HITACHI
Inspire the Next

Hitachi Storage Navigator Modular 2 Storage Features Reference Guide for AMS

Covers the following Program Products:

- Account Authentication
- Audit Logging
- Cache Partition Manager
- Cache Residency Manager
- Data Retention Utility
- LUN Manager
- Performance Monitor
- SNMP Agent Support

FastFind Links

Document Organization

Product Version

Getting Help

Table of Contents

Hitachi Data Systems

MK-97DF8148-07

Copyright © 2010 Hitachi Ltd., Hitachi Data Systems Corporation, 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 agreement(s). The use of Hitachi Data Systems products is governed by the terms of your agreement(s) 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.

Contents

1 Introduction 1-1

Account Authentication 1-2
User management 1-2
User authentication 1-2
Access control 1-2
Migrating from Password Protection to Account Authentication. 1-3
Key similarities 1-3
Key differences 1-3
Advanced Security Mode 1-3
Audit Logging. 1-4
Cache Partition Manager 1-5
Cache Residency Manager 1-6
Data Retention Utility 1-6
LUN Manager 1-7
Fibre Channel features 1-8
iSCSI features 1-9
iSCSI protocol 1-10
iSCSI network configuration 1-10
Performance Monitor 1-10
SNMP Agent Support. 1-11
Trap-issuing processing 1-12
Request processing 1-13
Modular Volume Migration 1-15
Usage guidelines 1-17
Environments 1-17
Requirements 1-17
Requirements for installing and enabling features 1-17
Audit Logging 1-18
Cache Partition Manager. 1-18
Modular Volume Migration 1-18
Requirements for uninstalling and disabling features 1-18
Account Authentication. 1-19

Cache Partition Manager 1-19
Data Retention 1-19
LUN Manager 1-19
Modular Volume Migration 1-19
SNMP Agent 1-19
Additional guidelines 1-20
Advanced Settings Java Applet 1-20

2 Installing and enabling storage features 2-1

Preinstallation information. 2-2
Environments 2-2
Storage feature requirements. 2-2
Requirements for installing and enabling features. 2-2
Audit Logging requirements 2-3
Cache Partition Manager requirements 2-3
Data Retention requirements 2-3
LUN Manager requirements 2-4
SNMP Agent requirements 2-4
Modular Volume Migration requirements 2-4
Additional guidelines 2-5
Installing storage features. 2-5
Enabling storage features 2-5
Disabling storage features. 2-5
Uninstalling storage features 2-6

3 Account Authentication 3-1

Account Authentication overview 3-2
Overview of Account Authentication 3-2
Accounts 3-3
Account types 3-4
Roles 3-5
Resources. 3-5
Session 3-7
Session types for operating resources 3-7
Warning banners. 3-8
Advanced Security Mode 3-8
Account Authentication procedures 3-10
Initial settings. 3-10
Managing accounts 3-10
Displaying accounts. 3-10
Adding accounts 3-11
Changing the Advanced Security Mode 3-13
Modifying accounts 3-14
Deleting accounts 3-15
Changing session timeout length 3-16

Forcibly logging out 3-17
Troubleshooting 3-17

4 Audit Logging. 4-1

Audit Logging overview 4-2
Audit Logging procedures 4-3
Initial settings 4-3
Optional operations 4-3
Enabling Audit Log data transfers 4-3
Viewing Audit Log data 4-5
Initializing logs 4-6
Configuring Audit Logging to an external Syslog server 4-6

5 Cache Partition Manager. 5-1

Cache Partition Manager overview 5-2
Cache Partition Manager 5-2
Partition Capacity. 5-3
Segment And Stripe Size 5-9
Restrictions 5-10
Cache Partition Manager settings 5-12
Initial settings 5-12
Stopping Cache Partition Manager 5-13
Working with cache partitions 5-13
Adding cache partitions 5-14
Deleting cache partitions 5-15
Assigning cache partitions 5-16
Setting a pair cache partition 5-16
Changing cache partitions. 5-17
Changing cache partition owner controllers 5-17
Installing SnapShot or TCE or Dynamic Provisioning under Cache Partition Manager 5-18

6 Cache Residency Manager 6-1

Cache Residency Manager overview 6-2
Termination Conditions 6-2
Disabling Conditions. 6-3
Equipment 6-3
Logical Unit Capacity 6-4
. 6-11
Restrictions 6-12
Cache Residency Manager operations 6-14
Initial settings 6-14
Stopping Cache Residency Manager 6-14
Setting and canceling residency logical units 6-14

. 6-14

7 Data Retention Utility 7-1

Data Retention Utility overview 7-2
Usage 7-3
Logical unit access attributes 7-3
Unified logical units 7-4
SnapShot and TCE. 7-4
SYNCHRONIZE CACHE command 7-4
Host Side Application example 7-4
Operating System (OS) Restrictions. 7-4
Logical units attributes set from the operating system. 7-4
Data Retention Utility operations 7-6
Initial settings 7-6
Optional operations 7-6
Opening the Data Retention window 7-6
Setting attributes. 7-8
Setting S-VOLs 7-8
Setting expiration locks 7-8

8 LUN Manager 8-1

LUN Manager overview 8-2
Design configurations and best practices 8-4
Fibre Channel configuration 8-4
Fibre Channel design considerations 8-6
iSCSI system design considerations 8-7
Assigning iSCSI targets and volumes to hosts. 8-15
Preventing unauthorized SAN access 8-17
Avoiding RAID Group Conflicts 8-18
SAN queue depth setting 8-20
Increasing queue depth and port sharing. 8-20
Increasing queue depth through path switching 8-21
LUN Manager operations. 8-23
Using Fibre Channel. 8-23
Using iSCSI. 8-23
Fibre Channel operations using LUN Manager. 8-24
Adding host groups 8-24
Enabling and disabling host group security. 8-24
Creating and editing host groups 8-26
Initializing Host Group 000 8-28
Deleting host groups 8-29
Changing nicknames 8-29
Deleting World Wide Names 8-29
Copy settings to other ports 8-30

iSCSI operations using LUN Manager 8-31
Creating an iSCSI target 8-32
Using the iSCSI Target Tabs 8-32
Setting the iSCSI target security 8-33
Editing iSCSI target nicknames.. 8-34
Adding and deleting targets 8-36
Editing target information 8-38
Editing authentication properties 8-38
Initializing Target 000 8-39
Changing a nickname 8-40
CHAP users 8-40
Adding a CHAP user 8-40
Changing the CHAP user 8-41

9 Performance Monitor 9-1

Performance Monitor overview 9-2
Performance Monitor operations 9-2
Initial settings 9-2
Optional operations 9-2
Optimizing system performance 9-3
Obtaining information 9-3
Using graphic displays 9-3
Working with the Performance Monitor Tree View 9-6
More About Tree View Items in Performance Monitor 9-8
Using Performance Monitor with Dynamic Provisioning 9-12
Working with Graphing and Dynamic Provisioning 9-13
Explanation of Displayed Items 9-14
Determining the Ordinate Axis 9-18
Saving Monitoring Data 9-21
Exporting Performance Monitor Information 9-22
Enabling Performance Measuring Items 9-26
Working with Port Information 9-28
Working with RAID Group, DP Pool and Logical Unit Information 9-28
Working with Cache Information 9-28
Working with Processor Information 9-28
Troubleshooting Performance 9-29
Performance Imbalance and Solutions 9-29
. 9-30

10 SNMP Agent Support 10-1

SNMP Agent Support overview 10-2
Error status 10-3
Dual controller GET/TRAP specifications 10-4
SNMP functions 10-5

TRAP reporting 10-5
Extended TRAPs 10-6
Request processing 10-7
Additional SNMP environment requirements 10-8
SNMP Agent Support operations 10-8
Managing SNMP Agent Support 10-9
SNMP setup 10-9
Disk array-side setup 10-9
SNMP Manager-side setup 10-9
Checking the connection 10-9
Creating environmental information files 10-10
Environment setting file 10-10
Array name setting file 10-13
Registering SNMP environmental information 10-14
Referencing the SNMP environment information file 10-16
Verifying SNMP connections 10-16
Detecting failures 10-18

11 Modular Volume Migration 11-1

Modular Volume Migration overview 11-2
Environments and Requirements 11-4
Setting up Volume Migration 11-5
Setting Logical Units to be recognized by the host 11-5
VxVM 11-7
MSCS 11-8
AIX 11-8
Windows 2000/Window Server 2003/Windows Server 2008 11-8
Linux and LVM 11-8
Windows 2000/Windows Server 2003/Windows Server 2008 and Dynamic Disk 11-8
Performance 11-8
Using unified logical units 11-9
Using with the Data Retention Utility 11-11
Using with ShadowImage 11-12
Using with Cache Partition Manager 11-14
Modular Volume Migration operations 11-14
Managing Modular Volume Migration 11-16
Adding reserved logical units 11-16
Deleting reserved logical units 11-17
Changing copy pace 11-22
Confirming Volume Migration Pairs 11-23
Splitting Volume Migration pairs 11-24
Canceling Volume Migration pairs 11-25

A Appendix A — Logical unit expansion/reduction A-1

Overview A-2
Expanding capacity using LUN "Grow/Shrink" A-2
Requirements and conditions for expanding logical units A-2
Changing logical unit capacity A-4
Displaying unified logical unit properties A-5
Separating unified logical units A-5
Separating the last logical unit A-5
Separating all logical units A-6

Available unified LUN information A-7

Glossary

Index

Preface

This document provides facilities requirements for preparing and installing Hitachi Adaptable Modular Storage (AMS) 2100, 2300, and 2500 storage systems. In this document, these storage systems are referred to collectively as the Hitachi AMS 2000 Family storage systems. If information pertains to certain members of this family, those systems are identified.

Using this document, you will be able to prepare your site for the arrival and installation of your units. To determine the total components your shipment will include, please consult your Hitachi Data Systems representative.

This preface includes the following information:

- Document revision level
- Changes in this revision
- Intended audience
- Document organization
- Document conventions
- Related documents
- Regulatory information
- Getting help
- Comments

Document revision level

This section provides a history of the revision changes to this document.

Revision	Date	Description
MK-97DF8148-P	July 2008	Preliminary Release
MK-97DF8148-00	October 2008	Revision 00, supersedes and replaces MK-97DF8148-P
MK-97DF8148-01	December 2008	Revision 01, supersedes and replaces MK-97DF8148-00
MK-97DF8148-02	March 2009	Revision 02, supersedes and replaces MK-97DF8148-01
MK-97DF8148-03	April 2009	Revision 03, supersedes and replaces MK-97DF8148-02
MK-97DF8148-04	May 2009	Revision 04, supersedes and replaces MK-97DF8148-03
MK-97DF8148-05	August 2009	Revision 05, supersedes and replaces MK-97DF8148-04
MK-97DF8148-06	November 2009	Revision 06, supersedes and replaces MK-97DF8148-05
MK-97DF8148-07	April 2010	Revision 07, supersedes and replaces MK-97DF8148-06

Changes in this revision

- New informationabout password and roles of the built-in account in chapter 3, Account Authentication.
- New content about Changing Licenses.
- New steps for setting Account Authentication using the Security Mode Wizard in chapter 3, Account Authentication.
- New feature, Advanced Security Mode in chapter 3, Account Authentication.
- New steps involved with displaying, modifying, and changing passwords for account information in chapter 3, Account Authentication.
- New content for Volume Migratio in chapter 11, Volume Migration.
- New content for Performance Monitor in chapter 9, Performance Monitor.

Intended audience

This document is intended for personnel who will schedule, manage, and perform the tasks required to prepare your site for installing a Hitachi AMS 2000 Family storage systems.

Product version

This document applies to Hitachi AMS 2000 Family firmware version 0890/A or later.

Document organization

The following table provides an overview of the contents and organization of this document. Click the chapter title in the first column to go to that chapter. The first page of every chapter or appendix contains a brief list of the contents of that section of the manual, with links to the pages where the information is located.

Chapter/Appendix Title	Description
Chapter 1, Introduction	Describes features in the Navigator 2 environment.
Chapter 2, Installing and enabling storage features	Describes installing and enabling storage features.
Chapter 3, Account Authentication	Describes how to create permissions for selected users who will be authenticated when they attempt to access the storage system.
Chapter 4, Audit Logging	Describes how the Audit Log facility works and where to retrieve messages sent to the log.
Chapter 5, Cache Partition Manager	Describes how to segment the storage system into discrete partitions, and provides allowable increments.
Chapter 6, Cache Residency Manager	Describes how Cache Residency Manager works.
Chapter 7, Data Retention Utility	Describes how to retain data on the storage system and to create settings that determine how much data will be retained and to specify a retention interval.
Chapter 8, LUN Manager	Describes how to work logical unit numbers on the storage system.
Chapter 9, Performance Monitor	Describes how to monitor activity and responsiveness on the storage system using the Performance Monitor tool.
Chapter 10, SNMP Agent Support	Describes how to implement Simple Network Management Protocol on the storage system and to work with SNMP Set and Get commands.
Chapter 11, Modular Volume Migration	Describes how to work with Modular Volume Migration on the storage system.
Appendix A — Logical unit expansion/reduction	Describes how to work with logical unit expansion and reduction on the storage system.

Convention for storage capacity values

Physical storage capacity values (e.g., disk drive capacity) are calculated based on the following values:

Physical capaciy 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 (e.g., logical device capacity) are calculated based on the following values:

Logical capaciy 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

Document conventions

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 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**.
Italic	Indicates a variable, which is a placeholder for actual text provided by the user or system. Example: copy *source-file target-file* Angled brackets (< >) are also used to indicate variables.
screen/code	Indicates text that is displayed on screen or entered by the user. Example: `# pairdisplay -g oradb`
< > angled brackets	Indicates a variable, which is a placeholder for actual text provided by the user or system. Example: `# pairdisplay -g <group>` 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]

Accessing product documentation

The AMS 2000 Family user documentation is available on the Hitachi Data Systems 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.

This documentation set consists of the following documents.

Release notes

- Adaptable Modular Storage System Release Notes
- Storage Navigator Modular 2 Release Notes

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

Installation and getting started

The following documents provide instructions for installing an AMS 2000 Family storage system. They include rack information, safety information, site-preparation instructions, getting-started guides for experienced users, and host connectivity information. The symbol ☞ identifies documents that contain initial configuration information about Hitachi AMS 2000 Family storage systems.

☞ **AMS2100/2300 Getting Started Guide**, MK-98DF8152

Provides quick-start instructions for getting an AMS 2100 or AMS 2300 storage system up and running as quickly as possible.

☞ **AMS2500 Getting Started Guide**, MK-97DF8032

Provides quick-start instructions for getting an AMS 2500 storage system up and running as quickly as possible.

AMS 2000 Family Site Preparation Guide, MK-98DF8149

Contains initial site planning and pre-installation information for AMS 2000 Family storage systems, expansion units, and high-density expansion units. This document also covers safety precautions, rack information, and product specifications.

AMS 2000 Family Fibre Channel Host Installation Guide, MK-08DF8189

Describes how to prepare Hitachi AMS 2000 Family Fibre Channel storage systems for use with host servers running supported operating systems.

AMS 2000 Family iSCSI Host Installation Guide, MK-08DF8188

Describes how to prepare Hitachi AMS 2000 Family iSCSI storage systems for use with host servers running supported operating systems.

Storage and replication features

The following documents describe how to use Storage Navigator Modular 2 (Navigator 2) to perform storage and replication activities.

Storage Navigator 2 Advanced Settings User's Guide, MK-97DF8039

Contains advanced information about launching and using Navigator 2 in various operating systems, IP addresses and port numbers, server certificates and private keys, boot and restore options, outputting configuration information to a file, and collecting diagnostic information.

Storage Navigator Modular 2 User's Guide, MK-99DF8208

Describes how to use Navigator 2 to configure and manage storage on an AMS 2000 Family storage system.

AMS 2000 Family Dynamic Provisioning Configuration Guide, MK-09DF8201

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

Storage Navigator 2 Storage Features Reference Guide for AMS, MK-97DF8148 — this document

Contains concepts, preparation, and specifications for Account Authentication, Audit Logging, Cache Partition Manager, Cache Residency Manager, Data Retention Utility, LUN Manager, Performance Monitor, SNMP Agent, and Modular Volume Migration.

AMS 2000 Family Copy-on-write SnapShot User Guide, MK-97DF8124

Describes how to create point-in-time copies of data volumes in AMS 2100, AMS 2300, and AMS 2500 storage systems, without impacting host service and performance levels. Snapshot copies are fully read/write compatible with other hosts and can be used for rapid data restores, application testing and development, data mining and warehousing, and nondisruptive backup and maintenance procedures.

AMS 2000 Family ShadowImage In-system Replication User Guide, MK-97DF8129

Describes how to perform high-speed nondisruptive local mirroring to create a copy of mission-critical data in AMS 2100, AMS 2300, and AMS 2500 storage systems. ShadowImage keeps data RAID-protected and fully recoverable, without affecting service or performance levels. Replicated data volumes can be split from host applications and used for system backups, application testing, and data mining applications while business continues to operate at full capacity.

AMS 2000 Family TrueCopy Remote Replication User Guide, MK-97DF8052

Describes how to create and maintain multiple duplicate copies of user data across multiple AMS 2000 Family storage systems to enhance your disaster recovery strategy.

AMS 2000 Family TrueCopy Extended Distance User Guide, MK-97DF8054

Describes how to perform bi-directional remote data protection that copies data over any distance without interrupting applications, and provides failover and recovery capabilities.

AMS 2000 Data Retention Utility User's Guide, MK-97DF8019

Describes how to lock disk volumes as read-only for a certain period of time to ensure authorized-only access and facilitate immutable, tamper-proof record retention for storage-compliant environments. After data is written, it can be retrieved and read only by authorized applications or users, and cannot be changed or deleted during the specified retention period.

Storage Navigator Modular 2 online help

Provides topic and context-sensitive help information accessed through the Navigator 2 software.

Hardware maintenance and operation

The following documents describe how to operate, maintain, and administer an AMS 2000 Family storage system. They also provide a wide range of technical information and specifications for the AMS 2000 Family storage systems. The symbol ☞ identifies documents that contain initial configuration information about Hitachi AMS 2000 Family storage systems.

☞ **AMS 2100/2300 Storage System Hardware Guide**, MK-97DF8010

Provides detailed information about installing, configuring, and maintaining AMS 2100 and 2300 storage systems.

☞ **AMS 2500 Storage System Hardware Guide**, MK-97DF8007

Provides detailed information about installing, configuring, and maintaining an AMS 2500 storage system.

☞ **AMS 2000 Family Storage System Reference Guide**, MK-97DF8008

Contains specifications and technical information about power cables, system parameters, interfaces, logical blocks, RAID levels and configurations, and regulatory information about AMS 2100, AMS 2300, and AMS 2500 storage systems. This document also contains remote adapter specifications and regulatory information.

AMS 2000 Family Storage System Service and Upgrade Guide, MK-97DF8009

Provides information about servicing and upgrading AMS 2100, AMS 2300, and AMS 2500 storage systems.

AMS 2000 Family Power Savings User Guide, MK-97DF8045

Describes how to spin down volumes in selected RAID groups when they are not being accessed by business applications to decrease energy consumption and significantly reduce the cost of storing and delivering information.

Command and Control (CCI)

The following documents describe how to install the Hitachi AMS 2000 Family Command Control Interface (CCI) and use it to perform TrueCopy and ShadowImage operations.

AMS 2000 Family Command Control Interface (CCI) Installation Guide, MK-97DF8122

Describes how to install CCI software on open-system hosts.

AMS 2000 Family Command Control Interface (CCI) Reference Guide, MK-97DF8121

Contains reference, troubleshooting, and maintenance information related to CCI operations on AMS 2100, AMS 2300, and AMS 2500 storage systems.

AMS 2000 Family Command Control Interface (CCI) User's Guide, MK-97DF8123

Describes how to use CCI to perform TrueCopy and ShadowImage operations on AMS 2100, AMS 2300, and AMS 2500 storage systems.

Command Line Interface (CLI)

The following documents describe how to use Hitachi Storage Navigator Modular 2 to perform management and replication activities from a command line.

Storage Navigator Modular 2 Command Line Interface (CLI) Unified Reference Guide, MK-97DF8089

Describes how to interact with all Navigator 2 bundled and optional software modules by typing commands at a command line.

Storage Navigator 2 Command Line Interface Replication Reference Guide for AMS, MK-97DF8153

Describes how to interact with Navigator 2 to perform replication activities by typing commands at a command line.

Dynamic Replicator documentation

The following documents describe how to install, configure, and use Hitachi Dynamic Replicator to provide AMS Family storage systems with continuous data protection, remote replication, and application failover in a single, easy-to-deploy and manage platform.

Dynamic Replicator - Scout Release Notes, RN-99DF8211

Dynamic Replicator - Scout Host Administration Guide, MK-98DF8212

Dynamic Replicator - Scout Installation and Configuration Guide, MK-98DF8213

Dynamic Replicator - Scout Quick Start Guide, MK-98DF8214

Dynamic Replicator - Scout Host Troubleshooting Guide, MK-98DF8215

Dynamic Replicator DR-Scout ICAT Utility Guide, MK-98DF8216

Dynamic Replicator - Scout RX Server Deployment Guide, MK-98DF8217

Dynamic Replicator VX Solution for Oracle (Solaris), MK-98DF8218

Dynamic Replicator - Scout Solution for SharePoint 2007, MK-98DF8219

Dynamic Replicator - Scout Solution for MySQL (Windows), MK-98DF8220

Protecting Citrix XenServer Using Hitachi Dynamic Replicator - Scout, MK-98DF8221

Dynamic Replicator Quick Install/Upgrade Guide, MK-98DF8222

Dynamic Replicator - Scout Protecting MS SQL Server, MK-98DF8223

Dynamic Replicator - Scout - Protecting Microsoft Exchange Server, MK-98DF8224

Dynamic Replicator - Scout File Server Solution, MK-98DF8225

Dynamic Replicator - Scout ESX - Protecting ESX Server (RCLI), MK-99DF8226

Getting help

If you need to contact the Hitachi Data Systems support center, please provide as much information about the problem as possible, including:

- The circumstances surrounding the error or failure.
- The exact content of any messages displayed on the host system(s).
- The exact content of any messages displayed on Storage Navigator Modualr 2.
- The Storage Navigator Modular 2 configuration information. This information is used by service personnel for troubleshooting purposes.

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 Hitachi Data Systems 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 section(s) and paragraph(s) whenever possible.

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

1

Introduction

This chapter provides information on AMS 2000 Family storage features available from Hitachi Storage Navigator Modular 2 Graphical User Interface (GUI) and covers the following topics:

- Account Authentication
- Audit Logging
- Cache Partition Manager
- Cache Residency Manager
- Data Retention Utility
- LUN Manager
- Performance Monitor
- SNMP Agent Support
- Modular Volume Migration
- Usage guidelines
- Advanced Settings Java Applet

NOTE: Some storage features may require the Java Runtime Environment (JRE) on your computer.

Account Authentication

Account Authentication is a feature that ensures the security of the disk array by protecting it from attacks such as illegal break-in and illegal operation from the management LAN interface. This feature protects the information on the disk array configuration and user data. It authenticates users who access the disk array, providing access (for monitoring and configuration) from the disk array resources based on the account information that is registered on the disk array.

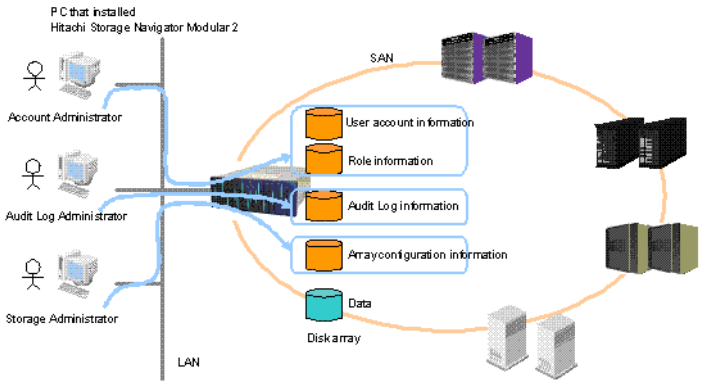

Figure 1-1: Account Authentication Outline

To ensure that unauthorized users cannot access the disk array, Account Authentication consists of user management, user authentication, and access control (see Figure 1-2 on page 1-2). Users are authenticated based on the account information registered with the array.

You can manage, authenticate, and control the access of users.

User management

This function registers the user account information in the disk array (for example, ID, password, role). The user must register their account information.

