message confirms that the settings are complete.
 9. Click **Close**.

Registering the SNMP environment information

After you register the SNMP information in the Hitachi modular storage array, refer to that information.

To refer to registered SNMP information

1. Start Navigator 2 and log in as a registered user.
2. From the Arrays page, check the check box in the left column that corresponds to the Hitachi modular storage array on which you will set up SNMP Agent Support Function.
3. At the bottom of the page, click **Show & Configure Array**.
4. Click the **SNMP Agent** icon in the Alert Settings of the tree view. The SNMP Agent dialog box appears, with the SNMP environment information displayed.




Figure 9-12: SNMP Agent window

Confirming your setup

After you set up the Hitachi modular storage system and SNMP manager, check for a connection between the storage system and SNMP manager.

To check for a connection between the storage system and SNMP manager

1. Check for a *Trap* connection by disabling the SNMP Agent Support function and then enabling it again (see [Hitachi SNMP Agent Support procedures on page 9-15](#)). Confirm that the standard trap *coldStart* was received at all SNMP managers that have been configured as trap receivers in the SNMP environment information file *Config.txt*.

If you cannot perform the above operation, set the SNMP environment information file again. Check that a standard trap, “warmStart,” has been received by all SNMP managers which have been set as a trap receiver in the SNMP environment information file (*Config.txt*).

2. Perform a *REQUEST* connection check by sending a MIB-supported *GET* request from all SNMP managers to the Hitachi modular storage array. Confirm that the array responds.

If the results of steps 1 and 2 succeed, it means all SNMP managers can communicate with the array via SNMP.

To respond to a failure of the above procedure

1. Obtain MIB information (*dfRegressionStatus*) periodically. This MIB value is set to 0 when there are no failures.
2. If an error occurs that results in a trap, the Hitachi modular storage array reports the error to the SNMP manager.

This trap lets you detect Hitachi modular storage array failures when they occur. The UDP protocol, however, may prevent the trap from being reported properly to the SNMP manager. Moreover, if a controller goes down, the *systemDown* trap may not be issued.

3. The MIB is configured to detect errors periodically, as noted in step 1. As a result, you will know when a failure occurs or a part fails, even if a trap described in step 2 is not reported, because the MIB value *dfRegressionStatus* in the event of failure is not 0.

Example: If a drive is blocked, *dfRegressionStatus* = 69

A request from the SNMP manager may receive no response if a controller is blocked. You can detect when a controller is blocked, even if a *systemDown* trap is not reported. However, the UDP protocol used with SNMP may cause requests from the SNMP manager to be ignored, even during normal operation. If continuous requests receive no response, it can indicate that a controller is blocked.




Table 9-5: SNMP Agent flow diagram

Operational guidelines

When using SNMP Agent Support Function, observe the following guidelines:

- Like other SNMP applications, SNMP Agent Support Function uses the UDP protocol. UDP might prevent error traps from being reported properly to the SNMP manager. Therefore, it is recommended that the SNMP manager acquire MIB information periodically.
- If the interval for collecting MIB information is set too short, it can adversely impact the Hitachi modular storage array's performance.
- If failures occur in a Hitachi modular storage array after the SNMP manager starts, the failures are not reported with a trap. In this case,

acquire the MIB objects *dfRegressionStatus* after starting the SNMP manager and check whether failures occur.

- The SNMP Agent Support Function stops if the controller is blocked and the SNMP managers receive no response.
- If a Hitachi modular storage array has two controllers, a failure of a hardware component, such as a fan, battery, power supply, or cache, between power-on and when the array becomes “Ready” are reported as traps from both array controllers. This includes failures that occurred at the last power off. Disk drive failures and failures that occur while an array is “Ready” are reported with a trap from only the controller that detects the failures.
- For Hitachi modular storage arrays with two controllers, SNMP manager must monitor both controllers. If only one of the controllers is monitored using the SNMP manager, traps are not reported on the unmonitored controller. In addition, observe the following considerations:
 - Monitor controller 0.
 - *dfRegressionStatus* of the MIB object is system failure information. Acquire *dfRegressionStatus* periodically from the SNMP Manager and check whether a failure is present.
 - If controller 0 becomes blocked, you cannot use the SNMP Agent Support Function.
 - If the acquisition of *dfRegressionStatus* of the MIB object fails, a controller blockage has occurred. Use Navigator 2 to check the status of the storage array.
- If the Hitachi modular storage array receives broadcasts or port scans on TCP port 199, response delays or time-outs can occur when the SNMP manager requests MIB information. In this case, check the network configuration to confirm that TCP port 199 of the storage array is not being accessed by other applications.
- The accumulated operating time (sysUpTime) of the SNMP agent is counted from the starting/restarting time of the SNMP agent. Also, the sysUpTime resets in the following cases:
 - Starting/restarting the SNMP agent. Starting the array, replacing the firmware, rebooting and others.
 - At the time of setting the SNMP to Enabled (including the “Disabled -> Enabled” operation).
 - When exceeding the upper limit (approximately 497 days) of the sysUpTime.

Table 9-6 details the connection status of the GET/TRAP specification.

Table 9-6: GET/TRAP specification

Connection Status	Controller Status	GET/TRAP Specification				Comments
		Controller 0		Controller 1		
Both controllers	1. Both controllers are normal	GET	YES	GET	YES	Master controller 0
		TRAP	YES	TRAP	*	
	2. Controller 1 is blocked	GET	YES	GET	NO	Master controller: 0 If controller 1 is recovered, the system goes to 1.
		TRAP	YES	TRAP	NO	
	3. Controller 0 is blocked	GET	NO	GET	YES	Master controller 1
		TRAP	NO	TRAP	YES	
	4. Controller 0 is recovered (board was replaced while power was on)	GET	YES	GET	YES	Master controller: 1 System goes to 0 when restarted (P/S ON).
		TRAP	*	TRAP	YES	
Controller 0 only	5. Both controllers are normal	GET	YES	GET	NO	Master controller 0
		TRAP	YES	TRAP	NO	
	6. Controller 1 is blocked	GET	YES	GET	NO	
		TRAP	YES	TRAP	NO	
	7. Controller 0 is blocked	GET	NO	GET	NO	Master controller 1
		TRAP	NO	TRAP	NO	
	8. Controller 0 is recovered (the board was replaced while the power is on)	GET	YES	GET	NO	Master controller: 1 System goes to ? when restarted (P/S ON).
		TRAP	*	TRAP	NO	

LEGEND:

YES = *GET* and *TRAP* are possible. Drive blockages and occurrences detected by the other controller in a dual-controller configuration are excluded.

NO = *GET* and *TRAP* are impossible.

* = A trap is reported only for its own controller blockade (drive extractions not included) detected by its own controller.



NOTE: A trap is reported for an error that is detected when a controller board is replaced while the power is on or when the power is turned on. Traps other than the above are also reported.

MIBs

Supported MIBs

Table 9-7 shows the MIBs that the Hitachi modular storage arrays support. The *GetResponse* of *noSuchName* is returned in response to the *GetRequest* or *SetRequest* issued to an unsupported object.

Table 9-7: Supported MIBs

MIB		Supported?	Relevant RFC
MIB II	system group	YES	RFC 1213
MIB II	interface group	Partially	RFC 1213
MIB II	at group	NO	RFC 1213
MIB II	ip group	Partially	RFC 1213
MIB II	icmp group	NO	RFC 1213
MIB II	tcp group	NO	RFC 1213
MIB II	udp group	NO	RFC 1213
MIB II	egp group	NO	RFC 1213
MIB II	snmp group	YES	RFC 1213
Extended MIB		YES	—

MIB access mode

The access mode for all community MIBs should be read-only.

The *GetResponse* of *noSuchName* is returned in response to each SNMP manager's *Set* request.

OID assignment system

Figure 9-13 on page 9-30 through Figure 9-15 on page 9-32 show the OID assignment system.




Figure 9-13: OID assignment system (1 of 3)

dfraidLanExMib(2)	
—dfSystemParameter(1)	
—dfSystemProductName(1)	Product name
—dfSystemMicroRevision(2)	Firmware Rev. No.
—dfSystemSerialNumber(3)	DFxxx Serial No.
—dfWarningCondition(2)	
—dfRegressionStatus(1)	Warning failure information
—dfPreventiveMaintenanceInformation(2)	Preventive maintenance information
—dfWarningReserve1(3)	Not used
—dfWarningReserve2(4)	Not used
—dfCommandExecutionCondition(3)	
—dfCommandTable(1)	
—dfCommandEntry(1)	
—dfLun(1)	Logical unit No.
—dfReadCommandNumber(2)	Number of read command receptions
—dfReadHitNumber(3)	Number of cache read hits
—dfReadHitRate(4)	Cache read hit rate
—dfWriteCommandNumber(5)	Number of write command receptions
—dfWriteHitNumber(6)	Number of cache write hits
—dfWriteHitRate(7)	Cache write hit rate
—dfCacheLoadCondition(4)	
—dfWriteDataRate(1)	Darty segment rate
—dfLUNS(5)	LUN security information
—dfLUNSSwitch(1)	Security switch
—dfLUNSSwitchEntry(1)	
—dfSwitchSerialNumber(1)	DFxxx serial number
—dfSwitchPortID(2)	Port number
—dfSwitchOnOff(3)	Valid/invalid status
—dfSwitchControlStatus(4)	Control flag
—dfLUNSWWN(2)	WWN information
—dfLUNSWWNEEntry(1)	
—dfWWNSerialNumber(1)	DFxxx serial number
—dfWWNPortID(2)	Port number
—dfWWNControlIndex(3)	Control index
—dfWWNWWN(4)	WWN (Port Name)
—dfWWNID(5)	WWN number
—dfWWNNickname(6)	Nickname
—dfWWNUseNickname(7)	Use/no use of nickname
—dfWWNControlStatus(8)	Control flag
—dfLUNSWWNGroup(3)	Not supported
—dfLUNSLUN(4)	LUN information
—dfLUNSLUNEntry(1)	
—dfLUNSerialNumber(1)	DFxxx Serial number
—dfLUNPortID(2)	Port number
—dfLUNLUN(3)	LUN
—dfLUNWWNSecurity(4)	WWN access permission
—dfLUNWWNGroupSecurity(5)	WWN group access permission
—dfLUNControlStatus(6)	Control flag
—dfLUNSLUNGroup(5)	Not supported

Continue to the next page (1)

Figure 9-14: OID assignment system (2 of 3)

Continued from the previous page (1)




Figure 9-15: OID assignment system (3 of 3)

Supported traps and extended traps

Table 9-8 on page 9-33 lists standard traps the SNMP agent supports, and Table 9-9 on page 9-33 lists extended traps. If the Hitachi modular storage array is used as a local TrueCopy or TrueCopy Extended array, these traps are issued if both paths are blocked after the remote array restarts. In addition, if the local array starts or restarts and becomes ready before the remote disk array becomes ready, both paths are blocked and this trap is issued.

For trap-issuing opportunities, see the extended traps in Table 9-9 on page 9-33.

Table 9-8: Supported standard traps

Generic Trap Code	Trap	Description	Supported?
0	coldStart	Reset from power-off. (P/S on) The SNMP agent started online.	YES
1	warmStart	Management module restarted. The SNMP information file was reset.	YES
2	linkDown	Link goes down	NO
3	linkUp	Link goes up	NO
4	authenticationFailure	Illegal SNMP accessed	YES
5	egpNeighborLoss	EGP error is detected	NO
6	enterpriseSpecific	Enterprise extended trap	YES

Table 9-9: Supported extended traps

Trap Code	Trap	Meaning
1	systemDown	Array down occurred. If a controller is blocked, the array issues TRAPs that show the blockage. The array may recover from a controller blockade automatically, depending on the cause of the failure.
2	driveFailure	Drive blocking occurred.
3	fanFailure	Fan failure occurred.
4	powerSupplyFailure	Power supply failure occurred.
5	batteryFailure	Battery failure occurred.
6	cacheFailure	Cache memory failure occurred.
7	upsFailure	UPS failure occurred.
10	otherControllerFailure	Other controller failure occurred. If a controller is blocked, the array issues TRAPs that show the blockage. The array may recover from a controller blockade automatically, depending on the cause of the failure.
11	warning	Warning occurred. The array warning status can be set automatically in the warning information via preventive maintenance, periodic part replacement, or field work conducted by Hitachi service personnel.
12	SpareDriveFailure	Spare drive failure occurred.
14	interfaceBoardFailure	Interface board failure.
16	pathFailure	Path failure occurred.
20	hostConnectorFailure	Host connector failure occurred.
250	interfaceBoardFailure	Interface board failure.

Table 9-9: Supported extended traps

Trap Code	Trap	Meaning
254	hostIoModuleFailure	Host I/O module failure occurred.
255	driveIoModuleFailure	Drive I/O module failure occurred.
256	managementModuleFailure	Management module failure occurred.
257	recoverableControllerFailure	Recoverable CTL alarm by the maintenance procedures of the blocked component
300	psueShadowImage	Failure occurred [ShadowImage].
301	psueSnapShot	Failure occurred [SnapShot].
302	psueTrueCopy	Failure occurred [TrueCopy]
303	psueTrueCopyExtendedDistance	Failure occurred [TrueCopy Extended Distance]
304	psueModularVolumeMigration	Failure occurred [Modular Volume Migration].
307	cycleTimeThresholdOver	Cycle time threshold over occurred.
308	luFailure	Data pool no free.
309	replaceAirFilterBezel	Replace the Air Filter of DC power supply.
310	dpPoolEarlyAlert	DP Pool consumed capacity early alert
311	dpPoolDepletionAlert	DP Pool consumed capacity depletion alert
312	dpPoolCapacityOver	DP Pool consumed capacity over
313	overProvisioningWarningThreshold	Over Provisioning Warning Threshold
314	overProvisioningLimitThreshold	Over Provisioning Limit Threshold
319	replicationDepletionAlert	Over replication depletion alert threshold
320	replicationDataReleased	Over replication data released threshold
321	ssdWriteCountEarlyAlert	SSD write count early alert
322	ssdWriteCountExceedThreshold	SSD write count exceeds threshold
323	sideCardFailure	Side Card failure occurred
324	pageRelocationFailure	The page relocation failed due to the destage timeout of the pool management information.
325	arrayRebootRequestForDPPPoolInvalid	Array reboot is requested because the invalid DP pool management information was detected due to the power shutdown.
326	dpPoolInformationInvalid	The invalid DP pool management information was detected due to the power shutdown.
327	fmdWriteCountEarlyAlert	FMD write count early alert.
328	fmdWriteCountExceedThreshold	FMD write count exceeds threshold
329	fmdBatteryLifeEarlyAlert	The FMD battery lifetime rate reached the set threshold value.
330	pduConnectionError	A problem occurred in the PDU operation.
331	pduHealthCheckError	A connection check with PDU failed.

MIB installation

This section provides installation specifications for MIBs supported by Hitachi modular storage arrays. The following conventions are used in this section:

- Standard = the standard shown on the subject standard document.
- Content: = the content of the subject extended MIB.
- Access = shows whether the item read/write (RW), read only (R), or not accessible (N/A).
- Installation - the specifications for mounting the subject MIB in the array.
- Supported status = can be YES, Partially, or NO.

MIB II

mgmt OBJECT IDENTIFIER :: = {iso(1) org(3) dod(6) internet(1) 2}

mib-2 OBJECT IDENTIFIER :: = {mgmt 1}

system group

system OBJECT IDENTIFIER :: = { mib-2 1 }

This section describes the *system* group of MIB-II.

Table 9-10 details the object identifier of the system group.

Table 9-10: system group

No.	Object Identifier	Access	Installation Specification	Support?	Comments
1	sysDescr {system 1}	R	[Standard] Name or version No. of hardware, OS, network OS [Installation] Fixed character string (Fibre connection for AMS) : HITACHI DF600F Verxxxxxxxxx (Same as inquiry information)	YES	
2	sysObjectID {system 2}	R	[Standard] Object ID indicating the agent vendor product identification number. [Installation] Value is fixed. 1.3.6.1.4.1.116.3.11.1.2	YES	
3	sysUpTime {system 3}	R	[Standard] Accumulated time since the SNMP agent software was started in units of 10 ms. [Installation] Value is fixed as 0 by default. Accumulated time in the case of SET SYSUPTIME line is added in the SNMP setting.	YES	
4	sysContact {system 4}	R	[Standard] agent manager's name and items for contact (manager, managing department, and extension number) [Installation] User-specified ASCII character string (within 255 characters). No default value (NULL).	YES	Should be Read_Only in the array. Data should be entered from the operation environment setting file.
5	sysName {system 5}	R	[Standard] A name given to the agent for management, namely, domain name. [Installation] User-specified ASCII character string (within 255 characters). No default value (NULL).	YES	Should be Read_Only in the array. Data should be entered from the operation environment setting file.

Table 9-10: system group (Continued)

No.	Object Identifier	Access	Installation Specification	Support?	Comments
6	sysLocation {system 6}	R	[Standard] Installation place of the agent [Installation] User-specified ASCII character string (within 255 characters). No default value (NULL).	YES	Should be Read_Only in the array. Data should be entered from the operation environment setting file.
7	sysServices {system 7}	R	[Standard] Service value [Installation] Value is fixed as 8.	YES	

interfaces OBJECT IDENTIFIER ::= { mib-2 2 }

This section describes the *interfaces* group of MIB-II.

Table 9-11 details the object identifiers of the interfaces group.

Table 9-11: interfaces group

No.	Object Identifier	Access	Installation Specification	Support?	Comments
1	ifNumber {interface 1}	R	[Standard] Number of network interfaces provided by this system. [Installation] Value is fixed as 1.	YES	
2	ifTable {interface 2}	N/A	[Standard] Information on each interface is presented in tabular form. The number of entries depends on the ifNumber value. [Installation] Same as the standard. (Refer to the lower hierarchical level.)	Partially	
2.1	ifEntry {ifTable 1}	N/A	[Standard] Each interface information comprising the entries shown below. [Installation] Same as the standard. (Refer to the lower hierarchical level.)	Partially	
2.1.1	ifIndex {ifEntry 1}	R	[Standard] Interface identification number. [Installation] Value is fixed as 1.	YES	(index)
2.1.2	ifDescr {ifEntry 2}	R	[Standard] Interface information [Installation] Fixed character string for each interface type. Ethernet Auto	YES	
2.1.3	ifType {ifEntry 3}	R	[Standard] Interface type ID number [Installation] Fixed value. ethernetCsmacd	YES	
2.1.4	ifMtu {ifEntry 4}	R	[Standard] Maximum sendable/receivable frame length in bytes. MTU (Max Transfer Unit) value [Installation] - (Not installed)	NO	
2.1.5	ifSpeed {ifEntry 5}	R	[Standard] Transfer rate in units of bit/s. [Installation] 100000000	YES	

Table 9-11: interfaces group (Continued)

No.	Object Identifier	Access	Installation Specification	Support?	Comments
2.1.6	ifPhysAddress {ifEntry 6}	R	[Standard] Interface physical address [Installation] MAC Address	YES	
2.1.7	ifAdminStatus {ifEntry 7}	RW	[Standard] Interface set status • 1: Operation • 2: Stop • 3: Test [Installation] - (Not installed)	NO	
2.1.8	ifOperStatus {ifEntry 8}	R	[Standard] Current interface status • 1: Operating • 2: Stopped • 3: Testing [Installation] - (Not installed)	NO	
2.1.9	ifLastChange {ifEntry 9}	R	[Standard] sysUpTime assumed when the subject interface ifOperStatus is changed last [Installation] - (Not installed)	NO	
2.1.10	ifInOctets {ifEntry 10}	R	[Standard] Total number of bytes (including synchronous bytes) in the frame received by the subject interface [Installation] - (Not installed)	NO	
2.1.11	ifInUcastPkts {ifEntry 11}	R	[Standard] Number of subnetwork unicast packets reported to the host protocol [Installation] - (Not installed)	NO	
2.1.12	ifInNUcastPkts {ifEntry 12}	R	[Standard] Number of broadcast or multicast packets reported to the host protocol [Installation] - (Not installed)	NO	
2.1.13	ifInDiscards {ifEntry 13}	R	[Standard] Number of received packets discarded due to insufficient buffer space, even if normal [Installation] - (Not installed)	NO	
2.1.14	ifInErrors {ifEntry 14}	R	[Standard] Number of received erred packets [Installation] - (Not installed)	NO	
2.1.15	ifInUnknownProtos {ifEntry 15}	R	[Standard] Number of received packets discarded due to incorrect or unsupported protocol [Installation] - (Not installed)	NO	

Table 9-11: interfaces group (Continued)

No.	Object Identifier	Access	Installation Specification	Support?	Comments
2.1.16	ifOutOctets {ifEntry 16}	R	[Standard] Total number of bytes (including synchronizing characters) in transmitted frames [Installation] - (Not installed)	NO	
2.1.17	ifOutUcastPkts {ifEntry 17}	R	[Standard] Number of packets (including those not sent) requested unicast from the upper layer. [Installation] - (Not installed)	NO	
2.1.18	ifOutNUcastPkts {ifEntry 18}	R	[Standard] Number of packets (including those discarded and not sent) requested broadcast or multicast from the upper layer. [Installation] - (Not installed)	NO	
2.1.19	ifOutDiscards {ifEntry 19}	R	[Standard] Number of packets discarded due to insufficient transmit buffer space, etc. [Installation] - (Not installed)	NO	
2.1.20	ifOutErrors {ifEntry 20}	R	[Standard] Number of packets not sent due to errors. [Installation] - (Not installed)	NO	
2.1.21	ifOutQLen {ifEntry 21}	R	[Standard] Sent frame queue length (indicated in number of packets) [Installation] - (Not installed)	NO	
2.1.22	ifSpecific {ifEntry 22}	R	[Standard] Object identifier number for defining the MIB specific to interface media [Installation] Value is fixed as 0.0.	YES	

at group

at OBJECT IDENTIFIER :: = { mib-2 3 }

The *at* group of MIB-II is not supported.

ip group

ip OBJECT IDENTIFIER ::= { mib-2 4 }

This section describes the *ip* group of MIB-II.

Table 9-12 details the object identifiers of the ip group.

Table 9-12: ip group

No.	Object Identifier	Access	Installation Specification	Support?	Comments
1	ipForwarding {ip 1}	R	[Standard] Specifies whether received IP packets are transferred as IP gateways. 1: Transfer 2: No transfer [Installation] - (Not installed)	NO	
2	ipDefaultTTL {ip 2}	R	[Standard] Default value to be set in TTL (Time to live: packet life) in IP header. [Installation] - (Not installed)	NO	
3	ipInReceives {ip 3}	R	[Standard] Total number of received IP packets, including erred ones [Installation] - (Not installed)	NO	
4	ipInHdrErrors {ip 4}	R	[Standard] Number of packets discarded due to IP header errors. Errors: Check sum error, version mismatch, or other format error, TTL value out of limits, IP header option error, etc. [Installation] - (Not installed)	NO	
5	ipInAddrErrors {ip 5}	R	[Standard] Number of packets discarded, since the address in IP header is illegal. [Installation] - (Not installed)	NO	
6	IpForwDatagrams {ip 6}	R	[Standard] Number of packets transferred to the last address. If not operated as an IP gateway, indicates the number of packets transferred successfully by source routing. [Installation] - (Not installed)	NO	
7	ipInUnknownProtos {ip 7}	R	[Standard] Number of discarded packets of received IP packets due to unknown or unsupported protocol. [Installation] - (Not installed)	NO	

Table 9-12: ip group (Continued)

No.	Object Identifier	Access	Installation Specification	Support?	Comments
8	ipInDiscards {ip 8}	R	[Standard] Number of IP packets discarded due to internal trouble such as insufficient buffer space. (Does not include packets discarded while waiting for Re_assembly.) [Installation] - (Not installed)	NO	
9	ipInDelivers {ip 9}	R	[Standard] Number of packets transferred to an IP user protocol (host protocol including ICMP) [Installation] - (Not installed)	NO	
10	ipOutRequests {ip 10}	R	[Standard] Number of IP packets requested by a local IP user protocol (including ICMP). (ipForwDatagrams is not included.) [Installation] - (Not installed)	NO	
11	ipOutDiscards {ip 11}	R	[Standard] Number of IP packets discarded due to insufficient buffer space, etc.; IP packets have no error. (IP packets discarded by ipForwDatagrams according to a send request are included.) [Installation] - (Not installed)	NO	
12	ipOutNoRoutes {ip 12}	R	[Standard] Number of packets discarded due to no route to destination. This is the number of packets that could not be transferred because the default gateway was down (including discarded IP packets that intended to be transferred with ipForwDatagrams because the router was unknown). [Installation] - (Not installed)	NO	
13	ipReasmTimeout {ip 13}	R	[Standard] Maximum time waiting for all IP packets to be assembled when receiving fragmented IP packets. [Installation] - (Not installed)	NO	
14	ipReasmReqds {ip 14}	R	[Standard] Number of received fragmented IP packets to be assembled with an entity. [Installation] - (Not installed)	NO	
15	ipReasmOKs {ip 15}	R	[Standard] Number of fragmented IP packets received and assembled successfully [Installation] - (Not installed)	NO	

Table 9-12: ip group (Continued)

No.	Object Identifier	Access	Installation Specification	Support?	Comments
16	ipReasmFails {ip 16}	R	[Standard] Number of fragmented IP packets received but failed to be assembled due to time-out, etc. [Installation] - (Not installed)	NO	
17	ipFragOKs {ip 17}	R	[Standard] Number of packets fragmented successfully with this entity [Installation] - (Not installed)	NO	
18	ipFragFails {ip 18}	R	[Standard] Number of IP packets discarded without fragmenting because the "No Fragment" flag was set - or some other reason - although they must be fragmented with this entity. [Installation] - (Not installed)	NO	
19	ipFragCreates {ip 19}	R	[Standard] Number of fragmented IP packets created by the fragment with this entity. [Installation] - (Not installed)	NO	
20	ipAddrTable {ip 20}	N/A	[Standard] Address information table for each IP address of this entity [Installation] Same as standard. (Refer to the lower hierarchical level.)	YES	
20.1	ipAddrEntry {ipAddrTable 1}	N/A	[Standard] IP address information [Installation] Same as standard. (Refer to the lower hierarchical level.)	YES	
20.1.1	ipAdEntAddr {ipAddrEntry 1}	R	[Standard] IP address of this entity [Installation] Same as standard. A system parameter set by users.	YES	(index)
20.1.2	ipAdEntIfIndex {ipAddrEntry 2}	R	[Standard] Interface identification number corresponding to this IP address. Same as ifIndex. [Installation] Same as standard. Value is fixed as 1.	YES	
20.1.3	ipAdEntNetMask {ipAddrEntry 3}	R	[Standard] Subnetwork mask value related to this IP address. [Installation] Same as standard.	NO	
20.1.4	ipAdEntBcastAddr {ipAddrEntry 4}	R	[Standard] LSB value of IP broadcast address when IP broadcast sending. [Installation] Value is fixed as 1.	NO	

Table 9-12: ip group (Continued)

No.	Object Identifier	Access	Installation Specification	Support?	Comments
20.1.5	ipAdEntReasm Max-Size {ipAddrEntry 5}	R	[Standard] Maximum size of IP packets that can be assembled with this entity from fragmented IP packets received by this interface. [Installation] Value is fixed as 65535.	NO	
21	ipRouteTable {ip 21}	N/A	[Standard] IP routing table of this entity [Installation] - (Not installed)	NO	
21.1	ipRouteEntry {ipRouteTable 1}	N/A	[Standard] Route to a specific destination [Installation] - (Not installed)	NO	
21.1.1	ipRouteDest {ipRouteEntry 1}	RW	[Standard] Destination IP address of this route table [Installation] - (Not installed)	NO	(index)
21.1.2	ipRouteIfIndex {ipRouteEntry 2}	RW	[Standard] Interface identification number to send to the host next to this route. Same as ifIndex. [Installation] - (Not installed)	NO	
21.1.3	ipRouteMetric1 {ipRouteEntry 3}	RW	[Standard] Primary routing metric of this route [Installation] - (Not installed)	NO	
21.1.4	ipRouteMetric2 {ipRouteEntry 4}	RW	[Standard] Alternate routing metric [Installation] - (Not installed)	NO	
21.1.5	ipRouteMetric3 {ipRouteEntry 5}	RW	[Standard] Alternate routing metric [Installation] - (Not installed)	NO	
21.1.6	ipRouteMetric4 {ipRouteEntry 6}	RW	[Standard] Alternate routing metric [Installation] - (Not installed)	NO	
21.1.7	ipRouteNextHop {ipRouteEntry 7}	RW	[Standard] Next hop IP address of this route [Installation] - (Not installed)	NO	
21.1.8	ipRouteType {ipRouteEntry 8}	RW	[Standard] Routing type • other = 1, • invalid (invalid route) = 2, • direct (direct connection) = 3, • indirect (indirect connection) = 4 [Installation] - (Not installed)	NO	

Table 9-12: ip group (Continued)

No.	Object Identifier	Access	Installation Specification	Support?	Comments
21.1.9	ipRouteProto {ipRouteEntry 9}	R	[Standard] Learned routing mechanism • other = 1 • local = 2 • netmgmt = 3 • icmp = 4 • epg = 5 • ggp = 6 • hello = 7 • rip = 8 • is-is = 9 • es-is = 10 • ciscoIgrp = 11 • bbnSpfIgp = 12 • ospf = 13 • bgp = 14 [Installation] - (Not installed)	NO	
21.1.10	ipRouteAge {ipRouteEntry 10}	RW	[Standard] Elapsed time (in seconds) since the route was recognized last as the normal one. [Installation] - (Not installed)	NO	
21.1.11	ipRouteMask {ipRouteEntry 11}	RW	[Standard] Subnet mask value [Installation] - (Not installed)	NO	
21.1.12	ipRouteMetric5 {ipRouteEntry 12}	RW	[Standard] Alternate routing metric [Installation] - (Not installed)	NO	
21.1.13	ipRouteInfo {ipRouteEntry 13}	R	[Standard] Defined number of the MIB for the routing protocol used for this route. [Installation] - (Not installed)	NO	
22	ipNetToMediaTable {ip 22}	N/A	[Standard] IP address conversion table used to convert IP addresses to physical addresses. [Installation] - (Not installed)	NO	
22.1	ipNetToMediaEntry {ipNetToMedia- Table 1}	N/A	[Standard] Entry including an IP address corresponding to a physical address. [Installation] - (Not installed)	NO	
22.1.1	ipNetToMediaIf- Index {ipNetToMedia- Entry 1}	RW	[Standard] Interface identification number of this entry. The ifIndex value is used. [Installation] - (Not installed)	NO	(index)

Table 9-12: ip group (Continued)

No.	Object Identifier	Access	Installation Specification	Support?	Comments
22.1.2	ipNetToMedia-PhysAddress {ipNetToMedia-Entry 2}	RW	[Standard] Physical address depending on medium [Installation] - (Not installed)	NO	
22.1.3	ipNetToMedia-NetAddress {ipNetToMedia-Entry 3}	RW	[Standard] P address corresponding to the physical address of this entry. [Installation] - (Not installed)	NO	(index)
22.1.4	ipNetToMediaType {ipNetToMedia-Entry 4}	RW	[Standard] Address conversion method • other = 1 • invalid = 2 • dynamic (conversion) = 3 • static (conversion) = 4 [Installation] - (Not installed)	NO	
23	ipRoutingDiscards {ip 23}	R	[Standard] Total of valid routing information items discarded due to insufficient memory space, etc. [Installation] - (Not installed)	NO	

icmp group

icmpOBJECT IDENTIFIER :: = {mib-2 5}

The *icmp* group of MIB-II is not supported.

tcp group

tcpOBJECT IDENTIFIER :: = {mib-2 6}

The *tcp* group of MIB-II is not supported.

udp group

udpOBJECT IDENTIFIER :: {mib-2 7}

The *udp* group of MIB-II is not supported.

egp group

egpOBJECT IDENTIFIER :: = {mib-2 8}

The *egp* group of MIB-II is not supported.

snmp group

snmpOBJECT IDENTIFIER ::= { mib-2 11 }

This section describes the *snmp* group of MIB-II.