User authentication

This function authenticates users (based on the account information) when they access or log in the disk array, and restricts their permissions appropriately.

Access control

This function controls access to the disk array resources based on the role type assigned to the account.

Figure 1-2: Account Authentication Example

Migrating from Password Protection to Account Authentication

There are some important similarities and differences to note if you have been using the Password Protection security feature to protect other AMS or SMS storage systems and are now migrating to the Account Authentication security feature.

NOTE: Maintain a secure environment and change the "built-in" default "root" password after logging in for the first time under Account Authentication.

Key similarities

- Account Authentication restricts access at the storage system (array) level. As with
- Password Protection, user names and passwords must be configured on the secured array itself
- The "built-in" or default root user account should only be used to create user names and passwords.
- Assigning the same Navigator 2 login information (user name and password) when creating an account under Account Authentication provides seamless access to both Navigator to and the secured array.
- Enabling or disabling the Account Authentication feature immediately puts the user back into the main Navigator 2 Array List window. If enabled on a specific array, the first-time login requires the "built-in" default account information to access the array and create accounts

Key differences

- Account Authentication provides role-based permissions for user accounts. Account or security administrators should consider the role(s) or account types to be assigned to a user. For more information about account types and role assignments, see Account types on page 3-4.
- Password Protection and Account Authentication are mutually exclusive and cannot be enabled at the same time for a given array
- User name and password information is not inherited if you switch from Password Protection to Account Authentication. New accounts and role assignments must be created under Account Authentication

Advanced Security Mode

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

Table 1-1: Advanced Security Mode Specifications

Feature	Description	Specification
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 • **Encryption type:** The encryption is executed using SHA256 when it is enabled and MD5 when it is disabled.

You need a built-in account to perform Advanced Security Mode operations. The mode can be set only when the storage system runs firmware version 0890/A or greater and the management PC runs Navigator 2 version 9.00 or greater.

By changing the Advanced Security Mode, the storage system removes or initializes the following information. As necessary, check the information in advance, and set it again after changing the Advanced Security Mode.

All session during login (accounts during login are logged out).

All public accounts registered in the array.

Role and password of the built-in account.

Audit Logging

Audit Logging audits and defers inappropriate disk array actions by recording the user, the operation, the location, and then creating a log (see Figure 1-3). The audit log is sent to the Syslog server using port 514 using the User Diagram Protocol (UDP). The log can also be saved inside the disk array as backup information in case a network or the Syslog server fails.

A log is sent when:

- An operation occurs outside the disk array
- Starting and terminating the disk array

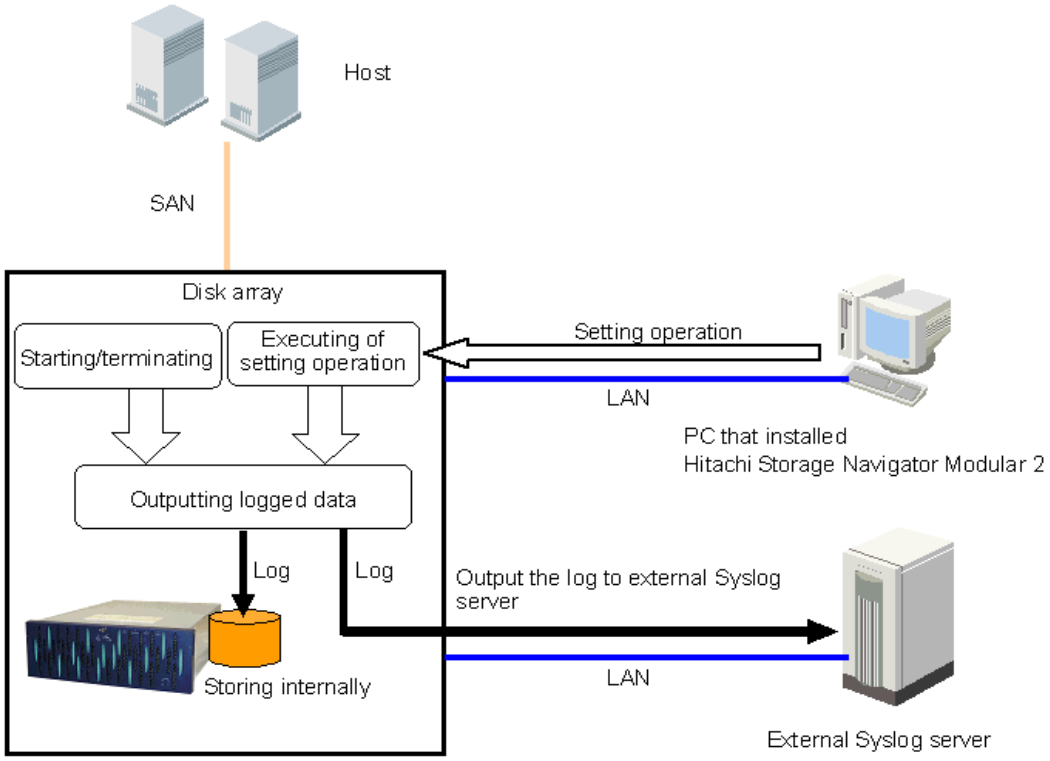

Figure 1-3: Audit Logging Example

For example, if user A accesses the disk array with Navigator 2 and creates a RAID group when setting an operation outside the disk array, the array creates a log where at x hours y minutes z seconds on m month d day in 2008, user A succeeded (or failed) creating a RAID group from a computer where Navigator 2 was operating and sends the log to the external Syslog server.

If the disk array enters the Ready status at the time of a status change (system event) inside the array, the array creates a log where "at x hours y minutes z seconds on m month d day in 2008, success of Subsystem Ready" and sends the log to the Syslog server (who, from, and where are not created because the status change is internal).

Cache Partition Manager

The cache memory on a disk array is a gateway for receiving/sending data from/to a host. In the array, the cache memory is divided into a system control area and a user data area. When sending and receiving data, the user data area is used.

The Cache Partition Manager divides the array user data area more finely, into partitions; then, a logical unit defined in the array is assigned to the partition.

A user can specify the partition and segment size (size of a data management unit). You can optimize the data reception/sending from/to a host by assigning the most suitable partition to a logical unit according to the data received from a host.

Cache Residency Manager

The Cache Residency Manager ensures that all the data in a logical unit is stored in cache memory. All read/write commands to the logical unit can be executed at a 100% cache hit rate without accessing the drive. Since a latency period is not needed to access the disk drive, the throughput is improved when this function is applied to a logical unit that contains data accessed frequently.

As shown in Figure 1-4 on page 1-6, part of the cache memory installed in the controller is used for the Cache Residency Manager. Cache memory utilizes a battery backup on both controllers, and the data is duplicated on each controller in case of a power failure, cache package failure, and so on.

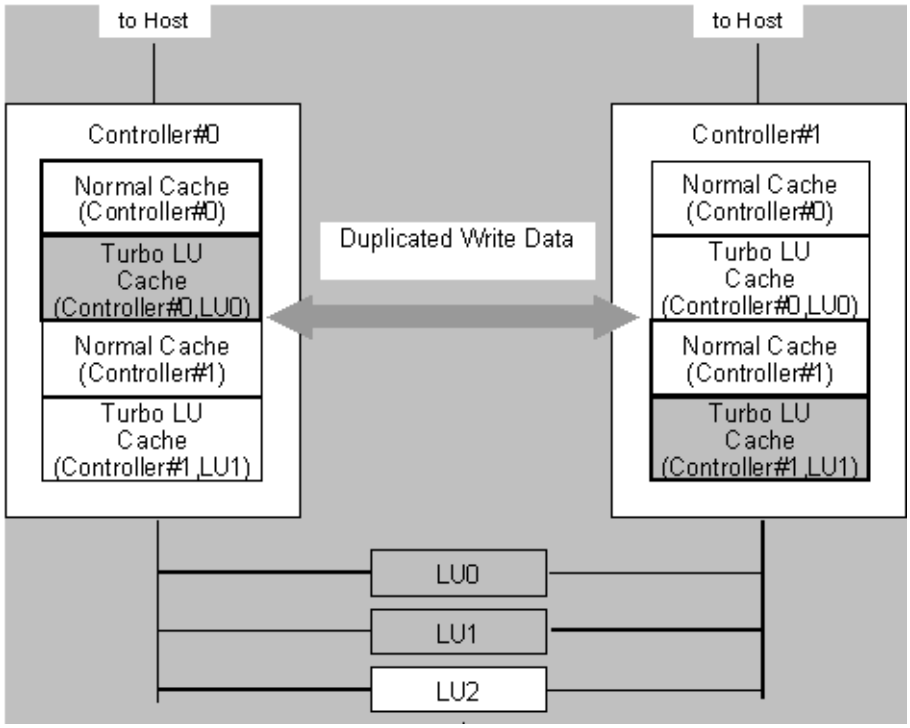

Figure 1-4: Cache Residency Manager Example

Data Retention Utility

The Data Retention Utility, which requires a separate license purchase, protects your disk array data and LUNs from input/output (I/O) operations performed by an open-systems hosts. It may help you comply to Federal mandates that certain day (files) be protected, for example:

- Emails/Email server data

- Health Records
- Banking transactions
- Brokerage transactions

Data Retention lets you assign an access attribute to each logical volume. You can use a logical volume as a read-only volume, and protect it from read and write operations.

NOTE: Logical volumes are sometimes referred to as logical devices (LDEVs). Also, logical volumes to be accessed by open-systems hosts are sometimes referred to as logical units or LUs.

Contact your sales representative for license information.

LUN Manager

LUN Manager, which is operated through Storage Navigator Modular 2, manages access paths between hosts and logical units, for each port in your array. Depending on your array model, LUN Manager can manage either fibre channel- (FC) or iSCSI-based host connections.

- For Fibre Channel, LUN Manager lets you set the option (host connection mode), Logical Unit (LU), and WWN (World Wide Name) parameters for each connected host so you can connect multiple hosts to the same port.
- For iSCSI, when setting up host connections in LUN Manager, for each host you specify the settings for Host Connection Mode and iSCSI Name. Each host can access a logical unit simulating a dedicated port to the host even if that host shares the port with other hosts.

With LUN Manager, each host can access a logical unit as if it was a dedicated port to the host, even if that host shares the port with other hosts.

NOTE: Although additional hosts can be connected to one port, this increases traffic on the port. When using LUN Manager, design the system configuration so that you evenly distribute traffic at the port, controller, and disk drive.

Figure 1-5 shows a port being shared by multiple host systems (setting access paths between hosts and logical units for Fibre Channel.)

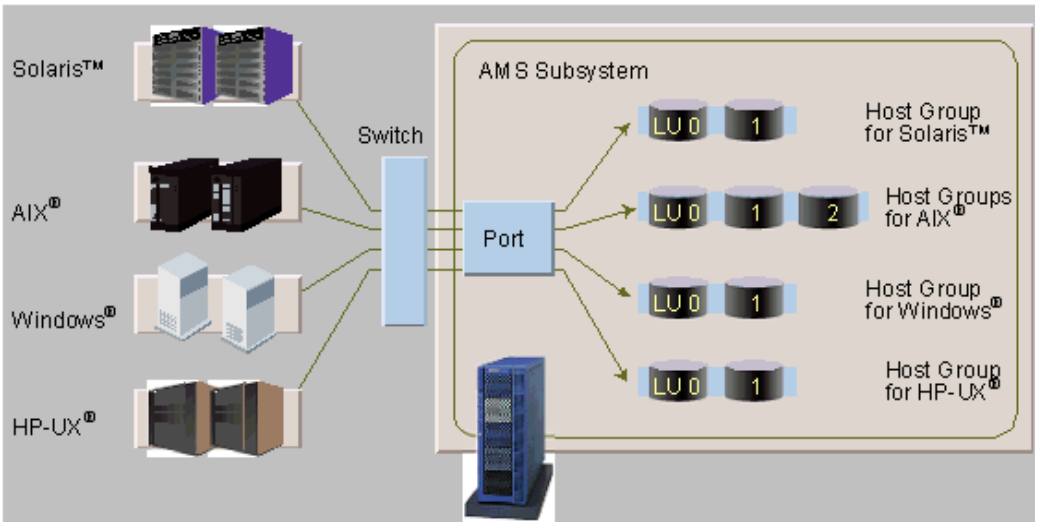

Figure 1-5: LUN Manager Fibre Channel Example

Fibre Channel features

Table 1-2 lists the LUN Manager features for fibre channel.

Table 1-2: LUN Manager Features for Fibre Channel

Feature	Description
Prevents illegal access from other hosts.	Logical units are grouped, and each group is registered in a port. LUN Manager specifies which host may access which logical unit, by assigning hosts and logical units to each host group.
The host connection mode can be set for each host connected.	The host connection mode can be set for each host group.
The logical unit mapping can be set for each connected host.	Logical unit numbers (H-LUN) recognized by a host can be assigned to each host group. Hosts that require LU0 can be connected to the same port.

iSCSI features

Table 1-3 lists the LUN Manager features for iSCSI.

Table 1-3: LUN Manager Features for iSCSI

Feature	Description
Connecting Hosts to Array Ports	You can connect more than one host to an array port. On an array with two ports, port A can connect to a Windows® and Solaris™ host, and port B can connect to another Windows®, AIX®, or HP-UX® host. When setting up host connections, specify the Host Connection Mode and iSCSI Name for each host. Each host can access a logical unit simulating a dedicated port to the host even if that host shares the port with other hosts
Mapping Logical Units to Hosts	You can map or assign your array logical units to the hosts on your network. You can share or restrict logical unit access among hosts.
Network Security	You can enable or disable Challenge Handshake Authentication Protocol (CHAP), a security protocol that requires users to enter a secret for access.

Figure 1-6 shows how multiple hosts on an iSCSI network can share ports on an array. Note that the volumes are grouped into targets. Each host is associated with one target that can contain one or more volumes. Hosts can share targets so that the hosts have access to the same volumes.

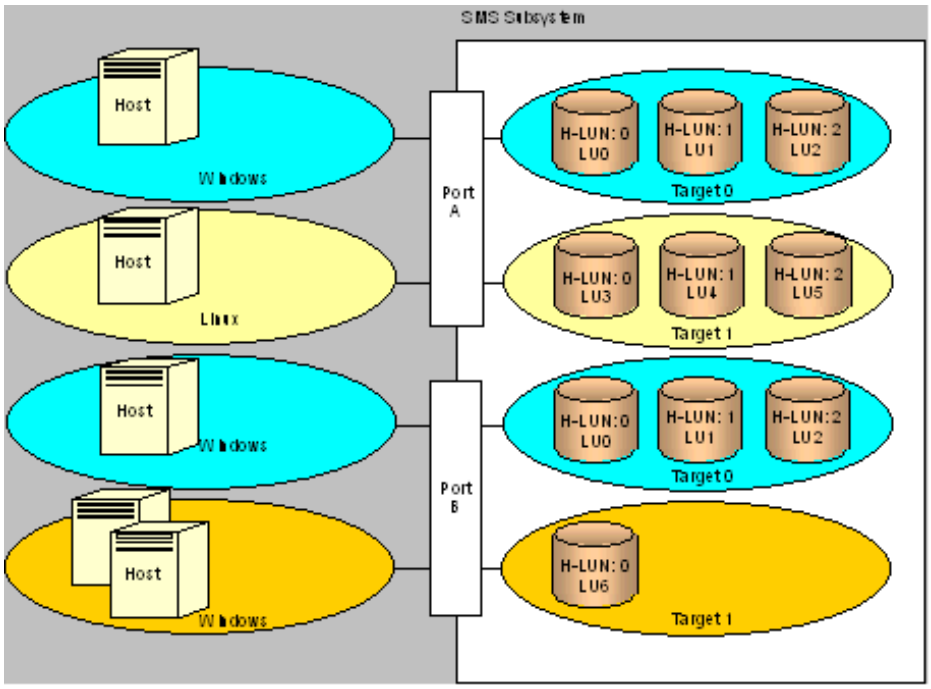

Figure 1-6: Targets (Volume Groups) Assigned to Hosts

iSCSI protocol

iSCSI is a network protocol standard that allows the SCSI protocol to be used over Transmission Control Protocol/Internet Protocol (TCP/IP) networks. Technically speaking, iSCSI is a transport-layer protocol in the SCSI-3 specifications framework. Other protocols in the transport layer include SCSI Parallel Interface (SPI), Serial Attached SCSI (SAS), and fibre channel.

For more information about iSCSI, refer to the *Hitachi iSCSI Resource and Planning Guide (MK-97DF8105).*

iSCSI network configuration

iSCSI makes it possible to build an IP-SAN by connecting hosts and arrays at a low cost. However, iSCSI increases the input/output (I/O) workload of the network and array. When using iSCSI, configure the network so that the workload among the network, port, controller, and drive is properly distributed.

Even though the Local Area Network (LAN) switches and Network Interface Cards (NICs) are the same, there are differences when you use iSCSI, particularly regarding the LAN connection. Note the following:

- iSCSI uses most of the Ethernet bandwidth, and can degrade the iSCSI traffic and LAN performance. Therefore, separate the iSCSI IP-SAN and the office LAN.
- Host I/O load affects iSCSI response time. The more I/O traffic, the lower the iSCSI performance.
- You must have a failover path between the host and the iSCSI, to update the firmware without stopping the system.

Performance Monitor

Performance Monitor obtains disk array performance and resource information (<h9Hyperlink9>Chapter 1, Performance Monitor Example). 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.

The resource use, such as loads on the disks and ports, can be measured and displayed with line graphs. The graphs appear after the data is collected and after you make a choice from that data. This data can be output to a comma-separated value (CSV) file.

When an issue such as a slow response occurs in a host, the system administrator can quickly determine the source of the problem by using Performance Monitor. Figure 1-7 shows a Performance Monitor example.

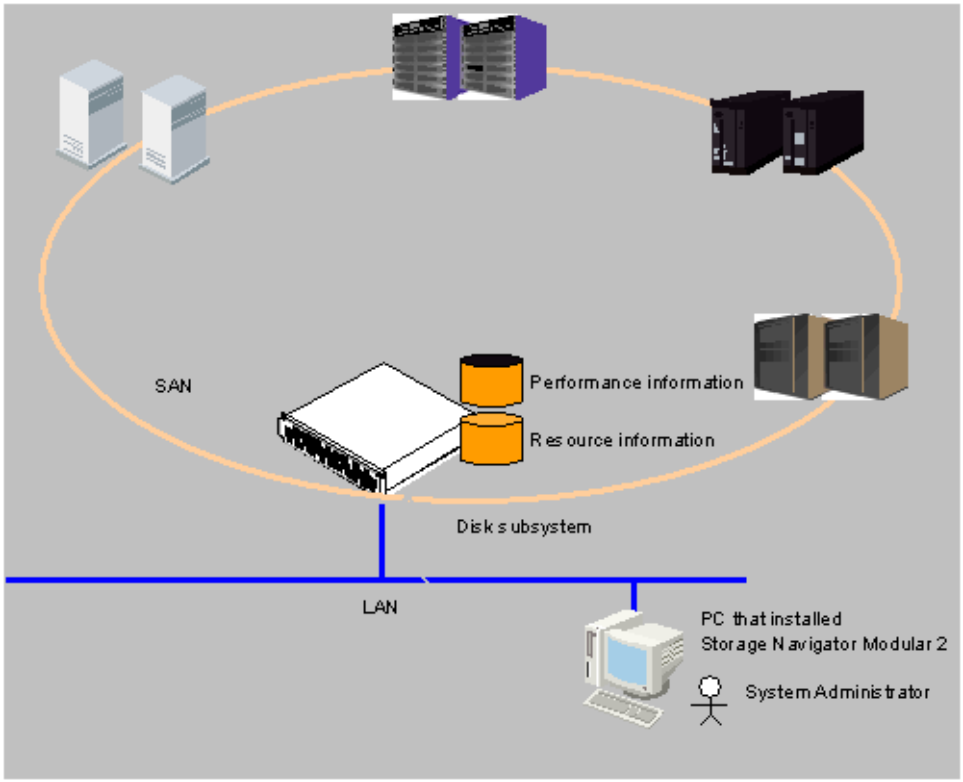

Figure 1-7: Performance Monitor Example

SNMP Agent Support

The SNMP Agent Support feature reports failures in the workstation for network monitoring to a properly configured SNMP manager application.

Command operating status (for example, number of commands received, number of cache hits, etc.) of the disk array is reported. This information can be used for performance tuning, since the command operating status, depending on the type of access from the host, can be referred to this function.

To use the SNMP Agent Support, you must have a LAN facility and a workstation in which the SNMP manager application (hereafter called SNMP manager) is installed.

Trap-issuing processing

A trap-issuing event in the disk array causes the array to issue a trap to the SNMP manager asynchronously and report the error once (see Figure 1-8).

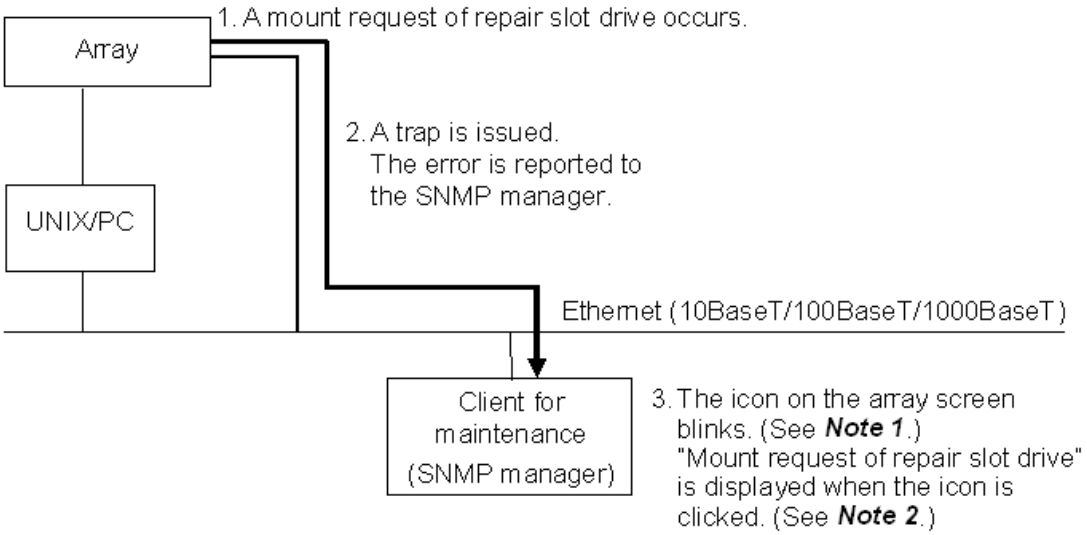

Figure 1-8: Drive Blockade and Trap Issue Example

The trap indicates an error and the relevant regressed site only. The trap does not identify its exact location, for example, the drive number.

NOTES: The action taken at the time the trap is received, and the display operation and specification of the trap codes depends on the specification of the SNMP manager. The display operation and the display specification of the trap codes depend on the specification of the SNMP manager used.

Request processing

This process returns the value of the Message Information Block (MIB) that the SNMP manager requested (Figure 1-9).

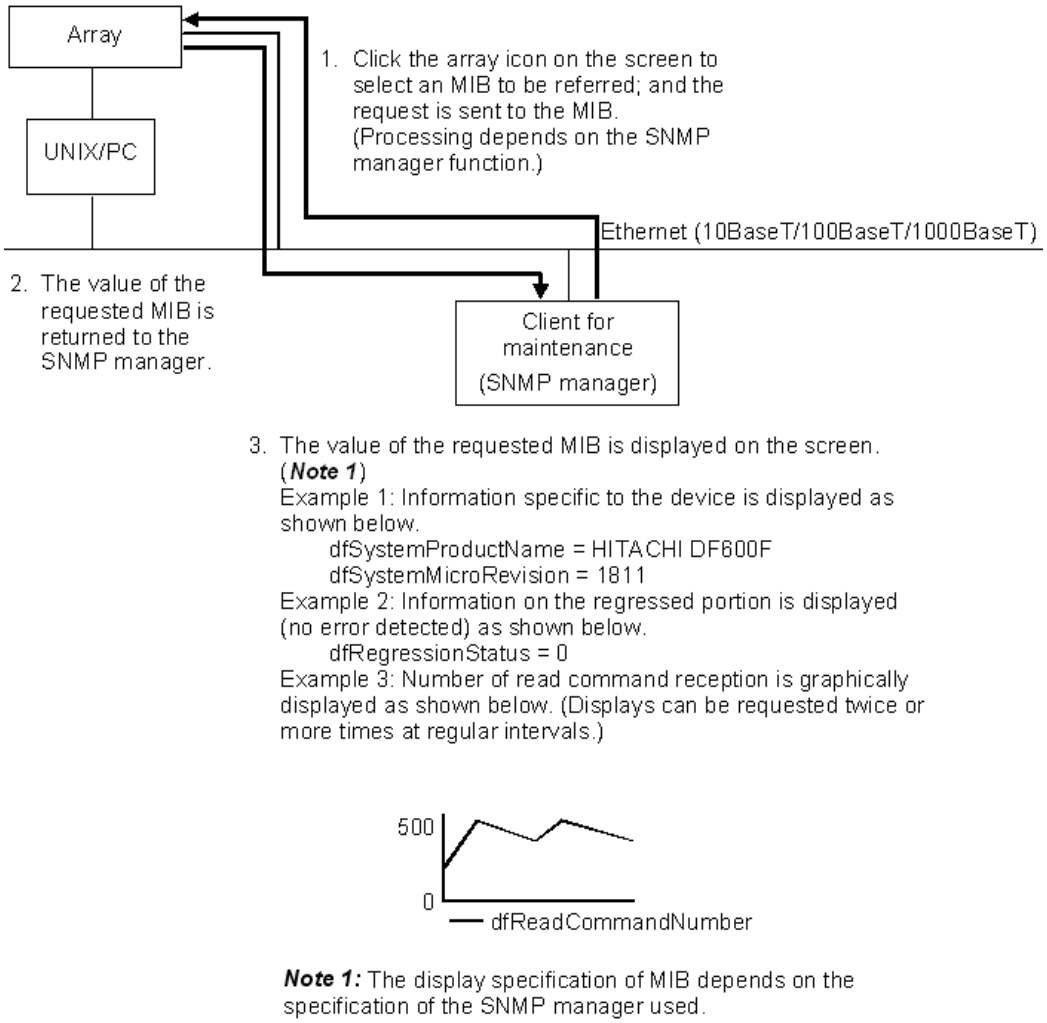

Figure 1-9: Request Processing Example

The regressed portion does not indicate the exact error location (for example, drive number). If the interval set for obtaining the MIB information is too short, the host command processing performance of the array can be affected.

The disk array cannot send/receive SNMP messages larger than 484 bytes, and in these cases, the message "tooBig" appears. Refer to Figure 1-10 on page 1-14 for more details on message specifications.

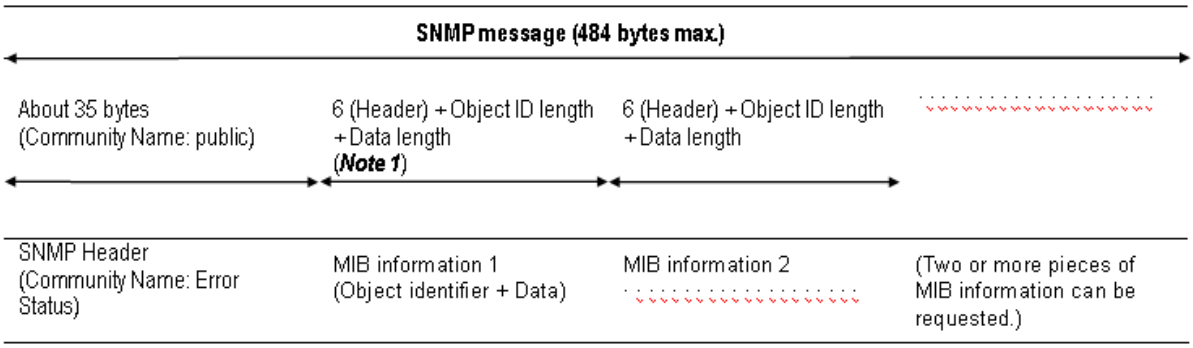

Figure 1-10: SNMP Message Management

NOTE: The action that occurs when a trap is received depends on the specifications of the SNMP manager being used. MIB information 1 becomes 6+8+10 = 24 bytes long. Be aware that header lengths vary with the data length, as shown in <h9Hyperlink9>Table 1-4.

Table 1-4: SNMP Data Length vs. Header Size

Data Length (Bytes)	Header Size (Bytes)
0 to 115	6
116 to 127	7
128 to 242	8
243 to 255	9
256+	10

Modular Volume Migration

Modular Volume Migration copies logical unit data to a logical unit in the other RAID group within the array. The host can continue the Read/Write operation even though data has migrated to another logical unit.

Figure 1-11 shows the status of the data migrated by Modular Volume Migration.

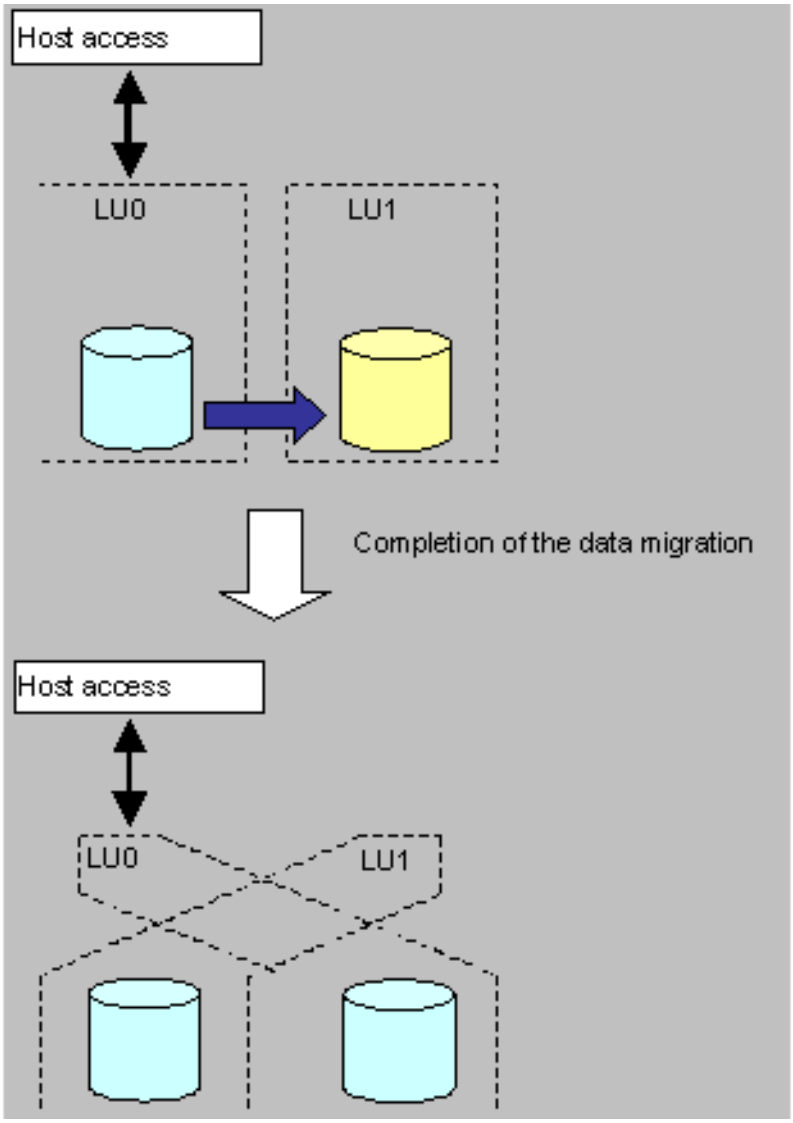

Figure 1-11: Modular Volume Migration Example

Modular Volume Migration requires a volume pair with a Primary Volume (P-VOL) which is the migration source of the data, and a Secondary Volume (S-VOL) which is the migration destination of the data, and a reserved logical unit. A typical Modular Volume Migration configuration appears in Figure 1-12.

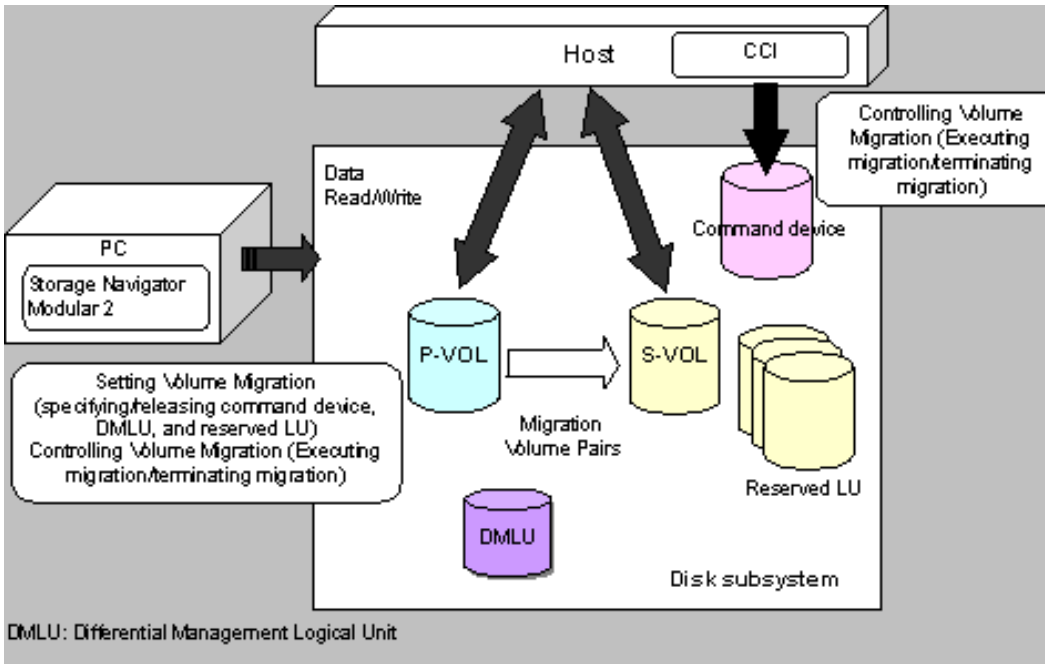

Figure 1-12: Modular Volume Migration Components

Usage guidelines

Before installing, uninstalling, enabling, or disabling your features, review the guidelines 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 **0850** or higher for the AMS 2100, 2300, or 2500 systems.
- Navigator 2 version 6.50 or higher for your computer.
- When using the command control interface (CCI), version 01-23-03/08 or higher is required for your computer.

Requirements

- Storage feature license key(s).
- 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 logical unit size must be 33 MB or more.
- Maximum of two Differential Management Logical Units (DMLUs). The DMLU size must be 10 GB or more. It is recommended that two DMLUs are set for different RAID groups.

The primary volume (P-VOL) size must equal the secondary volume (S-VOL) logical unit size.

Requirements for installing and enabling features

Before you install or enable your features, read the following notes.

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

Audit Logging

- 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

Cache Partition Manager

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 LUs 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.

For more information, refer to the following documents:

- *Hitachi AMS 2000 Family TrueCopy Extended Distance User's Guide* (MK-97DF8054*)*
- *Hitachi AMS Copy-on-Write SnapShot User's Guide (*MK-97DF8124)

Modular Volume Migration

- To install and enable the Modular Volume Migration license, follow the procedure provided in Installing storage features on page 2-5, and select the license **LU-MIGRATION**.

Requirements for uninstalling and disabling features

Before you uninstall or disable your features, read the following notes.

- 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.

Account Authentication

- 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

- Sub-partitions, except for the master partition, must be deleted.
- The capacity of the master partition must be the default partition size (see Table 5-1 on page 5-2).

Data Retention

- You must return the logical unit attributes the Read/Write setting.

LUN Manager

- The host group and target security on every port must be disabled.

Modular Volume Migration

- All the volume migration pairs must be released, including those with a Completed or Error status.
- You cannot have logical units registered as reserved.

SNMP Agent

- 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.

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 logical unit cache memory or partition.

Advanced Settings Java Applet

Users who access AMS arrays from Navigator 2 have an additional array tree item called **Advanced Settings** located under **[Array Name] > Settings**.

When you click **Advanced Settings**, a Java applet launches that provides additional system functionality. Some functions may require an additional license be installed and enabled.

You must have the proper Java Runtime Environment (JRE) loaded and the Java Console set properly on your system to view the Advanced Settings window. The requirements are as follows:

- JRE version required: **v1.6.0**
- Enter **-Xmx216m** to the Java Runtime Parameters field.

CAUTION! The Java applet window may time out after 20 minutes due to an automatic logout function. If this occurs, close the Web browser, stop the SNM2 Server and restart. Launch the SNM2 GUI and return to the array you want to manage.

For Linux and Solaris, perform the following steps:

1. Run the Java Control Panel from an XWindows terminal executing the <JRE installed directory> /bin/jcontrol.
2. Click **View** of the upper position in the Java tab.
3. Enter **-Xmx216m** to the Java Runtime Parameters field.
4. Click **OK**.
5. Click **OK** in the Java tab.

2

Installing and enabling storage features

This chapter describes how to install, enable, disable, and uninstall storage features.

This chapter covers the following topics:

- Preinstallation information
- Installing storage features
- Enabling storage features
- Disabling storage features
- Uninstalling storage features

Preinstallation information

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 **0850** or higher for the AMS 2100, 2300, or 2500 systems.
- Navigator 2 version 6.50 or higher for your computer.
- When using the command control interface (CCI), version 01-23-03/08 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(s).
- 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 logical unit size must be 33 MB or more.
- Maximum of two Differential Management Logical Units (DMLUs). The DMLU size must be 10 GB or more. It is recommended that two DMLUs are set for different RAID groups.

The primary volume (P-VOL) size must equal the secondary volume (S-VOL) logical unit 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.

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 LUs 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 5-1 on page 5-2).

For more information, refer to the following documents:

- *Hitachi AMS 2000 Family TrueCopy Extended Distance User's Guide* (MK-97DF8054*)*
- *Hitachi AMS Copy-on-Write SnapShot User's Guide (*MK-97DF8124)

Data Retention requirements

- If you uninstall or disable this storage feature, you must return the logical unit 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.

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 2-5, 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 logical units registered as reserved.

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 logical unit cache memory or partition.

Installing storage features

These instructions describe how to install your features for each array.

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 **Install License**.
3. In the Install License screen, 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.
4. Click **OK**.
5. Follow the on-screen instructions.

Enabling storage features

To enable 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. 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. On the De-Install License screen, enter the code in the **Key Code** box, and then click **OK**.
5. Follow the on-screen instructions.
6. Log out from the disk array.

Uninstalling of the feature is now complete.

3

Account Authentication

This chapter describes Account Authentication. The topics covered in this chapter are:

- Account Authentication overview
- Account Authentication procedures
- Troubleshooting

Account Authentication overview

The Account Authentication feature is pre-installed 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.Table 3-1 details the settings in the built-in default account.

Table 3-1: 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 20 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).

We recommend 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 3-7.

Overview of Account Authentication

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 informatoin) 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).

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 sotrage system, even in the case of users who have registered their accounts. Figure 3-1 provides an outline of the Account Authentication process.

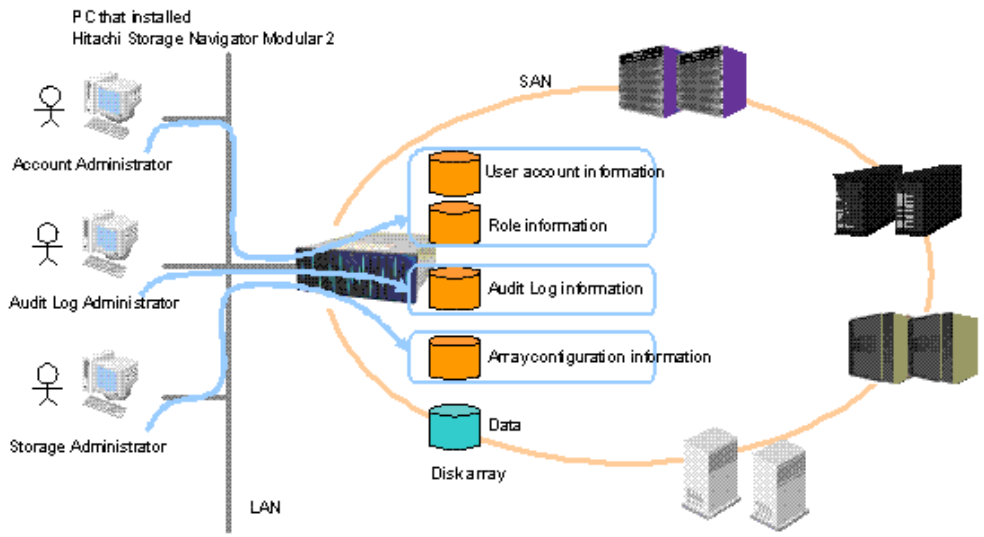

Figure 3-1: Account Authentication Outline

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 3-2: 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 3-5.
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 storage system, create a public account as the normally used account, and assign the necessary role to it. See Table 3-3 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 3-3: 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). To place restrictions, assign a role to an account.

Table 3-4: 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 an LU and to delete an account).

Table 3-5: 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 3-5: Resources (Continued)

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 Table 3-6. 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 3-6: Role and Resource Group Relationships

Resource Group Name (Repository) / Role Name	Role Definition	Key	Storage Resource	Account	Role Mapping	Account Setting	Audit Log Setting	Audit Log
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 window.

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 3-7: Session Types

Type	Operation	Maximum Number of Session IDs
Modify mode	View and modify (setting) array operations.	3 (0nly one log in for each role)
View mode	Only view the array setting information.	256

The built-in account 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.

NOTE: The built-in account is the root account and has all privileges. The Account Administrator (View and Modify or View Only) role can display account information, including which users have modification privileges.

Warning banners

The warning banner is a function that allows users with **User Management** privileges to post a pre-login message for all users who log in to Navigator 2. This function is available from the main Navigator 2 Explorer tree, **Administration > Security >Warning Banner**.

The basic specifications for banners are as follows:

- Maximum message length: 1000 characters
- Usable characters: **0 to 9, A to Z, a to z, " ! # $ % & ' ( ) * + - . / : ; < = > ? @ [\ ] ^ _ ` { | } ~**

The following table describes attributes added to the access request.

Table 1-4. Attribute Specifications Added to Access Request

Item	Specification
User-Name	Entered user ID
User-Password	Encrypted password in the specified format that is used when selecting PAP as the authentication protocol.
CHAP-Password	MD5 hash value processed in the specified format used when selecting CHAP for authentication protocol.
NAS-Identifier	Fixed value: Hitachi SMS/AMS RADIUS Client
NAS-IP-Address	IPv4 address of the array management port.
NAS-IPv6-Address	IPv6 address of the array management port.

This section describes how to modify Account Authentication accounts and log in operations.

CAUTION! The Account Authentication license is pre-installed and enabled. You will be required to enter the built-in (default) username and password information when you first log in to the array.

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 generatoin method which hs the 128-bit strength.

Table 3-1: 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 intialized. As necessary, check the set following information in advance, and set it again after changing the mode:

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

Account Authentication procedures

The following sections describe Account Authentication procedures.

Initial settings

1. Verify that you have the environments and requirements for Account Authentication (see Preinstallation information on page 2-2).
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 3-4).
5. Register an account (see Adding accounts on page 3-11).
6. Registering an account for the service personnel (see Adding accounts on page 3-11).

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 3-14.
- Delete accounts — see Deleting accounts on page 3-15.

Displaying accounts

To display accounts, you must have an Account Administrator (View and Modify or View Only) role. See Table 3-3 on page 3-4 for accounts types and permissions that may be created.

1. Start Navigator 2 and log in. The Arrays window 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 3-2 on page 3-11.

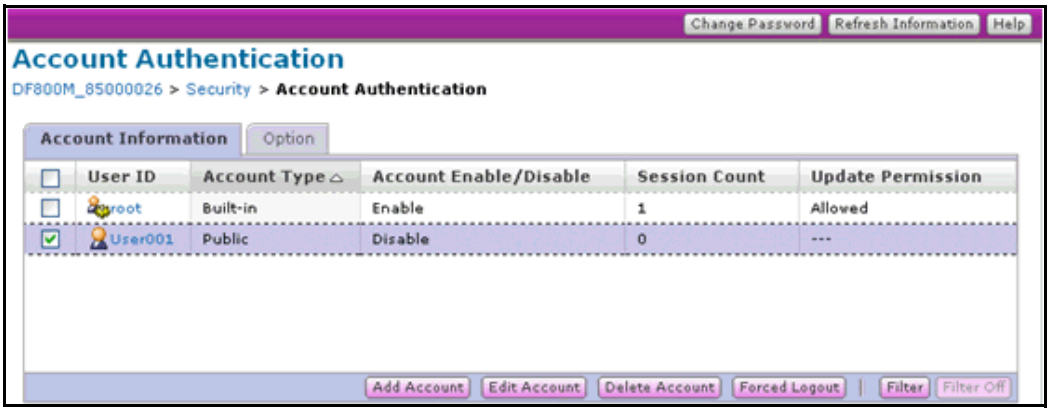

Figure 3-2: Account Information Screen

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, follow these steps.

1. Start Navigator 2 and log in. The Arrays window 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**. The Add Account screen is displayed. See Figure 3-3 on page 3-12.

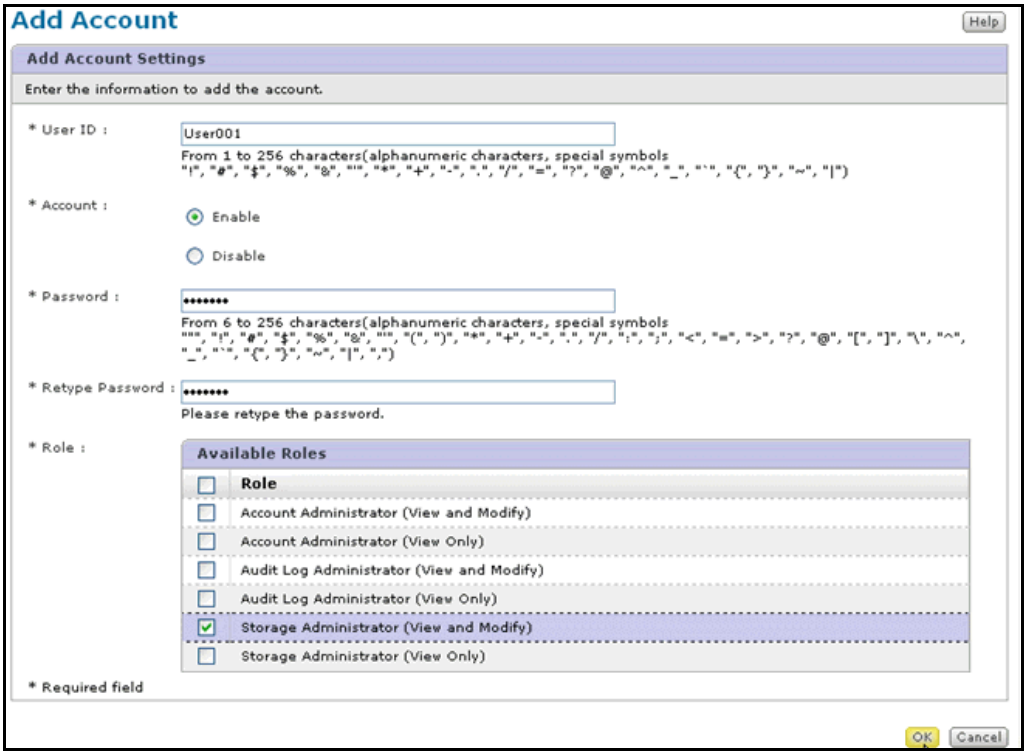

Figure 3-3: Add Account Window

6. Type a new username in the **User ID** field.
7. Select **Enable** in Account to enable the account.

Figure 3-4: User Registration Wizard - Change Root Password

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 as shown in Figure 3-5.