Table 9-13 details the object identifiers of the snmp group.

Table 9-13: snmp group

No.	Object Identifier	Access	Installation Specification	Support?	Comments
1	snmpInPkts {snmp 1}	R	[Standard] Total of SNMP messages received from a transport service [Installation] Same as standard.	YES	
2	snmpOutPkts {snmp 2}	R	[Standard] Total of SNMP messages requested to be transferred to the transport layer. [Installation] Same as standard.	YES	
3	snmpInBad-Versions {snmp 3}	R	[Standard] Total of received messages of an unsupported version. [Installation] Same as standard.	YES	
4	snmpInBad-CommunityNames {snmp 4}	R	[Standard] Total of received SNMP messages of an unused community. [Installation] Same as standard.	YES	
5	snmpInBad-CommunityUses {snmp 5}	R	[Standard] Total of received messages indicating operation disabled for the community. [Installation] Same as standard.	YES	
6	snmpInASNParse-Errs {snmp 6}	R	[Standard] Total of received messages of ASN.1 error [Installation] Same as standard.	YES	
8	snmpInTooBigs {snmp 8}	R	[Standard] Total of received PDUs of tooBig error status. [Installation] Same as standard.	YES	
9	snmpInNoSuchNames {snmp 9}	R	[Standard] Total of received PDUs of noSuchName error status. [Installation] Same as standard.	YES	
10	snmpInBadValues {snmp 10}	R	[Standard] Total of received PDUs of badValue error status. [Installation] Same as standard.	YES	
11	snmpInReadOnlys {snmp 11}	R	[Standard] Total of received PDUs with readOnly error status. [Installation] Same as standard.	YES	

Table 9-13: snmp group (Continued)

No.	Object Identifier	Access	Installation Specification	Support?	Comments
12	snmpInGenErrs {snmp 12}	R	[Standard] Total of received PDUs with genErr error status. [Installation] Same as standard.	YES	
13	snmpInTotalReq-Vars {snmp 13}	R	[Standard] Total of MIB objects for which MIB was gathered successfully. [Installation] Same as standard.	YES	
14	snmpInTotalSet-Vars {snmp 14}	R	[Standard] Total of MIB objects for which MIB was set successfully. [Installation] Same as standard.	YES	
15	snmpInGetRequests {snmp 15}	R	[Standard] Total of received GetRequest PDUs. [Installation] Same as standard.	YES	
16	snmpInGetNexts {snmp 16}	R	[Standard] Total of received GetNext Request PDUs. [Installation] Same as standard.	YES	
17	snmpInSetRequests {snmp 17}	R	[Standard] Total of received SetRequest PDUs. [Installation] Same as standard.	YES	
18	snmpInGet-Responses {snmp 18}	R	[Standard] Total of received GetResponse PDUs. [Installation] Same as standard.	YES	
19	snmpInTraps {snmp 19}	R	[Standard] Total of received TrapPDUs. [Installation] Same as standard.	YES	
20	snmpOutTooBigs {snmp 20}	R	[Standard] Total of transferred PDUs of tooBig error status. [Installation] Same as standard.	YES	
21	snmpOutNoSuch-Names {snmp 21}	R	[Standard] Total of transferred PDUs of noSuchName error status. [Installation] Same as standard.	YES	
22	snmpOutBadValues {snmp 22}	R	[Standard] Total of transferred PDUs of badValue error status. [Installation] Same as standard.	YES	
23	snmpOutBadValues {snmp 23}	R	[Standard] Total of transferred PDUs of badValue error status. [Installation] Same as standard.	YES	

Table 9-13: snmp group (Continued)

No.	Object Identifier	Access	Installation Specification	Support?	Comments
24	snmpOutGenErrs {snmp 24}	R	[Standard] Total of received PDUs of genErr error status. [Installation] Same as standard.	YES	
25	snmpOutGet-Requests {snmp 25}	R	[Standard] Total of transferred GetRequest PDUs. [Installation] Same as standard.	YES	
26	snmpOutGetNexts {snmp 26}	R	[Standard] Total of transferred GetNextRequest PDUs. [Installation] Same as standard.	YES	
27	snmpOutSet-Requests {snmp 27}	R	[Standard] Total of transferred SetRequest PDUs. [Installation] Same as standard.	YES	
28	snmpOutGet-Responses {snmp 28}	R	[Standard] Total of transferred GetResponse PDUs. [Installation] Same as standard.	YES	
29	snmpOutTraps {snmp 29}	R	[Standard] Total of transferred Trap PDUs. [Installation] Same as standard.	YES	
30	snmpEnable-AuthenTraps {snmp 30}	R	[Standard] This indicates whether an authentication-failure trap can be issued. - enabled = 1 - disabled = 2 [Installation] Fixed value 1 (enabled)	YES	Should be Read Only in array

Extended MIBs

EnterprisesOBJECT IDENTIFIER :: = {iso(1) org(3) dod(6) internet(1) 4}

Enterprises	OBJECT IDENTIFIER :: = {iso(1) org(3) dod(6) internet(1) 4}
hitachi	OBJECT IDENTIFIER :: = {enterprises 116}
systemExMib	OBJECT IDENTIFIER :: = {hitachi 5}
storageExMib	OBJECT IDENTIFIER :: = {systemExMib 11}
dfraidExMib	OBJECT IDENTIFIER :: = {storageExMib 1}
dfraidLanExMib	OBJECT IDENTIFIER :: = {dfraidExMib 2}

dfSystemParameter group

dfSystemParameterOBJECT IDENTIFIER :: {dfraidLanExMib 1}

This section describes the *dfSystemParameter* group of the Extended MIBs.

Table 9-14 details the object identifiers of the dfSystemParameter group.

Table 9-14: dfSystemParameter group

No.	Object Identifier	Access	Installation Specification	Support?	Comments
1	dfSystemProductName { dfSystemParameter 1}	R	[Content] Product name [Installation] (AMS): HITACHI DF600F (Same as inquiry information)	YES	
2	dfSystemMicro- Revision { dfSystemParameter 2}	R	[Content] Firmware revision number [Installation] Same as above	YES	
3	dfSystemSerialNumber { dfSystemParameter 2}	R	[Content] Disk array serial number [Installation] The eight digits of the manufacturing serial number	YES	

dfWarningCondition group

dfWarningConditionOBJECT IDENTIFIER ::= { dfraidLanExMib 2 }

This section describes the *dfWarningCondition* group of the Extended MIBs.

Table 9-15 details the object identifiers of the dfWarningCondition group.

Table 9-15: dfWarningCondition group

No.	Object Identifier	Access	Installation Specification	Support?	Comments
1	dfRegressionStatus { dfWarningCondition 1 }	R	[Content] Warning error information [Installation] Same as above. When normal, this is assigned to 0. (See Note 1)	YES	
2	dfPreventiveMaintenanceInformation { dfWarningCondition 2 }	R	[Content] Drive preventive maintenance information [Installation] Same as above. Value is fixed as 0.	YES	
3	dfRegressionStatus2 { d fWarningCondition 3 }	R	[Content] Warning error information [Installation] When normal, this is assigned to 0.	YES	
4	dfWarningReserve2 { dfWarningCondition 4 }	R	[Content] Reserved area [Installation] Not used. Value is fixed as 0.	YES	

Table 9-16 details the format of the dfRegressionStatus group.

Table 9-16: dfRegressionStatus format

Bit Byte	7	6	5	4	3	2	1	0
0	0	I/F board	0	Host connector	0	0	0	Cache
1	Managem ent Module	Host Module	0	Fan	0	0	PS	Battery
2	Recovera ble CTL	Drive Module	0	0	0	Path	0	UPS
3	CTL	Warning	0	0	ENC	D-Drive	S-Drive	Drive

Table 9-17: dfRegressionStatus2 Format

Bit Byte	7	6	5	4	3	2	1	0
0	0	0	0	0	0	0	0	0
1	0	0	0	0	0	0	0	0
2	0	0	0	0	0	0	0	0

Bit Byte	7	6	5	4	3	2	1	0
3	0	0	0	0	0	0	0	Side Card

Subject bits should be “on” if each part is in the regressed state. This value can be fixed as “0,” depending on the array type and the firmware revision.

[Table 9-18](#) shows this object value for each failure status.

Table 9-18: dfRegressionStatus value for each failure

No.	Bit Position		Object Value (Decimal)	Failed Component
	Byte	Bit		
1	—	—	0	Array normal status
2	3	0	1	Drive blocked
3	3	1	2	Drive (spare drive) blockade
4	3	2	4	Drive (data drive) blockade
5	3	3	8	ENC alarm
6	3	6	64	Warned array
7	3	7	128	Mate controller blocked
8	2	0	256	UPS alarm
9	2	1	—	—
10	2	2	1024	Path blocked
11	2	6	16384	Drive I/O module failure
12	2	7	32768	Controller failure by related parts
13	1	0	65536	Battery alarm
14	1	1	131072	Power supply failure
15	1	3	—	—
16	1	4	1048576	Fan alarm
17	1	5	—	—
18	1	6	4194304	Host I/O module failure
19	1	7	838608	Management module failure
20	0	0	16777216	Host connector alarm
21	—	1	—	—
22	0	2	—	—
23	0	3	—	—
24	0	4	268436456	Host connector alarm

Table 9-19: dfRegressionStatus2 Value for Each Failure

No.	Bit Position		Object Value (Decimal)	Failed Component
	Byte	Bit		
1	—	—	0	Array normal status
2	3	0	1	Side Card failure
3	3	1	-	-
4	3	2	-	-
5	3	3	-	-
6	3	6	-	-
7	3	7	-	-
8	2	0	-	-
9	2	1	-	-
10	2	2	-	-
11	2	6	-	-
12	2	7	-	-
13	1	3	-	-
14	1	4	-	-
15	1	5	-	-
16	0	0	-	-
17	0	1	-	-
18	1	6	-	-
19	1	7	-	-
20	0	4	-	-
21	0	5	-	-
22	0	6	-	-

If two or more components fail, the object value adds up each object value.

Example: When a failure occurs in the battery and the fan:

Object value: 1114112 (65536 + 1048576)

When a value of an object is converted into a binary number, it corresponds to the format in [Table 9-18](#).

Each *TRAP* signal (specific trap codes 2 to 6) is issued each time a warning failure in related component occurs (see [Figure 9-16 on page 9-54](#)). If a warning failure occurs, the bit of the related component of *dfRegistrationStatus* is turned on. The bit is turned off when the array recovers from the warning failure.




Figure 9-16: Relationship of traps and dfWarningCondition groups

dfCommandExecutionCondition group

dfCommandExecutionCondition OBJECT IDENTIFIER ::= { dfraidLanExMib 3 }

This section describes the *dfCommandExecutionCondition* group of the Extended MIBs.

[Table 9-20](#) details object identifiers in the *dfCommandExecutionCondition* group.


Table 9-20: dfCommandExecutionCondition group

No.	Object Identifier	Access	Installation Specification	Support?	Comments
1	dfCommandTable { dfCommandExecutionCondition 1 }	N/A	[Content] Command execution condition table [Installation] Same as above (Refer to the lower hierarchical level)	YES	
1.1	dfCommandEntry { dfCommandTable 1 }	N/A	[Content] Command execution condition entry [Installation] Same as above (Refer to the lower hierarchical level)	YES	
1.1.1	dfLun { dfCommandEntry 1 }	R	[Content] Volume number [Installation] Same as above - HUS110: 0 to 2,047 - Other HUS130/HUS150 models: 0 to 4,095	YES	(index)

Table 9-20: dfCommandExecutionCondition group (Continued)

No.	Object Identifier	Access	Installation Specification	Support?	Comments
1.1.2	dfReadCommandNumber {dfCommandEntry 2}	R	[Content] Number of read command receptions [Installation] Same as above	YES	
1.1.3	dfReadHitNumber {dfCommandEntry 3}	R	[Content] Number of cache read hits [Installation] Number of read commands whose host request range completely hits that of the cache	YES	
1.1.4	dfReadHitRate {dfCommandEntry 4}	R	[Content] Cache read hit rate (%) [Installation] (Number of cache read hits / Number of read command receptions) x 100	YES	
1.1.5	dfWriteCommandNumber {dfCommandEntry 5}	R	[Content] Number of write command receptions [Installation] Same as above	YES	
1.1.6	dfWriteHitNumber {dfCommandEntry 6}	R	[Content] Number of cache write hits [Installation] Number of write commands that were not restricted to write data (not made to wait for writing data) in cache by the dirty threshold value manager	YES	
1.1.7	dfWriteHitRate {dfCommandEntry 7}	R	[Content] Cache write hit rate (%) [Installation] Number of cache write hits / Number of write command receptions) x 100	YES	

The information of this group is updated every 10 seconds. The value accumulated in the previous ten seconds is set (see [Figure 9-17](#)).

**Figure 9-17: Accumulated values over time**

The *dfCommandExecutionCondition* group is updated every 10 seconds and is set to a value accumulated for individual 10 seconds. This interval time of 10 seconds can vary within an error span, depending on the command execution condition. In this case, the group is set to a value converted to every 10 seconds from an accumulated value.

Example: If an elapsed time is 11 seconds and the accumulated number of read command received for that time is 110, the *dfReadCommandNumber* is set to 100.

The number of hits (*dfReadHitNumber*, *dfWriteHitNumber*) can exceed the number of commands received (*dfReadCommandNumber*, *dfWriteCommandNumber*), depending on the timing of updating the *dfCommandExecutionCondition* group. The hit rate (*dfReadHitRate*, *dfWriteHitRate*) at this time is set to 100%.

The *dfCommandExecutionCondition* group indicates the information of the volumes that can be accessed from the host. If the unified volume is being used, this group indicates information of the unified volumes.

dfPort group

dfPortOBJECT IDENTIFIER ::= { dfraidLanExMib 4 }

This section describes the *dfPort* group of the Extended MIBs.

[Table 9-21](#) details object identifiers in the dPort group.

Table 9-21: dPort group

No.	Object Identifier	Access	Installation Specification	Support?	Comments
1	dfPortinf { dfPort 1 }	N/A	[Content] Port information table [Installation] Ditto. (See the lower layer.)	YES	
1.1	dfPortinf Entry { dfPortinf 1 }	N/A	[Content] Port information entry [Installation] Ditto. (See the lower layer.)	YES	
1.1.1	dfLUNSerialNumber { dfLUNSWWNEnt 1 }	R	[Content] Disk array serial number [Installation] The eight digits of the manufacturing serial number.	YES	(index)
1.1.2	dfPortID { dfPortinf Entry 2 }	R	[Content] Port number [Installation] Ditto. (0 to 15) See Table 9-22 on page 9-58 .	YES	(index)

Table 9-21: dPort group (Continued)

No.	Object Identifier	Access	Installation Specification	Support?	Comments
1.1.3	dfPortKind {dfPortinf Entry 3}	R	[Content] Port type [Installation] Ditto. See Port types on page 9-58 .	YES	
1.1.4	dfPortHostMode {dfPortinf Entry 4}	R	[Content] Host mode [Installation] Ditto.	YES	No Data
1.1.5	dfPortFibreAddress {dfPortinf Entry 5}	R	[Content] N_Port_ID of the port [Installation] Ditto. See Fibre address host mode on page 9-58 .	YES	
1.1.6	dfPortFibreTopology {dfPortinf Entry 6}	R	[Content] Topology information [Installation] Ditto. (1 to 4) See Table 9-24 on page 9-59 .	YES	
1.1.7	dfPortControlStatus {dfPortinf Entry 7}	R	[Content] Control flag [Installation] Ditto. (Fixed at 1.)	YES	• 1: Regular return value • 2: Request for setting
1.1.8	dfPortDisplayName {dfPortinf Entry 8}	R	[Content] Port name [Installation] Ditto. (0A to 0H, 1A to 1H) See Table 9-25 on page 9-60 .	YES	
1.1.9	dfPortWWN {dfPortinf Entry 9}	R	[Content] WWN of the port [Installation] Ditto. (8 bytes OCTET String) See Port WWN on page 9-60 .	YES	

[Table 9-22](#) details port display numbers.

Table 9-22: Port display numbers

Port Number	Controller Number	Fibre	Comments
0	0	0A	
1		0B	
2		0C	
3		0D	
4		0E	
5		0F	
6		0G	
7		0H	
8	1	1A	
9		1B	
10		1C	
11		1D	
12		1E	
13		1F	
14		1G	
15		1H	

Port types

Sets **Fibre** or **iSCSI**.

For ports other than those that are not applicable, **None** is set.

The item of the ports of a blocked controller is **None**.

Fibre address host mode

For Fibre-oriented ports, address translation is performed followed by setting. If the address is illegal, the value is 0.

For ports other than Fibre-oriented ones, the value is 0.

[Table 9-23](#) details port addresses and associated values.

Table 9-23: Port addresses and associated values

Value	Address	Value	Address	Value	Address	Value	Address
1	EF	33	B2	65	72	97	3A
2	E8	34	B1	66	71	98	39
3	E4	35	AE	67	6E	99	36
4	E2	36	AD	68	6D	100	35
5	E1	37	AC	69	6C	101	34

Table 9-23: Port addresses and associated values

Value	Address	Value	Address	Value	Address	Value	Address
6	E0	38	AB	70	6B	102	33
7	DC	39	AA	71	6A	103	32
8	DA	40	A9	72	69	104	31
9	D9	41	A7	73	67	105	2E
10	D6	42	A6	74	66	106	2D
11	D5	43	A5	75	65	107	2C
12	D4	44	A3	76	63	108	2B
13	D3	45	9F	77	5C	109	2A
14	D2	46	9E	78	5A	110	29
15	D1	47	9D	79	59	111	27
16	CE	48	9B	80	56	112	26
17	CD	49	98	81	55	113	25
18	CC	50	97	82	54	114	23
19	CB	51	90	83	53	115	1F
20	CA	52	8F	84	52	116	1E
21	C9	53	88	85	51	117	1D
22	C7	54	84	86	4E	118	1B
23	C6	55	82	87	4D	119	18
24	C5	56	81	88	4C	120	17
25	C3	57	80	89	4B	121	10
26	BC	58	7C	90	4A	122	0F
27	BA	59	7A	91	49	123	08
28	B9	60	79	92	47	124	04
29	B6	61	76	93	46	125	02
30	B5	62	75	94	45	126	01
31	B4	63	74	95	43	-	-
32	B3	64	73	96	3C	-	-

Table 9-24 details topology information.

Table 9-24: Topology information

Value	Meaning
1	Fabric (on) & FCAL
2	Fabric (off) & FCAL
3	Fabric (on) & Point to Point
4	Fabric (off) & Point to Point
5	Not Fibre

Table 9-25 details port display names.

Table 9-25: Port display names

Port Number	Controller Number	Fibre	Comments
0	0	*0A*	
1		*0B*	
2		*0C*	
3		*0D*	
4		*0E*	
5		*0F*	
6		*0G*	
7		*0H*	
8	1	*1A*	
9		*1B*	
10		*1C*	
11		*1D*	
12		*1E*	
13		*1F*	
14		*1G*	
15		*1H*	

Port WWN

For Fibre-oriented ports, the port identifier (WWN) is set.

For non-Fibre-oriented ports, the value is 0.

dfCommandExecutionInternalCondition group

dfCommandExecutionInternalConditionOBJECT IDENTIFIER :: =
{ dfraidLanExMib 7 }

This section describes the *dfCommandExecutionInternalCondition* group of the Extended MIBs.

Table 9-25 details object identifiers in the dfCommandExecutionInternalCondition group.

Table 9-26: dfCommandExecutionInternalCondition group

No.	Object Identifier	Access	Installation Specification	Support?	Comments
1	dfCommandInternalTable {dfCommandExecutionCon dition 1}	N/A	[Content] Command execution condition table [Installation] Same as above (Refer to the lower hierarchical level)	YES	
1.1	dfCommandInternalEntry {dfCommandTable 1}	N/A	[Content] Command execution condition entry [Installation] Same as above (Refer to the lower hierarchical level)	YES	
1.1.1	dfInternalLun {dfCommandEntry 1}	R	[Content] Volume number [Installation] Same as above - HUS110: 0 to 2,047 - Other HUS130/HUS150 models: 0 to 4,095	YES	(index)
1.1.2	dInternalReadCommand Number {dfCommandEntry 2}	R	[Content] Number of read command receptions [Installation] Same as above	YES	
1.1.3	dfInternalReadHitNumber {dfCommandEntry 3}	R	[Content] Number of cache read hits [Installation] Number of read commands whose host request range completely hits that of the cache	YES	
1.1.4	dfInternalReadHitRate {dfCommandEntry 4}	R	[Content] Cache read hit rate (%) [Installation] (Number of cache read hits / Number of read command receptions) x 100	YES	
1.1.5	dfInternalWriteCommand Number {dfCommandEntry 5}	R	[Content] Number of write command receptions [Installation] Same as above	YES	

Table 9-26: dfCommandExecutionInternalCondition group (Continued)

No.	Object Identifier	Access	Installation Specification	Support?	Comments
1.1.6	dfInternalWriteHitNumber {dfCommandEntry 6}	R	[Content] Number of cache write hits [Installation] Number of write commands that were not restricted to write data (not made to wait for writing data) in cache by the dirty threshold value manager	YES	
1.1.7	dfInternalWriteHitRate {dfCommandEntry 7}	R	[Content] Cache write hit rate (%) [Installation] Number of cache write hits / Number of write command receptions) x 100	YES	

Additional resources

For more information about SNMP, refer to the following resources and to the IETF Web site <http://www.ietf.org/rfc.html>.

SNMP Version 1

- [RFC 1155](#) — structure and identification of management information for TCP/IP-based internets.
- [RFC 1157](#) – simple protocol by which management information for a network element can be inspected or altered by logically remote users
- [RFC 1212](#) – format for producing MIB modules.
- [RFC 1213](#) — v2 of MIB-2 for network management of TCP/IP-based internets.
- [RFC 1215](#) – Trap-Trap macro for using experimental MIBs.

SNMP Version 2

- [RFC 2578](#) – adapted subset of OSI's Abstract Syntax Notation One, ASN.1 (1988) and associated administrative values.
- [RFC 2579](#) – initial set of textual conventions available to all MIB modules.
- [RFC 2580](#) – notation used to define the acceptable lower-bounds of implementation, along with the actual level of implementation achieved.
- [RFC 3416](#) – syntax and elements of procedure for sending, receiving, and processing SNMP PDUs.

SNMP Version 3

- [RFC 3410](#) – overview of SNMP v3.

- [RFC 3411](#) – vocabulary for describing SNMP Management Frameworks and an architecture for describing the major portions of SNMP Management Frameworks.
- [RFC 3412](#) – dispatching potentially multiple versions of SNMP messages to the proper SNMP Message Processing Models, and dispatching PDUs to SNMP applications.
- [RFC 3413](#) – five types of SNMP applications that use of an SNMP engine as described in STD 62, and MIB modules for specifying targets of management operations, notification filtering, and proxy forwarding.
- [RFC 3414](#) – elements of procedure for providing SNMP message level security and MIB for remotely monitoring/managing the configuration parameters for this Security Model.
- [RFC 3415](#) – view-based Access Control Model (VACM) for use in the SNMP architecture, and MIB for remotely managing the configuration parameters for the VACM.

Coexistence between SNMP standards

- [RFC 2576](#) — coexistence between SNMP v3, SNMP v2, and SNMP v1.

Virtualization

This chapter describes virtualization.

This chapter covers the following topics:

- ❑ [Virtualization overview](#)
- ❑ [Virtualization and applications](#)
- ❑ [A sample approach to virtualization](#)

Virtualization overview

Most data centers use less than 15 percent of available, compute, storage, and memory capacity. By underutilizing these resources, companies deploy more servers than necessary to perform a given amount of work. Additional servers increase costs and create a more complex and disparate environment that can be difficult to manage.

This scenario often results in reduced availability and failure to meet service-level agreements. To sustain an efficient data center environment with fast application deployment, predictable performance, and smooth growth, data centers must increase resource utilization while making sure of security to protect the infrastructure, applications, and data integrity.

Hitachi virtualization and tiered storage solutions, as part of Hitachi Data Systems Services Oriented Storage, enable organizations to strategically align business applications and storage infrastructure so that cost, performance, reliability and availability characteristics of storage can be matched to business requirements.

Tiered storage designs are a natural for both the enterprise Hitachi Universal Storage Platform™ family and the midrange Hitachi Adaptable Modular Storage systems with their ability to support a mix of drive types, sizes and speeds along with advanced RAID options. Solutions based around a Universal Storage Platform add the ability to virtualize both internal and external heterogeneous storage into a single pool with well defined tiers and the ability to transparently move data at will between them.

Virtualization features

The following are Virtualization features:

- **Premium storage reserved for critical applications** - Deploy premium storage for critical applications and data that need premium storage services
- **Cost prioritization model** - Assign lower cost, relatively slower storage for less critical data (like backups or archived data)
- **Data portability** - Move data across tiers as needed to meet application and business requirements

Virtualization task flow

The following is a task flow for Virtualization:

1. You determine that you need to virtualize some of your storage solutions.

2. You begin the process of virtualization using Hitachi Virtual Storage Platform. [Figure 10-1](#) details the approach to Virtualization.




Figure 10-1: Hitachi's Virtual Storage Platform

Virtualization benefits

The following are Virtualization benefits:

- **Basic task improvement** - Improves backup, recovery and archiving; utilization and availability.
- **Transparency** - Allows seamless transparent data volume movement among any storage systems attached to a Virtual Storage Platform
- **Data volume portability** - Enables movement of data volumes between custom storage tiers without requiring administrators to pause or halt applications
- **Complexity reduction** - Masks the underlying complexity of tiered storage data migration and does not require the administrator to master the operation of complex storage analysis
- **Cost and efficiency** – You can't keep throwing more storage as point solutions for each user or business need. You need to balance high business demands with low budgets, contain costs, and "do more with less". Virtualization helps you reclaim, utilize and optimize your storage assets.
- **Data and technology management** – You have more and more data to manage, and you're dealing with a multi-vendor environment as a result of data growth and business change. It's time to rein in all those assets and manage them to drive your business.
- **Improve customer service** – You're under pressure to meet SLAs, align IT with business strategy, and support users and customers.

Virtualizing enables you to deliver storage in right-sized, right-performing slices—slices of what you have now, but weren't maximizing before.

- **Stay competitive** – Business is always looking for ways to be better, faster, cheaper. Hitachi Storage Virtualization increases business agility and lets you do more with less so that you can ramp up fast to meet changing business needs.
- **Enhance performance** – The best way you can support your users and customers is to improve speed and access to their data. Virtualizing gives new life to your existing infrastructure because it lets you optimize all your multi-vendor storage and match storage to application requirements.

Virtualization and applications

The design of a virtualized, tiered storage system starts with the applications. It is the business needs and applications that drive the storage requirements, which in turn guide tier configuration. Most applications can benefit from a mix of storage service levels, using high performance where it is important and less expensive storage where it is not.

But operationally it is not efficient to configure unique tiers for each application. Individually configuring a unique scheme for each application leads to extra work, cost and provisioning delays. Instead, the recommended practice is to develop a catalog of tiers with pre-defined characteristics and then allocate storage to applications as needed from the catalog. Figure 1 outlines a four-tier model; your individual requirement may call for more or less.

Now that we have designed our tiers from a requirements standpoint, how do you configure a system to match? There are a variety of ways to configure tiered storage architectures. You can configure performance, but the bulk of the storage for the mailboxes themselves can be mapped to the less expensive but still performing "Lower Cost" Tier for business data. A small amount of storage space is also mapped in from "Less Critical" for development purposes. With stringent retention policies and an expanding amount of emails with large attachments, a large amount of "Archive" Tier storage is needed.

The NAS Head File and Print functions need some "Primary Tier" storage for several critical image processing applications. However, the bulk is file sharing used for shared directories within the company and print spooling and can use inexpensive "Low Critical" tier.

Additionally, the company's web server uses the "Lower Cost" Tier for business data for the core set of often accessed pages. The bulk of what is online is infrequently accessed and can be kept on "Less Critical" storage.

Storage Options

Now that we have designed our tiers from a requirements standpoint, how do you configure a system to match? There are a variety of ways to configure tiered storage architectures.

You can dedicate specific storage systems for each tier, or you can use different types of storage within a storage system for an "in-the-box" tiered storage system. The Hitachi best practice is to use the virtualization capabilities of the Hitachi Virtual Storage Platform (VSP) and the Hitachi Universal Storage Platform (USP) family to eliminate the inflexible nature of dedicated tiered storage silos and seamlessly combine both. This allows for the best overall solution possible.

For example, for the highest tier you could start with a VSP configured with Fibre Channel drives and a high performance RAID configuration. Here the highest levels of performance and availability for mission critical applications are required. As a second tier you could add the USP with Fibre Channel drives, which are configured at a RAID level that is more cost-effective and still highly reliable but with a little less performance.

The Hitachi storage virtualization architecture is differentiated by the way in which Hitachi storage virtualization maps its existing set of proven storage controller-based services, such as replication and migration, across all participating heterogeneous storage systems.

A sample approach to virtualization

The following sections describe the key components used in the Hitachi Data Systems lab when developing these best practice recommendations.

The Hitachi HUS systems are the only midrange storage systems with the Hitachi Dynamic Load Balancing Controller that provide integrated, automated hardware-based front to back end I/O load balancing. This eliminates many complex and time-consuming tasks that storage administrators typically face.

This type of approach this ensures I/O traffic to back-end disk devices is dynamically managed, balanced and shared equally across both controllers. The point-to-point backend design virtually eliminates I/O transfer delays and contention associated with Fibre Channel arbitration and provides significantly higher bandwidth and I/O concurrency.




Figure 10-2: View of a Hitachi HUS 110 in a controller

The active-active Fibre Channel ports mean the user does not have to consider with controller ownership. I/O is passed to the managing controller through cross-path communication.

Any path can be used as a normal path. The Hitachi Dynamic Load Balancing controllers assist in balancing microprocessor load across the storage systems. If a microprocessor becomes excessively busy, the volume management automatically switches to help balance the microprocessor load. [Table 10-1](#) lists some of the differences between the HUS 100 family storage systems.

Table 10-1: HUS 100 Family overview

Metric	HUS 110	HUS 130	HUS 150
Maximum number of disk drives supported	159	240	480
Maximum cache	8GB	32GB	32GB
Maximum attached hosts through Fibre Channel virtual ports	1,024	2,048	2,048

Table 10-1: HUS 100 Family overview

Metric	HUS 110	HUS 130	HUS 150
Host port options	• 8 Fibre Channel • 4 Fibre Channel • 4 Fibre Channel + 4 iSCSI	• 16 Fibre Channel • 8 Fibre Channel • 8 Fibre Channel + 4 iSCSI	• 16 Fibre Channel • 8 iSCSI • 8 Fibre Channel + 4 iSCSI
Back-end disk drive connections	8 x 3 Gb/s SAS links	16 x 3 Gb/s SAS links	32 x 3 Gb/s SAS links

Hitachi Dynamic Provisioning software

On HUS family systems, Hitachi Dynamic Provisioning software's thin provisioning and wide striping functionalities provide virtual storage capacity to eliminate application service interruptions, reduce costs and simplify administration, as follows:

- Optimizes or “right-sizes” storage performance and capacity based on business or application requirements.
- Supports deferring storage capacity upgrades to align with actual business usage.
- Simplifies and adds agility to the storage administration process.
- Provides performance improvements through automatic optimized wide striping of data across all available disks in a storage pool.

The wide-striping technology that is fundamental to Hitachi Data Provisioning software dramatically improves performance, capacity utilization and management of your environment. By deploying your virtual disks using DP-VOLs from Dynamic Provisioning pools on the HUS 100 family, you can expect the following benefits in your vSphere environment:

- A smoothing effect to virtual disk workload that can eliminate hot spots across the different RAID groups, reducing the need for VMFS workload analysis by the VM.
- Significant improvement in capacity utilization by leveraging the combined capabilities of all disks comprising a storage pool.

vSphere 4

This sample approach uses vSphere 4 as a Virtualization example. vSphere 4 is a highly efficient virtualization platform that provides a robust, scalable and reliable infrastructure for the data center. vSphere features provide an easy to manage platform. These features include

- Distributed Resource Scheduler
- High Availability
- Fault Tolerance

Use of ESX 4's round robin multipathing policy with the symmetric active-active controllers' dynamic load balancing feature distributes load across multiple host bus adapters (HBAs) and multiple storage ports. Use of VMware Dynamic Resource Scheduling (DRS) with Hitachi Dynamic Provisioning software automatically distributes loads on the ESX host and on the storage system's back end. For more information, see VMware's vSphere web site.

For more information, see the Hitachi Dynamic Provisioning data sheet.

Storage configuration

The following sections describe configuration considerations to keep in mind when optimizing a HUS 100 family storage infrastructure to meet your performance, scalability, availability, and ease of management requirements.

Redundancy

A high-performance, scalable, highly available and easy-to-manage storage infrastructure requires redundancy at every level.

To take advantage of ESX's built-in multipathing support, each ESX host needs redundant HBAs. This provides protection against both HBA hardware failures and Fibre Channel link failures.

Figure 10-1 shows that when one HBA is down with either hardware or link failure, another HBA on the host can still provide access to the storage resources. When ESX 4 hosts are connected in this fashion to a HUS 100 family storage system, hosts can take advantage of using round robin multipathing algorithm where the I/O load is distributed across all available paths. Hitachi Data Systems recommends a minimum of two HBA ports for redundancy.

Zone configuration

Zoning divides the physical fabric into logical subsets for enhanced security and data segregation. Incorrect zoning can lead to volume presentation issues to ESX hosts, inconsistent paths, and other problems. Two types of zones are available, each with advantages and disadvantages:

- **Port** — Uses a specific physical port on the Fibre Channel switch. Port zones provide better security and can be easier to troubleshoot than WWN zones. This might be advantageous in a smaller static environment. The disadvantage of this is ESX host's HBA must always be connected to the specified port. Moving an HBA connection results in loss of connectivity and requires rezoning.
- **WWN** — Uses nameservers to map an HBA's WWN to a target port's WWN. The advantage of this is that the ESX host's HBA can be connected to any port on the switch, providing greater flexibility. This might also be advantageous in a larger dynamic environment. However,

the disadvantage is the reduced security and adds more complexity in troubleshooting.

Zones can be created in two ways, each with advantages and disadvantages:

- **Multiple initiator** — Multiple initiators (HBAs) are mapped to one or more targets in a single zone. This can be easier to setup and reduce administrative tasks, but this can introduce interference caused by other devices in the same zone.
- **Single initiator** — Contains one initiator (HBA) with single or multiple targets in a single zone. This can eliminate interference but requires creating zones for each initiator (HBA).

When zoning, it's also important to consider all the paths available to the targets so that multipathing can be achieved. [Table 10-2](#) shows an example of a single-initiator zone with multipathing.

Table 10-2: Single-initiator zoning with multipathing

ESX 1	HBA 1 Port 1	ESX1_HBA1_1_A MS2K_0A_1A	0A
			1A
ESX 1	HBA 2 Port 1	ESX1_HBA2_1_A MS2K_0E_1E	0E
			1E
ESX 2	HBA 1 Port 1	ESX2_HBA1_1_A MS2K_0A_1A	0A
			1A
ESX 2	HBA 2 Port 1	ESX2_HBA2_1_A MS2K_0E_1E	0E
			1E
ESX 3	HBA 1 Port 1	ESX3_HBA1_1_A MS2K_0A_1A	0A
			1A
ESX 3	HBA 2 Port 1	ESX3_HBA2_1_A MS2K_0E_1E	0E
			1E

In this example, each ESX host has two HBAs with one port on each HBA. Each HBA port is zoned to one port on each controller with single initiator and two targets in one zone. The second HBA is zoned to another port on each controller. As a result, each HBA port has two paths and one zone. With a total of two HBA ports, each host has four paths and two zones.

Determining the right zoning approach requires prioritizing your security and flexibility requirements. With single initiator-zones, each HBA is logically partitioned in its own zone. Problems in the fabric caused by one HBA do not affect other HBAs. In a vSphere 4 environment, many storage targets are shared between multiple hosts. It is important to prevent the operations of one ESX host from interfering with other ESX hosts. Industry standard best practice is to use single-initiator zones.

Host Group configuration

Configuring host groups on the Hitachi Unified Storage family involves defining which HBA or group of HBAs can access a volume through certain ports on the controllers. The following sections describe different host group configuration scenarios.

One Host Group per ESX host, standalone host configuration

If you plan to deploy ESX hosts in a standalone configuration, each host's WWNs can be in its own host group. This approach provides granular control over volume presentation to ESX hosts. This is the best practice for SAN boot environments, because ESX hosts do not have access to other ESX hosts' boot volumes.

However, this approach can be an administration challenge because keeping track of which host has which volume can be difficult. In a scenario when multiple ESX hosts need to access the same volume for vMotion purposes, the volume must be added to each host group. This operation is error prone and might lead to confusion. If you have numerous ESX hosts, this approach can be tedious.

One Host Group per cluster, cluster host configuration

Many features in vSphere 4 require shared storage, such as vMotion, DRS, High Availability (HA), Fault Tolerance (FT) and Storage vMotion. Many of these features require that the same LUs are presented to all ESX hosts participating in these cluster functions. If you plan to use ESX hosts with these features, create host groups with clustering in mind.

Host Group options

On an HUS 100 family storage system, host groups are created using Hitachi Storage Navigator Modular 2 software. In the Available Ports box, select all ports. This applies the host group settings to all the ports that you select. Choose VMware from the Platform drop-down menu. Choose **Standard Mode** from the Common Setting drop-down menu. In the Additional Settings box, uncheck the check boxes. These settings automatically apply the correct configuration. Hitachi Dynamic Provisioning software with vSphere 4

The following sections describe best practices for using Hitachi Dynamic Provisioning Software with vSphere 4. *Dynamic Provisioning Space Saving and Virtual Disks*

Two of vSphere's virtual disk formats are thin-friendly, meaning they only allocate chunks from the Dynamic Provisioning pool as required. Thin and zeroedthick format virtual disks are thin-friendly, eagerzeroedthick format virtual disks are not. The eagerzeroedthick format virtual disk allocates 100 percent of the DP-VOLs space in the Dynamic Provisioning pool. While the

eagerzeroedthick format virtual disk does not give the benefit of cost savings by over provisioning of storage, it can still assist in the wide striping of the DP-VOL across all disks in the Dynamic Provisioning pool.

When using DP-VOLs to overprovision storage, follow these best practices:

- Create the VM template on a zeroedthick format virtual disk on non-VAAI enabled environment. When used with VAAI, create the VM template on an eagerzeroedthick format virtual disk. When deploying, select the **Same format as source** radio button in the vCenter GUI.
- Use eagerzeroedthick format virtual disk in VAAI environments.
- Use the default zeroe thick format virtual disk if the volume is not on VAAI-enabled storage.
- Using Storage vMotion when the source VMFS datastore is on a Dynamic Provisioning volume is a Dynamic Provisioning thin friendly operation.

Keep in mind that this operation does not zero out the VMFS datastore space that was freed by the Storage vMotion operation, meaning that Hitachi Dynamic Provisioning software cannot reclaim the free space.

Virtual Disk and Dynamic Provisioning performance

To obtain maximum storage performance for vSphere 4 when using the HUS 100 family storage, follow these best practices:

- Use eagerzeroedthick virtual disk format to prevent warm-up anomalies. Warm-up anomalies occur one time, when a block on the virtual disk is written to for the first time. Zeroedthick is fine for use on the guest OS boot volume where maximum write performance is not required.
- Use at least four RAID groups in the Dynamic Provisioning pool for maximum wide striping benefit.

Virtual disks on standard volumes

Zeroedthick and eagerzeroedthick format virtual period required by the zeroedthick virtual disk on standard LUs. Either virtual disk format provides similar throughput after some write latency.

When deciding whether to use zeroedthick or eagerzeroedthick format virtual disks, keep the

- If you plan to use vSphere 4 Fault Tolerance on a virtu eagerzeroedthick virtual disk format.
- If minimizing the time to create the virtual disk is m performance, use the zeroedthick virtual disk format.
- If maximizing initial write performance is more important than minimizing the time required to the virtual disk, use the eagerzeroedthick format. *Distributing Computing Resource and I/O*

Loads Hitachi Dynamic Provisioning software can balance I/O load in pools of RAID groups. VMware's Distributed Resource Scheduling (DRS) can balance computing capacity in CPU and memory pools. When you use Hmemory into a DRS resource pool and Hitachi Dynamic Provisioning groups into a Dynamic Provisioning pool. Figure 3 shows how resource pool.

Special functions

This chapter will provides details on Modular Volume Migration Manager, Volume Expansion, and Power Savings. The topics covered in this chapter are:

- ❑ [Modular Volume Migration overview](#)
- ❑ [Managing Modular Volume Migration](#)
- ❑ [Volume Expansion \(Growth not LUSE\) overview](#)
- ❑ [Power Savings overview](#)
- ❑ [Viewing volume information in a RAID group](#)

Modular Volume Migration overview

As data gets older and performance requirements decrease over time, Volume Migration can move data from higher performance SAS disk drives to lower cost disk drives. The available free SAS drives can now be used for higher performance data. Your organization can avoid provisioning additional costly SAS drives to satisfy your business needs.

Modular Volume Migration Manager features

The following are Modular Volume Migration Manager features:

- **Data fluidity** - Moves data between RAID groups. Enables you to move data online without host interruption.
- **Secure port mapping** - Security level mapping for SAN ports and virtual ports
- **Intersystem path mapping** - Mapping of data between storage systems.
- **Online volume migrations** - Seamless migration of data volumes.

Modular Volume Migration Manager benefits

The benefits of Modular Volume Migration Manager are:

- **Increased performance** - Does not require host resources to perform tasks so it does not hamper performance on the system. Removes performance bottlenecks.
- **Online configuration capability** - Enables tasks to execute without interruption to normal operation of storage system because of online configuration capability.

Modular Volume Migration task flow

The following is a task flow for Modular Volume Migration:

1. You determine that data on a SAS drive is aging and that needs for the data is not immediate.
2. You determine the old data can move be moved off the high performance SAS drive to a lower performance drive.
3. Select the primary volume.
4. Select the target disk by reserving it.
5. Create a Modular Volume Migration pair and select the pair.
6. Specify the primary volume (typically, a volume number).
7. Select the secondary volume enabling you to cross RAID levels and disk types.
8. You now have a choice of taking the content of a high performance disk (for example, a RAID 10 SAS) and migrating it to a lower performance disk (for example, a RAID 6 or RAID 5 SAS volume).

9. You then set the copy pace priority slower or normal.




Figure 11-1: Modular Volume Migration task flow

Modular Volume Migration Manager specifications

Table 11-1 lists the Modular Volume Migration specifications.

Table 11-1: Volume Migration specifications

Item	Description
Number of pairs	Migration can be performed for the following pairs per array, per system: • 1,023 (HUS 110) • 2,047 (HUS 130 and HUS 150) Note: The maximum number of the pairs is limited when using ShadowImage. For more information, see Using with ShadowImage on page 11-14.
Number of pairs whose data can be copied in the background	Up to two pairs per controller. However, the number of pairs whose data can be copied in the background is limited when using ShadowImage. For more information, see Using with ShadowImage on page 11-14 .
Number of reserved volumes	• 1,023 (HUS 100) • 2,047 (HUS 130 and HUS 150)
RAID level support	RAID 0 (2D to 16D), RAID 1 (1D+1D), RAID 5 (2D+1P to 15D+1P), RAID 1+0 (2D+2D to 8D+8D), RAID 6 (2D+2P to 28D+2P). We recommend using a P-VOL and S-VOL with redundant RAID level. Note that RAID 0 cannot be set for the SAS7.2K disk drive.
RAID level combinations	All combinations are supported.

Table 11-1: Volume Migration specifications (Continued)

Item	Description
Types of P-VOL/S-VOL drives	Volumes consisting of SAS drives can be assigned to any P-VOLs and S-VOLs. You can specify a volume consisting of SAS drives for the P-VOL and the S-VOL.
Host interface	Fibre Channel or iSCSI
Canceling and resuming migration	Migration cannot be stopped or resumed. When the migration is canceled and executed again, Volume Migration copies of the data again.
Handling of reserved volumes	You cannot delete volumes or RAID groups while they are being migrated.
Handling of volumes	You cannot format, delete, expand, or reduce volumes while they are being migrated. You also cannot delete or expand the RAID group. You can delete the pair after the migration, or stop the migration.
Formatting restrictions	You cannot specify a volume as a P-VOL or an S-VOL while it is being formatted. Execute the migration after the formatting is completed.
Volume restrictions	Data pool volume, DMLU, and command devices (CCI) cannot be specified as a P-VOL or an S-VOL.
Concurrent use of unified volumes	The unified volumes migrate after the unification. Using unified volumes on page 11-13.
Concurrent use of Data Retention	When the access attribute is not Read/Write, the volume cannot be specified as an S-VOL. The volume which executed the migration carries over the access attribute and the retention term. For more information, see Using with the Data Retention Utility on page 11-14.
Concurrent use of SNMP Agent	Available
Concurrent use of Password Protection	Available
Concurrent use of LUN Manager	Available
Concurrent use of Cache Residency Manager	The Cache Residency volume cannot be set to P-VOL or S-VOL.
Concurrent use of Cache Partition Manager	Available. Note that a volume that belongs to a partition and stripe size cannot carry over, and cannot be specified as a P-VOL or an S-VOL.
Concurrent use of Power Saving/Power Saving Plus	When a P-VOL or an S-VOL is included in a RAID group for which the Power Saving/Power Saving Plus has been specified, you cannot use Volume Migration. The reserved volumes can specify the Power Saving/Power Saving Plus. Also, the volumes included in the RAID group where the Power Saving/Power Saving Plus is specified can be specified as the reserved volumes. However, Volume Migration is impossible.

Table 11-1: Volume Migration specifications (Continued)

Item	Description
Concurrent use of ShadowImage	A P-VOL and an S-VOL of ShadowImage cannot be specified as a P-VOL or an S-VOL of Volume Migration unless their pair status is Simplex.
Concurrent use of SnapShot	A SnapShot P-VOL cannot be specified as a P-VOL or an S-VOL when the SnapShot volume (V-VOL) is defined.
Concurrent use of TrueCopy or TCE	A P-VOL and an S-VOL of TrueCopy or TCE cannot be specified as a P-VOL or an S-VOL of Volume Migration unless their pair status is Simplex.
Concurrent Use of Dynamic Provisioning	Available. The DP-VOLs created by Dynamic Provisioning and the normal volume can bet as a P-VOL, an S-VOL, or a reserved volume.
Failures	The migration fails if the copying from the P-VOL to the S-VOL stops. The migration also fails when a volume blockade occurs. However, the migration continues if a drive blockade occurs.
Memory reduction	To reduce the memory being used, you must disable Volume Migration and SnapShot, ShadowImage, TrueCopy, or TCE function.

Table 11-2 details reserved volume guard conditions.

Table 11-2: Reserved volume guard conditions

Item	Guard Condition
Concurrent use of ShadowImage	P-VOL or S-VOL.
Concurrent use of SnapShot	P-VOL or S-VOL.
Concurrent use of TrueCopy or TCE	P-VOL or S-VOL or TrueCopy or TCE
Concurrent use of Data Retention	Data Retention volume.
Concurrent use of Dynamic Provisioning	The DP-VOLs created by Dynamic Provisioning
Volume restrictions for special uses	Data pool volume, DMLU, command device (CCI).
Other	Unformatted volume. However, a volume being formatted can be set as reserved even though the formatting is not completed.

Requirements

Table 11-3 shows requirements for Modular Volume Migration Manager.

Table 11-3: Environments and requirements

Item	Description
Specifications	Number of controllers: 2 (dual configuration) Command devices: Max 128 (The command device is required only when CCI is used for the operation of Volume Migration. The command device volume size must be greater than or equal to 33 MB.) DMLU: Max. 1 (the DMLU size must be greater than or equal to 10 GB to less than 128 GB). Size of volume: The P-VOL size must equal the S-VOL volume size.

Supported capacity

Table 11-4 shows the maximum capacity of the S-VOL by the DMLU capacity. The maximum capacity of the S-VOL is the total value of the S-VOL capacity of ShadowImage, TrueCopy, and Volume Migration.

Table 11-4: Maximum S-VOL capacity and corresponding DMLU capacity

S-VOL Number	DMLU Capacity				
	10 GB	32 GB	64 GB	96 GB	128 GB
2	256 TB				
32	1,031 TB	3,411 TB	4,096 TB		
64	983 TB	3,363 TB	6,827 TB	7,200 TB	
128	887 TB	3,267 TB	6,731 TB	7,200 TB	
512	311 TB	2,691 TB	6,155 TB	7,200 TB	
1,024	N/A	1,923 TB	5,387 TB	7,200 TB	
4,096	N/A	N/A	779 TB	4,241 TB	7,200 TB



NOTE: The maximum capacity shown in [Table 11-3](#) is the value smaller than the pair creatable capacity displayed in Navigator 2. This condition is because the pair creatable capacity in Navigator 2 is treated not as the real capacity, but as the value rounded up by the 1.5 TB unit, not as the actual capacity when calculating the S-VOL capacity. The maximum capacity (the capacity of which the pair can be created) reduced by the capacity capable of rounding up by the number of S-VOLs becomes the capacity shown in [Table 11-3](#).

Setting up Volume Migration

This section explains guidelines to observe when setting up Volume Migration.

Setting volumes to be recognized by the host

During the migration, the data is copied to the destination logical volume (S-VOL), and the source logical volume (P-VOL) is not erased ([Figure 11-4 on page 11-11](#)). After the migration, the logical volume destination becomes a P-VOL, and the source logical volume becomes an S-VOL. If the migration stops before completion, the data that has been copied from source logical volume (P-VOL) remains in the destination logical volume (S-VOL). If you use a host configuration, format the S-VOL with Navigator 2 before making it recognizable by the host.

Volume Migration components

Volume Migration system components include:

- Volume Migration volume pairs (P-VOLs and S-VOLs).
- Reserved volume.
- DMLU




Figure 11-2: Components of Volume Migration

Volume Migration pairs (P-VOLs and S-VOLs)

The disk array controls the P-VOL, which is the migration source of the data, and the S-VOL which is the migration destination of the data in a pair. The pair of a P-VOL and an S_VOL is called a migration pair or simply a pair. The P-VOL can be read/written by a host whereas the S-VOL cannot.

Reserved Volume

Volume Migration registers the volume which is the migration destination of the data as a reserved volume before executing the migration in order to shut off the S-VOL from the Read/Write operation by a host beforehand. When executing the migration using Navigator 2. The volume that is selectable as an S-VOL is the reserved volume only. The reserved volume is a volume which is the migration destination of the data when the migration is executed, and data is not guaranteed.

DMLU

DMLU refers to Differential Management Logical Unit and a volume exclusive for storing differential information of a P-VOL and an S-VOL of a Volume Migration pair. To create a Volume Migration pair, you need to prepare one DMLU in the array.

The differential information of all Volume Migration pairs is managed by this singular DMLU. However, a volume that is set as the DMLU is not recognized by a host (it is hidden). The following table differentiates supportable platforms by the DMLU for both the AMS 2000 and SMS 100 product families and the HUS series.

Item	AMS 2000/SMS 100	HUS
Target feature	ShadowImage Copy on Write SnapShot TrueCopy remote replication' TrueCopy Extended Distance Modular Volume Migration	ShadowImage TrueCopy Remote Replication Modular Volume Migration
Assignable Number	2	1

As shown in [Figure 11-3](#), the array accesses the differential information stored in the DMLU and refers to and updates it in the copy processing to synchronize the P-VOL and the S-VOL and the processing to manage the difference of the P-VOL and the S-VOL.




Figure 11-3: Flow of operations using the DMLU

The createable pair capacity is dependent on the DMLU capacity. If the DMLU does not have enough capacity to store the pair differential information, the pair cannot be created. In this case, a pair can be added by expanding the DMLU. The DMLU capacity is a minimum of 10 GB and the maximum of 128 GB. Refer to the section that details the number of creatable pairs according to the capacity and the total capacity of the volume to be paired.

DMLU precautions

This section details DMLU precautions for setting, expanding, and removing.

Precautions for setting DMLUs include

- The volume belonging to RAID 0 cannot be set as a DMLU
- You cannot complete setting the unified volume as a DMLU if the capacity of each unified volume becomes less than 1 GB on average. For example, when setting a volume of 10 GB as a DMLU, if the volume consists of 11 sub-volumes, it cannot be set as a DMLU.
- The volume assigned to the host cannot be set as a DMLU.

Precautions for expanding DMLUs include

When expanding DMLUs, select a RAID group which meets the following conditions:

- The drive type and the combination are the same as the DMLU.
- A new volume can be created.
- A sequential free area for the capacity to be expanded exists.

Precautions for removing DMLUs include

- When either pair of ShadowImage, TrueCopy, or Volume Migration exists, the DMLU cannot be removed.
- The volume after the DMLU removing becomes the unformatted status. You can reset the DMLU as unformatted, but if you use it for another purpose, you need to format the volume.



NOTE: When the migration is completed or stopped, the latest data is stored in a logical volume (P-VOL).



NOTE: When formatting, format the S-VOL. If the P-VOL is formatted by mistake, some data may be lost.




Figure 11-4: Volume Migration host access

VxVM

Do not allow the P-VOL and S-VOL to be recognized by the host at the same time.

MSCS

Do not allow the P-VOL and S-VOL to be recognized by the host at the same time.

- Do not place the MSCS Quorum Disk in CCI.
- Shutdown MSCS before executing the CCI sync command.

AIX

- Do not allow the P-VOL and S-VOL to be recognized by the host at the same time.

Windows 2000/Window Server

- When specifying a command device in the configuration definition file, specify it as Volume GUID. For more information, see the Command Control Interface (CCI) Reference Guide.
- When the source volume is used with a drive character assigned, the drive character is taken to the migration volume. However, when both volumes are recognized at the same time, the drive character can be assigned to the S-VOL through a host restart.

Linux and LVM

Do not allow the P-VOL and S-VOL to be recognized by the host at the same time.

Windows 2000/Windows Server and Dynamic Disk

Do not allow the P-VOL and S-VOL to be recognized by the host at the same time.

UNMAP Short Length Mode

Enable UNMAP Short Length Mode when connecting to Windows 2012. If you not enable it, the UNMAP commands may not complete because of a time-out.

Performance

- Migration affects the performance of the host I/O to P-VOL and other volumes. The recommended Copy Pace is Normal, but if the host I/O load is heavy, select Slow. Select Prior to shorten the migration time; however, this can affect performance. The Copy Pace can be changed during the migration.
- The RAID structure of the P-VOL and S-VOL affects the host I/O performance. The write I/O performance concerning a VOL, which migrates from a disk area, consists of the SAS drives, the SAS7.2K drives or the SAS (SED) drives to a disk area is lower than that concerning a volume that consists of the lower cost drives.
- Do not concurrently migrate logical volumes that are in the same RAID group.
- Do not run Volume Migration from/to volumes that are in Synchronizing status with ShadowImage initial copy, or in resynchronization in the same RAID group. Additionally, do not execute ShadowImage initial copy or resynchronization in the case where volumes involved in the ShadowImage initial copy or resynchronization are from the same RAID group.

- It is recommended that Volume Migration is run during periods of low system I/O loads.

Using unified volumes

A unified logical volume can be used as a P-VOL or S-VOL as long as their capacities are the same (they can be composed of different number of volumes).

The number of volumes that can be unified as components of a P-VOL or S-VOL is 128 (Figure 11-5).




Figure 11-5: Unified volumes assigned to P-VOL or S-Vol (unification)

The volumes, including the unified volumes assigned to the P-VOL and S-VOL, cannot be on the same RAID level, or have the same number of disks (Figure 11-6 on page 11-13 and Figure 11-7 on page 11-13).




Figure 11-6: RAID level combination




Figure 11-7: Disk number combination

Do not migrate when the P-VOL and the S-VOL volumes belong to the same RAID group.




Figure 11-8: Volume RAID group combinations

Using with the Data Retention Utility

The volume that executed the migration carries the access attribute and the retention term set by Data Retention, to the destination volume. If the access attribute is not Read/Write, the volume cannot be specified as an S-VOL.

The status of the migration for a Read Only volume appears in [Figure 11-9 on page 11-14](#). When the migration of the Read Only VOL0 to the VOL1 is executed, the Read Only attribute is carried to the destination volume. Therefore, VOL0 is Read Only. When the migration pair is released and VOL1 is deleted from the reserved volume, a host can Read/Write to the VOL1.




Figure 11-9: Read only

Using with ShadowImage

The array limits the ShadowImage and Volume Migration pairs to 1,023 (AMS2100) and 2,047 (AMS2300). The numbers of migration pairs that can be executed are calculated by subtracting the number of ShadowImage pairs from the maximum number of pairs.

The number of copying operations that can be performed in the background is called the copying multiplicity. The array limits the copying multiplicity of the Volume Migration and ShadowImage pairs to 4 per controller. When Volume Migration is used with ShadowImage, the copying multiplicity of Volume Migration is 2 two per controller because Volume Migration and ShadowImage share the copying multiplicity.

Note that at times, copying does not start immediately ([Figure 11-10 on page 11-15](#) and [Figure 11-11 on page 11-15](#)).




Figure 11-10: Copy operation where Volume Migration pauses




Figure 11-11: Copy operation where ShadowImage operation pauses

Using with Cache Partition Manager

It is possible to use Volume Migration with Cache Partition Manager. Note that a volume that belongs to a partition cannot carry over. When a migration process completes, a volume belonging to a partition is changed to destination partition.

Concurrent Use of Dynamic Provisioning

Consider the following points when using Volume Migration and Dynamic Provisioning together. For the purposes of this discussion, the volume created in the RAID group is called a normal volume and the volume created in the DP pool that is created by Dynamic Provisioning is called a DP-VOL.

- When using a DP-VOL as a DMLU

Check that the free capacity (formatted) of the DP pool to which the DP-VOL belongs is 10 GB or more, and then set the DP-VOL as a DMLU. If the free capacity of the DP pool is less than 10- GB, the DP-VOL cannot be set as a DMLU.

- Volume type that can be set for a P-VOL or an S-VOL of Volume Migration

The DP-VOL created by Dynamic Provisioning can be used for a P-VOL or an S-VOL of Volume Migration. The following table shows a combination of a DP-VOL and a normal volume that can be used for a P-VOL or an S-VOL of Volume Migration. [Table 11-5](#) details the combination of a DP-VOL and a normal VOL.

Table 11-5: Combination of a DP-VOL and a normal VOL

Volume Migration P-VOL	Volume Migration S-VOL	Contents
DP-VOL	DP-VOL	Available.
DP-VOL	Normal VOL	Available. In this combination, the migration copy takes about the same time it takes when the normal volume is P-VOL.
Normal VOL	DP-VOL	Available. In this combination, executing initial copying, the DP pool of the same capacity as the normal volume (P-VOL) is used.



NOTE: When both the P-VOL and the S-VOL use DP-VOLs, a pair cannot be created by combining the DP-VOLs which have different setting of Enabled/ Disabled for Full Capacity Mode.

- Usable Combination of DP Pool and RAID Group

The following table shows a usable combination of DP Pool and RAID group. [Table 11-6](#) details the contents of a volume migration P-VOL and S-VOL.

Table 11-6: Contents of Volume Migration P-VOL and S-VOL

Volume Migration P-VOL and S-VOL	Contents
Same DP pool	Not available
Different DP pool	Available
DP pool and RAID group	Available
RAID group and DP pool	Available

- Pair status at the time of DP pool capacity depletion
When the DP pool is depleted after operating Volume Migration which uses the DP-VOL created by Dynamic Provisioning, the pair status of the pair concerned may be an error.
The following table shows the pair statuses before and after the DP pool capacity depletion. When the pair status becomes an error caused by the DP pool capacity depletion, add the DP pool capacity whose capacity is depleted, and execute Volume Migration again. [Table 11-7](#) details pair statuses before and after the DP pool capacity depletion.