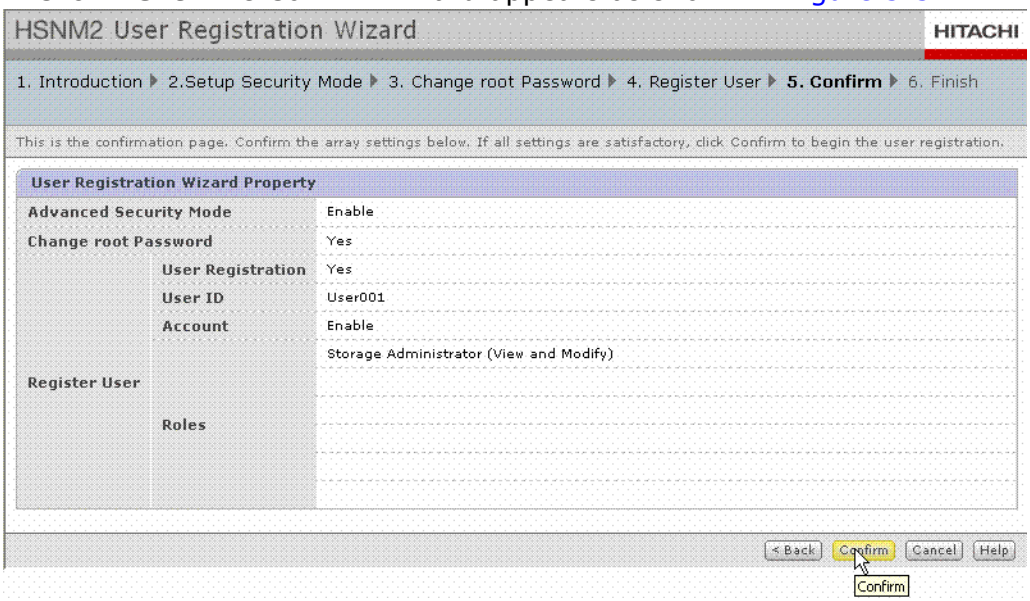

Figure 3-5: HSNM2 User Registration Wizard - Confirm

Changing the Advanced Security Mode

To add accounts, follow these steps.

1. Start Navigator 2 and log in. The Arrays window 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 **Change Security Mode**. The Change Security Mode screen displays as shown in Figure 3-6.

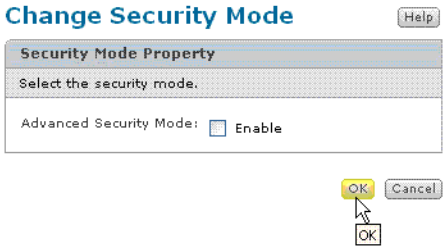

Figure 3-6: Change Security Mode dialog box

6. 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.

7. Click **OK**.
8. Observe any messages that display and click **Confirm** to continue. An example of a system message displays in Figure 3-7

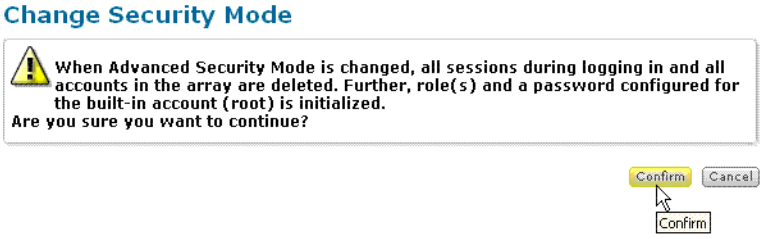

Figure 3-7: Change Security Mode system message

9. 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:

1. Start Navigator 2 and log in. The Arrays window appears
2. Select the appropriate array and click **Show & Configure Array**.
3. Log in as an Account Administrator (View and Modify).
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 you want to modify, and then click **Edit Account**. The Edit Account screen appears. The **Edit Account** window appears, as shown in Figure 3-8.

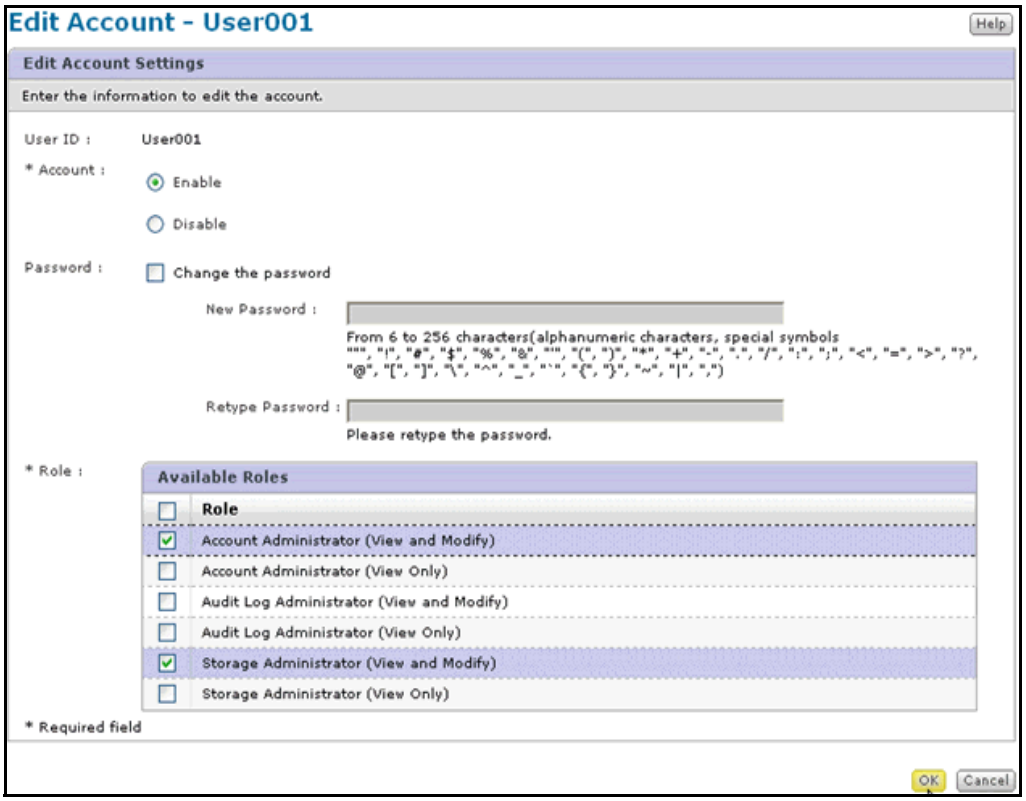

Figure 3-8: Edit Account Window

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, follow these steps:

1. Start Navigator 2 and log in. The Arrays window 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**.
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, follow these steps.

1. Start Navigator 2 and log in. The Arrays window 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.
6. Click **Change session timeout time**. The Change session timeout time screen appears as shown in Figure 3-9.

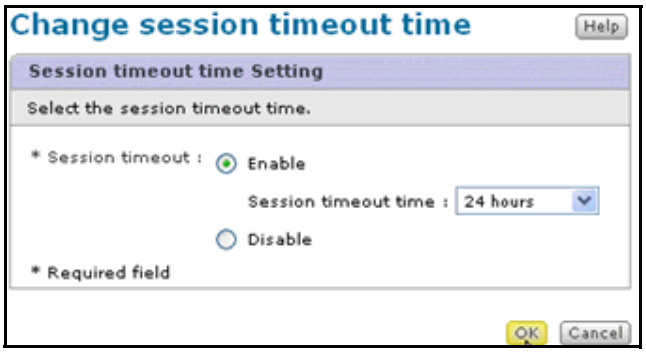

Figure 3-9: Change Session Timeout Window

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 a specific account:

1. Start Navigator 2 and log in. The Arrays window 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**.
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**.

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.

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.

4

Audit Logging

This chapter describes Audit Logging. The topics covered in this chapter are:

- Audit Logging overview
- Audit Logging procedures

Audit Logging overview

Table 4-1 describes specifications for Audit Logging.

Table 4-1: 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 *Hitachi Storage Navigator Modular 2 Command Line Interface (CLI) User's Guide* (MK-97DF8089).
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 logical unit). • 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.
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.

These events are not logged:

- Partial blockade of the array and recovery
- Reference/setting made from the array Web function
- Success/failure of the device authentication (iSCSI CHAP)

NOTE: CCI logs are not generated by Audit Logging because CCI outputs the log individually.

The audit log for an event has the format shown in Figure 4-1.

Figure 4-1: Audit Log Format

The output of an audit log is shown in Figure 4-2. Items are separated by commas. When there is no item to be output, nothing is output.

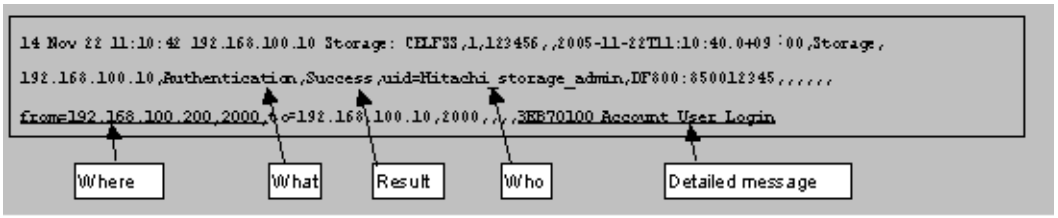

Figure 4-2: Log Example

For more details about Audit log format, see the *Hitachi Storage Navigator Modular 2 Command Line Interface (CLI) User's Guide* (MK-97DF8089).

Audit Logging procedures

The following sections describe the Audit Logging procedures.

Initial settings

1. Verify that you have the environments and requirements for Audit Logging (see Preinstallation information on page 2-2).
2. Set the Syslog Server (see Table 4-1 on page 4-2).

Optional operations

1. Export the internal logged data.
2. Initialize the internal logged data (see Initializing logs on page 4-6).

Enabling Audit Log data transfers

To transfer data to the Syslog server, these steps.

1. Start Navigator 2 and log in. The Arrays window 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 window is displayed.
5. Click **Configure Audit Log**. The Configure Audit Log window is displayed. See Figure 4-3.

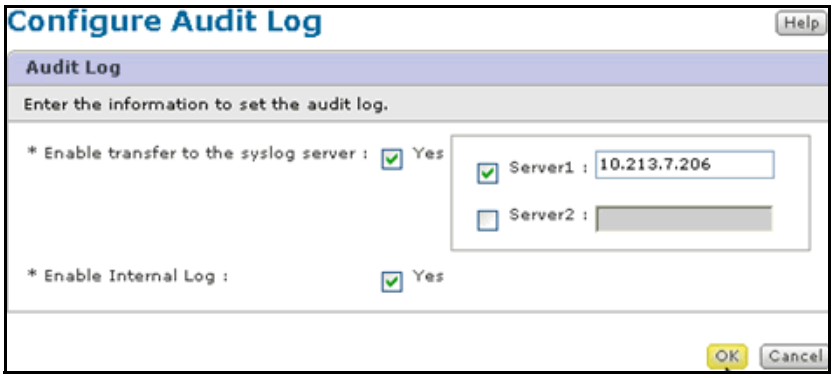

Figure 4-3: Audit Logging Window-Audit Log Tab

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 *Storage Navigator Modular 2 Command Line Interface (CLI) User's Guide* (MK-97DF8089) 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 window 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 window is displayed.
5. Click **Show Internal Log**. The Show Internal Log confirmation screen appears as shown in Figure 4-4.

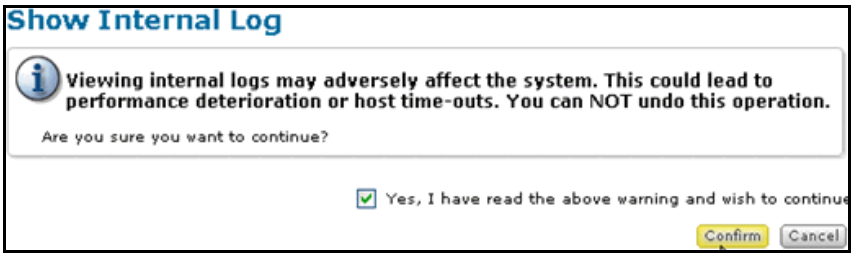

Figure 4-4: Show Internal Log Confirmation

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 4-5).

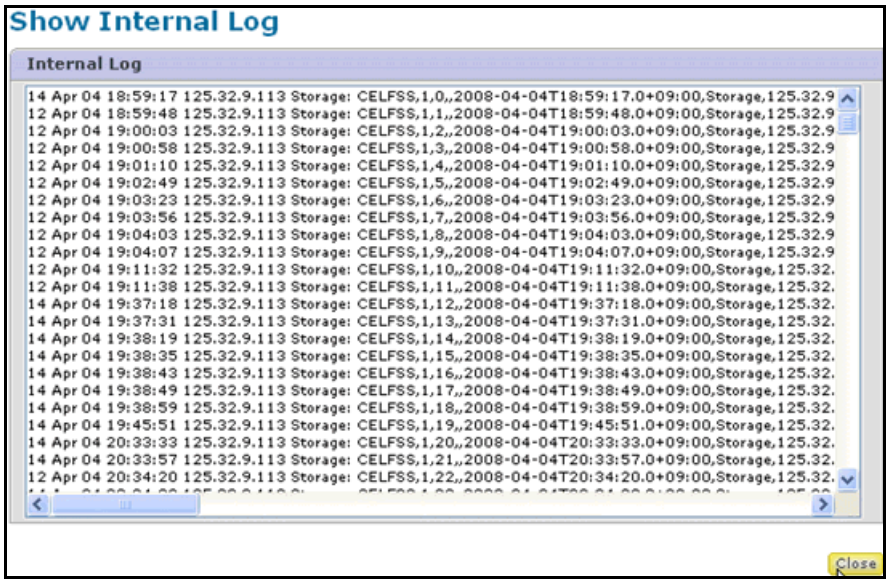

Figure 4-5: Internal Log Window

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 *Storage Navigator Modular 2 Command Line Interface (CLI) User's Guide* (MK-97DF8089).

To initialize logs:

1. Start Navigator 2 and log in. The Arrays window 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 window is displayed (see Figure 4-6).

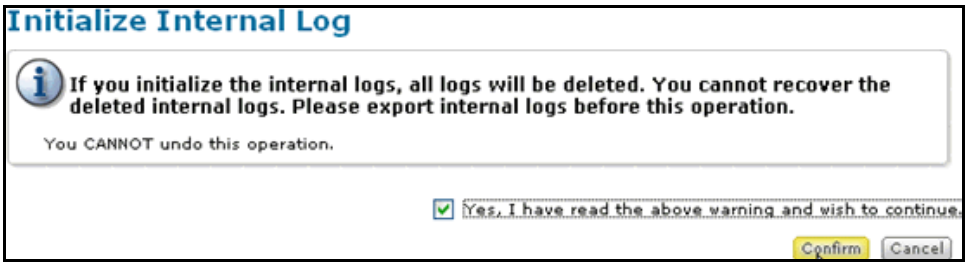

Figure 4-6: Initialize Internal Log Window

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.

5

Cache Partition Manager

This chapter describes Cache Partition Manager. The topics covered in this chapter are:

- Cache Partition Manager overview
- Cache Partition Manager settings

Cache Partition Manager overview

Table 5-1: Cache Partition Manager Specifications

Item	Description
Cache memory	• AMS2100: 2, 4 GB/controller • AMS2300: 2, 4, 8 GB/controller • AMS2500: 4, 6, 8, 10, 12, 16 GB/controller
Number of partitions	• AMS2100 (2 GB/controller): 2 to 6 • AMS2100 (4 GB/controller): 2 to 16 • AMS2300 (2 GB/controller): 2 to 6 • AMS2300 (4 GB/controller): 2 to 16 • AMS2300 (8 GB/controller): 2 to 32 • AMS2500 (4 GB/controller): 2 to 6 • AMS2500 (6 GB/controller): 2 to 12 • AMS2500 (8 GB/controller): 2 to 16 • AMS 2500 (10 GB/controller): 2 to 20 • AMS 2500 (12 GB/controller): 2 to 26 • AMS 2500 (16 GB/controller): 2 to 32 The number of partitions including the two master partitions, is shown. The maximum number of partitions varies depending on the capacity allocated to each partition.
Partition capacity	The partition capacity depends on the array and the capacity of the cache memory installed in the controller. For more information, see Cache Partition Manager settings on page 5-12.
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 5-16.
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 Restrictions on page 5-15.
Partition mirroring	Always On (it is always mirrored).

Cache Partition Manager

Table 5-2: Cache Partition Manager Specifications

Item	Description
Cache memory	• AMS2100: 2, 4 GB/controller • AMS2300: 2, 4, 8 GB/controller • AMS2500: 4, 6, 8, 10, 12, 16 GB/controller

Table 5-2: Cache Partition Manager Specifications

Item	Description
Number of partitions	• AMS2100 (2 GB/controller): 2 to 6 • AMS2100 (4 GB/controller): 2 to 16 • AMS2300 (2 GB/controller): 2 to 6 • AMS2300 (4 GB/controller): 2 to 16 • AMS2300 (8 GB/controller): 2 to 32 • AMS2500 (4 GB/controller): 2 to 6 • AMS2500 (6 GB/controller): 2 to 12 • AMS2500 (8 GB/controller): 2 to 16 • AMS 2500 (10 GB/controller): 2 to 20 • AMS 2500 (12 GB/controller): 2 to 26 • AMS 2500 (16 GB/controller): 2 to 32 The number of partitions including the two master partitions, is shown. The maximum number of partitions varies depending on the capacity allocated to each partition.
Partition capacity	The partition capacity depends on the array and the capacity of the cache memory installed in the controller. For more information, see Cache Partition Manager settings on page 5-12.
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 5-16.
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 Restrictions on page 5-15.
Partition mirroring	Always On (it is always mirrored).

Partition Capacity

The partition capacity depends on the user data area of the cache memory and the segment size.

User Data Area

The user data area depends on the array type, 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

Table 5-3 on page 5-4, Table 5-4 on page 5-6, and Table 5-7 on page 5-7 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 type. Sizes of partitions using all 4 KB or 8 KB segments must meet specific criteria for maximum partitions size of small segments.

Maximum partition size of small segments (MB) more than, or equal to, Size of partitions using all 4 KB segments (MB) × 3 + Size of partitions using all 8 KB segments (MB).

[(The size of partitions using all 4 KB segments in MB) + (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.

Table 5-3 on page 5-4 shows the supported capacity in case where SnapShot or TCE disabled in a dual-controller environment. Table 5-4 on page 5-6 shows the supported capacity in case where SnapShot, TCE, or Dynamic Provisioning is used in a dual-controller environment.

[(The size of partitions using all 4 KB segments in MB) + (Size of partitions using all 8 KB segments in MB/3)] has to be less or equal to Maximum partitions size of small segments (in MB) from the table.

Table 5-6 on page 5-6 shows the supported capacity in the case where SnapShot and TCE is disabled and Dynamic Provisioning is enabled in a dual-controller environment. Table 5-3 on page 5-4 shows the supported capacity for a single-controller environment.

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).

Table 5-3: Dual Controller and SnapShot, TCE, or Dynamic Provisioning Disabled Partition Capacity (MB)

AMS Equipment	Cache	User Data Area	Default Partition Size	Default Minimum Size	Default Maximum Size	Partition Capacity for Small Segment
2100	2 GB/CTL	1,520	760	200	1,320	1,120
2100	4 GB/CTL	3,520	1,760	200	3,320	1,520
2300	2 GB/CTL	1,440	720	200	1,240	1,040
2300	4 GB/CTL	3,280	1,640	200	3,080	2,880
2300	8 GB/CTL	7,160	3,580	200	6,960	3,280
2500	4 GB/CTL	2,960	1,480	400	2,560	2,160
2500	6 GB/CTL	4,840	2,420	400	4,440	4,040
2500	8 GB/CTL	6,740	3,370	400	6,340	5,940

Table 5-3: Dual Controller and SnapShot, TCE, or Dynamic Provisioning Disabled Partition Capacity (MB)

AMS Equipment	Cache	User Data Area	Default Partition Size	Default Minimum Size	Default Maximum Size	Partition Capacity for Small Segment
2500	10 GB/CTL	8,620	4,310	400	8,220	6,740
2500	12 GB/CTL	10,500	5,250	400	10,100	6,740
2500	16 GB/CTL	14,420	7,210	400	14,020	6,740

Table 5-4: Dual Controller and SnapShot or TCE Enabled and Dynamic Provisioning Disabled Partition Capacity (MB)

AMS Equipment	Cache	User Data Area	Default Partition Size	Default Minimum Size	Default Maximum Size	Partition Capacity for Small Segment
2100	2 GB/CTL	1,000	500	200	800	600
2100	4 GB/CTL	1,480	740	200	1,280	1,080
2300	2 GB/CTL	920	460	200	720	520
2300	4 GB/CTL	1,220	610	200	1,020	820
2300	8 GB/CTL	3,060	1,530	200	2,860	2,660
2500	4 GB/CTL	1,420	710	400	1,020	620
2500	6 GB/CTL	1,760	880	400	1,360	960
2500	8 GB/CTL	2,640	1,320	400	2,240	1,840
2500	10 GB/CTL	3,500	1,750	400	3,100	2,700
2500	12 GB/CTL	4,360	2,180	400	3,960	3,560
2500	16 GB/CTL	6,220	3,110	400	5,820	5,420

Table 5-5: Dual Controller and SnapShot or TCE and Dynamic Provisioning Enabled Partition Capacity (MB)

AMS Equipment	Cache	User Data Area	Default Partition Size	Default Minimum Size	Default Maximum Size	Partition Capacity for Small Segment
2100	2 GB/CTL	940	470	200	740	540
2100	4 GB/CTL	1,400	700	200	1,200	1,000
2300	2 GB/CTL	780	390	200	580	380
2300	4 GB/CTL	1,080	540	200	880	680
2300	8 GB/CTL	2,920	1,460	200	2,720	2,520
2500	4 GB/CTL	1,120	560	400	720	320
2500	6 GB/CTL	1,480	740	400	1,080	680
2500	8 GB/CTL	2,340	1,170	400	1,940	1,540
2500	10 GB/CTL	3,200	1,600	400	2,800	2,400
2500	12 GB/CTL	4,060	2,030	400	3,660	3,260
2500	16 GB/CTL	5,940	2,970	400	5,540	5,140

Table 5-6: Dual Controller and SnapShot and TCE Disabled and Dynamic Provisioning Enabled Partition Capacity (MB)

AMS Equipment	Cache	User Data Area	Default Partition Size	Default Minimum Size	Default Maximum Size	Partition Capacity for Small Segment
2100	2 GB/CTL	1,440	720	200	1,252	1,052
2100	4 GB/CTL	3,460	1,730	200	3,264	1,526

Table 5-6: Dual Controller and SnapShot and TCE Disabled and Dynamic Provisioning Enabled Partition Capacity (MB)

AMS Equipment	Cache	User Data Area	Default Partition Size	Default Minimum Size	Default Maximum Size	Partition Capacity for Small Segment
2300	2 GB/CTL	1,300	650	200	1,240	1,040
2300	4 GB/CTL	3,120	1,560	200	3,260	1,520
2300	8 GB/CTL	7,020	3,510	200	6,820	3,280
2500	4 GB/CTL	2,660	1,330	400	2,260	1,860
2500	6 GB/CTL	4,560	2,280	400	4,160	3,760
2500	8 GB/CTL	6,440	3,220	400	6,040	5,640
2500	10 GB/CTL	8,320	4,160	400	7,920	6,740
2500	12 GB/CTL	10,216	5,100	400	9,800	6,740
2500	16 GB/CTL	14,120	7,060	400	13,720	6,740

Table 5-7: Single Controller Partition Capacity (MB)

AMS Equipment	Cache	User Data Area	Default Partition Size	Default Minimum Size	Default Maximum Size	Partition Capacity for Small Segment
2100	2 GB/CTL	1,520	1,520	400	1,520	1,120
2100	4 GB/CTL	3,530	3,530	400	3,530	1,520
2300	2 GB/CTL	1,440	1,440	400	1,440	1,040
2300	4 GB/CTL	3,280	3,280	400	3,280	2,880
2300	8 GB/CTL	7,170	7,170	400	7,170	3,280

Table 5-8: Supported Partition Capacity (Dual Controller Configuration and SnapShot or TCE is Valid and Dynamic Provisioning is Invalid): Unit MB

AMS Equipment	Cache	User Data Area	Default Partition Size	Default Minimum Size	Default Maximum Size	Partition Capacity for Small Segment
2100	2 GB/CTL	880	440	200	680	480
2100	4 GB/CTL	1,300	650	200	1,100	900
2300	2 GB/CTL	820	410	200	620	420
2300	4 GB/CTL	1,060	530	200	880	660
2300	8 GB/CTL	2,840	1,420	200	2,640	2,440
2500	4 GB/CTL	1,260	620	400	840	440
2500	6 GB/CTL	1,600	800	400	1,200	800
2500	8 GB/CTL	2,360	1,180	400	1,960	1,500
2500	10 GB/CTL	3,200	1,600	400	2,800	2,400
2500	12 GB/CTL	3,820	1,910	400	3,420	3,020

AMS Equipment	Cache	User Data Area	Default Partition Size	Default Minimum Size	Default Maximum Size	Partition Capacity for Small Segment
2500	16 GB/CTL	5,880	2,940	400	5,480	5,080

Table 5-9: Supported Partition Capacity (Dual Controller Configuration and SnapShot or TCE is Valid and Dynamic Provisioning is Valid): Unit MB

AMS Equipment	Cache	User Data Area	Default Partition Size	Default Minimum Size	Default Maximum Size	Partition Capacity for Small Segment
2100	2 GB/CTL	800	400	200	600	400
2100	4 GB/CTL	1,240	620	200	1,040	840
2300	2 GB/CTL	680	340	200	480	280
2300	4 GB/CTL	900	450	200	700	500
2300	8 GB/CTL	2,680	1,340	200	2,480	2,280
2500	4 GB/CTL	940	470	400	540	N/A
2500	6 GB/CTL	1,300	650	400	900	500
2500	8 GB/CTL	2,060	1,030	400	1,660	1,260
2500	10 GB/CTL	2,920	1,460	400	2,520	2,120
2500	12 GB/CTL	3,540	1,770	400	3,140	2,740
2500	16 GB/CTL	5,580	2,790	400	5,180	4,780

NOTE: The sum total of capacities of all the partitions cannot be larger than the capacity of the user data area. The maximum partition capacity is the value that occurs when the capacity of the other partitions are at their minimum size when you configure only master partitions.