Table 11-7: Pair Statuses Before and after DP Pool capacity depletion

Pair Statuses before the DP Pool Capacity Depletion	Pair Statuses after the DP Pool Capacity Depletion belonging to P-VOL	Pair Statuses after the DP Pool Capacity Depletion belonging to S-VOL
Copy	Copy Error	Error
Completed	Completed	Completed
Error	Error	Error



NOTE: When write is performed to the P-VOL to which the capacity depletion DP pool belongs, the copy cannot be continued and the pair status becomes an error.

- DP pool status and availability of Volume Migration operation
When using the DP-VOL created by Dynamic Provisioning for a P-VOL or an S-VOL of Volume Migration, Volume Migration operation may not be executed depending on the status of the DP pool to which the DP-VOL belongs. The following table shows the DP pool status and availability of Volume Migration operation. When Volume Migration operation fails due

to the DP pool status, correct the DP pool status and execute Volume Migration operation again. [Table 11-8](#) details DP Pool statuses and availability of the volume migration operation.

Table 11-8: DP Pool statuses and availability of Volume Migration operation

Operation	Normal	Capacity in Growth	Capacity Depletion	Regressed	Blocked	DP in Optimization
Executing	○	X	○	○	X	○
Splitting	○	○	○	○	○	○
Canceling	○	○	○	○	○	○

Executing-Normal: Refer to the status of the DP pool to which the DP-VOL of the S-VOL belongs. If the status exceeds the DP pool capacity belonging to the S-VOL by Volume Migration operation, Volume Migration operation cannot be executed.

Executing-Capacity Depletion: Refer to the status of the DP pool to which the DP-VOL of the P-VOL belongs. If the status exceeds the DP pool capacity belonging to the P-VOL by Volume Migration operation, Volume Migration operation cannot be executed.

Also, When the DP pool was created or the capacity was added, the formatting operates for the DP pool. If Volume Migration is performed during the formatting, depletion of the usable capacity may occur. Since the formatting progress is displayed when checking the DP pool status, check if the sufficient usable capacity is secured according to the formatting progress, and then start Volume Migration operation.

Executing-DP in Optimization

- Operation of the DP-VOL during Volume Migration use

When using the DP-VOL created by Dynamic Provisioning for a P-VOL or an S-VOL of Volume Migration, any of the operations among the capacity growing, capacity shrinking, volume deletion, and Full Capacity Mode changing of the DP-VOL in use cannot be executed. To execute the operation, split the Volume Migration pair of which the DP-VOL to be operated is in use, and then execute it again.

- Operation of the DP pool during Volume Migration use

When using the DP-VOL created by Dynamic Provisioning for a P-VOL or an S-VOL of Volume Migration, the DP pool to which the DP-VOL in use belongs cannot be deleted. To execute the operation, split the Volume Migration pair of which the DP-VOL is in use belonging to the DP pool to be operated, and then execute it again. The attribute edit and capacity addition of the DP pool can be executed usually regardless of Volume Migration pair.

Concurrent Use of Dynamic Tiering

The considerations for using the DP pool or the DP-VOL whose tier mode is enabled by using Dynamic Tiering are described. For the detailed information related to Dynamic Tiering, refer to the *Hitachi Unified Storage Provisioning Configuration Guide* (MK-91DF8277).

When using a DP-VOL whose tier mode is enabled as a DMLU, check that the free capacity (formatted) of the Tier other than SSD/FMD of the DP pool to which the DP-VOL belongs is more than or equal to the DP-VOL used as a DMLU, and then set it. At the time of the setting, the entire capacity of the DMLU is assigned from the first tier. However, the tier configured by the SSD/FMD is not assigned to the DMLU. Furthermore, the area assigned to the DMLU is out of the relocation target.

Dirty Data Flush Limit number

The Dirty Data Flush Limit setting number determines whether to limit the number of executing the processing for flushing the dirty data in the cache to the drive at the same time. This setting is effective when Volume Migration is enabled. When all the volumes in the array are created in the RAID group of RAID 1 or RAID 1+0 configured of the SAS drives and in the DP pool.

If this setting is enabled, the dirty data flush number is limited even though Volume Migration is enabled. When the Dirty Data Flush number is limited, Write I/O time becomes long. To correct his condition, disable the Dirty Data Flush Number Limit setting when using Volume Migration.

When the Dirty Data Flush Number Limit setting is Enabled in ShadowImage, disable the setting once, and then execute Volume Migration. After it is completed, set the Dirty Data Flush Number Limit to **Enabled** again.

Load Balancing function

The Load Balancing function applies to a Volume Migration session. When the Load Balancing function is activated for a Volume Migration pair, the ownership of the P-VOL and S-VOL changes to the same controller. When the pair state is in the Synchronizing state, the ownership of the pair changes across the cores, but not across the controllers.

Contents related to the connection with the host

The VMware EXS clones the virtual machine. Although, the ESX clone function and Volume Migration can be linked, caution is required for the performance at the time of the execution.

When the volume which becomes the ESX clone destination is a Volume Migration P-VOL pair whose status is Copy, the data may be written to the S-VOL for writing to the P-VOL. Since the background data copy is executed,

the load on the drive becomes heavy. Therefore, the time required for a clone may become longer and the clone may be terminated abnormally in some cases.

To avoid abnormal termination, set the copy pace of the Volume Migration pair to Slow, or the operation to make a migration after executing the ESX clone. The same abnormal termination may occur when you execute functions such as migrating the virtual machine, deploying from the template, and inflating the virtual disk and Space Reclamation.

Hitachi recommends you enable UNMAP Short Length Mode when connecting to VMware. If you do not enable this feature, the UNMAP commands may not complete because of a time-out.

Modular Volume Migration operations

To perform a basic volume migration operation

1. Verify that you have the environments and requirements for Volume Migration (see [Preinstallation information for Storage Features on page 3-20](#)).
2. Set the DMLU (see [Adding reserved volumes on page 11-24](#)).
3. Create a volume in RAID group 1 and format it. The size of the volume must be same as the one you are migrating. When the volume that has already been formatted is to be the volume of the migration destination, it is not necessary to format it again.
4. Set volume X as a reserved volume (see [Adding reserved volumes on page 11-24](#)).
5. Migrate. Specify the VOL0 and the VOL1 for the P-VOL and the S-VOL, respectively.



NOTE: You cannot migrate while the reserved volume is being formatted.

6. Confirm the migration pair status. When the copy operation is in progress normally, the pair status is displayed as Copy and the progress rate can be referred to (see [Confirming Volume Migration Pairs on page 11-29](#)).
7. When the migration pair status is Completed, release the migration pair. The relation between the P-VOL/S-VOL of VOL0/VOL1 is released and the two volumes are returned to the status before the migration executing.



NOTE: When the pair status is displayed as Error, the migration failed because a failure occurred in the migration progress. When this happens, delete the migration pair after recovering the failure and execute the migration again.

8. When the migration is complete, VOL0 has been migrated to the RAID group 1 where VOL1 was created, and VOL1 has been migrated to the RAID group 0 where VOL0 was. If the migration fails, VOL0 is not migrated from the original RAID group 0 (see [Migrating volumes on page 11-26](#)).
9. The VOL1 migrated to the RAID group 0 can be specified as an S-VOL when the next migration is executed. If the next migration is not scheduled, delete VOL1 from the reserved volume. The LU1 deleted from the reserved volume can be used for the usual system operation as a formatted volume (see [Migrating volumes on page 11-26](#)).

Managing Modular Volume Migration

This section describes how to migrate volumes using the Modular Volume Migration tool.

Volume Migration runs under the Volume Migration menu under the Replication menu in the Navigation bar.

Pair Status of Volume Migration

Volume Migration can check the status of the migration pair using Navigator 2. The relation between the pair status changes of the Volume Migration and the operations of Volume Migration is shown in [Figure 11-12](#).




Figure 11-12: Volume Migration Pair Status Transitions

Setting the DMLU

Refer to the section [DMLU on page 11-8](#) for the description and setting related to the DMLU.

To designate the DMLU

1. Select the DMLU icon in the Setup tree view of the Replication tree view.
The Differential Management Logical Units list displays.
2. Click **Add DMLU**.

The Add DMLU screen displays as shown in [Figure 11-13](#).

Add DMLU

* Volumes to add as DMLU

VOL	Capacity	RAID Group	DP Pool	RAID Lev
0001	50.0GB	000	N/A	RAID5(4D)
1001	10.0GB	000	N/A	RAID5(4D)
1002	10.0GB	000	N/A	RAID5(4D)
1003	64.0GB	000	N/A	RAID5(4D)

* Required field

OK Cancel

Figure 11-13: Add DMLU window

3. Select one of the volumes you want to set as the DMLU and click **OK**.
A message displays. Select the checkbox and click **Confirm**.

Removing the designated DMLU

This section details how to remove the designated DMLU. Note that when Volume Migration, ShadowImage, and TrueCopy pairs exist, the DMLU cannot be released.

To remove the designated DMLU

1. Select the DMLU icon in the Setup tree view of the Replication tree view.
The Differential Management Logical Units list displays.
2. Select the volume you want to remove, and click Remove DMLU.
A message displays. Click **Close**.

Adding the designated DMLU

To add the designated DMLU

1. Select the DMLU in the Setup tree view of the Replication tree view.
The Differential Management Logical Units list displays.
2. Select the volume you want to add and click Add DMLU Capacity.

The Add DMLU Capacity screen displays as shown in [Figure 11-14](#).
Add DMLU Capacity - 1003

Differential Management Logical Unit Capacity Property Help

Enter the information for volume added to a DMLU.

VOL : 1003
 Total Capacity: 64.0GB
 RAID Group: 000
 RAID Level: RAID5(4D+1P)
 Drive Type: SAS

*New Capacity: 96 GB
 To 128GB

Use RAID Group: **Assignable RAID Groups**

Rows/Page: 25 | Page 1 of 1

RAID Group	RAID Level	Capacity	
		Total	Free
000	RAID5(4D+1P)	8.9TB	8.9TB

Filter Filter Off

* Required field

OK Cancel

Figure 11-14: Add DMLU Capacity window

3. Enter a capacity reflecting the expansion in units of GB to the New Capacity and click **OK**.
4. When the DMLU is a volume that belongs to the RAID group, select the RAID group that acquires the capacity to be expanded.
5. Select the RAID group that can acquire the capacity to be expanded in the sequential free area. A message displays. Click **Close**.

Adding reserved volumes

When mapping mode is enabled, the host cannot access the volume if it has been allocated to the reserved volume.



NOTE: When the mapping mode displays, the host cannot access the volume if it has been allocated to the reserved volume. Also when the mapping mode is enabled, the host cannot access the volume if the mapped volume has been allocated to the reserved volume.



WARNING! Stop host access to the volume before adding reserved volumes for migration.

To add reserved volumes for volume migration

1. Start Navigator 2 and log in. The Arrays window appears

- Click the appropriate array.
- Click **Show & Configure Array**.
- Select the **Reserve Volumes** icon in the Volume Migration tree view as shown in Figure 11-15.




Figure 11-15: Reserve Volumes window

- Click **Add Reserve Volumes**. The **Add Reserve Volumes** panel displays as shown in Figure 11-16.



Figure 11-16: Add Reserve Volumes panel

- Select the volume for the reserved volume and click **OK**.
- In the resulting message boxes, click **Confirm**.
- In the resulting message boxes, click **Close**.

Deleting reserved volumes

When canceling or releasing the volume migration pair, delete the reserve volume, or change the mapping. For more information, see [Table 11-1 on page 11-3](#) and [Setting up Volume Migration on page 11-7](#).



NOTE: Be careful when the host recognizes the volume that has been used by Volume Migration. After releasing the Volume Migration pair or canceling Volume Migration, delete the reserved volume or change the volume mapping.

To delete reserved volumes

1. From the Reserve Volumes dialog box, select the volume to be deleted and click **Remove Reserve Volumes** as shown in [Figure 11-17](#).

Reserve VOLS

HUS110_91100026 > Replication > Volume Migration > Reserve VOLS




Figure 11-17: Reserve Volumes window - volume selected for deletion

2. In the resulting message boxes, click **Confirm**.
3. In the resulting message boxes, click **Format VOL** if you want to format the removed reserved volumes. Otherwise click **Close**.

Migrating volumes

To migrate volumes

1. Start Navigator 2 and log in. The Arrays window appears
2. Click the appropriate array.
3. In the navigation tree view, click the **Replication** list and click the **Volume Migration**.
4. Click **Migration Pairs**. The Migration Pairs screen displays as shown in [Figure 11-18](#).

Migration Pairs

HUS110_91100026 > Replication > Volume Migration > Migration Pairs




Figure 11-18: Before pair creation

5. Click **Create Pair**.

The Create Migration dialog box displays as shown in [Figure 11-19](#).

Create Migration Pair




Figure 11-19: Create Volume Migration Pair window

6. Select the volume for the P-VOL, and click **OK**.
7. Select the volume for the S-VOL and Copy Pace, click **OK**.
8. Follow the on-screen instructions.

Changing copy pace

The pair copy pace can only be changed if it is in either Copy or Waiting status. There are three options for this feature:

- **Prior** - The copying pace from the previous copying session.
- **Normal** - The default copying pace.
- **Slow** - A copying pace that requires more time to complete than the default pace.



NOTE: Normal mode is the default for the Copy Pace. If the host I/O load is heavy, performance can degrade. Use the Slow mode to prevent performance degradation. Use the Prior mode only when the P-VOL is rarely accessed and you want to shorten the copy time.

To change the copy pace

1. Start Navigator 2 and log in. The Arrays window appears
2. Click the appropriate array.
3. In the navigation tree view, click the **Replication** list and click **Volume Migration**.
4. Click **Migration Pairs**. The Migration Pairs screen displays as shown in [Figure 11-20](#).
5. Select the pair whose copy pace you are modifying, and click **Change Copy Pace**.





Figure 11-20: Launching the Change Copy Pace window

6. The Change Copy Pace dialog box appears, as shown in [Figure 11-20](#).

7. Select the copy pace and click **OK**. The Change Copy Pace panel appears, as shown in [Figure 11-21](#).

Change Copy Pace - Pair0002:0001




The dialog box titled "Change Copy Pace" has a green header bar. Below the header, it says "Select the copy pace." There are three radio button options: "Fast (Copy is prioritized.)", "Medium", and "Slow (Host access is prioritized.)". The "Slow" option is selected. At the bottom right, there are "OK" and "Cancel" buttons. A mouse cursor is pointing at the "OK" button.

Figure 11-21: Change Copy Pace dialog box

8. In the resulting message box, click **OK**, as shown in.
9. Follow the on-screen instructions.

Confirming Volume Migration Pairs

[Figure 11-22](#) shows the pair migration status.



The "Migration Pairs" window shows a table with migration details. The breadcrumb path is "HUS110_91100026 > Replication > Volume Migration > Migration Pairs". The table has columns: Primary Volume, Secondary Volume, Capacity, Copy Pace, Owner, and Status. There is one row with a checkmark in the Primary Volume column, volume number 0002, secondary volume 0001, capacity 10.0GB, copy pace Medium, owner HSNM2, and status Copy(80%).

Primary Volume	Secondary Volume	Capacity	Copy Pace	Owner	Status
☒ 0002	0001	10.0GB	Medium	HSNM2	Copy(80%)

Figure 11-22: Migration Pairs window - P-VOL and S-VOL migration

- **P-VOL** - The volume number appears for the P-VOL.
- **S-VOL** - The volume number appears for the S-VOL.
- **Capacity** - The capacity appears for the P-VOL and S-VOL.
- **Copy Pace** - The copy pace appears.
- **Owner** - The owner of the migration appears. For Adaptable Modular Storage, this is Storage Navigator Modular 2. For any other, this is CCI.
- **Pair Status** - The pair status appears and includes the following items:
 - **Copy** - Copying is in progress.
 - **Waiting** - The migration has been executed but background copying has not started yet.
 - **Completed** - Copying completed and waiting for instructions to release the pair.

- **Error** - The migration failed because the copying was interrupted. The number enclosed in parentheses is the failure error code. When contacting service personnel, give them this error code.

Releasing Volume Migration pairs

A pair can only be split if it is in Completed or Error status.

To split Volume Migration pairs

1. Start Navigator 2 and log in. The Arrays window appears
2. Click the appropriate array.
3. In the navigation tree view, click the **Replication** list and click the **Volume Migration**.
4. Click **Migration Pairs**. The Migration Pairs screen displays as shown in Figure 11-23.



Figure 11-23: Migration Pairs - pair releasing

5. Select the migration pair to release, and click **Remove Pairs**.
6. Follow the on-screen instructions.

If you cancel the Migration Pair, you may have to wait up to five seconds before the following tasks can be performed:

- Creating a pair in ShadowImage, when the volume specified as the S-VOL of the canceled pair is an S-VOL.
- Creating a a pair in TrueCopy where the volume specified is the S-VOL of the canceled pair.
- Volume Migration where the volume specifies is the S-VOL of the canceled pair.
- Deleting the volume specified that is the S-VOL of the canceled pair.
- Removing the DMLU.
- Expanding the capacity of the DMLU.

Canceling Volume Migration pairs

A pair can only be canceled if it is in the **Copy** or **Waiting** status.



NOTE: When the migration starts, it cannot be stopped. If the migration is canceled, the data is copied again when you start over.

To cancel a migration

1. Start Navigator 2 and log in. The Arrays window appears
2. Click the appropriate array.
3. In the navigation tree view, click the **Replication** list and click **Volume Migration**.
4. Click **Migration Pairs**. The Migration Pairs screen displays as shown in [Figure 11-24](#).
5. Select the Volume Migration pair that you want to cancel and click **Cancel Migrations**.




Figure 11-24: Migration Pairs - pair cancellation

6. Follow the on-screen instructions.

Note that if you cancel the migration pair, you will not be able to perform any of the following tasks for up to five seconds after the cancel operation.

- Create a ShadowImage pair which specifies the volume as the S-VOL of the canceled pair.
- Create a TrueCopy pair which specifies the volume as the S-VOL of the canceled pair.
- Create a Migration pair which specifies the S-VOL of the canceled pair.
- Delete a volume which specifies the S-VOL of the canceled pair.
- Shrink a volume which specifies the S-VOL of the canceled pair.
- Remove the DMLU.
- Expand DMLU capacity.

If you cancel the migration pair, you will not be able to do any tasks related to migration pairs up to five minutes.

Volume Expansion (Growth not LUSE) overview

This section provides information to guide you through the procedure to increase the size of an existing volume on the storage array by adding one or more existing volumes to it. It also includes procedures to remove volumes that have been added.

The Volume Expansion feature provides the capability of combining two or more existing volumes into a single unit. The procedure includes designating a main volume and then adding other ("sub") volumes to it. The expanded volume is called a unified volume.

This feature is different from the volume "grow" (expand) feature, which allows you to expand the size of an existing volume using available free space in a RAID group to which it belongs.

The Volume Expansion can be reversed by removing the last volume combined with the main volume. The unified volume can also be separated back into the original volumes.

Volume Expansion features

The following are Volume Expansion features:

- **Enables volume combination** - Enables you to combine two or more existing volumes into a single unit.
- **Tiered volume creation** - Enables you to designate a main volume and a sub volume.
- **Volume Expansion reversal** - Volume Expansion can be reversed by removing the last volume combined with the main volume.

Volume Expansion benefits

By combining volumes, you reduce the number of objects the storage system firmware has to inspect. This creates the following benefits:

Volume management ease - Reduces the number of volumes, the storage system firmware has to track, in turn, easing management of objects on the system.

Increased performance - Reduces the time required to track and manage your storage system, in turn, improving overall system performance.

Volume Expansion task flow

Do not skip any steps and make sure that you follow the instructions carefully. If you do not execute a procedure correctly, data in the array can be lost and the unified volume will not be created. The following process details general steps in Volume Expansion.

1. Back up the unified volumes before modifying them.

2. Format the unified volumes to delete the volume label which the operating system adds to volumes.
3. Create unified volumes only from volumes within the same array.
4. You must format a volume that is undefined before you can use it.

Displaying Unified Volume Properties

To display the list of unified volumes

1. In the Arrays window, select the array whose date and time you want to update, and either click **Show and Configure Array**, or double click the name of the array.

The Array window and the Explorer tree are displayed.

2. In the Explorer tree, expand the Groups menu.
3. Click the **Volumes** tab. The Volumes window is displayed.
4. Click **Change VOL Capacity**.

Selecting new capacity

To select new capacity

1. In the Volume Expansion window, click **Create Unified Volume**. The Create Unified Volume dialog box is displayed. It shows the list of main and sub volumes that are available to create unified volumes.
2. In the Create Unified Volume dialog box, select the main volume and sub-volume units and click **OK**. A warning message is regarding the mismatch of RAID levels and hard drive types displayed.
3. To create the designated unified volume, click the to agree that you have read the warning, and then click **Confirm**. Navigator 2 creates the designated unified volume and displays a confirmation message that the unit has been created, and then displays the Volume Expansion window as described above.
4. Click **Close** to exit the message box and return to the Volume Expansion window.
5. Observe the Volume Expansion window and verify that the designated unified volume is listed correctly.

Modifying a unified volume

To modify a unified volume

1. Click the name of the unified volume. The Volume Expansion window is replaced with a window that displays the properties of the selected unified volume and the sub-volume(s) it contains. The properties are described in the table above.

In addition to the properties tables, the window contains the following function buttons:

- Add Volumes
- Separate Last Volume

- Separate All Volumes
2. Click the function button needed to accomplish the desired task. Each button displays a dialog box for the selected function. In addition to the information below, the dialog box for each function has its own help page.

Add Volumes

To add a volume to a unified volume

1. In the unified volume properties window, click **Add Volumes**. The Add Volumes dialog box is displayed. The dialog box includes a table that displays the parameters of the selected unified volume, and a table that lists the available volumes that can be added to the existing unified volume.
2. Click the to the left of the name of the volume that you want to add to the unified volume.
3. Click **OK**. A warning message regarding RAID levels and drive types is displayed. The warning message also includes information that the data in the volume that is added will be destroyed.
4. To add the selected volume to the unified volume, click the to agree that you have read the warning message, and then click **Confirm**. A message box confirming that the volume has been added is displayed.
5. Click **Close** to exit the message box and return to the unified volume properties window.
6. Observe the contents of the window and verify that the volume has been added.

Separate Last Volume

This process is the reverse of adding a volume to a unified volume.

To separate the last volume

1. In the Arrays window, select the array whose date and time you want to update, and either click **Show and Configure Array**, or double click the name of the array.

The Array window and the Explorer tree are displayed.

2. In the Explorer tree, expand the Settings menu to show the list of available functions.
3. In the expanded menu, select Volume Expansion. The Volume Expansion window is displayed.

It shows the list of unified volumes in the array and a set of parameters for each listed unit.

4. In the Volume Expansion window, click the volume that you want to separate.
5. In the unified volume property window, click **Separate Last Volumes**. A confirmation dialog box is displayed.

6. In the confirmation dialog box, click the to agree that you have read the warning message, and then click **Confirm**. A message box stating that the volume has been successfully separated is displayed.
7. Click **Close** to exit the message box and return to the unified volume properties window.
8. Observe the contents of the window and verify that the volume was separated from the unified volume.

Separate All Volumes

To separate a unified volume into the original volumes that were used to create it

1. In the Arrays window, select the array whose date and time you want to update, and either click **Show and Configure Array**, or double click the name of the array.
The Array window and the Explorer tree are displayed.
2. In the Explorer tree, expand the Settings menu to show the list of available functions.
3. In the expanded menu, select **Volume Expansion**. The Volume Expansion window is displayed.
It shows the list of unified volumes in the array and a set of parameters for each listed unit.
4. In the LUN Expansion window, click the volume that you want to separate.
5. In the unified volume property window, click **Separate All LUs**. A confirmation dialog box is displayed.
6. In the confirmation dialog box, click the to agree that you have read the warning message, and then click **Confirm**.
7. Click **Close** to exit the message box and return to the unified volume properties window.
8. Observe the contents of the window and verify that the volume was separated from the unified volume.

Power Savings overview

Information technology (IT) executives are increasingly aware of how energy usage and costs affects their company and the environment. For example, many companies are running the power equipment in their data centers at maximum capacity, which is needed to run and cool computing gear.

Excessive power and cooling demands can lead to failures, and as many data centers are running at dangerous levels of power consumption, they are at risk of failing due to a power shortfall. The Hitachi Unified Storage systems enable companies to reduce energy consumption and significantly reduce the cost of storing and delivering information.

The Power Saving feature, which can be invoked on an as-needed basis, reduces rising energy and cooling costs, and strengthens your security infrastructure. Power Saving reduces electricity consumption by powering down the spindles of unused drives (stopping the rotation of unused disk drives) that configure a redundant array of independent disks (RAID) group. The drives can then be powered back up quickly when the application requires them.

Power Saving is particular useful for businesses that have large archived data or virtual tape library applications where the data is accessed infrequently or for a limited period of time.

In keeping with the Hitachi commitment to environmental responsibility without compromising availability or reliability, the Power Savings Service is available on Fibre Channel (FC) disk drives on all HUS systems.

Power Saving features

- **Spin down** - Slowing or halting volumes and RAID groups in any selected RAID group when they are not being accessed by an application.
- **Spin IP**- Quick restarting of volumes when required.
- **Support for broad portfolio** - Occurs on the SAS disk drives. It also supports both Fibre Channel and iSCSI host interfaces.
- **Automatic power cycles** - Power cycles implemented automatically with no user intervention required.
- **High number of cycles** - Disk drives used in the systems of the HUS family are rated for at least 50,000 contact start-stop cycles.
- **Disk drive safety** - While some power saving processes can damage a disk drive, Hitachi Power Savings is designed in a way to protect drives from degradation.
- **Persistent disk drive integrity** - Disk drives spin up monthly for a six-minute health check.
- **Server-based software command execution** - Spin down and spin up commands occur directly on the server where the integrated application resides.

Power Saving benefits

The following are Power Saving benefits:

- **Reduce power utilization immediately** - Disk drive spin up and spin down cycles are integrated into applications that are scheduled to run infrequently.
- **Increase total data storage capacity** - Data storage capacity can be increased by as much as 50 percent.
- **General power consumption reduction (GB/kWh)** - Power consumption can be reduced substantially.
- **Transparency to user** - Power cycles are implemented automatically, with no user intervention required.
- **Power reduction by spin down/up** - Disk drives that are spun down in power savings mode consume very little or no power.
- **Cost reduction** - Power reduction decreases cost involved in having active system running.
- **Environmental benefit** - Assists in creating an environment-friendly environment to meet organization and government requirements.

Power Saving task flow

The following steps detail the task flow of the Account Authentication configuration process.

1. You determine that your storage system is consuming too much power and decide to implement Power Savings to bring down cost and increase performance on your system.

2. You launch the Power Savings feature on your storage system.




Figure 11-25: Power Savings task flow

Power Saving specifications

Table 11-9 lists Power Saving specifications.

Table 11-9: Power Saving specifications

Item	Specification
RAID level	Any RAID level supported by the array.
Start of the spin-down operation	When spinning down the drives, instruct the spin-down to the RAID group from Navigator 2. Specify the command monitoring time also at the time of instructing the spin-down. According to the instructed command monitoring time, monitor the command or monitor the I/O issuance from the host or the application to the RAID group to which the spin-down was instructed. The spin-down is done when no command is issued during the command monitoring. When a command is issued during the command monitoring, the disk array and RAID group are judged to be in use and therefore the spin-down fails.
Command monitoring time	Command monitoring time: Can be set in the range of 0 to 720 minutes in units of minute. The default of the command monitoring time is one minute. If you can manage the operation for using the RAID group and want to spin down immediately, specify the command monitoring time to 0 minute. The command monitoring is terminated immediately and migrated to the spin-down processing. Even if the command monitoring time is specified as 0 minute, when the uncompleted command remains in the array for the target RAID groups, the spin-down fails. When a drive fails occurred, the spin-down executed after a drive reconstruction completed.
When an instruction to spin down is issued to two or more RAID groups at the same time	RAID groups are spun down in ascending order of the RAID group numbers. The command monitoring is done for specified minute for the first RAID group. For the second and following RAID groups, the command monitoring is done until the spin down occurs.

Table 11-9: Power Saving specifications (Continued)

Item	Specification
Instructing spin-down during command monitoring	• If the spin-down is instructed during the command monitoring, reset the command monitoring time according to the instructed command monitoring time, and monitor the command again. • When the RAID group status is Normal (Command Monitoring), do not turn OFF the array. If the power is turned OFF while the RAID group status is Normal (Command Monitoring), even the power is turned ON, the command monitoring is considered to be suspended by the power-OFF and the RAID group status becomes Normal (Spin Down Failure: PS OFF/ ON), and it does not spin down. To spin down, instruct the spin-down again. • If a controller failure or a failure between the host and array has occurred during the command monitoring time, the command is issued from the host to the array, and it may be the spin-down cancellation. Moreover, if the controller failure or the failure between the host and the array is restored during the command monitoring time, the command is also issued to the array, and it may be the spin-down cancellation.
How to cancel the command monitoring	To cancel the command monitoring, instruct the target RAID group to spin up or instruct the command monitoring time by the short time such as 0, and instruct the spin-down
RAID groups which cannot issue the instruction to spin down	• The RAID group that includes the system drives (drives #0 to #4 of the basic cabinet for AMS2100/ AMS2300, drives #0 to #4 of the first expansion cabinet for AMS2500). The system drive is the drive where the firmware is stored. • The RAID group configuring the SSDs. • The RAID group for ShadowImage, TrueCopy, or TCE including a P-VOL or an S-VOL in a pair status other than following Simplex, Split, Takeover • The RAID group including a volume whose pair is not released during the Volume Migration or after the Volume Migration is completed • The RAID group including a volume being formatted • The RAID group including a volume to which the parity correction is being performed • The RAID group including a volume for data pool • The RAID group including a volume for DMLU • The RAID group including a volume for command device • The expanding RAID group • The RAID group that the drive firmware is being replaced • When Turbo LU Warning is enabled by specifying the System Parameter option, for the RAID group including the volume using Cache Residency Manager, the de-staging does not proceed, and the spin-down may fail. Disable Turbo LU Warning and instruct the spin-down again.

Table 11-9: Power Saving specifications (Continued)

Item	Specification
Items that will restrain the operation during the spin-down or command monitoring	• I/O command from a host • The ShadowImage pair operation including a copy process Creating pairs, re-synchronizing pairs, restoring pairs • The SnapShot pair operation including a copy process Restoring pairs • The TrueCopy or TCE pair operation including a copy process Creating pairs (including no copy), re-synchronizing pairs, swapping pairs (pair status changes to Takeover) • Executing Volume Migration • Creating a volume • Deleting the RAID group • Formatting a volume • Executing the parity correction of a volume • Setting a volume for DP • Setting a volume for DMLU • Setting a volume for command device • Expansion of a RAID group • Volume growth
Number of times the same RAID group is spun down	Up to seven times a day.
Two or more instructions to the same RAID group	The last instruction is enabled. If the spin-down is instructed during the command monitoring, according to the instructed command monitoring time, and monitor the command again. To cancel the command monitoring, instruct the RAID group to spin up.
Scheduling function	An instruction to spin down or spin up can be issued using a scheduling function provided by JP1, etc.

Table 11-9: Power Saving specifications (Continued)

Item	Specification
Action to be taken for the long time spin-down (health check)	In order to prevent the drive heads from sticking to the disk surfaces, a RAID group which has been kept spun down for 30 days is spun up for six minutes. It is then spun down again. Although the drives are spun up temporarily, no host I/O can be accepted in this period. The opportunity to update the start-up date of spin-down is the time when the spin-down and the health check instructed by Navigator 2 are completed. Neither of the following is included in the spin-down completion opportunities for the update: • Completion of the spin-down of a RAID group, which has been kept spun down, after it is rebooted following the planned shutdown, or • powering off with or without data volatilization; completion of the spin-down of a RAID group, which was spun up when it had been spun down for the purpose of recovery from a failure, after it was waiting for the completion of the recovery from the failure. The RAID group accepts an instruction to spin up given by Navigator 2 during the health check and it enters the status of spin-up. The RAID group does not enter a status of spin-down immediately after it accepts the instruction but it continues the operation, undergoes the health check for six minutes, and then it spins down again. When the planned shutdown is done during the health check, the health check is performed again for six minutes after the power is turned on.
Action to be taken for powering off or on the disk array	The information on the set spin-down is taken over even if the disk array is powered off and then powered on. When restarting the array, the drives which were in the spin-down status once spin up, but they spin down again. However, when the RAID group status is normal (command monitoring), do not turn OFF the array. If the power is turned OFF while the RAID group status is normal (command monitoring), even the power is turned ON, the command monitoring is considered to be suspended by the power-OFF and the RAID group status becomes (spin-down failed: PS OFF/ON), and it does not spin down. To spin down, instruct the spin-down again.
Time required for the spin-up of one RAID group	The time required for the spin-up of one RAID group varies depending on the number of drives that configure the RAID group. The normal spin-up time is as shown below. • 2 to 15 drives: Within 45 seconds normally • 16 to 30 drives: Within 90 seconds normally • 31 or more drives: (Number of drives) ÷ 15x45 seconds Example: When the number of drives configuring the RAID group is 80. The time required for the spin-up = $80 \div 15 \times 45$ seconds = 240 seconds

Table 11-9: Power Saving specifications (Continued)

Item	Specification
Unified Volume	The unified volume is put in the same status as being spun down if one of the configured RAID groups has been spun down, so that the same restrictions with the VOL in the spun down status are applied to the operation to prevent a host I/O, etc.



NOTE: When you refer to the Power Saving Modes and Normal (Spin Up) appears, the power-up is completed. If the host uses a volume, it must mount it.

Table 11-10 details Power Saving effects. Note that the percent of the saving of electric power consumption and value varies by drive type.

Table 11-10: Power Saving effects

Expansion TrayType	During input/output (I/O) operation Unit: validation authority (VA)	During Power Saving (Unit: VA)	Number of Drives Spun Down	Effect: Percentage of the saving of the electric power consumption
Drive tray for 2.5 inch drive	320	140	24 of 24	60% to 70%
Drive tray for 3.5 inch drive	280	90	12 of 12	
Dense drive tray for 3.5 inch drive	1,000	420	48 of 48	

Estimated Spin-Up time

Drives spin up from the Power Saving state in up to three phases in units of drive box. Spin-up time may vary depending on the layout of drives that comprise a RAID group even if the same RAID level and the same number of drives are used.

Table 11-11 provides detail for DBS/DBX drive boxes.

Table 11-11: Estimated Spin Up Time in Drive Box (DBS/DBX)

Number of Drives to Be Spun Up in Drive Box	Spin Up from Spin Down
1 to 2 drives	About 20 seconds
3 to 8 drives	About 40 seconds
9 to 24 drives	About 60 seconds

Table 11-12 provides detail for DBL drive boxes.

Table 11-12: Estimated Spin Up Time in Drive Box (DBL)

Number of Drives to Be Spun Up in Drive Box	Spin Up from Spin Down
1 drive	About 20 seconds
2 to 4 drives	About 40 seconds
5 to 12 drives	About 60 seconds

Power down best practices

You can power down the following:

- ShadowImage drive groups involved in backup to tape
- Virtual tape library (VTL) drive groups involved in backups
- Local or internal backups
- Drive groups within archive storage
- Unused drive groups

You can deliver savings by doing the following:

- Reduce electrical power consumption of idled hard drives
- Reduce cooling costs related to heat generated by the hard drives
- Extend the life of your hardware

Power saving procedures

To use Power Saving, you must have a RAID group in the array. For the target RAID groups that cannot issue the power down instruction, see [Power saving requirements on page 11-45](#).



NOTE: When a fibre channel HDD is in power down status, the LED blinks every 4 seconds. When a serial ATA HDD is in power down status, the LED is off and does not blink.

Power down

To power down

1. Make sure every volume is unmounted.
2. When LVM is used for the disk management, deport the volume or disk groups.
3. Using Navigator 2, power down the RAID group.
4. Using Navigator 2, confirm the RAID group status for specified minutes after powering down.

Power up

To power up

1. Using Navigator 2, power up the RAID group.
2. Using Navigator 2, confirm the RAID group status for several minutes after the powering up.
3. When you refer to the Power Saving Status and see that Normal (Spin Up) is displayed after a while, the power up is completed. Make a host mount the volume included in the RAID group (if the host uses the volume).

This section covers the following key topics:

- ☐ [Power saving requirements](#)
- ☐ [Power saving requirements](#)
- ☐ [Operating system notes](#)

Power saving requirements

This section describes what is required for Power Saving.

Start of the power down operation

The HUS system monitors commands when it receives a power down instruction from a host or a program. The power down can fail if the system detects commands within one minute from the initial power down

instruction. When issuing the power down instruction to multiple RAID groups, each RAID groups is spun down respectively. However, the monitoring continues until all RAID groups are spun down.

RAID groups that cannot power down

- The RAID group that includes the system drives (drives 0 to 4 of the basic cabinet)
- The RAID group that includes the SCSI Enclosure Service (SES) drives of the fibre channel drives (drives 0 to 3 of each extended cabinet)
- The RAID group for ShadowImage, TrueCopy, or TCE, including a primary volume (P-VOL) or a S-VOL in a pair status other than SMPL and PSUS
- The RAID group for SnapShot, including a V-VOL
- The RAID group, including a volume whose pair is not released during the Volume Migration, or is released after the Volume Migration is completed
- The RAID group, including a volume that is being formatted
- The RAID group, including a volume to which the parity correction is being performed
- The RAID group, including a volume for POOL
- The RAID group, including a volume for the differential management volume (DM-LU).
- The RAID group, including a volume for the command device.
- The RAID group, including a system volume for the network-attached storage (NAS).

Things that can hinder power down or command monitoring

- The instruction to power down cannot be issued while the microcode is replaced
- The I/O command from the host
- The paircreate, paircreate -split, pairresync, or pairresync -restore command of ShadowImage
- The pairresync -restore command of SnapShot
- The paircreate, paircreate -nocopy, pairresync, pairresync -swaps, or pairresync -swapp command of TrueCopy
- The paircreate, paircreate -nocopy, or pairresync command of TrueCopy Extended (TCE)
- Executing Volume Migration
- Creating a volume
- Deleting the RAID group
- Formatting a volume
- Executing the parity correction of a volume
- Setting a volume for POOL

- Setting a volume for DM-LU
- Setting a volume for the command device
- Setting a system volume for NAS
- Setting a user volume for NAS

Number of times the same RAID group can be powered down

The same RAID group can be powered down up to seven times a day.

Extended power down (health check)

To prevent the drive heads from sticking to the disk surface, RAID groups that are powered down for 30 days are powered up for 6 minutes, and then powered down again. Although the drives are powered up temporarily, no host I/O can be accepted in this period.

When the power down and the health check instructed by Navigator 2 are completed, you can change the date when the RAID groups are powered down.

The RAID groups accept instructions to power up from Navigator 2 during the health check, and enter power up status. The RAID groups do not enter power down status immediately after they accept the instruction. Instead, they continue the operation, undergo the health check for 6 minutes, and then power down.

When the planned power down is done during the health check, the health check is performed again for 6 minutes after the power is turned on.

If the RAID groups are powered down for 30 days, they are powered up and a health check is performed. After the health check is completed and no problems occur, the system powers the RAID groups down again. This happens every time the RAID groups are powered down for 30 days.

Turning off of the array

The power down information is still valid even if the array is turned off and then on. When the array is turned on, all the installed drives are spun up one time, and the drives that were spun down when the array was turned off remain spun down.

When you restart the array or perform the planned shutdown, execute the power down after verifying that the command monitoring is not being performed.

Time required for powering up

The power up time of RAID groups depends on the number of drives that configure the RAID group. Typical power up times are shown below.

- 2 to 15 drives: 45 seconds
- 16 to 30 drives: 90 seconds
- 31 or more drives: (Number of drives) / 15 X 45

For example, if the number of drives configuring the RAID group is 80, the power up time is 240 seconds, because 80 divided by 15 and then multiplied by 45, is 373.



NOTE: A system drive is the drive where the firmware is stored. An SES (SCSI Enclosure Service) drive is where the information in each extended cabinet is stored. When the command monitoring is operating, the power down fails; the operation instructed by the command is suppressed in the power down status.

Operating system notes

This section describes notes for each operating system.

Advanced Interactive eXecutive (AIX)

- If the host reboots while the RAID group is spun down, the Ghost Disks occurs. When using the volume concerned, delete the Ghost Disks and validate the defined disks after completing the power up of the RAID group concerned.
- When the LVM is used, after making the volume group of LVM including a volume of the RAID group to be spun down offline, power down the RAID group.

Linux

- When the LVM is used, power down the volume group after making the volume group offline and exporting it. When the LVM is not used, power down the volume group after unmounting it.
- When middleware such as Veritas Storage Foundation for Windows is used, specify power down after deporting the disk group.

Hewlett Packard UNIX (HP-UX)

After making the volume group of LVM including a volume of the RAID group to be spun down offline, power down the RAID group.

Windows

- Mount or unmount the volume using the command control interface (CCI) command.

For example:

pairedisplay -x umount D:\hd1

- When middleware such as Veritas Storage Foundation for Windows is used, deport the disk group. Do not use the mounting or unmounting function of CCI.

Solaris

- When Sun Volume Manager is used, perform the power down after releasing the disk set from Solaris.

- When middleware such as Veritas Storage Foundation for Windows is used, specify power down after deport the volume group.



NOTE: For more information, see the Hitachi Adaptable Modular Storage and Workgroup Modular Storage Command Control Interface (CCI) User and Reference Guide, and the Hitachi Simple Modular Storage Command Control Interface (CCI) User's Guide.


Installing, uninstalling, enabling, and disabling Power Saving is set for each array. Before installing and uninstalling, make sure the array is operating correctly. If a failure such as a controller blockade has occurred, you cannot install or uninstall Power Saving.

Viewing Power Saving status

The disk drive information displayed by an operating system or a program when the disk drive is spun down and spun up may be different because reading or writing to a disk drive cannot be performed in power down status.

To view Power Saving status

1. Start Navigator 2.
2. Register the array which you are displaying information for, and connect to it. When you connect to this array.
3. Click the **Logical Status** tab.
4. Log in as a registered user.
5. Select the HUS system where you are enabling or disabling Power Saving.
6. Click **Show & Configure Array**.
7. Click **Energy Saving** in the Navigation bar and click **RG Power Saving**. The power saving information appears (Figure 11-26).



The screenshot shows the 'Arrays' section of the Hitachi Storage Navigator. The left navigation pane lists various components, with 'Energy Saving' and 'RG Power Saving' highlighted. The main area displays the 'RG Power Saving' status for the array HUS110_91100026. It includes a table with columns for RAID Group, Remaining I/O Monitoring Time, and Power Saving Status. The table shows data for RAID Groups 000 through 007. Below the table are buttons for 'Execute Spin Down', 'Execute Spin Up', 'Filter', and 'Filter Off'.

RAID Group	Remaining I/O Monitoring Time	Power Saving Status
000	N/A	Normal(Spin Up)
001	60minutes	Normal(Command Monitoring)
002	N/A	Power Saving(Executing Spin Down)
003	N/A	Power Saving(Spin Down)
004	N/A	Power Saving(Executing Spin Up)
005	N/A	Power Saving(Recovering)
006	N/A	Power Saving(Health Checking)
007	N/A	Power Saving(Spin Down Failure:Error)

Figure 11-26: Power Saving information

Table 11-13 describes power saving details.

Table 11-13: Power Saving details

Items	Contents
RAID Group	The RAID group appears.
I/O Link	Fixed to N/A.
Spin Down Remaining I/O Monitoring Time	The remaining time to spin down or drive power OFF since the last host command is received. In I/O link mode, if a host command is received when the power saving state is "Normal (Command Monitoring)", the remaining I/O monitoring time is reset.
Spin Down I/O Monitoring Time	The I/O monitoring time to spin down that was specified at the power saving request.
Power Off Remaining I/O Monitoring Time	Fixed to N/A.
Power Off I/O Monitoring Time	Fixed to N/A.
Remaining Power Saving Count	Fixed to N/A.
Remaining I/O Monitoring Time	The remaining time of the command monitoring is displayed. The N/A display is exempt.

Table 11-13: Power Saving details

Items	Contents
Power Saving Status	The power saving information appears. Normal (Spin Up): The status in which the drive is operating (being operated). Normal (Command Monitoring): The status in which an issue of a host command is monitored before the drive is spun down. Power Saving (Executing Spin Down): The status in which the spin-down processing of a drive is being executed. Power Saving (Spin Down): The status in which the drive is being spun down. Power Saving (Spin Up Executing): The status in which the spin-up processing of a drive is being executed. Power Saving (Recovering): The status in which the completion of a failure recovery processing is being waited. Power Saving (Health Checking): The status in which the drive has been spun up in order to prevent its head from sticking the disk surface. Normal (Spin Down Failure: Error): The status in which the spin-down processing failed because of a failure. Normal (Spin Down Failure: Host Command): The status in which the spin-down processing failed because of an issue of a host command. Normal (Spin Down Failure: Non-Host Command): The status in which the spin-down processing failed because of an issue of a command other than a host command. Normal (Spin Down Failure: Host Command/Non-Host Command): The status in which the spin-down processing failed because of an issue of a host command and a command other than a host command. Normal (Spin Down Failure: PS OFF/ON): The status in which the spin-down processing failed due to turning OFF/ON the array.




NOTE: The Power Saving Mode includes the power up and down of the drives that configure the RAID group. The RAID group does not show the mode of each drive.

Powering down

For the RAID groups that are not available, see [Power saving requirements on page 11-45](#). You can specify more than one RAID group.

To power down


1. Start Navigator 2.
2. Log in as a registered user.
3. Select the system you want to view information about.
4. Click **Show & Configure Array**.
5. Select the **RG Power Saving** icon in the **Power Saving** tree view.
6. Select the RAID group that you will spin down and click **Execute Spin Down**. The **Spin Down Property** screen displays.
7. Enter an I/O monitoring time in minutes (0 to 720) and click **OK**.



The dialog box is titled "Execute Spin Down - 000". It has a green header bar with the text "Spin Down Property". Below the header, it says "Enter the I/O monitoring time for spin down." There is a text input field for "I/O Monitoring Time" with the value "1" and the unit "minute(s)". Below the input field, it says "From 0 to 720". There is a note "* Required field". At the bottom right, there are "OK" and "Cancel" buttons.

Figure 11-27: Execute Spin Down - 000 dialog box

8. The volume information included in the specified RAID group is displayed. Verify that the spin-down does not cause a problem, and click **Confirm**.



The dialog box is titled "Execute Spin Down - 001". It has a yellow warning icon and the text "Are you sure you want to spin down the selected RAID group? Click Confirm to continue." Below this, there are two tables. The first table is titled "Specified RAID Groups" and has columns "RAID Group" and "VOL". It shows one row with "001" in the "RAID Group" column and "0001" in the "VOL" column. The second table is titled "Affected RAID Groups(Unified VOL)" and has columns "RAID Group" and "VOL". It shows "No Object" in the "RAID Group" column and "VOL" in the "VOL" column. At the bottom right, there are "Confirm" and "Cancel" buttons.

Figure 11-28: Specifying Spin Down window

9. Add a check mark to the box and click **Confirm**.
10. The resulting message appears. Click **Close**.
11. After you power down one RAID group, check the power saving status after the specified minutes have passed. When you power down two or more RAID groups, check the status after several minutes have passed. Refer to [Table 11-14](#) if a phrase other than Normal (Spin Down Failure: Host Command), Normal (Spin Down Failure: Non-Host Command), Normal (Spin Down Failure: Error), or Normal (Spin Down Failure: PS OFF/ON) is displayed.

Table 11-14: Power down errors and recommended action

Cause	Recommended Action
Host Command	A command was issued by a host to a volume that is included in the RAID group to which an instruction to power down had been issued. Check if the RAID group instructed to power down is correct. When the RAID group is correct, instruct it to power down in the status in which no command has been issued.
Non-Host Command	Each volume in the paired state, such as PAIR, must be included in the RAID group instructed to power down. Check that the RAID group instructed to power down is correct, and reissue the instruction to power down.
Error	A failure has occurred in the RAID group that was instructed to power down. After a recovery from the failure is completed, issue an instruction to power down again.
PS OFF/ON	The spin-down is instructed to the RAID group, and the power of the array is turned OFF/ON in the status where the RAID group status is Normal (Command Monitoring). To change it to the spin-down status, instruct the spin-down to the RAID group again.

Notes

- Only one power down instruction per minute can be issued. Before powering down, make sure that all volumes are unmounted. After powering down the LVM volume group offline, power down the RAID group.
- Do not use RAID group volumes that are going to be powered down.
- If there is a mounted volume, unmount it.
- When the logical volume manager (LVM) is used for the disk management, (for example, Veritas) unmount the volume or disk groups.
- Before issuing a power down instruction, verify that all previously issued power down instructions are completed. If the power down fails,

verify that the RAID group you want to power down is not in use, and then power it down again.

- When issuing a power down instruction, if a command is issued by a host or a program during the command monitoring, the power down fails. When the array restarts or performs the planned shutdown during the command monitoring, the monitoring continues after the array restarts.
- If a host or a program issues a command after the array restarts, the power down fails.
- In power down status, data reading or writing in a RAID group volume cannot be done. Instruct the RAID group to power up, verify that the Power Saving Mode in the operation window of Navigator is Normal (Power Up), and then perform the data reading/writing.
- An instruction to power down in the middle of the power up cancels the original instruction. Only the final instruction occurs.

Powering up

Power up a RAID group after it has been powered down. You can specify more than one RAID group.

To power up

1. Start Navigator 2.
2. Register the array where you are powering up the RAID group, and connect to it.
3. Click the **Logical Status** tab.
4. Log in as a registered user.
5. Select the system and RAID group you want to power up.
6. Click **Show & Configure Array**.
7. Select the **RG Power Saving** icon in the Energy Saving tree view.
8. Select the RAID group that you will remove power saving (spin up).
9. Click **Remove Power Saving & Execute Spin Up**.
10. The volume information included in the specified RAID group is displayed. Verify that the spin-up does not cause a problem, and click **Confirm**.

Execute Spin Up - 002

Help



Are you sure you want to spin up the selected RAID group? Click Confirm to continue.

Specified RAID Groups	
RAID Group ▲	VOL
002	0002

Affected RAID Groups(Unified VOL)	
RAID Group ▲	VOL
No Object	

Confirm Cancel
Confirm

Figure 11-29: Specifying Spin Up window

11. The resulting message appears. Click **Close**.

Notes

- Depending on the status of the array, more time may be required to complete the power up.
- An instruction to power up in the middle of the power down cancels the original instruction. Only the final instruction occurs.



NOTE: When you refer to the Power Saving Mode and Normal (Spin Up) appears, the power up is completed. If the host uses a volume, it must mount to it.

Viewing volume information in a RAID group

This section describes how to view volume information for a RAID group.

To view volume information for a RAID group

1. Start Navigator 2.
2. Log in as a registered user.
3. Select the system you want to view information on.
4. Click **Show & Configure Array**.
5. Click **Energy Saving** in the Navigation bar and click **RG Power Saving**. The power saving information appears. Click the **Power Saving** icon.
6. Click **Volume Information**.

7. When you are done, click **Close**.

This chapter provides information to help you identify and resolve problems when using Power Saving.

Failure notes

- If the Power Saving function is enabled, a copy back session occurs in the following two cases even if Spare Drive Operation Mode has been set to default mode where copy back typically does not run.

Table 11-15: Target Spare Drive Response to License Key States

License Key Status			Target Spare Drive	
		Source Data Drive	System Drive	Non System Drive
Power Saving	Disabled	System Drive	As specified	As specified
		Non System Drive	As specified	As specified
	Enabled	System Drive	As specified	Copy back
		Non System Drive	Copy back	As specified



NOTE: System drives correspond to drives 0 to 4 in CBS/CBSL.CBXSS/CBXSL, drives 0 to 4 in DBX of unit ID 0 connected to CBL, drives 0 to 4 of Unit 0 in DBSD/DBLD connected to CBLD.

- Spin-up time may vary depending on the layout of drives the comprise a RAID Group even if the same RAID level and the same number of drives are used. If Spare Drive Operation Mode is set to Variable, spin-up time from the Power Saving state may vary because the layout of the drives in a RAID group changes when drives recover from a failure. If you configure a RAID Group considering spin-up time from the Power Saving state, we recommend setting Spare Drive Operation Mode to Fixed.
- When the system or the spare drive at the position of the FC SES drive is used, you must perform the backup in the same way as that the Spare Drive Operation Mode Fixed, even if the Spare Drive Operation Mode is set to Variable.
- When a failure occurs during the power down in a RAID group other than RAID 0, the array lets the RAID group power up and then makes it power down after restoring the failure. However, if a failure occurs while a RAID group is spun down, the drives being spun down are spun

up and the power down fails. The drives are not spun down automatically after the failed drive is replaced.

- The drives in the power down status in the cabinet where a FC SES failure occurs are spun up. After the SENC failure is restored, the RAID group that has been instructed to power down is spun down.

This section provides use case examples when implementing Power Saving in the Hitachi Data Protection Suite (HDPS) using the Navigator 2 CLI and Account Authentication for a Windows and UNIX environment.

These use cases are only examples, and are only to be used as reference. Your particular use case may vary.

- ☐ [Overview](#)
- ☐ [Security](#)
- ☐ [HDPS AUX-Copy plus aging and retention policies](#)
- ☐ [HDPS Power Saving vaulting](#)
- ☐ [HDPS sample scripts](#)

Overview

These use cases focus on integrating Power Saving with HDPS by creating a power up and power down script which is called by the application before and after executing a disk-to-disk backup.

Power Saving implementations require the following:

- Detailed knowledge of the data environment; Service Level agreements; policies and procedures
- Knowledge in developing storage scripts
- An HUS array
- Storage Navigator GUI and CLI
- Power Savings feature enabled on the array
- Account authentication feature enabled on the array
- Volume Mapping
- Power up script
- Power down script

Power Saving powers down and powers up hard disk drives (HDDs) that contain volumes. You must be aware of where the target data is located, which applications access the data, and how often and what happens if the data is not available. Storage layout is critical. Target Power Saving storage should have a minimal number of application access (preferably only one application). Data availability service level agreements (SLAs) must be understood and modified if required.

To simplify the implementation of Power Saving, Hitachi provides sample scripts. These sample scripts are provided as a learning tool only and are not intended for production use. You must be familiar with script writing and the Navigator 2 CLI.

Security

This use case provides two levels of security. The first level is the array built-in security provided by Hitachi Account Authentication. Account authentication is required, and provides role based array security for the Navigator GUI and protection from rogue scripts.

The second level of security is provided by the HDPS (CommVault) console. Only authorized users can login to the CommVault console and schedule backups.

Account authentication requires that external scripts obtain the appropriate credentials (usernames/passwords). After the appropriate credentials are obtained, the scripts run in the context of that user. The scripts are stored on the MediaAgent and their permissions are dictated by the host operating system.

Set the account authentication password by using the simple network manager (SNM) CLI to specify the following environment parameters and commands.

```
%set STONAVM_ACT=on
```

set User ID and password with the auaccountenv command

[Manual operation] Only once at setting-up account authentication.

```
% accountenv -set -uid xxxxxx (xxxxxx: User ID)
```

Are you sure you want to set the account information? (y/n [n]): y

Please input password. password: yyyyyyy (where yyyyyyy is the password).

To bypass having to answer the confirmation questions: Confirming Command Execution (% set STONAVM_RSP_PASS=on)

HDPS AUX-Copy plus aging and retention policies

AUX-Copy is an HDPS feature that copies a data set which can then be powered down.

In [Figure 11-30](#), HDPS is copying data from the P-VOL to the S-VOL using the auxiliary copy function. After the data is copied, Power Saving powers down the RAID group.




Figure 11-30: HDPS AUX-Copy plus aging and retention

HDPS Power Saving vaulting

Figure 11-31 and Figure 11-32 show the HDPS Power Saving vaulting process.

Power Saving & HDPS Vaulting Process Flow




Figure 11-31: HDPS With Power Saving process flow (1/2)




Figure 11-32: HDPS with Power Saving process flow (2/2)

HDPS sample scripts

This section provides examples of how Power Saving scripts can be written and used in a HDPS Windows and UNIX environment. These are only snapshots of sample scripts, and do not include the whole script. Sample scripts are included in the installation CD. For customized scripts, contact your service delivery team.

```
echo off

setlocal

if not defined GALAXY_BASE set GALAXY_BASE=C:\Program Files\CommVault
Systems\Galaxy\Base

#####
##RUN POWER ON SCRIPT HERE
#####

set PATH=%PATH%;%GALAXY_BASE%

set tmpfile="aux_script.bat.tmp"

qlogin -cs "gordon.marketing.commvault.com" -u "cvadmin" -p "jhN;0w7" > c:\loginerr.txt

if %errorlevel% NEQ 0 (
    echo Login failed. > c:\cmdlog.txt
    goto :EOF )

qoperation auxcopy -af "c:\aux_script.bat.input" > %tmpfile%

if %errorlevel% NEQ 0 (
    for /F "tokens=1* usebackq" %%i in (%tmpfile%) do echo %%i %%j
    echo Failed to start job.
    goto end )
```

Windows scripts

This is only a snapshot of a sample Power Saving script for Windows, and does not include the whole script.

Power down and power up

This is a snapshot of the sample script when powering down and up in Windows.

```
'*++
'Copyright (c) Hitachi Data Systems Corporation
'@Module Name:
' hds-ps-script.vbs
'@Description:
' Script to power up and power down raid groups for a given set of volumes.
'@Revision History:
' 08/07/2007 (HDS)
' v1.0 - Initial script version
'--*/

'////////////////////////////////////
'//
'//Customer specific setting
'Set the SNM User Name / password / CLI directory
const HDS_DFUSER=""
const HDS_DFPASSWD=""
const HDS_STONAVM_HOME="C:\Program Files\Storage Navigator Modular CLI"
```

Using a Windows power up and power down script

The following example details how to use a script when setting up Power Saving for Windows.

To use a script when setting up Power Savings for Windows

1. Create a single volume on a RAID group. The Raid Group can be any size and type.
2. Install SNM CLI on the host where the scripts are going to run (Media Server).
3. Register the arrays with SNM CLI. Refer the Storage Navigator Modular CLI User Guide for command details.

```
auunitadd -unit <name> -LAN -ctl0 <ip of ctl0> -ctl1 <ip of ctl1>
```

4. Create a user account ID that HDPS (Hitachi Data Protection Suite) will use to power down the drives using the SNM CLI.

```
auaccount -unit <name> -add -uid <userid> -account enable -rolepattern 000001
```

5. Install the scripts in the same directory where SNM CLI is installed.
 - a. Copy the script files **hds-ps-app.exe** and **hds-ps-script.vbs** to the SNM CLI directory.

The hds-ps-app.exe is a stand-alone executable used by the Windows power saving script to obtain Windows volume ID information and HUS array information (for example, the array serial number and volume number).

The power saving script captures the output of the **hds-ps-app.exe** file when performing various script actions.

hds-ps-app.exe -volinfo <volume drive letter or mount point> displays the Windows volume ID information.

hds-ps-app.exe -diskextents <volume drive letter or mount point> displays the Windows disk mapping information for the volume.

hds-ps-app.exe -psluinfo <volume drive letter or mount point> displays all the volume information required by the power saving script.

- b. Set these variables in the script under Customer specific setting.

-HDS_STONAVM_HOME

set to install the SNM CLI directory (specify the complete path. For example C:\Program Files\Storage Navigator Modular CLI.

HDS_DFUSER

set to the user ID you defined when you created your account.

HDS_DFPASSWD

set to the password you defined when you created your account.

6. Log files: The script files generate a log file (pslog.txt) under the directory <SNM CLI path>\PowerSavings.
7. Map files: The script generates a volume map file (.psmap) under the directory <SNM CLI path>\PowerSavings.



CAUTION! Do not delete *.psmap files under the PowerSavings directory because they are required by the script to power up raid groups.

8. Error codes: The script returns the following error codes.
 - 0 - The script completed successfully.
 - 1 - Invalid argument/parameter passed to the script.
 - 2 - The specified volume is not valid.
 - 3 - The unmount volume operation failed.
 - 4 - The mount volume operation failed.
 - 5 - Power down failed.
 - 6 - The customer specific settings in the script are not valid.

Powering down

This is an example of how to use the sample script when powering down in Windows.

This amounts the list of volumes (separated by a space) and powers down the raid group that supports it. The list of volumes can be drive letters or mount points.

```
cscript -nologo hds-ps-script.vbs -powerdown <list of volumes>
```

For example:

```
cscript -nologo hds-ps-script.vbs -powerdown y: c:\mount
```

Powering up

This is an example of how to use the sample script when powering up in Windows.

This mounts the list of volumes (separated by space) and powers up the raid group that supports it. The list of volumes can be drive letters or mount points.

```
cscript -nologo hds-ps-script.vbs -powerup <list of volumes>
```

For example:

```
cscript -nologo hds-ps-script.vbs -powerup y: c:\mount
```

UNIX scripts

This is only a snapshot of a Power Saving sample script for UNIX, and does not include the whole script.