[(The size of partitions using all 4 KB segments in MB) + (The size of partitions using all 8 KB segments is shown in MB/3)] has to be less or equal to the maximum partition size of small segments (in MB) from the table.

Table 5-10: Supported Partition Capacity (Dual Controller Configuration and SnapShot and TCE is Invalid and Dynamic Provisioning is Valid, and DP Capacity Mode is Maximum Capacity): Unit MB

AMS Equipment	Cache	User Data Area	Default Partition Size	Default Minimum Size	Default Maximum Size	Partition Capacity for Small Segment
2100	2 GB/CTL	1,180	590	200	980	780
2100	4 GB/CTL	3,140	1,570	200	2,940	1,520
2300	2 GB/CTL	1,020	510	200	820	620
2300	4 GB/CTL	2,780	1,390	200	2,580	2,380
2300	8 GB/CTL	6,600	3,300	200	6,400	3,280
2500	4 GB/CTL	2,200	1,100	400	1,800	1,400

AMS Equipment	Cache	User Data Area	Default Partition Size	Default Minimum Size	Default Maximum Size	Partition Capacity for Small Segment
2500	6 GB/CTL	4,100	2,050	400	3,770	3,300
2500	8 GB/CTL	5,880	2,940	400	5,480	5,080
2500	10 GB/CTL	7,760	3,880	400	7,360	6,740
2500	12 GB/CTL	9,400	4,700	400	9,000	6,740
2500	16 GB/CTL	13,500	6,750	400	13,100	6,740

Table 5-11: Supported Partition Capacity (Single Controller Configuration): Unit MB

AMS Equipment	Cache	User Data Area	Default Partition Size	Default Minimum Size	Default Maximum Size	Partition Capacity for Small Segment
2100	2 GB/CTL	1,380	1,380	400	1,280	980
2100	4 GB/CTL	3,360	3,360	400	3,360	1,520
2300	2 GB/CTL	1,340	1,340	400	1,340	940
2300	4 GB/CTL	3,100	3,100	400	3,100	2,700
2300	8 GB/CTL	6,940	6,940	400	6,940	3,280

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.

Segment And Stripe Size

A logical unit stripe size depends on the segment size of the partition, as shown in Table 5-20 on page 5-15. The default stripe size is 256 KB.

Table 5-12: 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

Restrictions

Table 5-13: Cache Partition Manager 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 logical unit 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, SnapShot, or TCE on page 5-17.
Concurrent use of SnapShot	When SnapShot is enabled, the partition status is initialized. When using SnapShot, see Using ShadowImage, SnapShot, or TCE on page 5-17.
Concurrent use of a unified LU	All the default partitions of the logical unit must be the same partition.
LU Expansion	You cannot expand LUs while making changes with the Cache Partition Manager.
Concurrent use of RAID group Expansion	• You cannot change the Cache Partition Manager configuration for logical units 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 logical unit belongs to, cannot be changed.
Concurrent use of Volume Migration	A logical unit that belongs to a partition cannot carry over. When the migration is completed, the logical unit 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 logical units 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 LU. 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 logical units belong to) increases, the use capacity of the Cache also increases. When a logical unit 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 5-22 on page 5-17.

Table 5-14: Partition Capacity when Changing Segment Size

Segment Size	Partition Capacity	
	AMS2100/2300	**AMS2500**
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 logical unit 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 logical unit partition, invalidate the load balance.

Using ShadowImage, SnapShot, or TCE

The recommended segment size of the ShadowImage S-VOL, SnapShot or TCE Data Pool volume 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 LUs:

- The P-VOL and S-VOL (V-VOL in the case of SnapShot) belong to the master partition (partition 0 or 1).
- The DP-VOL
- The LU 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 LU 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 AMS 2000 Family ShadowImage In-system Replication User's Guide* (MK-97DF8129) or *Hitachi AMS 2000 Family Copy-on-Write SnapShot User's Guide* (MK-97DF8124).

The P-VOL and S-VOL/V-VOL partitions that you want to specify as LUs must be controlled by the same controller. See page 4 17 for more information.

After creating the pair, monitor the partitions for each LU to ensure they are controlled by the same controller.

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 settings

The following sections describe Cache Partition Manager settings.

Initial 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.

1. Verify that you have the environments and requirements for Cache Partition Manager (see Preinstallation information on page 2-2).
2. Change the partition size of the master partition (*Note 1*).
3. Add a sub partition (*Note 1*).
4. Change the partition the logical unit belongs to (*Note 1*).
5. Restart the array (*Note 1*).
6. Create a logical unit (*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 logical unit with the partition you created, determine the partition beforehand. Then, add the logical unit 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.

1. In the master partition, change logical unit 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 LUs.
- Selectable stripe size — Helps increase performance by customizing the disk access size.

NOTE: The following will occur when you restart an array on the remote side of TrueCopy or TCE after setting, deleting, or changing of Cache Partition Manager:

- Both paths of TrueCopy or TCE are blocked. When a path is blocked, an SNMP TRAP is sent if you are using the SNMP Agent Support feature. Be sure to inform the storage administrators prior to restarting the array. The paths are recovered after the array has restarted.
- When the pair status of TrueCopy or TCE is either **Paired** or **Synchronizing**, the array status changes to **Failure**.

We recommend you change the pair status of TrueCopy or TCE to **Split** before making changes using Cache Partition Manager.

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(s) 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.

The Cache Partition Manager runs under the Java applet used for some of the storage features. Please see Advanced Settings Java Applet on page 1-20 for more information on JRE and Java console settings.

After making changes to cache partitions, you must restart the array.

Adding cache partitions

To add cache partitions, follow these steps.

1. Start Navigator 2 and log in. The Arrays window appears
2. Click the appropriate array.
3. Expand the **Settings** list, and click **Advanced Settings**.
4. Click **Open Advanced Settings**. After some minutes, the Array Unit window appears.
5. Click **Cache Partition**. Figure 5-1 appears.

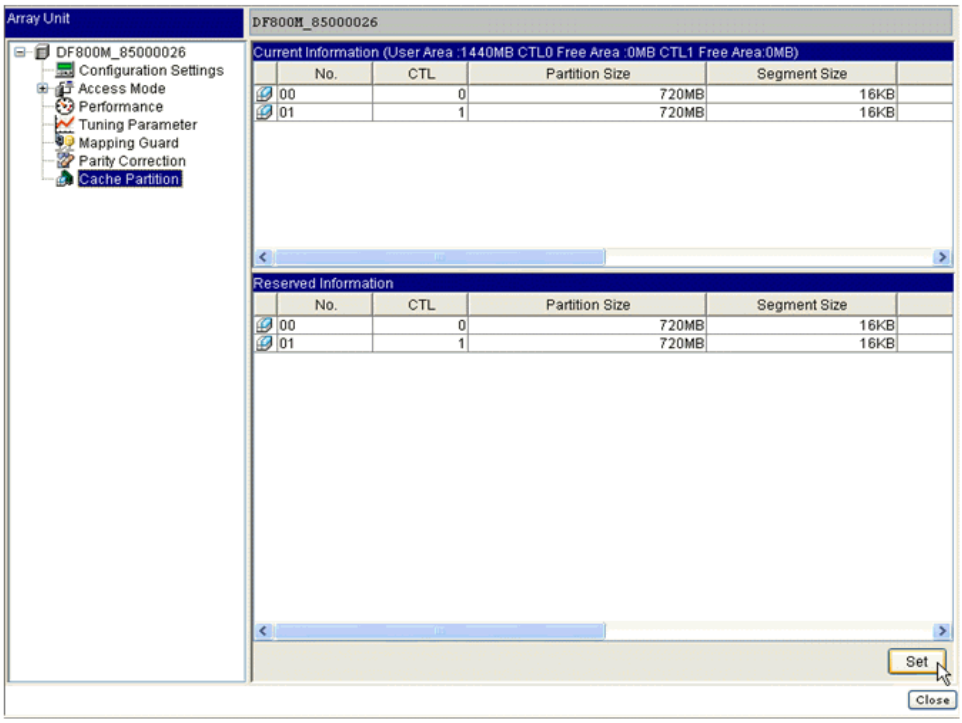

Figure 5-1: Cache Partition

6. Click **Set**. The Cache Partition dialog box appears, as shown in Figure 5-2.

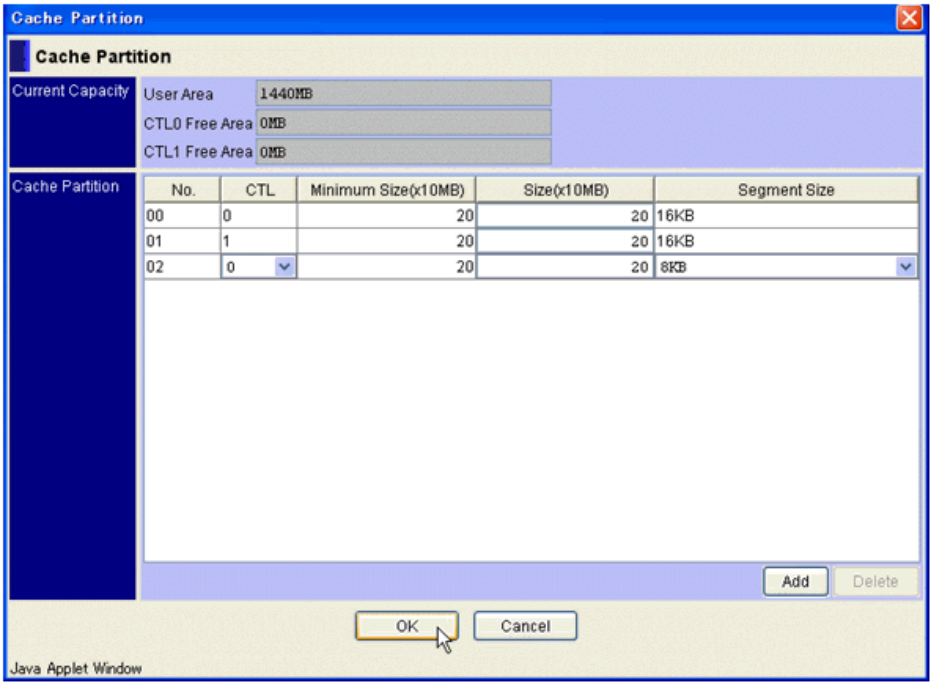

Figure 5-2: Cache Partition Dialog Box

7. Select cache partition **00** or **01**, and click **Add**. Cache partition 02 is added.
8. 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 settings on page 5-12 for more information about supported partition sizes.
9. Click **OK** and follow the on-screen instructions.

Deleting cache partitions

Before deleting a cache partition, move the logical unit 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. Expand the **Settings** list, and click **Advanced Settings**.
4. Click **Open Advanced Settings**. After some minutes, the Array Unit window appears.
5. Click **Cache Partition**. Figure 5-1 on page 5-14 appears.

6. Click **Set**. The Cache Partition dialog box appears, as shown in Figure 5-2 on page 5-15.
7. Select the cache partition number that you are deleting, and click **Delete**.
8. Click **OK** and follow the on-screen instructions.

Assigning cache partitions

If you do not assign a logical unit to a cache partition, it is assigned to the master partition. Also, note that the controllers for the logical unit and pair cache partitions must be different.

To assign cache partitions:

1. Start Navigator 2 and log in.
2. Select the appropriate array.
3. Click **Show & Configure Array**.
4. Under Arrays, click **RAID Groups**.
5. Click the **Logical Units** tab. Figure 5-3 appears.

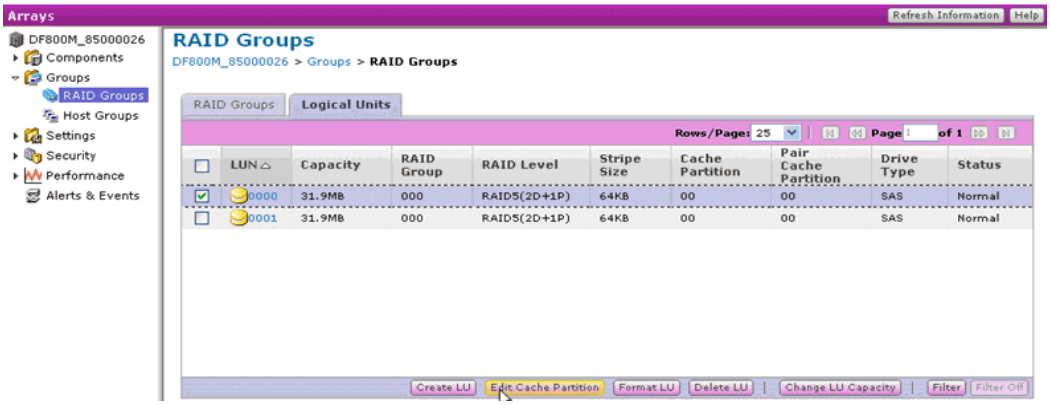

Figure 5-3: Logical Units Tab

6. Select a logical unit from the **LUN** list, and click **Edit Cache Partition**.
7. Select a partition number from the **Cache Partition** drop-down menu, and click **OK**.
8. Follow the on-screen instructions.

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 LU 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 **RAID Groups**.
5. Click the **Logical Units** tab. (See Figure 5-3 on page 5-16)
6. Select a logical unit from the LUN 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. Expand the **Settings** list, and click **Advanced Settings**.
4. Click **Open Advanced Settings**. After some minutes, the Array Unit window is displayed.
5. Click **Cache Partition**. Figure 5-1 on page 5-14 appears.
6. Click **Set**. The Cache Partition dialog box appears, as shown in Figure 5-2 on page 5-15.
7. To change capacity, double-click the **Size (x10MB)** field and make the desired change.
8. To change the segment size, select **segment size** from the drop-down menu to the left of **Segment Size**.
9. Follow the on-screen instructions.

Changing cache partition owner controllers

The controller that processes the I/O of a LUN is referred to as the owner controller.

To change cache partitions owner controllers:

1. Start Navigator 2 and log in. The Arrays window appears

2. Click the appropriate array.
3. Expand the **Settings** list, and click **Advanced Settings**.
4. Click **Open Advanced Settings**. After some minutes, the Array Unit window appears.
5. Click **Cache Partition**. Figure 5-1 on page 5-14 appears.
6. Click **Set**. The Cache Partition dialog box appears, as shown in Figure 5-2 on page 5-15.
7. Select the Cache Partition number and the controller (CTL) number (0 or 1) from the drop-down menu and click **OK**.
8. Follow the on-screen instructions.
9. The **Automatic Pair Cache Partition** Confirmation is displayed.

   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 array.

   Click **OK** to continue. The **Restart Array** message is displayed. You must restart the array to validate the settings, however, you do not have to do it at this time.
10. To restart now, click **OK**. 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 Provisioning under Cache Partition Manager

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 5-15 on page 5-13 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 logical units 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.

Examples where Cache Partition Manager is used without and together with SnapShot, TCE, and Dynamic Provisioning are shown in Figure 5-4 and Figure 5-5 respectively.

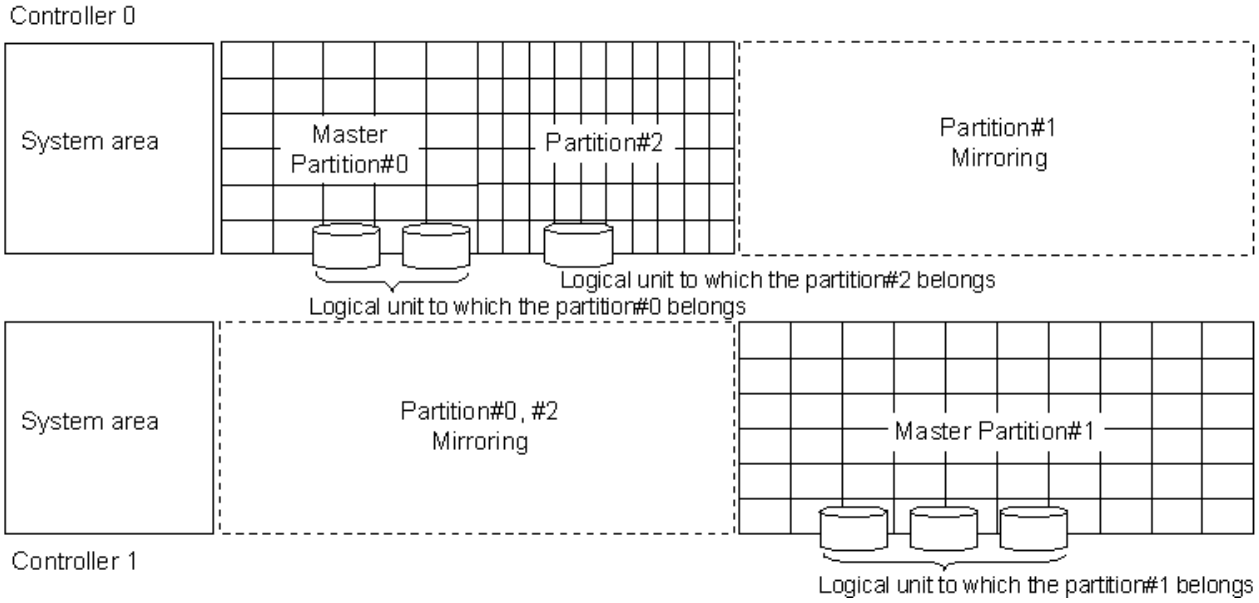

Figure 5-4: Cache Partition Manager without SnapShot, TCE, or Dynamic Provisioning

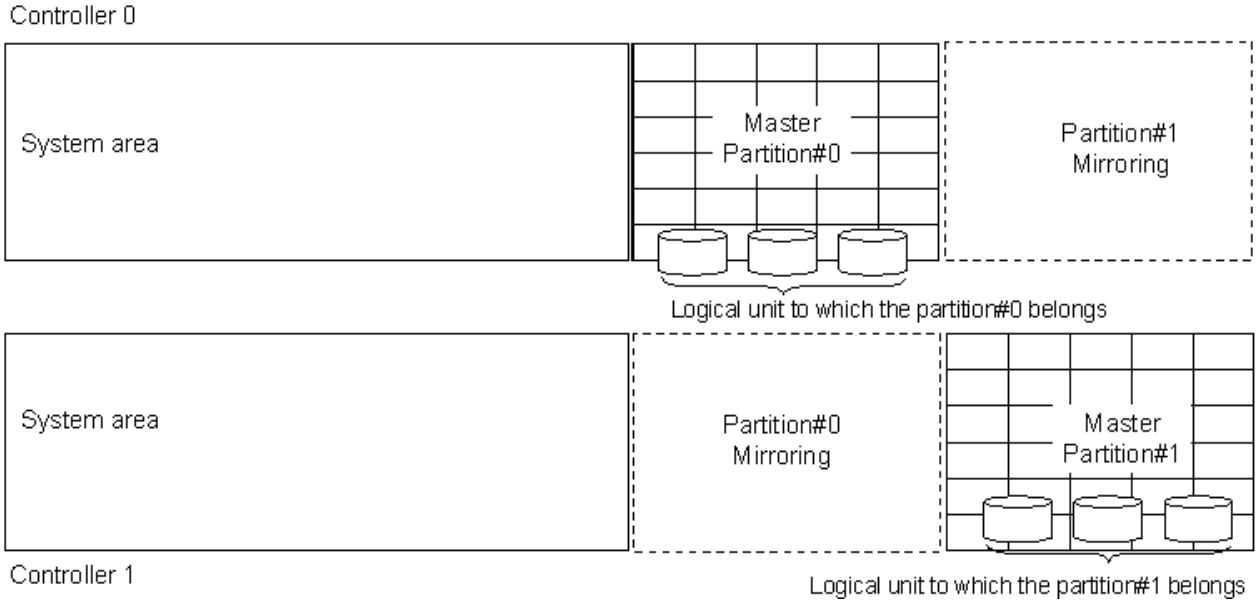

Figure 5-5: Cache Partition Manager with SnapShot, TCE, or Dynamic Provisioning

See Advanced Settings Java Applet on page 1-20 for information about the Java Runtime settings required to access the Cache Partition Manager from the Advanced Settings window.

6

Cache Residency Manager

This chapter describes the Cache Residency Manager.

This chapter covers the following topics:

- Cache Residency Manager overview
- Cache Residency Manager operations

Cache Residency Manager overview

The controller executes read/write commands to the logical unit using the Cache Residency Manager as follows:

- 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.
- Write data from the host is stored in the cache memory, and not written to the disk drives until the array is turned off.
- The cache memory utilizes a battery backup and the write data is duplicated (stored in the cache memory on both controllers).
- 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 logical units, except that the read/write command to the logical unit 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 logical units.

- The command execution may wait for the completion of commands issued to other logical units.
- 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 logical unit.
- The command execution may wait for the completion of processing for internal operation such as data reconstruction, etc.

Termination Conditions

The following conditions terminate Cache Residency Manager. Cache Residency Manager restarts when the failures are corrected.

Table 6-1: 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 logical unit size	Cache uninstallation.
A controller failure	Failure.
The battery alarm occurs	Failure.
A battery backup circuit failure	Failure.

Table 6-1: Cache Residency Manager Termination (Continued)

Condition	Description
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

The following conditions disable Cache Residency Manager.

Table 6-2: 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 logical unit or RAID group is deleted	Caused by the user.
The controller configuration is changed (Dual/Single)	Caused by the user.

NOTE: When the controller configuration is changed from single to dual after setting up the Cache Residency logical unit, the Cache Residency logical unit is cancelled. You can open the Cache Residency Manager in single configuration, but neither setup nor operation can be performed.

Equipment

The following equipment is required fore Cache Residency Manager.

Table 6-3: 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 logical unit belonging to a master partition.
Number of logical units with the Cache Residency function	1/controller (2/arrays)

Logical Unit Capacity

The maximum size of the Cache Residency Manager logical unit depends on the cache memory. Note that the Cache Residency logical unit is only assigned a master partition.

The capacity varies with Cache Partition Manager and SnapShot or TCE. There are three scenarios:

- Cache Partition Manager and SnapShot/TCE/Dynamic Provisioning re disabled
- Cache Partition Manager is disabled, while SnapShot/TCE/Dynamic Provisioning is enabled
- Cache Partition Manager is enabled, while SnapShot/TCE/Dynamic Provisioning is enabled or disabled
- Only when Dynamic Provisioning is valid

Note the following restrictions:

- When the hardware revision is 0200 and Dynamic Provisioning is valid, a supported capacity changes the DP Capacity Mode setting.
- When Cache Partition Manager, SnapShot/TCE/Dynamic provisioning are invalid (when the hardware revision is 0100).

Table 6-4: Supported Capacity of Cache Residency LU with Cache Partition Manager, SnapShot/TCE/Dynamic Provisioning Disabled

AMS Equipment	Cache	Logical Unit Capacity
2100	1 GB/CTL	181,440 blocks (approximately 88 MB).
2100	2 GB/CTL	1,128,960 blocks (approximately 551 MB).
2100	4 GB/CTL	3,144,960 blocks (approximately 1,535 MB).
2300	1 GB/CTL	100,800 blocks (approximately 49 MB).
2300	2 GB/CTL	1,048,320 blocks (approximately 511 MB).
2300	4 GB/CTL	2,903,040 blocks (approximately 1,417 MB).
2300	8 GB/CTL	6,814,080 blocks (approximately 3,327 MB).
2500	2 GB/CTL	342,720 blocks (approx. 167 MB).
2500	4 GB/CTL	2,177,280 blocks (approx. 1,063 MB).
2500	6 GB/CTL	4,072,320 blocks (approx. 1,988 MB).
2500	8 GB/CTL	5,987,520 blocks (approx. 2,923 MB).
2500	10 GB/CTL	7,882,560 blocks (approx. 3,848 MB).
2500	12 GB/CTL	9,777,600 blocks (approx. 4,774 MB).
2500	16 GB/CTL	13,728,960 blocks (approx. 6,703 MB).