Power down

This is a snapshot of the sample script when powering down in UNIX.

```
#!/bin/ksh
# PowerOff.ksh
# Arguments:
# 1 - Mount Point to issue Power Saving OFF function
# Prerequisites:
# 1 - Mountpoint is set in /etc/vfstab file
# Version History:
# v1.0 - HDS.com : Initial Development
##### Only change these variables #####
# Set STONAVM_HOME to where Storage Navigator Modular is installed
export STONAVM_HOME=/opt/snm7.11
# Set SNMUserID to the userid create in Account Authentication
SNMUserID=jpena
```

```

# Set SNMPasswd to the password for the userid set as SNMUserID
SNMPasswd=sac1sac1

## Don't change anything below ##

# Assign mount point parameter to variable
if [[ "$1" = "" ]] then
    echo Usage: $0 "<Mount_Point>"
    echo Example: $0 /backup01
    exit 1
fi

MntPoint=$1

# Check to see if Mount Point is currently mounted
RC=`mount -p | grep " $MntPoint " | wc -l`
if [[ $RC -eq 0 ]] then
    echo Mount Point \"$MntPoint\" is not currently mounted
    exit 2

```

Power up

This is a snapshot of the sample script when powering up in UNIX.

```

#!/bin/ksh

# PowerOn.ksh

# Arguments:

# 1 - Mount Point to issue Power Saving ON function

# Prerequisites:

# 1 - Mountpoint is set in /etc/vfstab file

# Version History:

# v1.0 - Joe.Pena@HDS.com : Initial Development

##### Only change these variables #####

# Set STONAVM_HOME to where Storage Navigator Modular is installed
export STONAVM_HOME=/opt/snm7.11

# Set SNMUserID to the userid create in Account Authentication
SNMUserID=jpena

# Set SNMPasswd to the password for the userid set as SNMUserID
SNMPasswd=sac1sac1

## Don't change anything below ##

# Assign mount point parameter to variable
if [[ "$1" = "" ]] then
    echo Usage: $0 "<Mount_Point>"

```

```

        echo Example: $0 /backup01

        exit 1

    fi

    MntPoint=$1

    # Check to see if Mount Point is currently mounted

    RC=`mount -p | grep " $MntPoint " | wc -l`

    if [[ $RC -ne 0 ]] then

        echo Mount Point \"$MntPoint\" is currently mounted

        exit 2
    fi

```

Using a UNIX power down and power up script

This is an example of how to use the sample script when setting up Power Saving for UNIX.

1. Create a single LDEV (LU) on a Raid group. The Raid Group can be any size and type.
2. Install SNM CLI on the host where the scripts are going to run (Media Server).
3. Register the arrays with SNM CLI.

```
auunitadd -unit <name> -LAN -ctl0 <ip of ctl0> -ctl1 <ip of ctl1>
```

4. Create a user account ID that HDPS (Hitachi Data Protection Suite) will use to power down the drives using the SNM CLI.

```
auaccount -unit <name> -add -uid <userid> -account enable -rolepattern 000001
```

5. Install the scripts in the same directory where SNM CLI is installed.
 - a. **PowerOn.ksh**, **PowerOff.ksh**, and **inqraid.exe**. Make sure all have a permission of -r-x----- and are owned by the root. The **inqraid** command tool confirms and displays details of the HDD connection between the array and the host computer. For more information, see the Command Control Interface (CCI) User's and Reference Guide.
 - b. Set the variables in the script.

```
STONAVM_HOME
```

set to install the SNM CLI directory.

```
SNMUserID
```

set to the userid you defined when you created your account.

```
SNMPasswd
```

set to the password you defined when you created your account.

6. Make sure that all the file systems that are going to be mounted and unmounted are in the mount tab file for your operating system. For example:

```
Solaris - /etc/vfstab
```

Powering down

This is an example of how to use the sample script when powering down in UNIX. This unmounts the file system and powers down the raid group that supports it.

```
PowerOff.ksh
```

For example:

```
PowerOff.ksh /backup01
```

Powering up

This is an example of how to use the sample script when powering up in UNIX. This mounts the file system and powers up the raid group.

```
PowerOn.ksh
```

For example:

```
PowerOn.ksh /backup01
```

Power Saving Plus

The Power Saving Plus feature reduces electricity consumption by spinning down the drives (stopping the rotation of the disk drives) or drive power OFF (stopping drive power feeding) that configure a RAID group.

You can request a spin-down operation either in non-host I/O link mode or host I/O link mode, or drive power OFF in addition to non host I/O-linked spin down supported by Power Saving.

With Power Saving Plus you can reduce the disk array electricity consumption and also lighten the load on the air-conditioning equipment of a data center in addition to reducing electricity consumption by the power saving mechanism. Effective application of Power Saving Plus reduces electricity wasted consumption by instructing an unused RAID group to spin down or drive power OFF.

You can also have host I/O-linked spin down or drive power OFF enabled to a RAID group accessed less frequently such as for backup. In this case, these drives automatically spin down or power OFF when no host I/O is requested and spun up upon I/O request, efficiently reducing power consumption.

Example of Power Saving Plus

To use Power Saving Plus, one or more target RAID groups, to which Power Saving Plus can be applied, must exist in the disk array. For example, assume there is an operation where host I/O-linked spin down is requested for each RAID group with its host I/O monitoring time set to 60 minutes. If you do not request a host I/O for 60 minutes to the RAID group, drives in the RAID group spin down automatically.

Table 11-16: Differences between Power Saving and Power Saving Plus

Support Function	Feature	Power Saving	Power Saving Plus
Non host I/O link	Spin Down	Supported	Supported
	Drive Power Off	Not Supported	Not Supported
Host I/O link	Spin Down	Not Supported	Supported
	Drive Power Off	Not Supported	Supported

Preparing to Use Power Saving Plus

This chapter describes preparations for using Power Saving Plus. This section covers the following topics:

- Specifications
- Power Saving
- Layout in RAID Group

Power Saving Plus Specifications

Table 11-17 shows Power Saving Plus specifications.

Instructing to spin down or drive power OFF is referred to as a power saving request hereafter. Also, the spin down status or the drive power OFF status is referred to as being in the power saving state hereafter.

Table 11-17: Power Saving Plus Specifications

Item	Description
Environment requires	Firmware: Version 0940/A or more is required. Hitachi Storage Navigator Modular 2 (called Navigator 2 hereafter): Version 24.00 or more is required for management PC. License key for Power Saving Plus.
Supported model	HUS150/HUS130/HUS110.
RAID level	All RAID levels supported by the disk array.
Power saving function	The following two power saving functions are supported. • Spin down (stopping the rotation of the disk drives) • Drive power OFF (stopping drive power feeding) Note that the drive power OFF function is only available when the RAID group is configured in a DBW.
Request target	Requesting per RAID group in Navigator 2
Power saving method	Host I/O is monitored (I/O monitoring or command monitoring) and the power saving state become effective if no host I/O is requested for a specified time. The host I/O monitoring time can be specified by the user as I/O monitoring time.
Spin down	I/O link can be disabled or enabled. Host I/O monitors the command issued from host I/O or application to RAID group requesting the power saving with the specified host I/O monitoring time at requesting the power saving (Command monitoring I/O monitoring). After requesting power saving, drives spin down if no host I/O is requested during the I/O monitoring time. Non I/O link mode: If a host I/O is requested during the command monitoring, the monitoring is directed to an active RAID group and the spin-down fails. Even if a host I/O is requested while the drives spin down, they do not respond during the host I/O monitoring time. I/O link mode: If a host I/O is requested during the command monitoring, the monitoring resets and command monitoring continues. If a host I/O is requested while the drives spin down, they spin up to respond. After that, command monitoring continues.

Table 11-17: Power Saving Plus Specifications

Item	Description
Drive power OFF	Only I/O link mode can be specified. Host I/O monitors the command issued from the host I/O or application to the RAID group requesting the power saving with the specified host I/O monitoring time at requesting the power saving (command monitoring I/O monitoring) After the power saving is requested, the drives go into a drive power OFF state if no host I/O is requested during the I/O monitoring time. If a host I/O is requested during the command monitoring, the monitoring resets and command monitoring continues. If a host I/O is requested while in the drive power OFF, they spin up to respond. After that, I/O monitoring continues. Note: Drive power OFF state cannot be requested if the drives in a RAID group are configured only with DBS/DBL/DBX. A RAID group configured with drives in DBW and CBSL can be requested to be powered OFF, but only drives in DBW are powered OFF and only the drives in CBSL spin down.
Command monitoring time	Non I/O link mode: The I/O monitoring time can be specified between 0 and 720 iminutes (Spin down default: 1 minute). If you want to spin down immediately, specify 0 to the command monitoring time. After that, the command monitoring immediately completes and the state changes to the spin down state. However, even if you set the command monitoring time to 0, if a non-completed command to the target RAID group remains in the array, the spin down fails. I/O link mode: The I/O monitoring time can be specified between 0 and 720 minutes (Spin down default: 30 minutes). If the drive failure occurs, the drive spins down after completing a failure recovery.
Concurrent use of spin down and drive power OFF	In I/O link mode, both the spin-down and drive power OFF operations can be used together. The following restriction applies to the I/O monitoring time. Spin down < Drive power OFF

Table 11-17: Power Saving Plus Specifications

Item	Description
Request a RAID group that has already been power saving again	Only command monitoring time can be changed in a RAID group that has already requested power saving. If the power saving request settings change, release the power saving request and request it again after spin up. For a power saving request while commands are monitored: - The command monitoring time is newly set with the specified value and commands are monitored again. For a power saving request while in the power saving state: - The requested settings are reflected with the power saving state unchanged. The newly set command monitoring time takes effect when commands start to be monitored after the drives spin up by receiving a host I/O.
How to release the power saving state	To release the power saving state, request the target RAID group.
System condition which cannot issue the power saving request	During the microprogram replacement. During the drive firmware replacement (Note 1).
RAID groups which cannot issue the power saving request	The RAID group that includes the system drives. System drives: - Drives #0 to #4 of the CBXSS/CBXSL/CBSS/CBSL. - Drives #0 to #4 of the DBS/DBL/DBW corresponding to the unit ID #0 connected to the CBL. - Drives #0 to #4 of the DBSD/DBLD corresponding to the unit ID #0 connected to the CBLD. - Drive #A0 to #A4 of the DBX. - Drive #0 to #4 of the DBW. - The RAID group configuring the SSDs. - The RAID group for ShadowImage, SnapShot, TrueCopy, or TCE including a P-VOL or an S-VOL in a pair status other than following Simplex, Split, Takeover - The RAID group including a volume whose pair is not released during the Volume Migration or after the Volume Migration is completed - The RAID group including a volume being formatted - The RAID group including a volume to which the parity correction is being performed - The RAID group including a volume for DP pool - The RAID group including a volume for DMLU - The RAID group including a volume for command device - The expanding RAID group - The RAID group including an unified LU (Host I/O link mode)

Table 11-17: Power Saving Plus Specifications

Item	Description
Items to restrain the operation while in the power saving request	• The ShadowImage pair operation including a copy process (Note 2) (Note 3) Creating pairs, re-synchronizing pairs, restoring pairs • The SnapShot pair operation including a copy process (Note 2) (Note 3) Restoring pairs • The TrueCopy or TCE pair operation including a copy process (Note 2) (Note 3) Creating pairs (including no copy), re-synchronizing pairs, swapping pairs (pair status changes to Takeover) • Executing Volume Migration (Note 2) (Note 3) • Creating a volume • Deleting the RAID group • Formatting a volume • Executing the parity correction of a volume • Setting a volume for DMLU • Setting a volume for command device • Expansion of a RAID group • Expansion of a volume • United LU • Drive firmware replacement (Note 1)
RAID group spin down count (the remaining power saving count)	We recommend limiting power saving requests to about 7 times per day to prevent drive failure caused by repeated transitions to the power saving state. Particularly in I/O link mode, the drives are automatically spun up or spun down (or drive power OFF) according to host I/Os, they are limited to 7 times per day. However, because the remaining power saving count on the previous day is added to 7 times of that on the day after midnight, the remaining power saving count becomes The remaining power saving count on the previous day + 7 times (up to 200 times). In I/O link mode, if the remaining power saving count is 0 when spin down (or drive power OFF) is requested, it is not executed with the power saving state remaining Normal (Command Monitoring) and the Remaining I/O Monitoring Time remaining 1 minute. It is triggered when the remaining power saving count becomes 1 or more.
Health check (Action to be taken for the long time spin down)	A RAID group that have been in the power saving state for about 30 days is spun up about 6 minutes for drive health check and return to the original power saving state.

Table 11-17: Power Saving Plus Specifications

Item	Description
Power OFF or ON of the disk array	Non I/O link mode: The RAID group status when the array is turned OFF is recovered when it is turned ON. If the array is rebooted, the drives that were spun down is spun up and then spun down after the array becomes Ready. However, if the array is turned OFF when the RAID group is in Normal (Command Monitoring), the command monitoring is assumed to be canceled by turning OFF the array, the drives are not spun down with the RAID group status becoming Normal (Spin Down Failure: PS OFF/ON). If you want to spin down the drives, request it again. I/O link mode: The RAID group that had been requested power saving when the array is turned OFF, its status becomes command monitoring after the array is turned ON.
Spin-up of one RAID group	Spin up process from the power saving state can handle up to CBL: 50 RAID groups, CBSL/CBSS: 30 RAID groups or CBXSL/CBXSS: 20 RAID groups parallelly. If more than the maximum number of RAID groups are concurrently spun up, it may take a long time to spin up, which may be up to about 5 minutes. Spin up time may vary depending on the layout of drives that comprises a RAID group even if the same RAID level and the same number of drives are used. If Spare Drive Operation Mode is set to Variable, spin up time from the power saving state may vary because the layout of drives in a RAID group is changed when drives are recovered from a failure. If you configure a RAID group considering spin up time from the power saving state, we recommend setting Spare Drive Operation Mode to Fixed.
Unified volume	If at least one RAID group that belongs to a unified volume spins down by a non I/O-linked request, the unified volume goes into the state identical to the spin-down state, and the same restrictions in spun down volumes apply to it about operations that suppress host I/Os, etc. You can not unify volumes if this operation causes a RAID group in non I/O link mode and one in I/O link mode to coexist.

About Power Savings Plus

A system drive is the drive that firmware is stored. The drive firmware replacement (download) may operate after upgrading the firmware version of the array. When an error such as **The process cannot be performed because the drive firmware is being replaced.** occurs with the instruction to spin down, wait for about one hour and execute the instruction to spin down again.

Drive firmware updating can be performed only if the state is in **Normal (Spin Up)**, **Normal (Spin Down Failure: Host Command)**, **Normal (Spin Down Failure: Error)**, **Normal (Spin Down Failure: Non-Host**

Command), or **Normal (Command Monitoring)** in I/O-linked mode. In other states, drive firmware updating can not be performed (it is performed after spin up).

If drive firmware updating is being performed to one of the drives in the system, spin down (drive power OFF) does not execute. For this reason, the spin-down (or drive power OFF) process is not executed even if the I/O monitoring time passes in the I/O-linked mode, with the power saving state remaining **Normal (Command Monitoring)** and the remaining monitoring time remaining being 1 minute. The spin-down (or drive power OFF) operation occurs when drive firmware updating completes

When the command monitoring is in operation, the spin-down operation fails, and the operation requested by the command is suppressed in the spin-down status.

Effect of Power Saving

The following table details the effects of Power Saving:

Table 11-18: Effects of Power Saving

Expansion Tray Type (Number of total disks spun down)	During Idle (Unit: VA)	During Spin Down (Unit: VA)	During Drive Power Off (Unit: VA)	Effect: (Compare with During Idle)
24 of 24	320	140	-	During Spin Down: 50 percent Note 6 During Drive Power Off: 70 percent (Percent of the saving of the electric power consumption and value).
12 of 12	280	90	-	
48 of 48	1,000	420	-	
84 of 84	1,260	600 Note 4	340 (84 disk drives of 84 are Drive Power Off)	

Drive Layout in RAID Group

The drive layout in a RAID Group consists of different types.

DBS/DBL/DBX

In DBS/DBL/DBX, drives in the RAID Group are spun up from the power saving state with up to 3 phases. Spin up time may vary depending on the layout of drives that comprise a RAID group even if the same RAID level and the same number of drives are used.

Table 11-19: Estimated Spin Up Time in a Drive Box (DBS/DBX)

Number of drives to be spun up in a drive box	Spin up from Spin down
1 to 2 drives	Around 20 seconds
3 to 8 drives	Around 40 seconds
9 to 24 drives	Around 60 seconds

Table 11-20: Estimated Spin Up Time in a Drive Box (DBL)

Number of drives to be spun up in a drive box	Spin up from Spin down
1 drive	Around 20 seconds
2 to 4 drives	Around 40 seconds
9 to 24 drives	Around 60 seconds

The spin-up process from the power saving state can handle the following number of RAID groups depending on what platform is deployed:

- CBL: 50 RAID groups
- CBSL/CBSS: 30 RAID groups
- CBXSL/CBXSS: 20 RAID groups in a parallel fashion.

If more than the maximum number of RAID groups are concurrently spun up, it may take a long time to spin up, which may be up to about 5 minutes.

In DBW, drives in the RAID Group are spun up from the power saving state with up to every 3 drives in 5 phases in sets of 14 drives, each set being HDU 0 to 13, HDU 14 to 27, HDU 28 to 41, HDU 42 to 55, HDU 56 to 69, and HDU 70 to 83. Spin-up time may vary depending on the layout of drives that comprise a RAID group even if the same RAID level and the same number of drives are used.

[Table 11-21](#) details the estimated spin up time in a set of 14 drives

Table 11-21: Estimated Spin Up Time in a Set of 14 Drives

Number of drives to be spun up in a drive box	Spin up from Spin down	Spin up from drive power OFF
1 to 3 drives	Around 20 seconds	Around 25 seconds
4 to 6 drives	Around 40 seconds	Around 50 seconds
7 to 9 drives	Around 60 seconds	Around 75 seconds
10 to 12 drives	Around 80 seconds	Around 100 seconds
13 to 14 drives	Around 100 seconds	Around 125 seconds

[Table 11-22](#) details the estimated spin up time for 84 drives.

Table 11-22: Estimated Spin Up Time

Number of drives to be spun up in a drive box	Spin up from Spin down	Spin up from drive power OFF
84 drives	Around 125 seconds	Around 150 seconds

The spin-up process from the power saving state can handle up to 50 RAID groups in a parallel fashion. If 51 or more RAID groups are concurrently spun up, it may take a long time to spin up, which may be up to about 5 minutes.guration (Spin up in 1 phase (1))

The spin-up process can be performed in 1 phase because three drives are the targets of spin up in each horizontal row of 14 drives of HDU 0 to 13, HDU 14 to 27, HDU 28 to 41, HDU 42 to 55, HDU 56 to 69 and HDU 70 to 83. Spin up process can be performed in 1 phase because three drives are the targets of spin up in each horizontal row of 14 drives of HDU 0 to 13, HDU 14 to 27, HDU 28 to 41, HDU 42 to 55, HDU 56 to 69, HDU 70 to 83.

Spin up process requires 2 phases because three or few drives are the targets of spin up in each horizontal row of 14 drives of HDU 0 to 13, HDU 28 to 41, HDU 42 to 55, HDU 56 to 69, and HDU 70 to 83, but the targets are 6 drives in HDU 14 to 27.

General Power Saving Plus details

The following details are common to I/O Link and non I/O Link modes.

If you use the Power Saving Plus function, perform elaborate testing with OS and applications to be used with and verify that no problems occur.

Recommended maximum power saving count is 7 times per day. If transitions to power saving state occur nearly 7 times, review when the power saving request is made or how log commands are monitored.

If a failure occurs in the controller or between host and array, the array may receive a command from the host. This is also true when the array is recovered from a failure in the controller or between host and array. Spin down cancellation or drive spin up from the power saving state may occur.

If LU Cache Warning is enabled in the System Parameter option, spin down may fail in a RAID group that contains a volume using Cache Residency Manager because of prolonged destaging. Disable LU Cache Warning and request power saving again.

If a failure occurs in the controller or between host and array, spin down may be cancelled due to a command issued from host to array. This is also true when the array is recovered from a failure in the controller or between host and array.

You can not request spin down when drive firmware updating is being performed.

The following details pertain to Power Saving in Non I/O Link mode.

Because no reading/writing from/to a disk drive can be done in the spin-down status, the disk drive information displayed by an OS or an application program when the disk drive is spun down and spun up may be different from each other.

When an application program using the disk drives concerned is used in the spin-down status, the operation of the application program may be affected because no reading/writing from/to the disk drives can be done in the spin-down status. Use the Power Saving Plus function after examining the effect thoroughly.

Stop the use of all volumes included in the RAID group to be spun down.

① If there is any volume that is mounted, un-mount volume.

✎ When a function for managing logical volumes (LVM: Logical Volume Manager, etc.) is used for the disk management, export/deport volume groups or disk groups.

When you issued an instruction to spin down in non I/O link mode, check if the spin-down has been completed. The spin-down may fail when a host or an application program issues a command. When the spin-down fails, check that the RAID group that you instructed to spin down is not in use and that the spin-down causes no problem, and then issue the instruction to spin down again.

If the power is turned OFF while the RAID group status is **Normal (Command Monitoring)**, even the power is tuned ON, the command monitoring is considered to be suspended by the power-OFF and the RAID group status becomes **Normal (Spin Down Failure: PS OFF/ON)**, and it does not spin down. To spin down, instruct the spin-down again. When restarting the disk array or performing the planned shutdown, do it after checking that the command monitoring is not being done. When the disk array restarts or performs the planned shutdown during the command monitoring, and the spin-down fails after the restart, issue the instruction to spin down again.

When you use a volume, spin up the RAID group. It should be in the power saving mode when it is not considered to be used.

The following details pertain to Power Saving in I/O Link mode.

If you use a volume, the RAID group is spun up automatically by a host I/O.

If a RAID group that has been requested to be in the power saving state does not transition to the power saving state, an application used may be issuing I/Os. Review the environment.

A RAID group that has been requested to be in the power saving state automatically spins up from the power saving state according to host I/Os. If you are using AIX/VMware, transition to the power saving state fails or spin up occurs soon after the transition to the power saving state even if no

host I/O is requested by the user because Read accesses are requested periodically even if a volume is recognized by a host. For this reason, the I/O-linked power saving is not available in a RAID group when AIX/VMware is used.

Power Saving details by operating system

Table 11-23 describes considerations for each OS when power saving is used in non I/O link mode.

Table 11-23: Operating system notes

Operating System	Notes
AIX	When host reboots while the RAID group is being spun down, the Ghost Disks occurs. When you use the volume concerned, it is required to perform operations to delete the Ghost Disks and validate the defined disks after completing the spin-up of the RAID group concerned. When the LVM is used, after making the volume group of LVM including a volume of the RAID group to be spun down offline, spin down the RAID group.
Linux	When the LVM is used, spin down the volume group after making the volume group offline and exporting the volume group. When the LVM is not used, spin down the volume group after un-mounting the volume group. When middleware such as Veritas Storage Foundation™ for Windows is used, specify spin down after deport the disk group.
HP-UX	After making the volume group of LVM including a volume of the RAID group to be spun down offline, spin down the RAID group.
Windows	Mount or un-mount the volume using the CCI command. (Note) Example: pairedisplay -x umount D:\hd1 When middleware such as Veritas Storage Foundation for Windows is used, do not use the mounting or un-mounting function of RAID Manager but deport the disk group instead.
Solaris	When Sun Volume Manager is used, perform the spin-down after releasing the disk set from Solaris. When middleware such as Veritas Storage Foundation™ for Windows is used, specify spin down after deport the volume group.

Notes on Failures

If the Power Saving function is enabled, copy back is performed in the following two cases even if Spare Drive Operation Mode has been set to the default mode, which is copy back less.

Table 11-24: Copy Back Cases

License Key Status		Source Data Drive	Target Spare Drive	
			System Drive	Non System Drive
Power Saving Plus	Disabled	System Drive	As specified	As specified
		Non-System Drive	As specified	As specified
	Enabled	System Drive	As specified	As specified
		Non-System Drive	Copy back	As specified

- Drives #0 to #4 of the CBXSS/CBXSL/CBSS/CBSL.
- Drives #0 to #4 of the DBS/DBL/DBW corresponding to the unit ID #0 connected to the CBL.
- Drives #0 to #4 of the DBSD/DBLD corresponding to the unit ID #0 connected to the CBLD.
- Drive #A0 to #A4 of the DBX.
- Drive #0 to #4 of the DBW.
- If the Drive restoration to the Spare Drive operates between the Drives of CBSL and DBW at the time of the Drive failure restoration of the RAID group in the host I/O-linked power saving request, the copy-back-less function does not operate and the copy-back function operates surely after replacing the drives.
- Even if a RAID groups have the same RAID level and the number of drives, the spin up time may differ depending on the drive positions which configure the RAID groups. When setting the spare drive operation mode to **Variable**, the drive positions which configure the RAID groups change due to the drive failure recovery and so that the spin up processing time from the power saving state may change. When the RAID groups are configured considering the spin up time from the power saving state, it is recommended to set the spare drive operation mode to **Fixed**.

When a failure occurs in a RAID group other than RAID 0 during the spin-down, the disk array let it spin up automatically and then makes a RAID group spin down after restoring the failure. However, if a failure occurs while a RAID 0 group is being spun down, the drives being spun down are spun up and the spin-down results in a failure.

Notes on Hosts

- If a volume in the I/O-linked power saving state receives a host I/O, it spins up to respond. Response time is about 20 to 300 seconds,

depending on the RAID group configuration. If you use the power saving function, use a long timeout period in the OS/path switching software.

- If a failure occurs in the controller or between host and array, a command may be issued from host to array. This is also true when the array is recovered from a failure in the controller or between host and array. Spin down cancellation or drive spin up from the power saving state may occur.
- Because of the periodical health check from a host, a RAID group that has been requested power saving may not be in the power saving state. You should address this by extending the interval of the health check, etc.

Operations Example

This section provides examples of operations in I/O link mode and non I/O link mode of Power Saving Plus.

Example of operation in non I/O link mode

The following process details the Power Saving request operation:

1. If there is any volume that is mounted, unmount the volume.
2. When Veritas Storage Foundation is used for the disk management, deport volume groups.
3. Spin down the specified RAID group in Navigator 2.
4. Confirm the RAID group status in Navigator 2 for specified minutes after the system issues instructions for a spin down.

The following process details the Power Saving removal operation (spin up):

1. Spin up the specified RAID group in Navigator 2.
2. Confirm the RAID group status in Navigator 2 for several minutes after you request the spin-up operation.
3. Mount when the spin-up operation completes.

Example of operation in I/O link mode

The following process details a Power Saving request operation.

1. Extend the command timeout period in the operating system and middleware.
2. Issue the Power Saving request to the specified RAID Group in Navigator 2.
3. Confirm the RAID Group status in Navigator 2 for the specified number of minutes after the Power Saving request is issued.

Removing Power Saving

The following process details a Power Saving removal operation.

1. Spin up the specified RAID Group in Navigator 2.
2. Confirm the RAID group status in Navigator 2 for several minutes after the spin-up request is issued.

Operations of Power Saving Plus

This section details the following Power Saving Plus operations:

- [Displaying Power Saving Information on page 11-84](#)
- [Requesting Non I/O-linked Spin Down on page 11-86](#)
- [Requesting I/O-linked Spin Down on page 11-87](#)
- [Requesting I/O-linked Drive Power OFF on page 11-87](#)
- [Requesting I/O-linked Spin Down with Drive Power OFF on page 11-88](#)
- [Requesting Remove Power Saving \(Spin Up\) on page 11-89](#)

Displaying Power Saving Information

To display the specified power saving information in RAID group:

1. Start Navigator 2.
2. Log in as registered user to Navigator 2.
3. Select the array in which you will reference the power saving information.
4. Click **Show & Configure Array**.
5. Select the **RG Power Saving** icon in the **Energy Saving** tree view.
6. The power saving information is displays.

Table 11-25: Referencing Contents of Power Saving Information

Operating System		Notes
RAID Group		The created RAID Group displays.
I/O Link		Whether the I/O Link mode is enabled or disabled.
Spin Down or Power Off	Renaming I/O Monitoring Time	The remaining time to spin down or drive power off since the last host command is received. In I/O Link mode, if a host command is received when the Power Saving Status is Normal (Command Monitoring), the remaining I/O monitoring time resets.
	I/O Monitoring Time	The I/O monitoring time to spin down that was specified at the Power Saving request.
Remaining Power Saving Council.		The remaining Power Saving count displays.
Power Saving Status		The power saving state is displayed. Note: Normal (Spin Up): The status in which the drive is operating (being operated). Normal (Command Monitoring): The status in which an issue of a host command is monitored before the drive is spun down. Power Saving (Spin Down Executing): The status in which the spin-up processing of a drive is being executed. Power Saving (Spin Down): The status in which the drive is being spun down. Power Saving (Power OFF): The status in which the drive is being turned OFF. Power Saving (Spin Up Executing): The status in which the spin-up processing of a drive is being executed. Power Saving (Recovering): The status in which the completion of a failure recovery processing is being waited. Power Saving (Health Checking): The status in which the completion of a failure recovery processing is being waited. Normal (Spin Down Failure: Error): The status in which the spin-down processing failed because of a failure. Normal (Spin Down Failure: Host Command): The status in which the spin-down processing failed because of an issue of a host command. Normal (Spin Down Failure: Non-Host Command): The status in which the spin-down processing failed because of an issue of a command other than a host command. Normal (Spin Down Failure: Host Command/Non-Host Command): The status in which the spin-down processing failed because of an issue of a host command and a command other than a host command. Normal (Spin Down Failure: PS OFF/ON): The status where the spin-down processing failed due to turning OFF/ON the array.

The **Power Saving Status** shows the **state of the power saving** including the spin-up/spin-down of the drives that configure the RAID group. RAID group does not show the status of each drive.

Requesting Non I/O-linked Spin Down

If there is any volume that is mounted, un-mount volume. In the case of Windows, un-mount the volume using the RAID Manager command. After making the volume group of LVM including a volume of the RAID group to be spun down offline, spin down the RAID group.

For the not available RAID groups, see the appropriate section. A procedure for spinning down a specified RAID group is shown below. (One or more RAID groups can be specified.)

To spin down the specified RAID group:

1. Start Navigator 2.
2. Log in as registered user to Navigator 2.
3. Select the array in which you will spin down the RAID group.
4. Click **Show & Configure Array**.
5. Select the **RG Power Saving** icon in the **Energy Saving** tree view.
6. Select the RAID group that you will spin down and click **Execute Spin Down**.

The **Power Saving Properties** screen appears.

7. Select the **Disable** radio button in **I/O Link**. And then check the mark to the **Spin Down** check box, and then enter an I/O monitoring time (0 to 720) in **Spin Down**. Click **OK**.
8. Because the volume information included in the specified RAID group displays, verify that the spin-down does not cause a problem and click **Confirm**.
9. Confirm the message that appears, check the mark to the check box and click **Confirm**.
10. The Result message appears, click **Close**.
11. When you instruct to spin down one RAID group, check the power saving state after specified minutes or more passed. When you instruct to spin down two or more RAID groups, check the status after several minutes passed.

Take appropriate actions indicated in the previous sections if a phrase other than **Normal (Spin Down Failure: Host Command)**, **Normal (Spin Down Failure: Non-Host Command)**, **Normal (Spin Down Failure: Error)**, **Normal (Spin Down Failure: PS OFF/ON)** or is displayed.

Operating System	Notes
Host Command	A command was issued by a host to a volume that is included in the RAID group to which an instruction to spin down had been issued. Check if the RAID group instructed to spin down is correct. When the RAID group is correct, instruct it to spin down in the status in which no command has been issued.

Operating System	Notes
Non-Host Command	It is considered that a volume in the pair status such as PAIR is included in the volumes included in the RAID group instructed to spin down. Check if the RAID group instructed to spin down is correct. When the RAID group is correct, release the cause of the error and instruct it to spin down again.
Error	A failure has occurred in the RAID group that was instructed to spin down. After a recovery from the failure is completed, issue an instruction to spin down again.
PS OFF/ON	The spin-down is instructed to the RAID group, and the power of the array is turned OFF/ON in the status where the RAID group status is Normal (Command Monitoring). To change it to the spin-down status, instruct the spin-down to the RAID group again.

Requesting I/O-linked Spin Down

The following instructions describe how to request I/O-linked spin down. A procedure for requesting I/O-linked spin down to a specified RAID group is shown below. (One or more RAID groups can be specified.)

1. Start Navigator 2.
2. Log in as registered user to Navigator 2.
3. Select the array in which you will spin down the specified RAID group.
4. Click **Show & Configure Array**.
5. Select the **RG Power Saving** icon in the **Energy Saving** tree view.
6. Select the RAID group that you will request I/O-linked spin down. Click **Execute Spin Down**.
7. The **Power Saving Properties** screen appears.
Select **Enable** radio button in **I/O Link**. And then check the mark to the **Spin Down** check box, and then enter an I/O monitoring time (10 to 720) in **Spin Down**. Click **OK**.
8. Because the volume information included in the specified RAID group is displayed, check that the spin-down causes no problem and click **Confirm**.
9. Confirm the message that appears, check the mark to the check box and click **Confirm**.
10. The Result message appears, click **Close**.

Requesting I/O-linked Drive Power OFF

The following instructions describe how to request I/O-linked drive power OFF. This section describes how to request drive power OFF to RAID groups. (One or more RAID groups can be specified.)

Drive power OFF can not be requested if the drives in a RAID group are configured only with DBS/DBL/DBX. A RAID group that is configured with drives in DBW and CBSL can be requested to be powered OFF, but only drives in DBW are powered OFF and drives in CBSL are only spun down.

1. Start Navigator 2.
2. Log in as registered user to Navigator 2.
3. Select the array in which you will spin down the specified RAID group.
4. Click **Show & Configure Array**.
5. Select the **RG Power Saving** icon in the **Energy Saving** tree view.
6. Select the RAID group that you will request I/O-linked drive power OFF. Click **Execute Spin Down**.
7. The **Power Saving Properties** screen appears.
Select **Enable** radio button in **I/O Link**. And then check the mark to the **Power Off** check box, and then enter an I/O monitoring time (10 to 720) in **Spin Down**. Click **OK**.
8. Because the volume information included in the specified RAID group is displayed, check that the spin-down causes no problem and click **Confirm**.
9. Confirm the message that appears, check the mark to the check box and click **Confirm**.
10. The Result message appears, click **Close**.

Requesting I/O-linked Spin Down with Drive Power OFF

A procedure for requesting I/O-linked spin down with drive power OFF to a specified RAID group is shown below. (One or more RAID groups can be specified.)

1. Start Navigator 2.
2. Log in as registered user to Navigator 2.
3. Select the array in which you will spin down the specified RAID group.
4. Click **Show & Configure Array**.
5. Select the **RG Power Saving** icon in the **Energy Saving** tree view.
6. Select the RAID group that you will request I/O-linked spin down with drive power OFF.
Click **Execute Spin Down**.
7. The **Power Saving Properties** screen appears.
Select **Enable** radio button in **I/O Link**. And then check the mark to the **Spin Down** and **Power Off** check box, and then enter an I/O monitoring time (10 to 720) in **Spin Down** and **Power Off**. Click **OK**.
Be sure to make the I/O monitoring time for drive power OFF longer than the command monitoring time for spin down.
8. Because the volume information included in the specified RAID group is displayed, check that the spin-up causes no problem and click **Confirm**.
9. Confirm the message that appears, check the mark to the check box and click **Confirm**.

10. The Result message appears, click **Close**.

Requesting Remove Power Saving (Spin Up)

The following instructions describe how to request the remove power saving (spin up).

Request the remove power saving (spin up) to spin up a RAID group that was spun down or powered OFF to be available.

A procedure for requesting the remove power saving (spin up) to a specified RAID group is shown below. (One or more RAID groups can be specified.)

1. Start Navigator 2.
2. Log in as registered user to Navigator 2.
3. Select the array in which you will request the remove power saving (spin up) to the specified RAID group.
4. Click **Show & Configure Array**.
5. Select the **RG Power Saving** icon in the **Energy Saving** tree view.
6. Select the RAID group that you will request the remove power saving (spin up).
Click **Execute Spin Up**.
7. Because the volume information included in the specified RAID group is displayed, check that the remove power saving causes no problem and click **Confirm**.
8. The Result message that appears and click **Close**.

If you see **Normal (Spin Up)** in the power saving state after a while, spin up is complete. If a volume in the RAID group is used by a host, mount it in the host.

Specifications

This appendix provides specifications for Navigator 2.

This appendix includes the following:

- ❑ [Navigator 2 specifications](#)

Navigator 2 specifications

The following sections details specifications for various operating systems for Navigator 2:

- Windows
- Red Hat Linux
- Solaris
- HP UX

The following sections detail specifications for various aspects of Navigator 2.

System requirements

This section describes system requirements for your environment.

Windows server

Windows® XP (with SP2 or SP3), Windows Server™ 2003 (with SP1 or SP2), Windows Server™ 2003 (R2) (with or without SP2), Windows Server™ 2003 R2 (x64) (with or without SP2), Windows® Vista (with SP1 or SP2), or Windows Server™ 2008 (x86, x64). The 64-bit Windows is not supported except Windows Server™ 2003 R2 (x64), Windows Server™ 2008 (x64), Windows Server™ 2012 (x64) (Non SP) or Windows 7 (X86, X64) (with or without SP1). Intel® Itanium® is not supported. Windows 6 (x86) Non SP; Windows 8 (x64) Non SP.

CPU: Pentium®

Memory: 256 MB minimum

Disk capacity: 60 MB minimum

Network adapter

Virtual memory: 128 MB

The following table shows the supported Windows.

Operating System	Service Pack
Windows XP (x86)	SP2, SP3
Windows Server 2003 (x86)	SP1, SP2
Windows Server 2003 R2 (x86)	Non SP, SP2
Windows Server 2003 R2 (x64)	Non SP, SP2
Windows Vista (x86)	SP1
Windows Server 2008 (x86)	Non SP, SP2
Windows Server 2008 (x64)	Non SP, SP2
Windows 7 (x86)	Non SP, SP1
Windows 7 (x64)	Non SP, SP1
Windows Server 2008 R2 (x64)	Non SP, SP1

Operating System	Service Pack
Windows Server 2012 (x64)	Non SP
Windows 8 (x86)	Non SP
Windows 8 (x64)	Non SP
Windows 8.1 (x86)	Non SP
Windows 8.1 (x64)	Non SP

Virtual OS

The following table shows the supported Windows versions for various virtual operating system hosts.

Host Operating System	Guest Operating System
VMware ESX Server 3.x	Windows XP
	Windows Server 2003 R2
VMware 4.1	Windows Server 2008 SP2 (x64)
	Windows Server 2008 R2 (x64)
	Windows 7 SP1 (x64)
VMware 5.0	Windows Server 2008 R2 (SP1 (x64)
Windows Server 2008 R2 (x64) (Hyper-V2)	Windows Server 2008 R2 (x64)
Windows Server 2012 (x64) (Hyper-V3)	Windows Server 2012 R2 (SP1)
Windows Server 2012R2 (x64) (Hyper-V3.1)	
VMware 5.1 update1	Windows Server 2008 R2 (SP1)
	Windows Server 2012
VMware 5.5	Windows Server 2012
VMware 5.5 update1	Windows Server 2012
VMware 5.5 update2	Windows Server 2012

Windows client settings

- Browser: IE6.0 (SP1, SP2, SP3) or IE7.0. The 64-bit IE6 (SP1, SP2, SP3) on Windows Server 2003 R2 (x64) and the 64-bit-IE7.0 on windows Server 2008 (x64) is supported.
- Only IE10.0 (x86, x64) is supported on Windows Server 2012 (x64), Windows 7 (x86, x64).
- IE 11.0 is only supported by Windows 8 (x86, x64), Windows 7 (x86, x64), Windows Server 2012 R2 and Windows Server 2008 R2 (SP1). When using IE11.0, use an account other than "administrator" for the Windows login user. If the login user is "administrator", the window may not be displayed correctly.
- CPU: (1 GHz or more is recommended)
- JRE: JRE 1.8.0_31, JRE 1.8.0_25, JRE 1.8.0_20, JRE 1.8.0_5, JRE 1.7.0_75, JRE 1.7.0_71, JRE 1.7.0_67, JRE 1.7.0_60, JRE 1.7.0_51, JRE 1.7.0_45, JRE 1.6.0_45, JRE 1.6.0_43, JRE 1.6.0_41, JRE

1.6.0_37, JRE 1.6.0_33, JRE 1.6.0_31, JRE 1.6.0_30, JRE 1.6.0_25, JRE 1.6.0_22, JRE 1.6.0_20, JRE 1.6.0_15, JRE 1.6.0_13, JRE 1.6.0_1.

- Only JRE 1.7.0_45 is supported on Windows Server 2008 (SP2) (x86, x64), Windows 7 (SP1) (x86, x64), Windows 8 (SP1) (x86, x64), Windows Server 2008 RS (SP1) (X64), and Windows Server 2012 (x64).
- Only JRE 1.8.0_20, JRE 1.8.0_5, JRE 1.7.0_67, and JRE 1.7.0_60 are supported on Windows 7 (SP1) (x86, x64), Windows 8 (SP1) (x86, x64), and windows Server 2012 (x64).
- The JRE download from
<http://java.sun.com/products/archive>

and then install JRE. When installing the JRE, check the setting in the Client Side section of [Starting Navigator 2 host and client configuration on page 3-27](#). For more installation about JRE, refer to the Java download page.

- Memory: 1 GB or more (2 GB or more is recommended)
When using Hitachi Storage Navigator Modular 2 and other software products together, the memory capacity totaling the value of each software product is required.
- Available disk capacity: A free capacity of 800 MB or more is required.
- Monitor: Resolution 800 × 600, 1,024 × 768 or more is recommended, 256 color or more.
- Virtual OS: VMware ESX Server 3.x: Windows XP, Windows Server 2003 R2, Windows Server 2008, SP2 (x64), Windows Server 2008 R2 (x64); VMware 5.0: Windows Server 2008 R2 (SP1) (x64); Windows Server 2008 R2 (x64) (Hyper-V2): Windows Server 2008 R2 (x64). Windows 8 (x86) Non SP; Windows 8 (x64) Non SP; VMware 5.1 update1: Windows Server 2008 R2 (SP1), Windows Server 2012; VMware 5.5: Windows Server 2012. Windows 8.1 (x86) Non SP; Windows 8.1 (x64) Non SP

Solaris™ (SPARC)

Solaris 8, 9, 10

CPU: UltraSPARC or higher

Memory: 256 MB minimum

Disk capacity: 100 MB minimum

Network adapter

CPU: SPARC minimum 1 GHz (2 GHz or more is recommended)

Solaris 10 (x64): Minimum 1.8 GHz (2 GHz (2 GHz or more is recommended))

Not supported x86 processor as like Opteron.

Solaris 10 (x64) is supported 64 bits kernel mode on Sun Fire x64 server family only. Do not change the kernel mode to other than 64 bits after installing Hitachi Storage Navigator Modular 2.

- Memory: 1 GB or more (2 GB or more is recommended)
When using Hitachi Storage Navigator Modular 2 and other software products together, the memory capacity totaling the value of each software product is required.
- Available disk capacity: A free capacity of 1.5 GB or more is required.
- JDK: JDK1.5.0 is required.

Client

- OS: Solaris 8
Solaris 9 (SPARC)
Solaris 10 (SPARC)
Solaris 10 (x86), or
Solaris 10 (x64)
- CPU: SPARC minimum 1 GHz (2 GHz or more is recommended)
Solaris 10 (x64): Minimum 1.8 GHz (GHz or more is recommended)
Not supported x86 processor as like Opteron.
Solaris 10 (x64) is supported 64-bit kernel mode on Sun Fire x64 server family only. Do not change the kernel mode to other than 64 bits after installing Hitachi Storage Navigator Modular 2.
- JRE: JRE 1.7.0_75, JRE 1.7.0_71, JRE 1.7.0_67, JRE 1.7.0_60, JRE 1.7.0_51, JRE 1.6.0_45, JRE 1.6.0_43, JRE 1.6.0_41, JRE 1.6.0_37, JRE 1.6.0_33, JRE 1.6.0_31, JRE 1.6.0_30, JRE 1.6.0_25, JRE 1.6.0_22, JRE 1.6.0_20, JRE 1.6.0_15, JRE 1.6.0_13, JRE 1.6.0_1. The 64-bit JRE is not supported. For more installation about JRE, refer to java download page.
- Only JRE 1.7.0_45 is supported on Solaris 10.
- Browser: Mozilla 1.7, Firefox 2
- CPU: (1 GHz or more is recommended)
- Memory: 1 GB or more (2 GB or more is recommended)
When using Hitachi Storage Navigator Modular 2 and other software products together, the memory capacity totaling the value of each software product is required.
- Available disk capacity: A free capacity of 100 MB or more is required.
Monitor: Resolution 800 × 600, 1,024 × 768 or more is recommended, 256 color or more.

HP-UX

HP-UX 11.0, 11i, 11i v2.0, 11i v3.0

CPU: PA8000 or higher (HP-UX 11i v2.0 operates in Itanium® 2 environment)

Memory: 256 MB minimum

Disk capacity: 110 MB minimum

Network adapter

AIX

AIX 5.1, 5.2, 6.1, or 7.1

CPU: PowerPC/RS64 II or higher

Memory: 256 MB minimum

Disk capacity: 90 MB minimum

Network adapter

Remise program: install the patch of IY33524 if needed after VisualAge C++ Runtime 6.0.0.0. Download from the IBM® Web site.

Linux

Host

- OS: Red Hat Enterprise Linux AS 4.0 (x86) update1 or
 - Red Hat Enterprise Linux AS 4.0 (x86) update5
 - Red Hat Enterprise Linux 5.3 (x86) (excluding SELinux)
 - Red Hat Enterprise Linux 5.4 (x86) (excluding SELinux)
 - Red Hat Enterprise Linux 5.4 (x64) (excluding SELinux)
 - Red Hat Enterprise Linux 5.5 (x86) (excluding SELinux)
 - Red Hat Enterprise Linux 5.5 (x64) (excluding SELinux)
 - Red Hat Enterprise Linux 5.6 (x86) (excluding SELinux)
 - Red Hat Enterprise Linux 5.6 (x64) (excluding SELinux)
 - Red Hat Enterprise Linux 5.7 (x86) (excluding SELinux)
 - Red Hat Enterprise Linux 5.7 (x64) (excluding SELinux)
 - Red Hat Enterprise Linux 5.8 (x86) (excluding SELinux)
 - Red Hat Enterprise Linux 5.8 (x64) (excluding SELinux)
 - Red Hat Enterprise Linux 6.1 (x86) (excluding SELinux)
Premise patch:
kernel-2.6.32-220.4.2.el6.i686.rpm or its inheritor
kernel-firmware-2.6.32-220.4.2.el6.noarch.rpm or its inheritor
- Red Hat Enterprise Linux 6.1 (x64) (excluding SELinux)
Premise patch:
glibc-2.12-1.25.el6.i686.rpm or its inheritor
nss-softokn-freebl-3.12.9-3.el6.i686.rpm or its inheritor
libgcc-4.4.5-6.el6.i686.rpm
libstdc++-4.4.5-6.el6.i686.rpm
kernel-2.6.32-220.4.2.el6.x86_64.rpm or its inheritor
kernel-firmware-2.6.32-220.4.2.el6.noarch.rpm or its inheritor
- Red Hat Enterprise Linux 6.2 (x86) (excluding SELinux)
Premise patch: kernel-2.6.32-220.4.2.el6.i686.rpm or its inheritor
kernel-firmware-2.6.32-220.4.2.el6.noarch.rpm or its inheritor
- Red Hat Enterprise Linux 6.2 (x64) (excluding SELinux)
Premise patch: glibc-2.12-1.47.el6.i686.rpm or its inheritor
nss-softokn-freebl-3.12.9-11.el6.i686.rpm or its inheritor
libgcc-4.4.6-3.el6.i686.rpm
libstdc++-4.4.6-3.el6.i686.rpm
- Red Hat Enterprise Linux 6.5 (x86) (excluding SELinux)

- Red Hat Enterprise Linux 6.5 (x64) (excluding SELinux)
Premise patch: glibc-2.12-1.132.el6.i686.rpm or its inheritor
nss-softokn-freebl-3.14.3-9.el6.i686.rpm or its inheritor
libgcc-4.4.7-4.el6.i686.rpm



NOTE: An update from Red Hat Enterprise Linux AS 4.0 is not supported.

- CPU: Minimum 1 GHz (2 GHz or more is recommended)
- Physical memory: 1 GB or more (2 GB or more is recommended)
When using Hitachi Storage Navigator Modular 2 and other software products together, the memory capacity totaling the value of each software product is required.
- Available disk capacity: A free capacity of 1.5 GB or more is required.

Client

- OS: Red Hat Enterprise Linux AS 4.0 (x86) update1 or
 - Red Hat Enterprise Linux AS 4.0 (x86) update5
 - Red Hat Enterprise Linux 5.3 (x86) (excluding SELinux)
 - Red Hat Enterprise Linux 5.4 (x86) (excluding SELinux)
 - Red Hat Enterprise Linux 5.4 (x64) (excluding SELinux)
 - Red Hat Enterprise Linux 5.5 (x86) (excluding SELinux)
 - Red Hat Enterprise Linux 5.5 (x64) (excluding SELinux)
 - Red Hat Enterprise Linux 5.6 (x86) (excluding SELinux)
 - Red Hat Enterprise Linux 5.6 (x64) (excluding SELinux)
 - Red Hat Enterprise Linux 5.7 (x86) (excluding SELinux)
 - Red Hat Enterprise Linux 5.7 (x64) (excluding SELinux)
 - Red Hat Enterprise Linux 5.8 (x86) (excluding SELinux)
 - Red Hat Enterprise Linux 5.8 (x64) (excluding SELinux)
 - Red Hat Enterprise Linux 6.1 (x86) (excluding SELinux)
 - Red Hat Enterprise Linux 6.1 (x64) (excluding SELinux)
- Premise patch:
 - glibc-2.12-1.25.el6.i686.rpm or its inheritor
 - nss-softokn-freebl-3.12.9-3.el6.i686.rpm or its inheritor
 - libgcc-4.4.5-6.el6.i686.rpm
 - libstdc++-4.4.5-6.el6.i686.rpm
- Red Hat Enterprise Linux 6.2 (x86) (excluding SELinux)
- Red Hat Enterprise Linux 6.2 (x64) (excluding SELinux)
Premise patch:
 - glibc-2.12-1.47.el6.i686.rpm or its inheritor
 - nss-softokn-freebl-3.12.9-11.el6.i686.rpm or its inheritor
 - libgcc-4.4.6-3.el6.i686.rpm
 - libstdc++-4.4.6-3.el6.i686.rpm

- Red Hat Enterprise Linux 6.1 (x86) (excluding SELinux)
- Red Hat Enterprise Linux 6.1 (x64) (excluding SELinux)
- Red Hat Enterprise Linux 6.2 (x86) (excluding SELinux)
- Red Hat Enterprise Linux 6.2 (x64) (excluding SELinux)
- Red Hat Enterprise Linux 6.3 (x86) (excluding SELinux)
- Red Hat Enterprise Linux 6.3 (x64) (excluding SELinux)
Premise patch: glibc-2.12-1.80.el6.i686.rpm or its inheritor
nss-softokn-freebl-3.12.9-11.el6.i686.rpm or its inheritor
libgcc-4.4.6-4.el6.i686.rpm
libstdc++-4.4.6.4.el6.i686.rpm
- Red Hat Enterprise Linux 6.4 (x86) (excluding SELinux)
- Red Hat Enterprise Linux 6.4 (x64) (excluding SELinux)
Premise patch: glibc-2.12-1.107.el6.i686.rpm or its inheritor
nss-softokn-freebl-3.12.9-11.el6.i686.rpm or its inheritor
libgcc-4.4.7-3.el6.i686.rpm
libstdc++-4.4.7.3.el6.i686.rpm
- Red Hat Enterprise Linux 6.5 (x86) (excluding SELinux)
- Red Hat Enterprise Linux 6.5 (x64) (excluding SELinux)
Premise patch: glibc-2.12-1.132.el6.i686.rpm or its inheritor
nss-softokn-freebl-3.14.3-9.el6.i686.rpm or its inheritor
libgcc-4.4.7-4.el6.i686.rpm
libstdc++-4.4.7-4.el6.i686.rpm



NOTE: An update from Red Hat Enterprise Linux AS 4.0 is not supported.

- Browser: Mozilla 1.7, firefox-17.0.9, firefox-10.0.12, firefox-10.0.5, firefox-3.6.16, firefox-3.6.17, firefox-3.6.24, firefox-16.0.2
 - Only Red Hat Enterprise Linux 6.2 (x86, x64) is supported on firefox-3.6.24 and firefox-16.0.2.
 - Only Red Hat Enterprise Linux 5.8 (x86, x64) is supported on firefox-3.6.16.
 - Only Red Hat Enterprise Linux 6.1 (x86, x64) is supported on firefox-3.6.17.
 - Only Red Hat Enterprise Linux 6.3 (x86, x64) is supported on firefox-10.0.5.
 - Only Red Hat Enterprise Linux 6.4 (x86, x64) is supported on firefox-10.0.12.
 - Only Red Hat Enterprise Linux 6.3 (x86, x64) and Red Hat Enterprise Linux 6.4 (x86, x64) are supported on firefox-17.0.9.
- JRE: JRE 1.8.0_31, JRE 1.8.0_25, JRE 1.8.0_20, JRE 1.8.0_5, JRE 1.7.0_75, JRE 1.7.0_71, JRE 1.7.0_67, JRE 1.7.0_60, JRE 1.7.0_51, JRE 1.7.0_45, JRE 1.6.0_45, JRE 1.6.0_43, JRE 1.6.0_41, JRE 1.6.0_37, JRE 1.6.0_33, JRE 1.6.0_31, JRE 1.6.0_30, JRE 1.6.0_25, JRE 1.6.0_22, JRE 1.6.0_20, JRE 1.6.0_15, JRE 1.6.0_13, JRE 1.6.0_10

- Only JRE 1.7.0_45 is supported on Linux 5.8 (x86), Linux 6.1 (x86), and Linux 6.2 (x86).
- Only JRE 1.8.0_20, JRE 1.8.0_5, JRE 1.7.0_67, and JRE 1.7.0_60, are supported on Linux 6.4 (x86).
- The JRE download from <http://java.sun.com/products/archive>, and then install JRE.
- For more installation about JRE, refer to java download page. The 64-bit JRE is not supported.
- When you operate HUS100, AMS2000 and SMS, installation of JRE is not necessary.
- When you operate 9500V and AMS/WMS, install JRE.
- CPU: (1 GHz or more is recommended)
- Physical memory: 1 GB or more (2 GB or more is recommended)
When using Hitachi Storage Navigator Modular 2 and other software products together, the memory capacity totaling the value of each software product is required.
- Available disk capacity: A free capacity of 100 MB or more is required.
- Monitor: Resolution 800 R 600, 1,024 768 or more is recommended, 256 color or more.

IPv6 supported platforms

Table A-1 shows the IPv6 supported platforms.

Table A-1: IPv6 Supported Platforms

Vendor	Operating System		IPv6 Supported
	Name	Service Pack	
SUN	Solaris 8 (SPARC)	–	Supported
	Solaris 9 (SPARC)	–	Supported
	Solaris 10 (SPARC)	–	Supported
	Solaris 10 (x86)	–	Supported
	Solaris 10 (x64)	–	Supported
Microsoft	Windows Server 2003 (x86)	SP1	Supported
	Windows Server 2003 (x86)	SP2	Supported
	Windows Server 2003 R2 (x86)	Without SP, With SP2	Supported
	Windows Server 2003 R2 (x64)	Without SP	Supported
	Windows Vista (x86)	SP1	Supported
	Windows Server 2008 (x86)	SP1, SP2	Supported
	Windows Server 2008 (x64)	SP1, SP2	Supported
	Windows Server 2008 R2 (x64)	SP, With SP1	Supported
	Windows Server 2012 (x64)	Without SP	Supported
	Windows 7 (x86)	Without SP, With SP1	Supported
	Windows 7 (x64)	Without SP, With SP1	Supported
Red Hat	Red Hat Enterprise Linux 4.0 Update1	–	Address searching function is not supported on the server.
	Red Hat Enterprise Linux 4.0 Update5	–	Address searching function is not supported on the server.
	Red Hat Enterprise Linux 5.3	-	Address searching function is not supported on the server.
	Red Hat Enterprise Linux 5.4	-	Address searching function is not supported on the server.
	Red Hat Enterprise Linux 5.5	-	Address searching function is not supported on the server.

Table A-1: IPv6 Supported Platforms

Red Hat Enterprise Linux 5.6	-	Address searching function is not supported on the server.
Red Hat Enterprise Linux 5.7 (x86)	-	Address searching function is not supported on the server.
Red Hat Enterprise Linux 5.7 (x64)	-	Address searching function is not supported on the server.
Red Hat Enterprise Linux 5.8 (x86)	-	Address searching function is not supported on the server.
Red Hat Enterprise Linux 5.8 (x64)	-	Address searching function is not supported on the server.
Red Hat Enterprise Linux 6.1 (x86)	-	Address searching function is not supported on the server.
Red Hat Enterprise Linux 6.1 (x64)	-	Address searching function is not supported on the server.
Red Hat Enterprise Linux 6.2 (x86)	-	Address searching function is not supported on the server.
Red Hat Enterprise Linux 6.2 (x64)	-	Address searching function is not supported on the server.
Red Hat Enterprise Linux 6.3 (x86)	-	Address searching function is not supported on the server.
Red Hat Enterprise Linux 6.3 (x64)	-	Address searching function is not supported on the server.
Red Hat Enterprise Linux 6.4 (x86)	-	Address searching function is not supported on the server.
Red Hat Enterprise Linux 6.4 (x64)	-	Address searching function is not supported on the server.

Table A-1: IPv6 Supported Platforms

Red Hat Enterprise Linux 6.5 (x86)	-	Address searching function is not supported on the server.
Red Hat Enterprise Linux 6.5 (x64)	-	Address searching function is not supported on the server.

Considerations at Time of Operation

The following sections detail recommended formatting sizes to ensure the best performance for given configurations.

Volume formatting

The total size of volumes that can be formatted at the same time has restrictions. If the configuration exceeds the possible formatting size, the firmware of the array does not execute the formatting (error messages are displayed). Moreover, if the volumes are expanded, the expanded volume unit size is automatically formatted and the size becomes the restriction target that permits which entities can be formatted at the same time.

Note that the possible formatting size differs depending on the array type. Format the total size of volumes by the recommended batch formatting size or less as shown in [Table A-3](#).

Table A-2: Batch formatting size by platform

Array Type	Recommended Batch Formatting Size		
HUS 100	359 TB (449 GB x 800)	308 TB (193 GB x 1,600)	208 TB (65 GB x 3,200)
HUS 130	287 TB (449 GB x 640)	247 TB (193 GB x 1,280)	166 TB (65 GB x 2,560)
HUS 150	179 TB (449 GB x 400)	154 TB (193 GB x 800)	104 TB (65 GB x 1,600)

The formatting is executed in the following three operations. However, it has no effect on the DP volumes using the Dynamic Provisioning function [Table A-3](#) details formatting capacity operation.

Table A-3: Formatting capacity by operation

Operation	Formatting Capacity
Volume creation (format is specified)	Size of volumes to create
Volume format	Size of volumes to format
Volume expansion	Size of volumes to expand

The restrictions of the possible formatting size becomes the size of totaling three operations. Perform it so that the total of each operation becomes the recommended batch formatting size or less.

When the above-mentioned operation is executed and the restrictions of the possible formatting size are exceeded, the following messages are displayed. [Table A-4](#) details messages that display when formatting size is exceeded.

Table A-4: Messages for exceeded size

Operation	Formatting Capacity
Volume creation (format is specified)	DMED100005: The quick format size is over maximum value. Please retry after that specified quick format size is decreased or current executed quick format is finished.
Volume format	
Volume expansion	DMED0E0023: The quick format size is over maximum value. Please retry after that specified quick format size is decreased or current executed quick format is finished.

(1) Volume creation (format is specified):

If the volume creation (format is specified) becomes an error, the volumes are created, but the formatting is not executed and the Status of the Volumes tab becomes Unformat. After checking that the Status of volumes which are already executing the other formatting or expansion the other volumes becomes Normal, execute only the formatting for the volumes which performed the creation of volumes.

(2) Volume format:

If the formatting of volume s becomes an error, the formatting is not executed and the Status of the Volumes tab is still kept as before the execution. After checking that the Status of volumes which are already executing the other formatting or expansion the other volumes becomes Normal, execute the formatting again.

(3) Volume expansion:

If the expansion of volumes becomes an error, the expansion is not executed and the Status of the Volumes tab is still kept as before the execution. After checking that the Status of volumes which are already executing the other formatting or expansion the other volumes becomes Normal, execute the expansion again.

Constitute array

Configurations set successfully.

Regardless of the configuration file that you specified, the cache partition number is specified to 0 or 1 and Full Capacity Mode s specified to disabled. When the result is different from your expectation, change the configurations. It cannot set configurations for optional storage features and you need to specify them manually.

Recording Navigator 2 Settings

This appendix contains a table where you can record your configuration settings for future reference. We recommend that you make a copy of the following table and record your Navigator 2 configuration settings for future reference.