A case exists when Cache Partition Manager and Dynamic Provisioning are invalid, and SnapShot/TCE is valid when the hardware revision is 0200.

Table 6-5 on page 6-5 provides maximum capacity values for instancesw when Cache Partition Manager and Dynamic Provisioning are invalid, and SnapShot/TCE are valid.

Table 6-5: Supported Capacity of Cache Residency LU with Cache Partition Manager/Dynamic Provisioning Disabled and SnapShot/TCE Enabled

AMS Equipment	Cache	Logical Unit Capacity
2100	1 GB/CTL	SnapShot is not available.
2100	2 GB/CTL	483,840 blocks (approx. 236 MB)
2100	4 GB/CTL	907,200 blocks (approx. 442 MB)
2300	1 GB/CTL	SnapShot is not available.
2300	2 GB/CTL	423,360 blocks (appox. 206 MB)
2300	4 GB/CTL	665,280 blocks (approx. 324 MB)
2300	8 GB/CTL	2,459,520 blocks (approx. 1,200 MB)
2500	2 GB/CTL	SnapShot cannot be used together with Cache Residency Manager. The Cache Residency LU will be canceled.
2500	4 GB/CTL	423,360 blocks (approx. 206 MB)
2500	6 GB/CTL	766,080 blocks (approx. 374 MB)
2500	8 GB/CTL	1,532,160 blocks (approx. 748 MB)
2500	10 GB/CTL	2,399,040 blocks (approx. 1,171 MB)
2500	12 GB/CTL	3,044,160 blocks (approx. 1,486 MB)
2500	16 GB/CTL	5,120,640 blocks (approx. 2,500 MB)

A case exists when only Dynamic Provisioning is valid (when the hardware revision is 0100).

Table 6-6 on page 6-5 details the maximum capacity of Cache Residency LU when only Dynamic Provisioning is valid.

Table 6-6: Supported Capacity of Cache Residency LU With Dynamic Provisioning Enabled

AMS Equipment	Cache	Logical Unit Capacity
2100	1 GB/CTL	100,800 blocks (approx. 49 MB)
2100	2 GB/CTL	927,360 blocks (approx. 452 MB)
2100	4 GB/CTL	2,903,040 blocks (approx. 1,417 MB)
2300	1 GB/CTL	Dynamic Provisioning cannout be used together with Cache Residency Manager. the Cache Residency LU will be canceled.
2300	2 GB/CTL	806,400 blocks (approx. 393 MB)
2300	4 GB/CTL	2,580,480 blocks (approx. 1,260 MB)
2300	8 GB/CTL	6,431,040 blocks (approx. 3,140 MB)
2500	2 GB/CTL	60,480 blocks (approx. 817 MB)
2500	4 GB/CTL	1,673,280 blocks (approx. 817 MB)
2500	6 GB/CTL	3,568,320 blocks (approx. 1,742 MB)
2500	8 GB/CTL	5,362,560 blocks (approx. 2,618 MB)
2500	10 GB/CTL	7,277,760 blocks (approx. 3553 MB)

Table 6-6: Supported Capacity of Cache Residency LU With Dynamic Provisioning Enabled

AMS Equipment	Cache	Logical Unit Capacity
2500	12 GB/CTL	8,951,040 blocks (approx. 4,370 MB)
2500	16 GB/CTL	13,083,840 blocks (approx. 6,388 MB)

Note the following:

- A case exists when Cache Partition Manager, SnapShot/TCE/Dynamic Provisioning are invalid (when the hardware revision is 200).
- When Cache Partition Manager, snapShot/TCE/Dynamic Provisioning are invalid, the maximum capacity of Cache Residency LU is as follows:

Table 6-7: Supported Capacity of Cache Residency LU (Cache Partition Manager, SnapShot/TCE/Dynamic Provisioning are Invalid)

AMS Equipment	Cache	Logical Unit Capacity
2100	1 GB/CTL	Not available.
2100	2 GB/CTL	Approx. 726 MB
2100	1 GB/CTL	Not Available
2100	2 GB/CTL	Approx. 726 MB
2100	4 GB/CTL	Approx. 1,732 MB
2300	1 GB/CTL	Not available.
2300	2 GB/CTL	Approx. 651 MB
2300	4 GB/CTL	Approx. 1,527 MB
2300	8 GB/CTL	Approx. 3,512 MB
2500	2 GB/CTL	Not available.
2500	4 GB/CTL	Approx. 1,335 MB
2500	6 GB/CTL	Approx. 2,280 MB
2500	8 GB/CTL	Approx. 3,225 MB
2500	10 GB/CTL	Approx. 4,167 MB
2500	12 GB/CTL	Approx. 5,108 MB
2500	16 GB/CTL	Approx. 7,066 MB

A case exists when Cache Partition Manager is invalid, and SnapShot/TCE/Dynamic Provisioning are valid (when the hardware revision is 0100).

Table 6-8 on page 6-7 details the maximum capacity of Cache Residency LU when Cache Partition Manager and Dynamic Provisioning are disabled, and SnapShot and TCE are enabled.

Table 6-8: Supported Capacity of Cache Residency LU With Cache Partition Manager Disabled and SnapShot/TCE/Dynamic Provisioning Enabled

AMS Equipment	Cache	Logical Unit Capacity
2100	1 GB/CTL	SnapShot is not available
2100	2 GB/CTL	Approx. 236 MB
2100	4 GB/CTL	Approx. 442 MB
2300	1 GB/CTL	SnapShot is not available
2300	2 GB/CTL	Approx. 206 MB
2300	4 GB/CTL	Approx. 324 MB
2300	8 GB/CTL	Approx 1,200 MB
2500	2 GB/CTL	SnapShot cannot be used together with Cache Residency Manager.
2500	4 GB/CTL	Approx. 206 MB
2500	6 GB/CTL	Approx. 374 MB
2500	8 GB/CTL	Approx. 748 MB
2500	10 GB/CTL	Approx. 1,171 MB
2500	12 GB/CTL	Approx. 1,486 MB
2500	16 GB/CTL	Approx. 2,500 MB

Note the following:

- When only Dynamoic Provisioning is valid (when the hardware revision is 0200).
- When only Dynamic Provisioning is valid, the maximum capacity of Cache Residency LU is as follows:

Table 6-9: Supported Capacity of Cache Residency LU (When only Dynamic Provisioning is Valid)

AMS Equipment	Cache	Logical Unit Capacity
2100	1 GB/CTL	Approx. 49 MB
2100	2 GB/CTL	Approx. 452 MB
2100	4 GB/CTL	Approx. 1,417 MB
2300	1 GB/CTL	Dynamic Provisioning cannot be used together with Cache Residency Manager
2300	2 GB/CTL	Approx. 393 MB
2300	4 GB/CTL	Approx. 1,260 MB
2300	8 GB/CTL	Approx. 3,140 MB
2500	2 GB/CTL	Approx. 29 MB
2500	4 GB/CTL	Approx. 817 MB
2500	6 GB/CTL	Approx. 1,742 MB
2500	8 GB/CTL	Approx. 2,618 MB
2500	10 GB/CTL	Approx. 3,553 MB

AMS Equipment	Cache	Logical Unit Capacity
2500	12 GB/CTL	Approx. 4,370 MB
2500	16 GB/CTL	Approx. 6,388 MB

Table 6-10: Supported Capacity of Cache Residency LU (When only Dynamic Provisioning is Valid and DP Capacity Mode is Maximum Capacity)

AMS Equipment	Cache	Logical Unit Capacity
2100	1 GB/CTL	Not available.
2100	2 GB/CTL	383 MB)
2100	4 GB/CTL	1,348 MB)
2300	1 GB/CTL	Not available
2300	2 GB/CTL	305 MB)
2300	4 GB/CTL	1,171 MB)
2300	8 GB/CTL	3,051 MB)
2500	2 GB/CTL	Not available
2500	4 GB/CTL	679 MB
2500	6 GB/CTL	1,604 MB
2500	8 GB/CTL	2,480 MB
2500	10 GB/CTL	3,405 MB
2500	12 GB/CTL	4,232 MB
2500	16 GB/CTL	6,250 MB

Note the following:

When Cache Partition Manager is invalid, and Snapshot/TCE/Dynamic Provisioning are valid)when the hardware revision is 0200).

When Cache Partition Manager is invalid, and SnapShot/TCE/Dynamic Provisioning are valid, the maximum capacity of Cache Residency LU is as follows:

Table 6-11: Supported Capacity of Cache Residency LU (When Cache Partition Manager is Invalid, and SnapShot/TCE/Dynamic Provisioning are Valid)

AMS Equipment	Cache	Logical Unit Capacity
2100	1 GB/CTL	SnapShot is not available.
2100	2 GB/CTL	Approx. 196 MB
2100	4 GB/CTL	Approx. 413 MB
2300	1 GB/CTL	SnapShot is not available.
2300	2 GB/CTL	Approx. 137 MB
2300	4 GB/CTL	Approx. 246 MB
2300	8 GB/CTL	Approx 1,122 MB

AMS Equipment	Cache	Logical Unit Capacity
2500	2 GB/CTL	SnapShot cannot be used together with Cache Residency Manager.
2500	4 GB/CTL	Approx. 59 MB
2500	6 GB/CTL	Approx. 236 MB
2500	8 GB/CTL	Approx. 610 MB
2500	10 GB/CTL	Approx. 1,033 MB
2500	12 GB/CTL	Approx. 1,348 MB
2500	16 GB/CTL	Approx. 2,352 MB

Table 6-12: Supported Capacity of Cache Residency LU (When Cache Partition Manage3r is Invalid, and SnapShot/TCE/ Dynamic Provisioning are Valid, and DP Capacity Mode is Maximum Capacity)

AMS Equipment	Cache	Logical Unit Capacity
2100	1 GB/CTL	SnapShot is not available
2100	2 GB/CTL	Approx. 127 MB
2100	4 GB/CTL	Approx. 344 MB
2300	1 GB/CTL	SnapShot is not available
2300	2 GB/CTL	Approx. 49 MB
2300	4 GB/CTL	Approx. 167 MB
2300	8 GB/CTL	Approx. 1,043 MB
2500	2 GB/CTL	Not available
2500	4 GB/CTL	Not available
2500	6 GB/CTL	Approx. 96 MB
2500	8 GB/CTL	Approx. 462 MB
2500	10 GB/CTL	Approx. 885 MB
2500	12 GB/CTL	Approx. 1,210 MB
2500	16 GB/CTL	Approx. 2,214 MB

Note the following:

- When Cache Partition Manager is valid:
- When Cache Partition Manager is valid, supported capacity of Cache Residency LU is not concerned Snapshot/TCE/Dynamic Provisioning is valid or invalid, the maximum capacity is decided by the capacity of a master partition.

Table 6-13: Supported Capacity of Cache Residency LU with Cache Partition Manager Enabled

AMS Equipment	Cache	Cache Residency Logical Unit Capacity
2100	1 GB/CTL	Cache Partition Manager is not available.

Table 6-13: Supported Capacity of Cache Residency LU with Cache Partition Manager Enabled

AMS Equipment	Cache	Cache Residency Logical Unit Capacity
2100	2 GB/CTL	(The master partition size (MB) See Note 1 - 200 MB) x 2,016 (Blocks)
2100	4 GB/CTL	
2300	1 GB/CTL	Cache Partition Manager is not available.
2300	2 GB/CTL	(The master partition size (MB) See Note 1 - 200 MB) x 2,016 (Blocks)
2300	4 GB/CTL	
2300	8 GB/CTL	
2500	2 GB/CTL	Cache Partition Manager is not available.
2500	4 GB/CTL	(The master partition size (MB) See Note 1 - 400 MB) x 2,016 (Blocks)
2500	6 GB/CTL	
2500	8 GB/CTL	
2500	10 GB/CTL	
2500	12 GB/CTL	
2500	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.

Table 6-14: Supported Capacity of Cache Residency LU (When Cache Partition Manager is Invalid, and SnapShot/TCE/Dynamic Provisioning are Valid, and DP Capacity Mode is MaximumCapacity)

AMS Equipment	Cache	Logical Unit Capacity
2100	1 GB/CTL	SnapShot is not available.
2100	2 GB/CTL	Approx. 127 MB
2100	4 GB/CTL	Approx. 344 MB
2300	1 GB/CTL	SnapShot is not available
2300	2 GB/CTL	Approx. 49 MB
2300	4 GB/CTL	Approx. 167 MB
2300	8 GB/CTL	Approx. 1,043 MB
2500	2 GB/CTL	Not available
2500	4 GB/CTL	Not avaiable
2500	6 GB/CTL	Approx. 98 MB
2500	8 GB/CTL	Approx. 462 MB

AMS Equipment	Cache	Logical Unit Capacity
2500	10 GB/CTL	Approx. 885 MB
2500	12 GB/CTL	Approx. 1,210 MB
2500	16 GB/CTL	Approx. 2,214 MB

The following is true:

- When Cache Manager is Valid:
- When Cache Partition Manager is valid, supported capacity of Cache Residency LU is not concerned Snapshot/TCE/Dynamic Provisioning is valid or invalid, the maximum capacity is decided by the capacity of a master partition.

AMS Equipment	Cache	Cache Residency Logical Unit Capacity
2100	1 GB/CTL	Cache Partition Manager is not available.
2100	2 GB/CTL	The master partition size (MB) Note 1 - 200 MB) x 2,016 (Blocks)
2100	4 GB/CTL	
2300	1 GB/CTL	Cache Partition Manager is not available.
2300	2 GB/CTL	(The master partition size (MB) Note 1 - 200 MB) x 2,016 (blocks)
2300	4 GB/CTL	
2300	8 GB/CTL	
2500	2 GB/CTL	Cache Partition Manager is not available
2500	4 GB/CTL	(The master partition size (MB) Note 1 - 400 MB) x 2,016 (blocks)
2500	6 GB/CTL	
2500	8 GB/CTL	
2500	10 GB/CTL	
2500	12 GB/CTL	
2500	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

The following sections provide Cache Residency Manager restrictions.

Table 6-15: Cache Residency Manager Restrictions

Item	Description
Concurrent use of SnapShot	The Cache Residency Manager logical unit (logical unit cache residence) cannot be set to P-VOL, V-VOL, or data Pool volume. When using SnapShot, the logical unit capacity that can be specified as a cache residency is limited. For more information, see Table 2-24 on page 2-22.
Concurrent use of Cache Partition Manager	You cannot change a partition affiliated with the Cache Residency logical unit. After you cancel the Cache Residency logical unit, you must set it up again.
Concurrent use of Volume Migration	The Cache Residency Manager logical unit (logical unit cache residence) cannot be set to P-VOL or S-VOL. After you cancel the Cache Residency logical unit, you must set it up again.
Concurrent use of Power Saving	A RAID group logical unit that has powered down can be specified as the Cache Residency logical unit. However, if a host accesses a Cache Residency RAID group logical unit that has powered down, and error occurs.
Concurrent use of TCE	The LU specified for Cache Residency Manager (LU cache residence) cannot be set to P-VOL, S-VOL, or a data pool volume. When using TCE concurrently, LU capacity is limited.
Concurrent use of LUN Expansion	The unified LU cannot be set to the Cache Residency LU. The Cache Residency LU cannot be used as a unified LU.
Concurrent use of RAID group expansion	You cannot configure an LU as a Cache Residency LU while executing a RAID group expansion. You cannot execute a RAID group expansion for a RAID group that contains a Cache Residency LU.
LU Expansion	You cannot configure an LU as a Cache Residency LU if that LU has been expanded. growing as a Cache Residency LU. You cannot expand LUs that have been configured as Cache Residency LUs.
LU Reduction (shrinking)	You can specify the LU after the LU reduction as a Cache Residency LU. However, you cannot execute an LU reduction for a Cache Residency LU.
Load balancing	The LU 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.

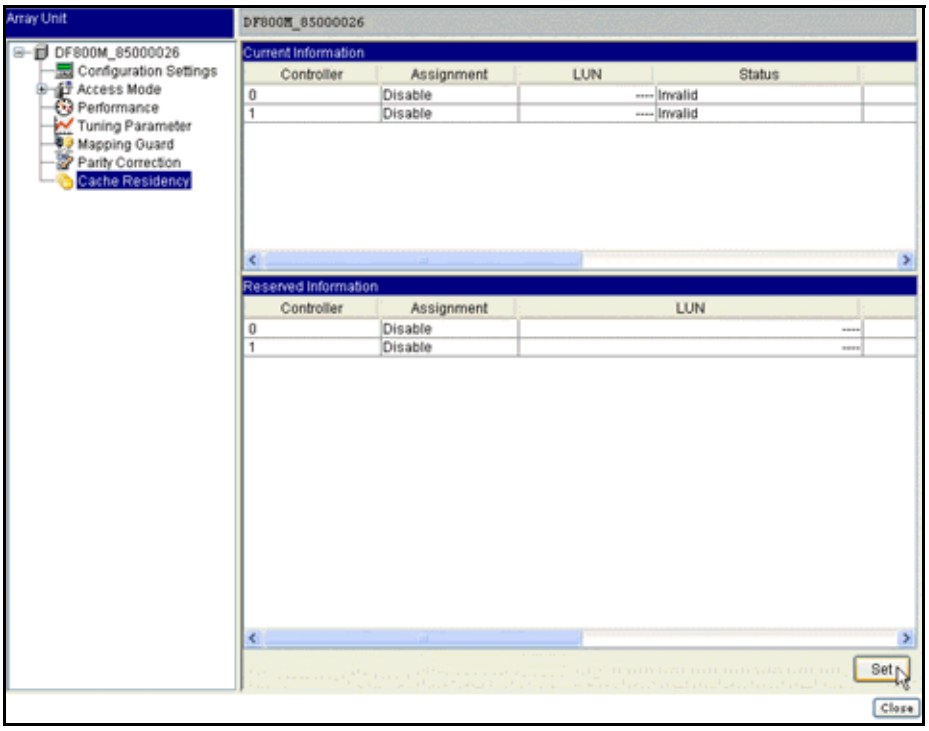

Figure 6-1: Cache Residency Manager

1. Click **Set**. The Cache Residency dialog box appears, as shown in Figure 6-2.

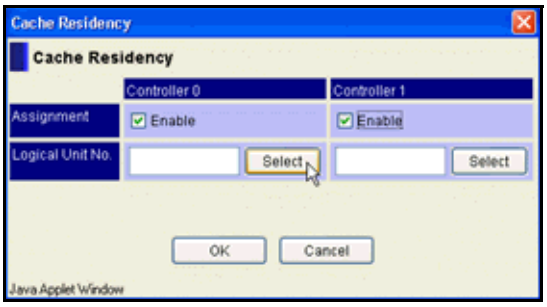

Figure 6-2: Cache Residency Dialog Box

2. In the **Assignment** field, do one of the following:
   - Select the **Enable** checkbox to set the residency logical unit.
   - Clear the **Enable** checkbox to cancel the residency logical unit.
3. In **Logical Unit No.** field, specify the logical unit number where you are setting or canceling the cache residency information. You can also click **Select** and specify the logical unit number.
4. Click **OK**.
5. Follow the on-screen instructions.

NOTE: Before you restart the array, be sure the host is not accessing data.

Cache Residency Manager operations

The procedure for Cache Residency Manager appears below.

Initial settings

1. Verify that you have the environments and requirements for Cache Residency Manager (see Preinstallation information on page 2-2).
2. Set the Cache Residency Manager (see Setting and canceling residency logical units on page 6-14).

Stopping Cache Residency Manager

1. Cancel the logical unit (see Setting and canceling residency logical units on page 6-14).
2. Disable Cache Residency Manager (see Setting and canceling residency logical units on page 6-14).

Before managing cache residency logical units, make sure that they have been defined.

Setting and canceling residency logical units

1. Start Navigator 2 and log in. The Arrays window appears
2. Click the appropriate array.
3. Expand the **Settings** list, and click **Advanced Settings**.
4. Click **Open Advanced Settings**. After some minutes, the Array Unit window appears.
5. Click **Cache Residency**. Figure 6-1 appears.

7

Data Retention Utility

This chapter describes the Data Retention utility.

This chapter covers the following topics:

- Data Retention Utility overview
- Data Retention Utility operations

Data Retention Utility overview

Table 7-1 describes the Data Retention Utility specifications.

Table 7-1: Data Retention Specifications

Item	Description
Logical unit setting	Set each logical unit. However, the expiration Lock is set for each array.
Number of logical units you can set	AMS2100: 2,048 logical units. AMS2300 and AMS2500: 4,096 logical units.
Access attributes	Read/Write (default). S-VOL Disable. Read Only. Protect. Read Capacity 0 (can only be set with CCI). Invisible from Inquiry Command (can only be set with CCI).
Protection against access attribute changes	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 on.
Unsupported logical units	Command devices. DMLUs. Sub-logical units (of a unified logical unit). Unformatted logical units. Logical units set as a Pool volumes in SnapShot.
Relation with ShadowImage/ SnapShot	When setting S-VOL Disable for a logical unit, the pair formation using the logical unit as an S-VOL (Pool volume) is suppressed. Setting of the S-VOL Disable of a volume that has already become an S-VOL (V-VOL or Pool volume) is not suppressed when the pair status is Split. When the S-VOL Disable is set for a P-VOL, you cannot restore SnapShot and ShadowImage.
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 when a controller detaches.
Relation with drive restoration	A correction copy, dynamic sparing, and copy back are performed as a usual logical unit.
Logical unit detachment	An access attribute that has been set for a logical unit is retained even when the logical unit is detached.
Firmware replacement	For logical units whose access attributes are not Read/ Write or S-VOL Disable, the initial setting up and initialization of settings (Configuration Clear) are suppressed.
Access attribute setting	The following operations cannot be performed for logical units whose access attributes are not Read/Write, and for RAID groups that include logical units: • Logical unit deletion • Logical unit formatting • RAID group deletion

Table 7-1: Data Retention Specifications (Continued)

Item	Description
Setting by Navigator 2	When Navigator 2 sets an access attribute, it can only be set for one logical unit at a time.
Unified logical units	A unified logical unit whose access level is not Read/Write, cannot be composed or dissolved.
Deleting, expanding, or reducing LU	An LU that has been configured for data retention where an access attribute has been set cannot be deleted, expanded, or reduced. You may set an access attribute after an LU has been expanded or reduced.
Cache Residency Manager	A logical unit whose access attribute is set can be used for the Cache Residency Manager. Conversely, an access attribute can be set for a logical unit in the Cache Residency Manager.
Concurrent use of LUN Manager	Yes.
Concurrent use of Volume Migration	Yes. The logical unit which executed the migration carries over the access attribute and retention term set by Data Retention, to the logical unit of the migration destination. When the access attribute is not Read/Write, the logical unit cannot be specified as an S-VOL of Volume Migration.
Concurrent use of SNMP Agent	Yes.
Concurrent use of Cache Partition Manager	Yes.
Concurrent Use of Dynamic Provisioning	Available. However, the DP-VOLs created by Dynamic Provisioning cannot be used. The Data Retention Utility can be executed for the normal logical unit.
Setting range of Retention Term	Unlimited.

Usage

This section provides notes on using Data Retention.

Logical unit access attributes

Do not modify logical unit access attributes while operations are performed on the data residing on the logical unit, or the operation may terminate abnormally.

You cannot change access attributes for the following logical volumes:

- A logical unit assigned to command device
- A logical unit assigned to a DMLU
- An uninstalled logical unit
- A unformatted logical unit

Unified logical units

You cannot combine logical volumes that do not have a Read/Write attribute. Unification of a unified logical unit, whose access attribute is not Read/Write, cannot be dissolved.

SnapShot and TCE

A logical unit whose access attribute is not Read/Write, cannot be assigned to a data pool. Additionally, an access attribute that is not Read/Write cannot be set for a logical unit that has been assigned to a data 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.

Logical units attributes set from the operating system

If you set access attributes from the OS, you must do so before mounting the logical unit. If the access attributes are set after the logical unit is mounted, the system may not operate properly.

When a command (create partition, format, etc.) is issued to a logical unit 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 logical unit.