Table B-1: Recording configuration settings

Field	Description
Storage System Name	
Management console static IP address (used to log in to Navigator 2)	
Email Notifications	
Email Notifications	?Disabled ?Enabled (record your settings below)
Domain Name	
Mail Server Address	
From Address	
Send to Address Address 1: Address 2: Address 3:	
Reply To Address	
Management Port Settings	
Controller 0	
Configuration	?Automatic (Use DHCP) ?Manual (record your settings below)
IP Address	

Table B-1: Recording configuration settings (Continued)

Field	Description
Subnet Mask	
Default Gateway	
Controller 1	
Configuration	?Automatic (Use DHCP) ?Manual (record your settings below)
IP Address	
Subnet Mask	
Default Gateway	
Data Port Settings	
Controller 0/ Port A	
IP Address	
Subnet Mask	
Default Gateway	
Negotiation	
Controller 0/ Port B	
IP Address	
Subnet Mask	
Default Gateway	
Negotiation	
Controller 1/ Port A	
IP Address	
Subnet Mask	
Default Gateway	
Negotiation	
Controller 1/ Port B	
IP Address	
Subnet Mask	
Default Gateway	
Negotiation	
VOL Settings	
RAID Group	
Free Space	
VOL	
Capacity	
Stripe Size	

Table B-1: Recording configuration settings (Continued)

Field	Description
Format the Volume	? Yes ? No



Glossary

This glossary provides definitions for replication terms as well as terms related to the technology that supports your Hitachi storage system. Click the letter of the glossary section to display the related page.

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

A

Arbitrated loop

A Fibre Channel topology that requires no Fibre Channel switches. Devices are connected in a one-way loop fashion. Also referred to as FC-AL.

Array

A set of hard disks mounted in a single enclosure and grouped logically together to function as one contiguous storage space.

B

bps

Bits per second. The standard measure of data transmission speeds.

C

Cache

A temporary, high-speed storage mechanism. It is a reserved section of main memory or an independent high-speed storage device. Two types of caching are found in computers: memory caching and disk caching. Memory caches are built into the architecture of microprocessors and often computers have external cache memory. Disk caching works like memory caching; however, it uses slower, conventional main memory that on some devices is called a memory buffer.

Capacity

The amount of information (usually expressed in megabytes) that can be stored on a disk drive. It is the measure of the potential contents of a device. In communications, capacity refers to the maximum possible data transfer rate of a communications channel under ideal conditions.

CBL

3U controller box.

CBXS

Controller box. Two types of CBXS controller boxes are available:

- A 2U CBXSS Controller Box that mounts up to 24 2.5-inch drives.
- A 3U CBXSL Controller Box that mounts up to 12 3.5-inch drives.

CBS

Controller box. There are two types of CBS controller boxes available:

- A 2U CBSS Controller Box that mounts up to 24 2.5-inch drives.

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

Glossary–2

- A 3U CBSL Controller Box that mounts up to 12 3.5-inch drives.

CCI

See command control interface.

Challenge Handshake Authentication Protocol

An authentication technique for confirming the identity of one computer to another. Described in RFC 1994.

CHAP

See Challenge Handshake Authentication Protocol.

CLI

See command line interface.

Cluster

A group of disk sectors. The operating system assigns a unique number to each cluster and then keeps track of files according to which clusters they use.

Cluster capacity

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

Command devices

Dedicated logical volumes that are used only by management software such as CCI, to interface with the storage systems. Command devices are not used by ordinary applications. Command devices can be shared between several hosts.

Command line interface (CLI)

A method of interacting with an operating system or software using a command line interpreter. With Hitachi's Storage Navigator Modular Command Line Interface, CLI is used to interact with and manage Hitachi storage and replication systems.

CRC

Cyclic Redundancy Check. An error-correcting code designed to detect accidental changes to raw computer data.

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

D

Disaster recovery

A set of procedures to recover critical application data and processing after a disaster or other failure. Disaster recovery processes include failover and failback procedures.

Differential Management Logical Unit (DMLU)

The volumes used to manage differential data in a storage system. In a TrueCopy Extended Distance system, there may be up to two DM logical units configured per storage system. For Copy-on-Write and ShadowImage, the DMLU is an exclusive volume used for storing data when the array system is powered down.

DMLU

See Differential Management-Logical Unit.

Drive Box

Chassis for mounting drives that connect to the controller box. The following drive boxes are supported:

- DBS, DBL: 2U drive box
- DBX: 4U drive box

Drive I/O Module

I/O module for the CBL that has drive interfaces.

Duplex

The transmission of data in either one or two directions. Duplex modes are full-duplex and half-duplex. Full-duplex is the simultaneous transmission of data in two direction. For example, a telephone is a full-duplex device, because both parties can talk at once. In contrast, a walkie-talkie is a half-duplex device because only one party can transmit at a time.

E

Ethernet

A computer networking technology for local-area networks.

Extent

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

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

F

Fabric

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

Failover

The automatic substitution of a functionally equivalent system component for a failed one. The term failover is most often applied to intelligent controllers connected to the same storage devices and host computers. If one of the controllers fails, failover occurs, and the survivor takes over its I/O load.

Fallback

Refers to the process of restarting business operations at a local site using the P-VOL. It takes place after the storage systems have been recovered.

Fault tolerance

A system with the ability to continue operating, possibly at a reduced level, rather than failing completely, when some part of the system fails.

FC

See Fibre Channel.

FC-AL

See Arbitrated Loop.

FCOE

See Fibre Channel over Ethernet.

Fibre Channel

A gigabit-speed network technology primarily used for storage networking.

Fibre Channel over Ethernet

A way to send Fiber Channel commands over an Ethernet network by encapsulating Fiber Channel calls in TCP packets.

Firmware

Software embedded into a storage device. It may also be referred to as Microcode.

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

Full-duplex

Transmission of data in two directions simultaneously. For example, a telephone is a full-duplex device because both parties can talk at the same time.

G**Gbps**

Gigabit per second.

Gigabit Ethernet

A version of Ethernet that supports data transfer speeds of 1 gigabit per second. The cables and equipment are very similar to previous Ethernet standards. Abbreviated GbE.

GUI

Graphical user interface.

H**HA**

High availability.

Half-duplex

Transmission of data in just one direction at a time. For example, a walkie-talkie is a half-duplex device because only one party can talk at a time.

HBA

See Host bus adapter.

Host

A server connected to the storage system via Fibre Channel or iSCSI ports.

Host bus adapter

An I/O adapter located between the host computer's bus and the Fibre Channel loop that manages the transfer of information between the two channels. To minimize the impact on host processor performance, the host bus adapter performs many low-level interface functions automatically or with minimal processor involvement.

Host I/O Module

I/O module for the CBL that has host interfaces.

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

I

IEEE

Institute of Electrical and Electronics Engineers (read “I-Triple-E”). A non-profit professional association best known for developing standards for the computer and electronics industry. In particular, the IEEE 802 standards for local-area networks are widely followed.

I/O

Input/output.

I/O Card (ENC)

I/O Card (ENC) installed in a DBX Drive Box, with interfaces for the controller box or drive box.

I/O Module (ENC)

I/O Module (ENC) installed in a DBS/DBL Drive Box, with interfaces for the controller box or drive box.

IOPS

Input/output per second. A measurement of hard disk performance.

initiator

See iSCSI initiator.

IOPS

I/O per second.

iSCSI

Internet-Small Computer Systems Interface. A TCP/IP protocol for carrying SCSI commands over IP networks.

iSCSI initiator

iSCSI-specific software installed on the host server that controls communications between the host server and the storage system.

iSNS

Internet Storage Naming Service. An automated discovery, management and configuration tool used by some iSCSI devices. iSNS eliminates the need to manually configure each individual storage system with a specific list of initiators and target IP addresses. Instead, iSNS automatically discovers, manages, and configures all iSCSI devices in your environment.

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

L

LAN

Local-area network. A computer network that spans a relatively small area, such as a single building or group of buildings.

Load

In UNIX computing, the system load is a measure of the amount of work that a computer system is doing.

Logical

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

M

MIB

Message Information Block.

Microcode

The lowest-level instructions directly controlling a microprocessor. Microcode is generally hardwired and cannot be modified. It is also referred to as firmware embedded in a storage subsystem.

Microsoft Cluster Server

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

P

Pair

Refers to two volumes that are associated with each other for data management purposes (for example, replication, migration). A pair is usually composed of a primary or source volume and a secondary or target volume as defined by you.

Pair status

Internal status assigned to a volume pair before or after pair operations. Pair status transitions occur when pair operations are performed or as a result of failures. Pair statuses are used to monitor copy operations and detect system failures.

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

Parity

The technique of checking whether data has been lost or corrupted when it's transferred from one place to another, such as between storage units or between computers. It is an error detection scheme that uses an extra checking bit, called the parity bit, to allow the receiver to verify that the data is error free. Parity data in a RAID array is data stored on member disks that can be used for regenerating any user data that becomes inaccessible.

Parity groups

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

Point-to-Point

A topology where two points communicate.

Port

An access point in a device where a link attaches.

Primary or local site

The host computer where the primary data of a remote copy pair (primary and secondary data) resides. The term “primary site” is also used for host failover operations. In that case, the primary site is the host computer where the production applications are running, and the secondary site is where the backup applications run when the applications on the primary site fail, or where the primary site itself fails.

R

RAID

Redundant Array of Independent Disks. A storage system in which part of the physical storage capacity is used to store redundant information about user data stored on the remainder of the storage capacity. The redundant information enables regeneration of user data in the event that one of the storage system's member disks or the access path to it fails.

RAID group

A set of disks on which you can bind one or more volumes.

Remote path

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

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

S

SAN

See Storage-Area Network

SAS

Serial Attached SCSI. An evolution of parallel SCSI into a point-to-point serial peripheral interface in which controllers are linked directly to disk drives. SAS delivers improved performance over traditional SCSI because SAS enables up to 128 devices of different sizes and types to be connected simultaneously.

SAS (ENC) Cable

Cable for connecting a controller box and drive box.

Secure Sockets Layer (SSL)

A protocol for transmitting private documents via the Internet. SSL uses a cryptographic system that uses two keys to encrypt data - a public key known to everyone and a private or secret key known only to the recipient of the message.

Snapshot

A term used to denote a copy of the data and data-file organization on a node in a disk file system. A snapshot is a replica of the data as it existed at a particular point in time.

SNM2

See Storage Navigator Modular 2.

Storage-Area Network

A dedicated, high-speed network that establishes a direct connection between storage systems and servers.

Storage Navigator Modular 2

A multi-featured scalable storage management application that is used to configure and manage the storage functions of Hitachi storage systems. Also referred to as "Navigator 2."

Striping

A way of writing data across drive spindles.

Subnet

In computer networks, a subnet or subnetwork is a range of logical addresses within the address space that is assigned to an organization. Subnetting is a hierarchical partitioning of the network address space of

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

an organization (and of the network nodes of an autonomous system) into several subnets. Routers constitute borders between subnets. Communication to and from a subnet is mediated by one specific port of one specific router, at least momentarily. SNIA.

Switch

A network infrastructure component to which multiple nodes attach. Unlike hubs, switches typically have internal bandwidth that is a multiple of link bandwidth, and the ability to rapidly switch node connections from one to another. A typical switch can accommodate several simultaneous full link bandwidth transmissions between different pairs of nodes. SNIA.

T

Target

The receiving end of an iSCSI conversation, typically a device such as a disk drive.

TCP

Transmission Control Protocol. A common Internet protocol that ensures packets arrive at the end point in order, acknowledged, and error-free. Usually combined with IP in the phrase TCP/IP.

10 GbE

10 gigabit Ethernet computer networking standard, with a nominal data rate of 10 Gbit/s, 10 times as fast as gigabit Ethernet

U

URL

Uniform Resource Locator. A standard way of writing an Internet address that describes both the location of the resource, and its type.

W

World Wide Name (WWN)

A unique identifier for an open systems host. It consists of a 64-bit physical address (the IEEE 48-bit format with a 12-bit extension and a 4-bit prefix). The WWN is essential for defining the SANTinel™ parameters because it determines whether the open systems host is to be allowed or denied

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

access to a specified logical unit or a group of logical units.

Z

Zoning

A logical separation of traffic between host and resources. By breaking up into zones, processing activity is distributed evenly.

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

A

- access attribute restrictions
 - LU deletion [5-38](#)
 - LU formatting [5-38](#)
 - RAID group deletion [5-38](#)
- access attributes [5-37](#), [5-40](#)
 - invisible from inquiry command [5-37](#)
 - protect [5-37](#)
 - read capacity [5-37](#)
 - read only [5-37](#)
 - read/write [5-37](#)
 - read-only [5-51](#)
 - restrictions [5-38](#)
 - setting with SNM2 [5-38](#)
 - S-VOL Disable [5-37](#)
- access attributes, assigning to logical units [5-40](#)
- Account Authentication
 - account types [5-9](#)
 - adding accounts [5-17](#)
 - default account [5-9](#)
 - deleting accounts [5-21](#)
 - modifying accounts [5-19](#)
 - permissions and roles [5-9](#)–[5-12](#)
 - session timeout settings [5-22](#)
 - viewing accounts [5-15](#)
- account types [5-9](#)
- Activities in Navigator 2 [3-41](#)
- Add Array wizard [4-6](#)
- Anti-virus software [2-10](#)
- Applet dialog box [3-37](#)
- array
 - turning off [11-47](#)
 - turning on [11-47](#)
- Array dialog box [3-36](#)
- Arrays screen
 - help [3-38](#)
- attribute
 - setting [5-51](#)
- attributes, access [5-40](#)
- attributes, assigning [5-40](#)
- audit logging
 - external syslog servers [5-35](#)
 - initializing logs [5-35](#)
 - protocol compliance [3-21](#)
 - setup guidelines [5-32](#)
 - syslog server [3-21](#)
 - transferring log data [5-32](#)
 - viewing log data [5-34](#)
- Audit Logging. *See* audit logging

B

- browser
 - client [A-5](#)
- Button Panel [3-41](#)

C

- Cache Partition Manager

- setting a pair cache partition [7-18, 7-19](#)
 - setup guidelines [7-15](#)
- Cache Residency Manager [5-38](#)
 - setting residency LUs [7-34](#)–[7-35](#)
 - setup guidelines [7-33](#)–[7-34](#)
- client
 - browser [A-5](#)
 - CPU [A-3, A-5](#)
 - disk capacity [A-5](#)
 - JRE [A-5](#)
 - memory [A-4, A-5](#)
 - monitor [A-5](#)
 - operating system [A-5](#)
- Command Suite Common Components
 - installation [3-3](#)
- concurrent use
 - LUN Manager [5-38](#)
 - password protection [5-38](#)
 - SNMP agent [5-38](#)
 - volume migration [5-38](#)
- Configuring
 - email alerts [4-9](#)
 - host ports [4-12](#)
 - management ports [4-11](#)
 - spare drives [4-14](#)
 - system date and time [4-14](#)
- Connecting
 - management console [4-3](#)
 - to a host [4-20](#)
- connecting to host [3-2](#)
- controller detachment [5-37](#)
- copy speed (pace). *See* Modular Volume Migration
- correction copy
 - dynamic sparing [5-37](#)
- CPU
 - client [A-5](#)
- Create & Map Volume wizard [4-15](#)
 - connecting to a host [4-20](#)
 - defining host groups or iSCSI targets [4-19](#)
 - defining logical units [4-18](#)
- creating
 - Host Groups (FC) [6-32](#)
 - iSCSI targets [6-40](#)

D

- Data Execution Prevention [3-5](#)
- Data Retention Utility
 - Expiration Lock configuration [5-50](#)
 - setting attributes [5-50](#)
 - setup guidelines [5-48, 5-48](#)
 - S-VOL configuration [5-50](#)
- Defining
 - host groups or iSCSI targets [4-19](#)
 - Logical units [4-18](#)
- deleting accounts. *See* Account Authentication
- disk capacity
 - client [A-5](#)
 - host [A-5](#)

- drive restoration [5-37](#)
 - correction copy [5-37](#)
- drive restoratoin
 - copy back [5-37](#)
- Dynamic Provisioning
 - logical unit capacity [7-27](#)

E

- Email alerts [4-9](#)
- environment [5-50](#)
- Explorer Panel [3-40](#)

F

- Failed installation on Windows [3-15](#)
- fibre channel
 - adding host groups [6-30](#)
 - deleting host groups [6-35](#)
 - initializing Host Group 000 [6-35](#)
- fibre channel setup workflow. *See* LUN Manager
- Firewalls [2-10](#)

G

- ghost disks [11-48](#)
- GUI
 - Specifying
 - Spin Down [11-86](#), [11-87](#), [11-88](#), [11-89](#)

H

- Hardware considerations [4-3](#)
- help
 - Arrays screen [3-38](#)
 - individual screen [3-39](#)
- High Availability cluster software [5-47](#)
- Hitachi SNMP Agent Support
 - additional resources [9-62](#)
 - confirming setup [9-25](#)
 - frame types [9-13](#)
 - installing [9-13](#)
 - license key [9-13](#)
 - management information base
 - defined [9-6](#)
 - extended MIBs [9-50](#)
 - dfCommandExecutionCondition group [9-54](#), [9-60](#), [9-56](#), [9-50](#), [9-51](#)
 - installation [9-35](#)
 - MIB access mode [9-29](#), [9-35](#)
 - at group [9-40](#)
 - egp group [9-46](#)
 - icmp group [9-46](#), [9-38](#), [9-41](#)
 - snmp group [9-47](#), [9-36](#)
 - tcp group [9-46](#)
 - udp group [9-46](#)
 - OID system assignment [9-29](#)
 - supported [9-29](#), [9-32](#)

- object identifiers [9-6](#)
- operating environment file [9-16](#)
- operational guidelines [9-26](#)
- preparing
 - SNMP manager [9-15](#)
 - storage array [9-15](#)
- referencing SNMP environment [9-24](#)
- registering SNMP environment [9-22](#)
- SNMP command messages [9-7](#)
- SNMP manager and agents [9-6](#)
- SNMP overview [9-2](#)
- SNMP traps [9-9](#)
- SNMP versions [9-5](#), [9-62](#)
- storage array name file [9-22](#)
- supported configurations [9-12](#)
- theory of operation [9-6](#)
- Hitachi Storage Command Suite Common Components, preinstallation considerations [2-10](#)
- Host
 - connecting to in Create & Map Volume wizard [4-20](#)
- host
 - connecting SNM2 [3-2](#)
 - disk capacity [A-5](#)
 - JDK [A-5](#)
 - memory [A-5](#)
- Host groups, defining [4-19](#)
- Host port configuration [4-12](#)
- HP-UX [5-46](#)

I

- Initial (Array) Setup wizard [4-8](#)
 - email alert configuration [4-9](#)
 - host port configuration [4-12](#)
 - management port configuration [4-11](#)
 - spare drive configuration [4-14](#)
 - system date and time configuration [4-14](#)
- Inquiry command [5-41](#)
- Installation
 - types [3-12](#)
- installation
 - Command Suite Common Components [3-3](#)
 - firewall [3-4](#)
 - Linux [3-2](#)
 - preparation [3-2](#)
 - services operational status [3-3](#)
 - Solaris [3-2](#)
 - windows [3-2](#)
- Installation fails on Windows [3-15](#)
- Installing
 - Navigator 2 [3-13](#)
- installing
 - preparation [3-2](#)
- installing SNM2 [3-2](#), [3-12](#)
- InstallShield Wizard
 - uninstalling [3-26](#)
- Interface of Navigator 2 [3-40](#)
- invisible from inquiry command, access attribute [5-37](#)
- Invisible mode [5-41](#)
- IPv6

- supported platforms [A-11](#)
- iSCSI
 - adding targets [6-45](#)
 - creating a target [6-40](#)
 - creating iSCSI targets [6-40](#)
 - deleting targets [6-49](#)
 - editing authentication properties [6-51](#)
 - editing target information [6-50](#)
 - host platform options [6-48](#)
 - initializing Target 000 [6-52](#)
 - nicknames, changing [6-52](#)
 - system configuration [6-14](#)
 - Target 000 [6-49](#)
 - using CHAP [6-39](#), [6-53](#), [??-8-9](#)
- iSCSI setup workflow. *See* LUN Manager
- iSCSI targets, defining [4-19](#)

J

- JDK
 - host [A-5](#)
- JRE
 - client [A-5](#)

L

- Linux
 - Navigator 2 installation [3-16](#)
- Logging in to Navigator 2 [4-4](#)
- logical unit
 - deleting [5-38](#)
 - growing [5-38](#)
 - inhibiting assignment as secondary volume [5-42](#)
 - shrinking [5-38](#)
- Logical units
 - defining [4-18](#)
- logical units
 - assigning access attributes [5-40](#)
 - deleting, growing, shrinking [5-38](#)
 - number allowed [5-37](#)
 - settable [5-37](#)
- logical units, protecting [5-40](#)
- logical units, using [5-40](#)
- logical volume
 - with Protect attribute [5-41](#)
- logical volumes
 - protecting [5-42](#)
- LU detachment [5-37](#)
- LUN Manager [5-38](#)
 - adding host groups [6-30-6-35](#)
 - creating iSCSI targets [6-40](#)
 - fibre channel setup workflow [6-25](#)
 - Host Group 000 [6-35](#)
 - host group security, fibre channel [6-30](#)
 - iSCSI setup workflow [6-26](#)
- LVM [5-46](#)

M

- Management console

- connecting to storage system [4-3](#)
- Management port configuration [4-11](#)
- memory
 - host [A-5](#)
- Menu Panel [3-40](#)
- Microsoft Windows
 - Navigator 2 installation [3-14](#)
 - Navigator 2 installation fails [3-15](#)
- migrating volumes. *See* Modular Volume Migration
- Modular Volume Migration
 - copy pace, changing [11-28](#)
 - migration pairs, canceling. [11-31](#)
 - migration pairs, confirming [11-29](#)
 - migration pairs, splitting [11-30](#)
 - Reserved LUs, adding [11-22](#)
 - Reserved LUs, deleting [11-26](#)
 - setup guidelines [11-21](#)
- monitor
 - client [A-5](#)

N

- Navigator 2
 - activities [3-41](#)
 - hardware considerations [4-3](#)
 - logging in [4-4](#)
 - operating environment [2-9](#)
 - terms [2-8](#)
 - understanding the interface [3-40](#)
- Navigator 2 installation [3-13](#)
 - fails on Microsoft Windows [3-15](#)
 - Linux [3-16](#)
 - Microsoft Windows [3-14](#)
 - Solaris [3-15](#)
 - types of [3-12](#)
- Navigator 2 interface
 - Button Panel [3-41](#)
 - Explorer Panel [3-40](#)
 - Menu Panel [3-40](#)
 - Page Panel [3-41](#)
- notes
 - failure [11-57](#)
 - operating system [11-48](#)
 - power down [11-54](#)
 - power up [11-56](#)
- NTP, using SNMP [3-38](#)

O

- Operating environment [2-9](#)
- operating system
 - Advanced Interactive eXecutive (AIX) [11-48](#)
 - client [A-5](#)
 - Hewlett Packard UNIX (HP-UX) [11-48](#)
 - Linux [11-48](#)
 - notes [11-48](#)
 - Solaris [11-48](#), [A-5](#)
 - Windows [11-48](#)
- operations
 - expiration lock [5-52](#)

- retention term [5-52](#)
- overview [1-5](#), [5-36](#)
 - Advanced Settings [1-6](#)
 - alerts and events [1-7](#)
 - command devices [1-6](#)
 - component status [1-5](#)
 - components [1-5](#)
 - DMLU [1-5](#)
 - E-mail alerts [1-6](#)
 - error monitoring [1-7](#)
 - FC settings [1-6](#)
 - firmware [1-6](#)
 - groups [1-5](#)
 - host groups [1-6](#)
 - iSCSI targets [1-6](#)
 - LANs [1-6](#)
 - licenses [1-6](#)
 - performance [1-7](#)
 - RAID groups [1-5](#)
 - security [1-6](#)
 - settings [1-5](#)
 - spare drives [1-6](#)

P

- Page Panel [3-41](#)
- password, default. *See* account types
- Performance Monitor
 - exporting information [8-30](#)
 - obtaining system information [8-8](#)
 - performance imbalance [8-36–8-37](#)
 - troubleshooting performance issues [8-36](#)
 - using graphs [8-8–8-9](#)
- permissions. *See* Account Authentication
- power down [11-53](#)
 - notes [11-54](#)
 - number of times [11-47](#)
 - problems [11-46](#)
 - UNIX [11-66](#)
 - Windows [11-64](#)
- Power Saving
 - effects [11-45](#)
 - modes [11-50](#)
 - operations [11-45](#)
 - requirements [11-45](#)
 - start [11-45](#)
- Power Saving Plus
 - specifications [11-71](#)
- power up [11-55](#)
 - notes [11-56](#)
 - time required [11-47](#)
 - UNIX [11-67](#)
 - Windows [11-64](#)
- Preconfigured
 - LUN on AMS 2000 storage systems [6-5](#)
- Preinstallation
 - anti-virus software [2-10](#)
 - firewalls [2-10](#)
- Preinstallation information on page 2-2 [7-33](#)

- preparation
 - installation [3-2](#)
 - Linux [3-4](#)
- protect, access attribute [5-37](#)
- protecting logical volume [5-41](#)
- protecting logical volumes [5-42](#)

R

- RAID groups
 - cannot power down [11-46](#)
- read capacity, access attribute [5-37](#)
- read only, access attribute [5-37](#)
- read/write, access attribute [5-37](#)
- read-only access attribute [5-51](#)
- read-only attribute
 - assigned to a logical volume [5-40](#)
 - copying data from utilities [5-40](#)
- read-write operations
 - about [5-40](#)
 - copying data from utilities [5-40](#)
 - protecting [5-40](#)
 - restricting [5-40](#)
 - volumes with attribute [5-40](#)
 - with open systems volumes [5-40](#)
 - with ShadowImage [5-40](#)
 - with SnapShot [5-40](#)
 - with TCE [5-40](#)
 - with TrueCopy [5-40](#)
- Red Hat Linux
 - installation [3-2](#)
 - preparation [3-4](#)
 - setting kernel [3-5](#)
 - starting SNM2 [3-30](#)
- Remote Desktop [3-4](#)
- report zero read cap [5-41](#)
- restrictions
 - operating systems [5-46](#)
- retention terms [5-41](#)

S

- scripts
 - UNIX [11-66](#)
 - Windows [11-64](#)
- scripts, samples [11-63](#)
- security [11-59](#)
- security, setting iSCSI target [6-43](#), [6-44](#)
- setting and attribute [5-51](#)
- setting kernel
 - Red Hat Linux [3-5](#)
 - Solaris 10 [3-8](#)
 - Solaris 8 [3-7](#)
 - Solaris 9 [3-5](#)
- SnapShot [5-45](#)
- SNM2
 - Applet dialog box [3-37](#)
 - Applet screen [3-36](#)
 - array and SMS Array dialog boxes [3-36](#)
 - operations [3-35](#)

- SNM2 Server [3-37](#)
- SNMP
 - MIB information [3-22](#)
- SNMP manager, dual-controller environment [3-23](#)
- Solaris
 - installation [3-2](#)
 - Navigator 2 installation [3-15](#)
 - preparation [3-4](#)
 - setting kernel [3-7](#)
 - starting SNM2 [3-30](#)
- Solaris 8
 - setting kernel [3-7](#)
- Solaris 9
 - setting kernel [3-7](#)
- Spare drive configuration [4-14](#)
- specifications [11-71](#)
 - access attribute change [5-37](#)
 - access attribute restrictions [5-38](#)
 - access attributes [5-37](#)
 - Cache Partition Manager [5-38](#)
 - Cache Residency Manager [5-38](#)
 - controller detachment [5-37](#)
 - deleting, growing, shrinking LUs [5-38](#)
 - drive restoration [5-37](#)
 - dynamic provisioning [5-38](#)
 - firmware replacement [5-37](#)
 - logical unit detachment [5-37](#)
 - number of settable LUs [5-37](#)
 - password protection [5-38](#)
 - powering on/off [5-37](#)
 - retention term, range [5-38](#)
 - ShadowImage [5-37](#)
 - SnapShot [5-37](#)
 - SNMP agent [5-38](#)
 - TCE [5-37](#)
 - TrueCopy [5-37](#)
 - unified logical unit [5-38](#)
 - unit of setting [5-37](#)
 - unsupported logical units [5-37](#)
 - Volume Migration [5-38](#)
- specifications
 - expansion of RAID group [5-38](#)
- starting
 - SNM2 [3-27](#)
- starting SNM2
 - client [3-28](#)
 - host [3-27](#)
 - Linux, Solaris [3-30](#)
 - Windows [3-30](#)
- S-VOL Disable [5-37](#)
- S-VOL Disable, access attribute [5-37](#)
- syslog server. *See* audit logging
- system configuration [6-14](#)
- System date and time configuration [4-14](#)
- system requirements
 - AIX [A-6](#)
 - HP-UX [A-5](#)
 - IRIX [A-11](#)
 - Linux [A-7](#)

Solaris [A-4](#)

T

TCE [5-45](#)

Terms associated with Navigator 2 [2-8](#)

timeout length, changing [5-22](#)

Troubleshooting

 installation fails on Microsoft Windows system [3-15](#)

 installation fails on Windows [3-15](#)

Types of installation [3-12](#)

U

Understanding the Navigator 2 interface [3-40](#)

unified logical unit [5-38](#)

uninstalling

 add or remove programs [3-26](#)

 InstallShield Wizard [3-26](#)

unit of setting [5-37](#)

UNIX

 power down [11-66](#)

 power up [11-67](#)

 use cases [11-66](#)

Unix [5-46](#)

unsupported logical units

 command device [5-37](#)

 DMLU [5-37](#)

 LU as data pool in SnapShot/TCE [5-37](#)

 sub-LU, unified LU [5-37](#)

 unformatted LU [5-37](#)

use cases [11-63](#)

V

Volume Migration [5-38](#), [5-51](#)

W

Windows

 installation [3-2](#)

 power down [11-64](#)

 power up [11-64](#)

 starting SNM2 [3-28](#)

 use cases [11-64](#)

Windows 2000 [5-46](#)

Windows Server 2003 [5-46](#)

Windows Server 2008 [5-46](#)

Wizards

 Add Array [4-6](#)

 Create & Map Volume [4-15](#)

 Initial (Array) Setup [4-8](#)

Hitachi Data Systems

Corporate Headquarters

2845 Lafayette Street
Santa Clara, California 95050-2639
U.S.A.

www.hds.com

Regional Contact Information

Americas

+1 408 970 1000

info@hds.com

Europe, Middle East, and Africa

+44 (0)1753 618000

info.emea@hds.com

Asia Pacific

+852 3189 7900



MK-91DF8275-25