A OS may not recognize a logical unit when the logical unit number (LUN) is larger than the one on which Invisible mode was set.

Windows 2000

A logical unit with a Read Only access attribute cannot be mounted.

Windows Server 2003/Windows Server 2008

When mounting a logical unit 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, logical units are not correctly recognized.

Unix

When mounting a logical unit with a Read Only attribute, mount it as Read Only (using the mount –r command).

Hewlett Packard Unix (HP-UX)

If there is a logical unit with a Read Only attribute, host shutdown may not be possible. When shutting down the host, change the logical unit attribute from Read Only to Protect.

If there is a logical unit with Protect attribute, host startup time may be lengthy. When starting the host, change the logical unit attribute to Read Only, or make the logical unit unrecognizable from the host by using mapping functions.

If a write is completed on the logical unit 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 logical unit 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 logical unit must be temporarily suspended using the raidvchkset -vg command. Place the logical unit again in the status in which it is checked when the LVM configuration change is completed.

HA Cluster Software

At times, a logical unit 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.

Data Retention Utility operations

Initial settings

1. Verify that you have the environments and requirements for Data Retention (see Preinstallation information on page 2-2).
2. Set the command device using the CCI. Refer to the following documentation for more information on the CCI:
   - *Hitachi AMS 2000 Family Command Control Interface (CCI) Reference Guide* (MK-97DF8121)
   - *Hitachi AMS 2000 Family Command Control Interface (CCI) Installation Guide* (MK-97DF8122)
   - *Hitachi AMS 2000 Family Command Control Interface (CCI) User's Guide* (MK-97DF8123)
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 operations

1. Set an attribute (see Setting attributes on page 7-8).
2. Changing the retention term (see Setting attributes on page 7-8).
3. Set an S-VOL (see Setting S-VOLs on page 7-8).
4. Set the expiration lock (see Setting expiration locks on page 7-8).

Opening the Data Retention window

To open the Data Retention window, follow these steps.

1. Start Navigator 2 and log in. The Arrays window appears
2. Click the appropriate array.
3. Expand the **Settings** list, and click **Advanced Settings**.
4. Click **Open Advanced Settings**. After some minutes, the Array Unit window appears.

   The Java applet window may time out after 20 minutes due to an automatic logout function. If this occurs, close the Web browser, stop the SNM2 Server and restart. Launch the SNM2 GUI and return to the array you want to manage.
5. Click **Data Retention**. Figure 7-1 appears.

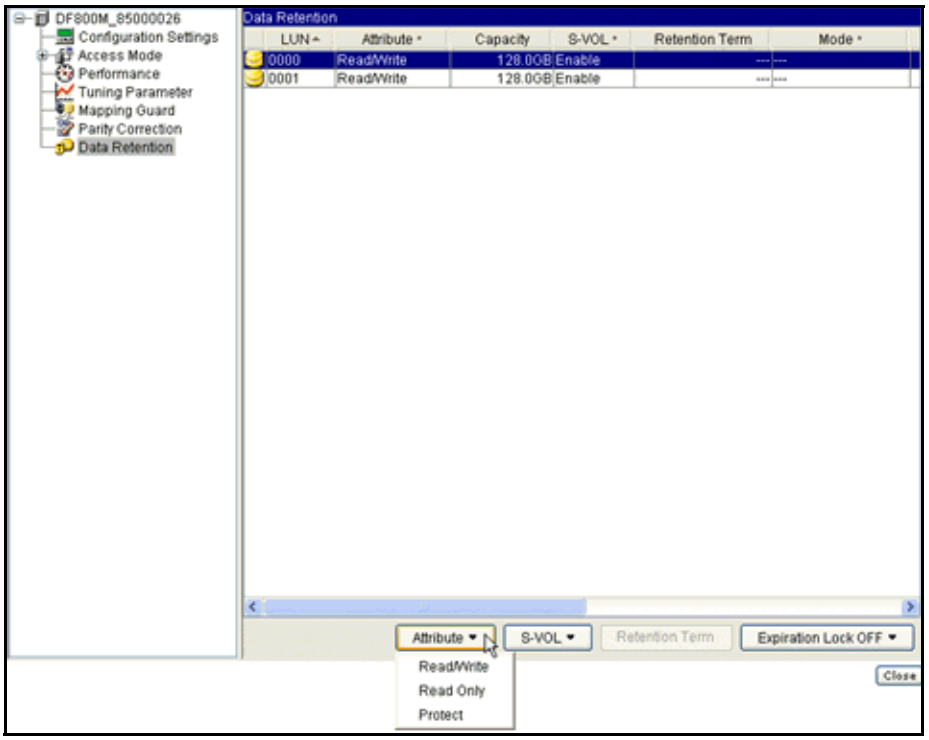

Figure 7-1: Data Retention Window

6. The following options are available:
   - **LUN** - logical unit number
   - **Attribute -** Read/Write, Read Only, Protect, or Can't Guard
   - **Capacity** - logical unit size
   - **S-VOL -** whether the logical unit 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.

7. Continue with the following sections to configure the desired Data Retention attributes.

Setting attributes

To set data retention attributes:

1. Select a LUN, and from the **Attribute** drop-down menu, specify the appropriate information. The **Term Setting** dialog box appears, as shown in Figure 7-2.

Figure 7-2: Term Setting Dialog Box

2. In the **Retention** Term field, specify the appropriate information and click **OK**.
3. Follow the on-screen instructions.

Setting S-VOLs

To set S-VOLs, follow these steps.

1. Select a LUN, and from the **S-VOL** drop-down menu, select **Disable**.
2. Follow the on-screen instructions.

Setting expiration locks

To set expiration locks, follow these steps.

1. From the **Expiration Lock** drop-down menu, select **ON**.
2. Follow the on-screen instructions.

8

LUN Manager

This chapter describes LUN Manager.

This chapter covers the following topics:

- LUN Manager overview
- LUN Manager operations

LUN Manager overview

This section provides the Fibre Channel and iSCSI specifications for LUN Manager.

Table 8-1: 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 logical unit 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)	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.
Logical Unit Mapping	Logical unit mapping can be set to the host group. 2,048 logical unit mappings can be set for a host group, and 16,384 can be set for a port.
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 logical unit, and 512 commands per port.

Table 8-2: 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 logical unit 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 8-2: 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	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	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.
Logical Unit Mapping	A logical unit can be set to the target. 2,048 logical unit mappings can be set for a target. Up to 16,384 logical unit 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 8-3: 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
Windows 2000®	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/.

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.

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 is connected to the array port.

Logical units 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.

Identify which logical units you want to use with a host, and then define a host group on that port for them (see Figure 8-1 on page 8-5).

:

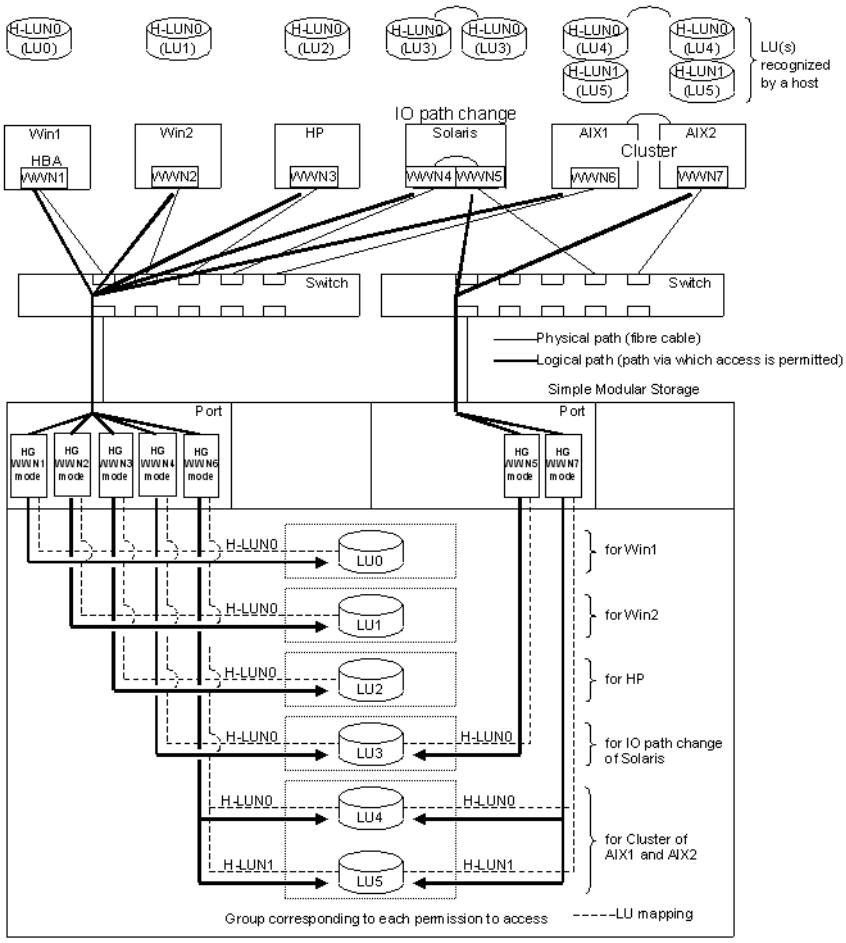

Figure 8-1: Fibre Channel System Configuration

Examples of configurations for creating host groups in multipathed and clustered environments appear in Figure 8-2 and Figure 8-3.

:

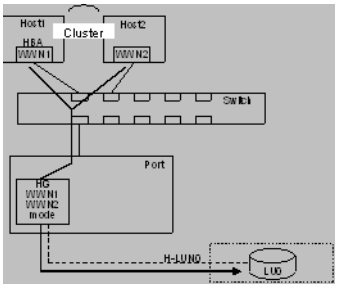

Figure 8-2: One Host Group Fibre Channel Configuration

.

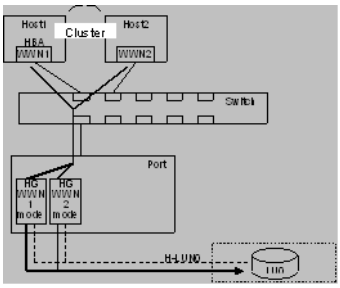

Figure 8-3: Two Host Groups Fibre Channel Configuration

Fibre Channel design considerations

When connecting multiple hosts to an array port, make sure you do the following.

Fibre Channel system design

- Assign logical units to hosts. Group logical units into host groups. Specify the host group when using LUN Manager.
- Assign logical units to RAID groups. A logical unit belongs to a RAID group. When two logical units (belonging to the same RAID group) are accessed at the same time, performance may decrease. When operating more than one host at the same time, assign logical units to separate RAID groups.
- Determine how to prevent unauthorized access. Determine input/output paths between hosts and logical units. The input/output path is a route through which access from the host is permitted.

  Set switch zoning to prevent interference from the other hosts that share the same switch. When the zoning is set, ports outside the zone do not affect ports within the zone.

  If you do not have enough ports, increase their number using the fibre channel switch.
- Determine queue depth. 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.

Fibre system configuration

To specify the input/output paths between hosts and logical units, 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
- Logical unit 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.

Overview

To set up and manage iSCSI network storage 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.

- Use Storage Navigator Modular 2 to set up logical units on the array.
- 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 logical units to targets.
  - Register CHAP users that are authorized to access the logical units.
  - Keep a record of the iSCSI names and related settings to simplify making any changes later.
- Physically connect the array to the network.
- 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.
- As a security measure, use LUN Manager in assignment mode to determine input/output paths between hosts and logical units. The input/output path is a route through which access from the host is permitted.
  - 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.
- Test host connections to the logical units on the array.
- Perform maintenance as needed: host and HBA addition, logical unit addition, HBA replacement, and switch replacement. Refer to your HBA vendor's documentation and Web site.

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 LU.

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/logical unit paths with LUN Manager

  Determine input/output paths between hosts and logical units 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.

:

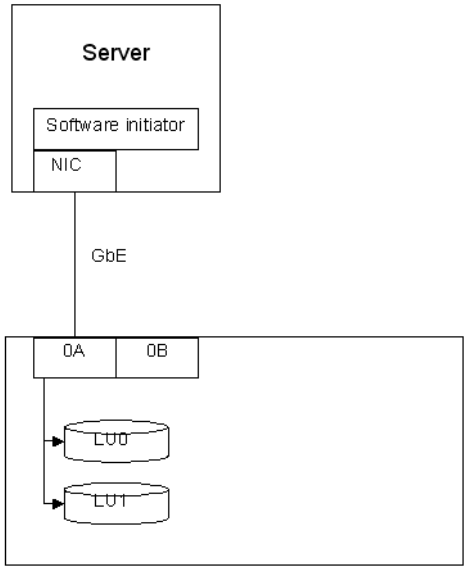

Figure 8-4: Direct Attached Type 1 for iSCSI

.

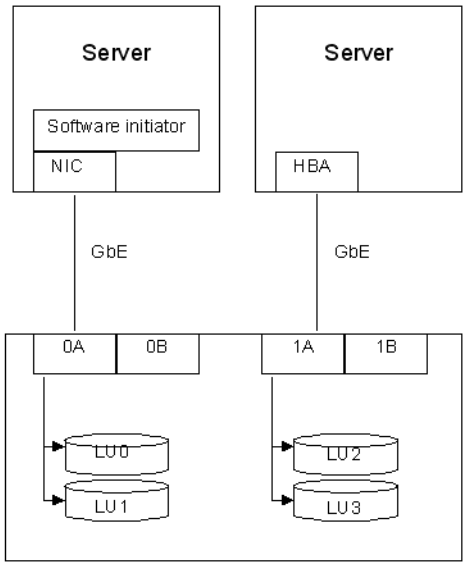

Figure 8-5: Direct Attached Type 2 for iSCSI

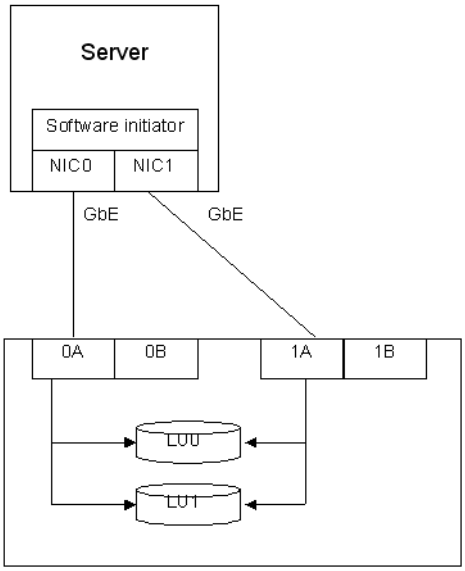

Figure 8-6: Direct Attached Type 3 for iSCSI

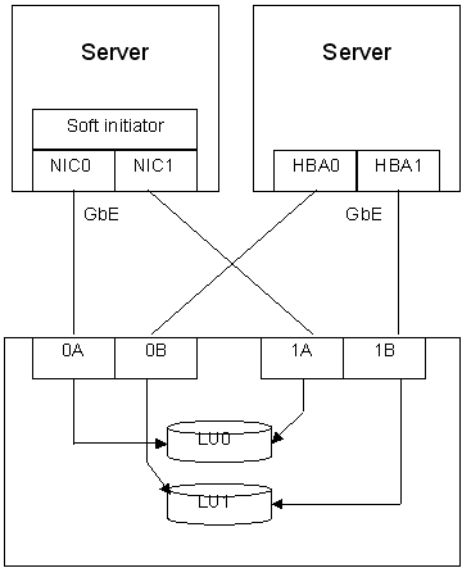

Figure 8-7: Direct Attached Type 4 for iSCSI

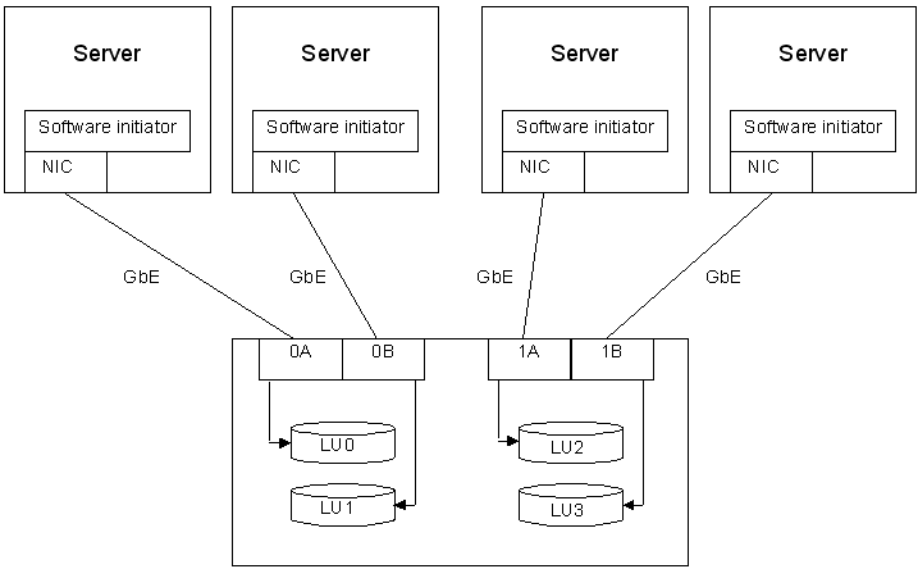

Figure 8-8: Direct Attached Type 5 for iSCSI

The following figures show possible topologies for switch-attached connections.

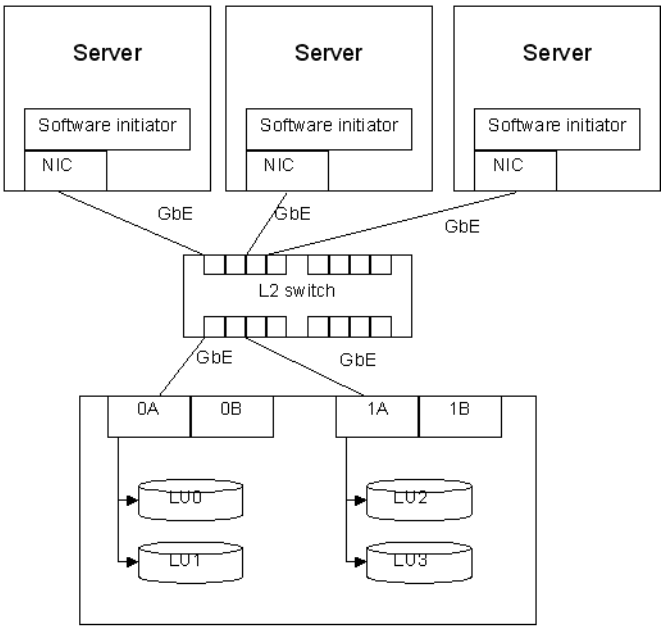

Figure 8-9: Switch Attached Type 1 for iSCSI

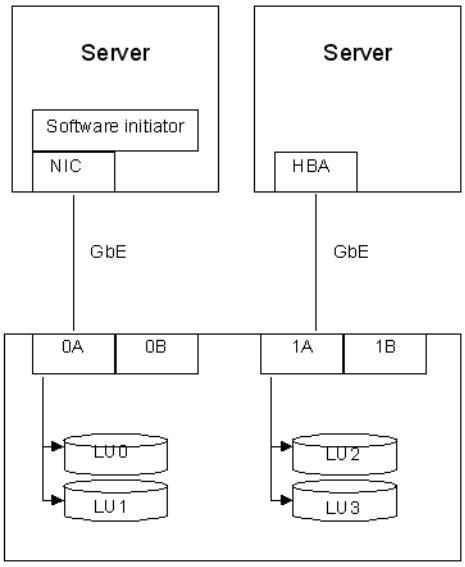

Figure 8-10: Switch Attached Type 2 for iSCSI

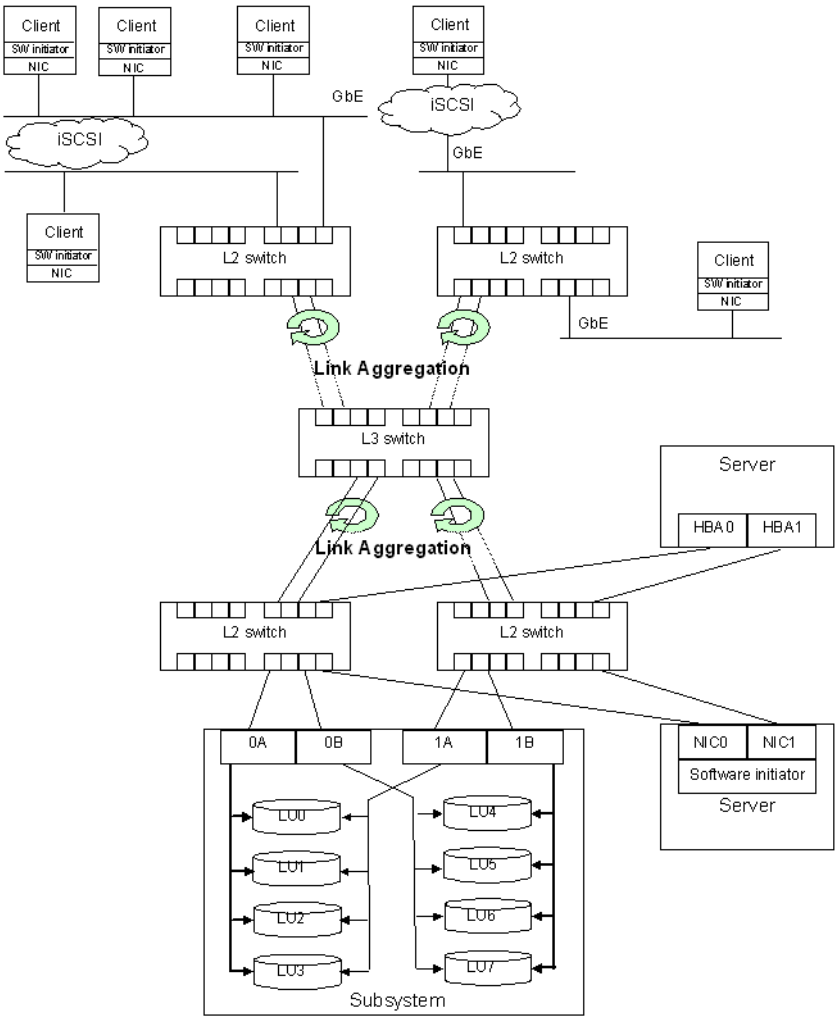

Figure 8-11: Switch Attached Type 3 for iSCSI

Assigning iSCSI targets and volumes to hosts

The host recognizes LUN(s) between H-LUN0 and H-LUN255. When you assign volumes of more than 256 logical units to the host, you must set the target logical unit mapping to be between H-LUN0 and H-LUN255.

- Up to 2,048 logical unit mappings can be set for a target.
- Up to 16,384 logical unit mappings can be set for a port.

:

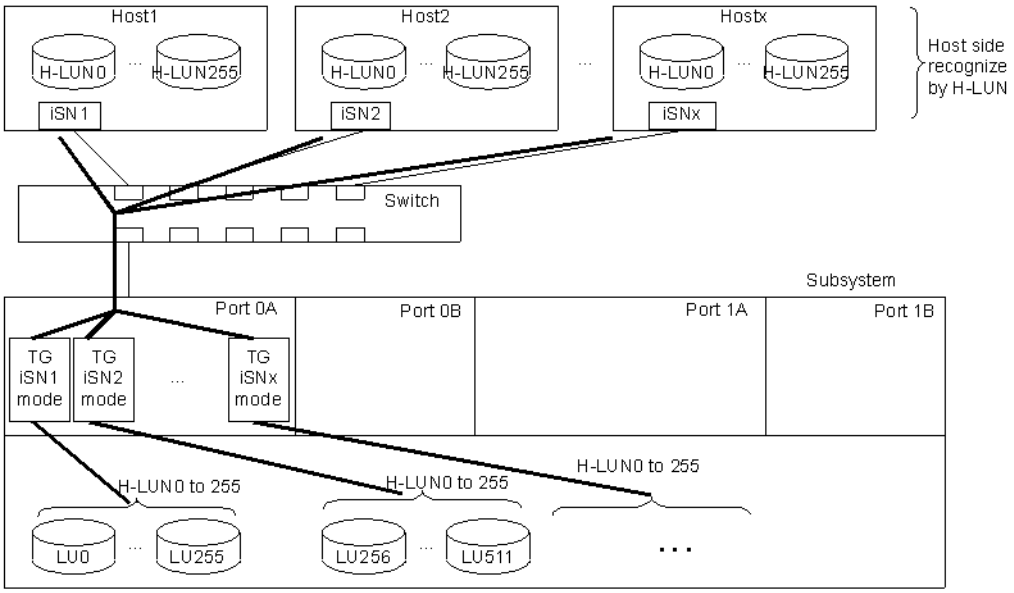

(iSN: iSCSI Name, TG: Target)

Figure 8-12: Mapping Volumes Between LU256-511 to the Host

When assigning LU3 to Host 1 and LU4 to Host 2, both hosts can access the same volume if the volume mapping is set alone as shown in Figure 8-13 on page 8-16. 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 8-14 on page 8-16.

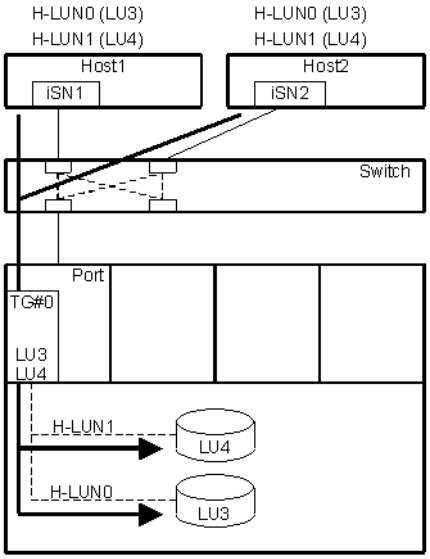

(TG: Target, iSN: iSCSI Name)

Figure 8-13: LUN Mapping—Different Hosts Can Access Volumes

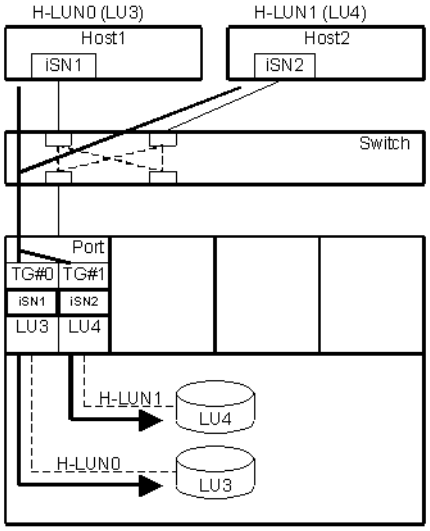

(TG: Target, iSN: iSCSI Name)

Figure 8-14: LUN 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 logical unit.

When assigning LU3 to Host 1 and LU4 to Host 2 as in Figure 8-15 on page 8-17, both hosts can access the same logical unit if the mapping is set separately.

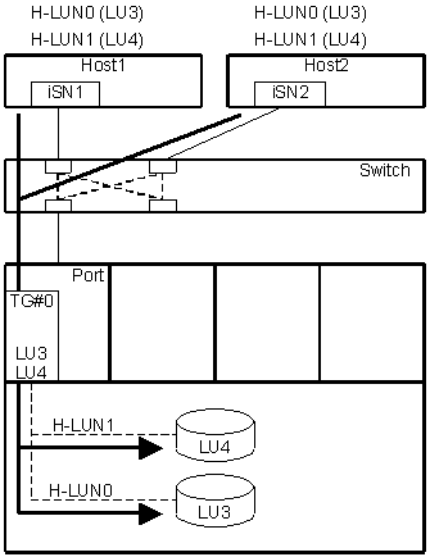

Figure 8-15: Volume Mapping—No Host Access Restrictions

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 8-16 on page 8-17.

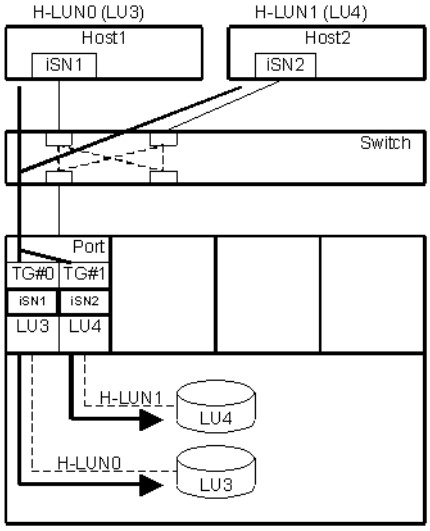

Figure 8-16: LUN Manager/CHAP—Restricted Host Access

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 8-17 on page 8-18.

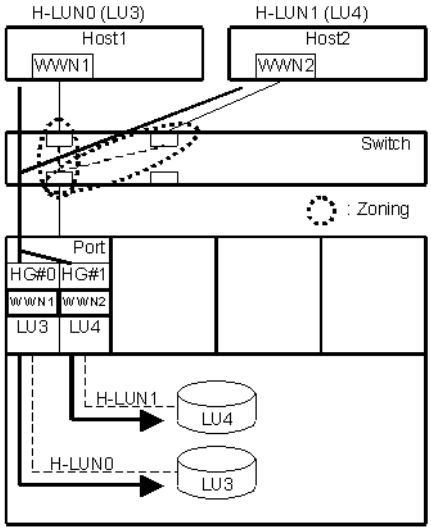

Figure 8-17: Switch Zoning

Avoiding RAID Group Conflicts

When multiple hosts are connected to an array and the logical units 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 LUNs 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 8-18 on page 8-19 and Figure 8-19 on page 8-19).

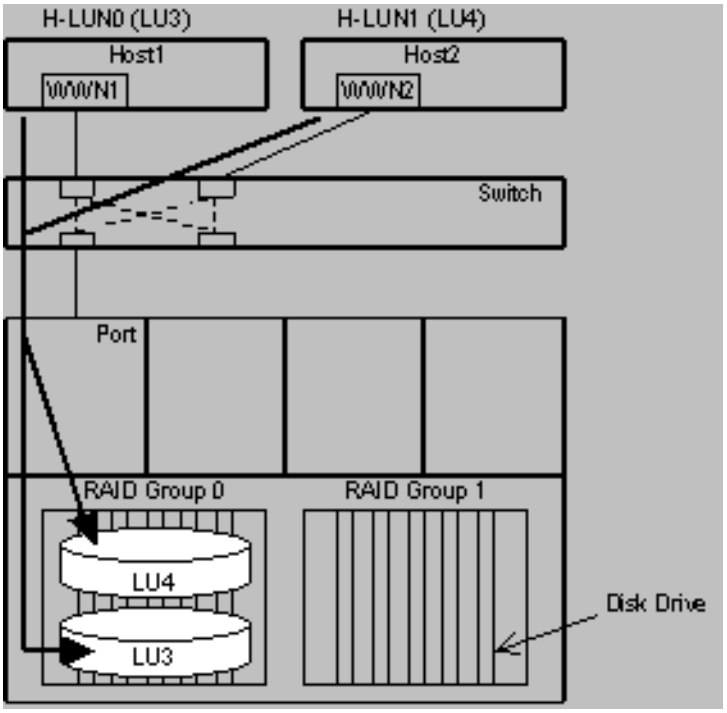

Figure 8-18: Hosts Connected to the Same RAID Group

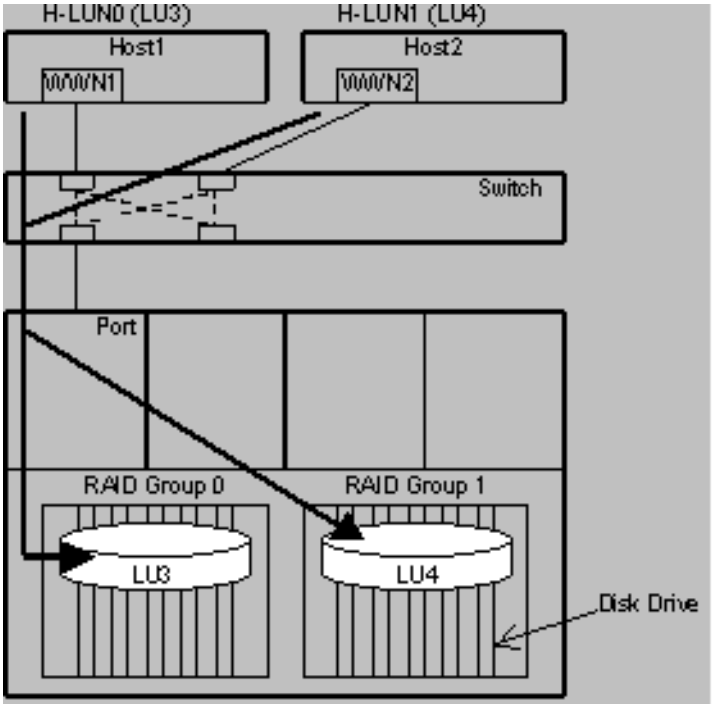

Figure 8-19: Hosts Connected to Different RAID Groups

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 • on page 8-20.

Table 2-9. 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.

Increasing queue depth and port sharing

Figure 8-20 on page 8-21 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.

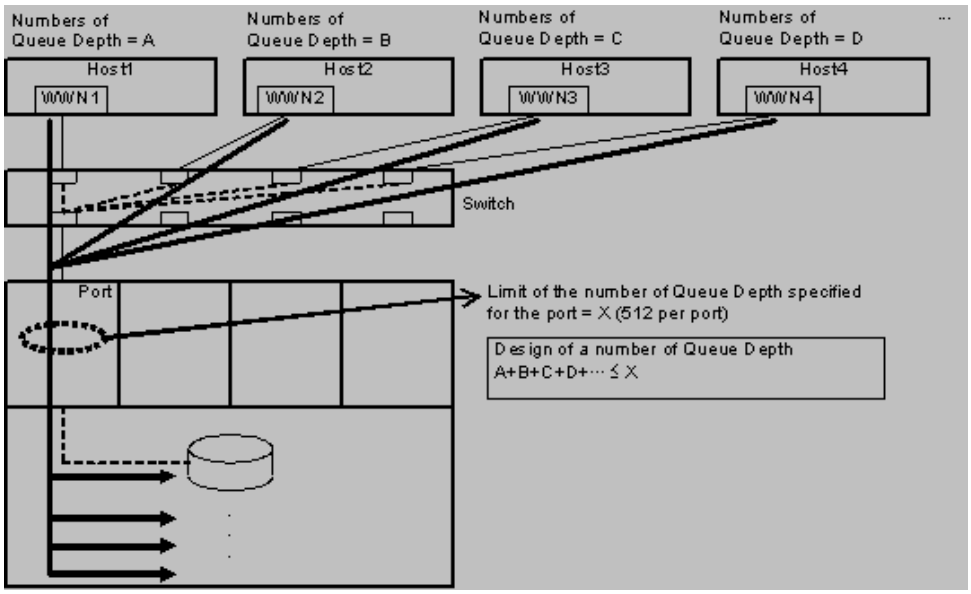

Figure 8-20: Queue Depth Does Not Exceed Port Limit

Increasing queue depth through path switching

Figure 8-21 on page 8-22 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 logical unit 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.

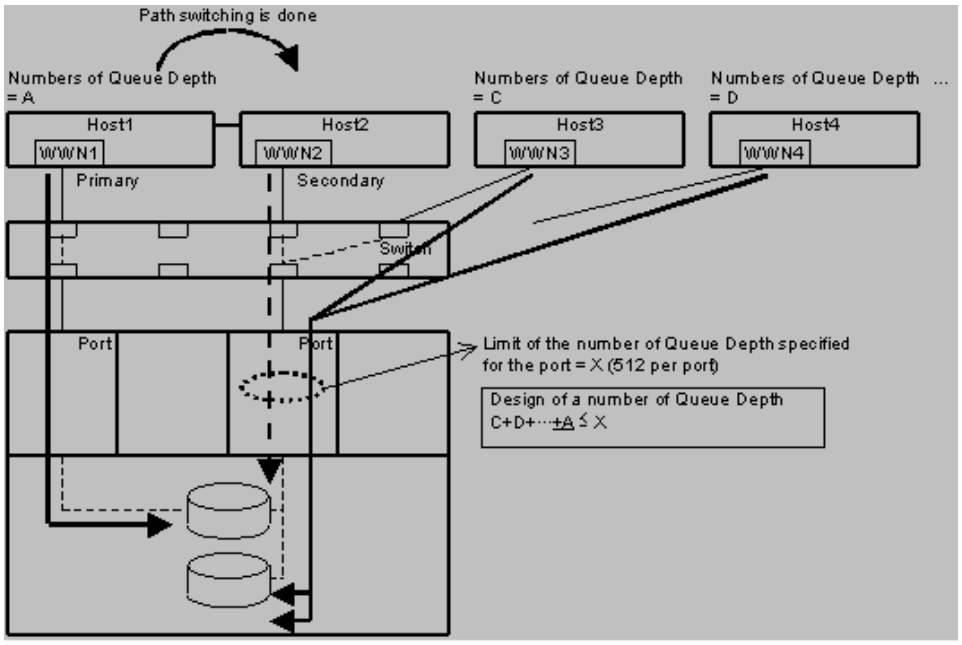

Figure 8-21: Queue Depth Increase From Path Switching

Queue depth allocation according to host job priority

Figure 8-22 on page 8-22 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.

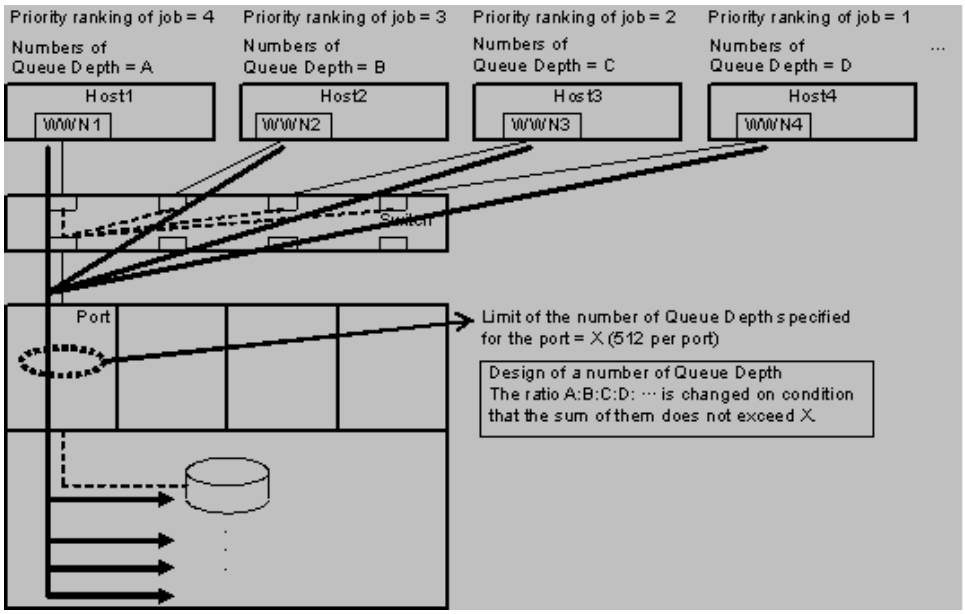

Figure 8-22: Host Job Priority

NOTE: We recommend that you execute any ping command tests when there is no I/O between hosts and controllers.

LUN Manager operations

This section describes LUN Manager operations for Fibre Channel and iSCSI.

Using Fibre Channel

1. Verify that you have the environments and requirements for LUN Manager (see Preinstallation information on page 2-2).

   For the array:

2. Set up a fibre channel port (see Fibre Channel operations using LUN Manager on page 8-24).
3. Create a host group (see Adding host groups on page 8-24).
4. Set the World Wide Name (WWN).
5. Set the host connection mode.
6. Create a logical unit.
7. Set the logical unit 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).

Using iSCSI

The procedure flow for iSCSI below. For more information, see the *Hitachi iSCSI Resource and Planning Guide* (MK-97DF8105).

1. Verify that you have the environments and requirements for LUN Manager (see Preinstallation information on page 2-2).

   For the array:

2. Set up the iSCSI port (see iSCSI operations using LUN Manager on page 8-31).
3. Create a target (see Adding and deleting targets on page 8-36).
4. Set the iSCSI host name (see Setting the iSCSI target security on page 8-33).

5. Set the host connection mode. For more information, see the *Hitachi iSCSI Resource and Planning Guide (MK-97DF8105).*
6. Set the CHAP security (see CHAP users on page 8-40).
7. Create a logical unit.
8. Set the volume mapping.
9. Set the network switch parameters. For more information, see the *Hitachi iSCSI Resource and Planning Guide (MK-97DF8105).*

   For the host:

10. Set the host bus adapter (HBA). For more information, see the *Hitachi iSCSI Resource and Planning Guide (MK-97DF8105).*
11. Set the HBA driver parameters. For more information, see the *Hitachi iSCSI Resource and Planning Guide (MK-97DF8105).*
12. Set the queue depth. For more information, see the *Hitachi iSCSI Resource and Planning Guide (MK-97DF8105).*
13. Set the CHAP security for the host (see CHAP users on page 8-40).
14. Create the disk partitions. For more information, see the *Hitachi iSCSI Resource and Planning Guide (MK-97DF8105).*

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 Word Wide Names
- Copy settings to other ports

Adding host groups

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

Enabling and disabling host group security

By default, the host group security is disabled for each port. To enable or disable host group security, follow these steps.

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 8-23).

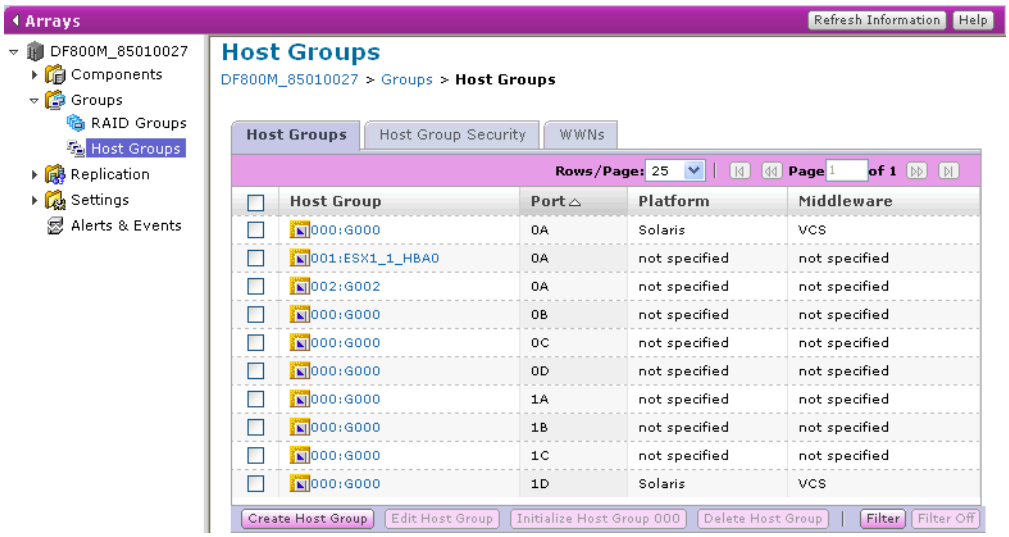

Figure 8-23: Host Groups Window

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. See Figure 8-24.
5. Select the port you want to configure and click **Change Host Group Security**.

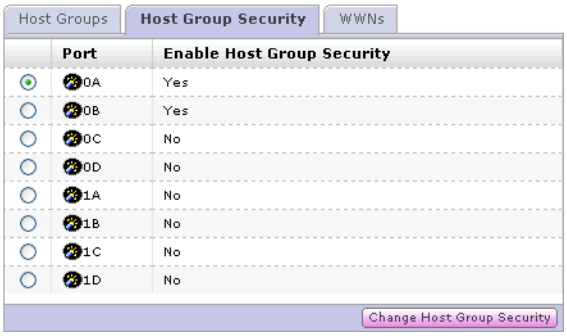

Figure 8-24: Host Group Security Tab — AMS System

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.

Creating and editing host groups

If you click **Create Host Group** without selecting a port, you can apply the same setting for multiple ports.

1. In the Host Groups tab, click **Create Host Group** or **Edit Host Group**. Figure 8-25 appears

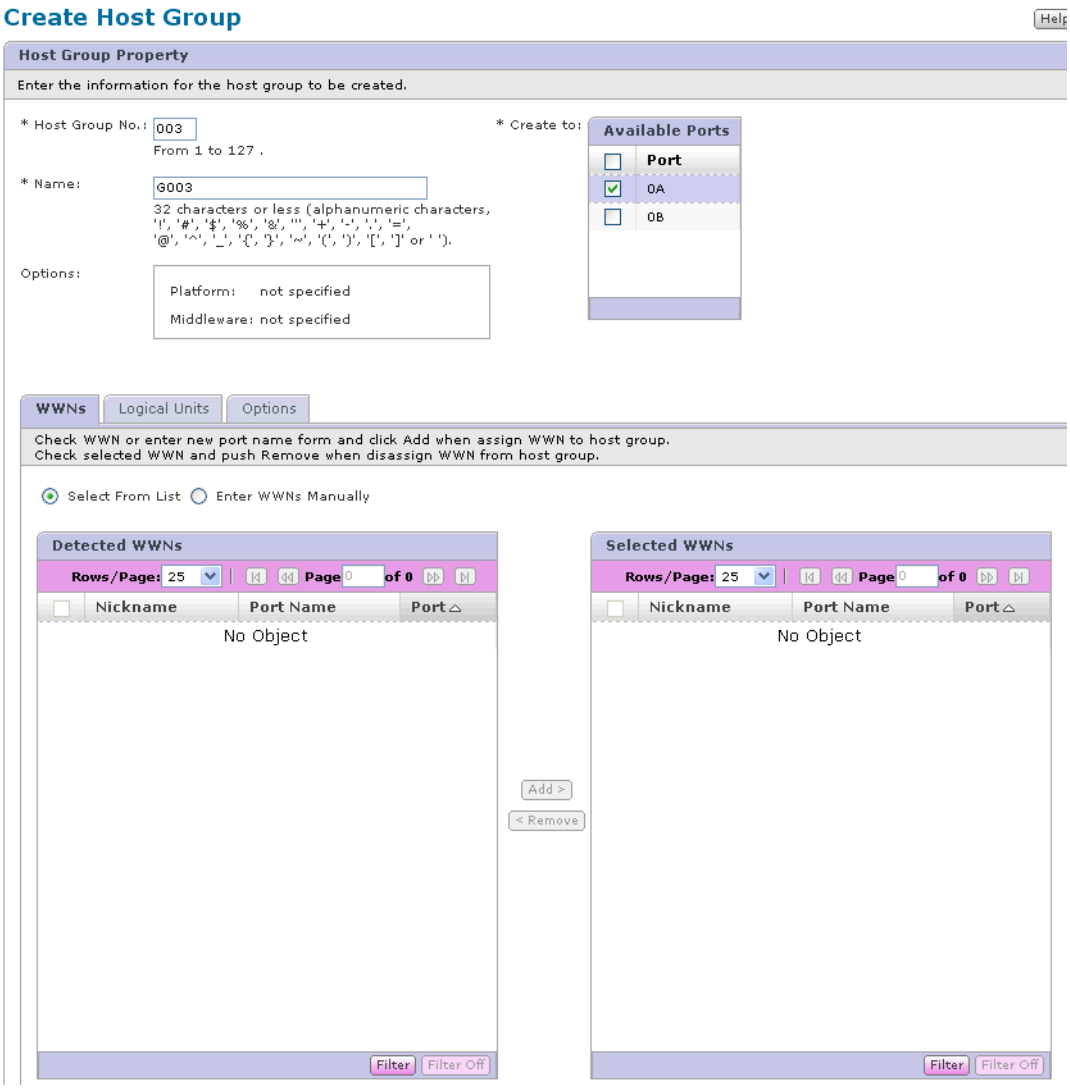

Figure 8-25: Host Group Property Window-WWNs Tab

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.
   - To specify the host information by selecting from a list, select **Select From List**, and click the appropriate WWN.
   - To specify the host information manually, select **Enter WWNs Manually**, and specify the port name that identifies the host (the port name must be 16 hexadecimal numerals).
   - **Port Name** is used to identify the host. Enter the Port Name using sixteen hexadecimal numerals.
4. Click **Add**. The added host information appears in the Selected WWNs pane.

NOTE: HBA WWNs are set to each host group, and are used for identifying hosts. When a port is connected to a host, the WWNs appear in the Detected WWNs pane and can be added to the host group. 128 WWNs can be assigned to a port. If you have more than 128 WWNs, delete one that is not assigned to a host group. Occasionally, the WWNs may not appear in the Detected WWNs pane, even though the port is connected to a host. When this happens, manually add the WWNs (host information).

5. Click the **Logical Units** tab. Figure 8-26 appears.

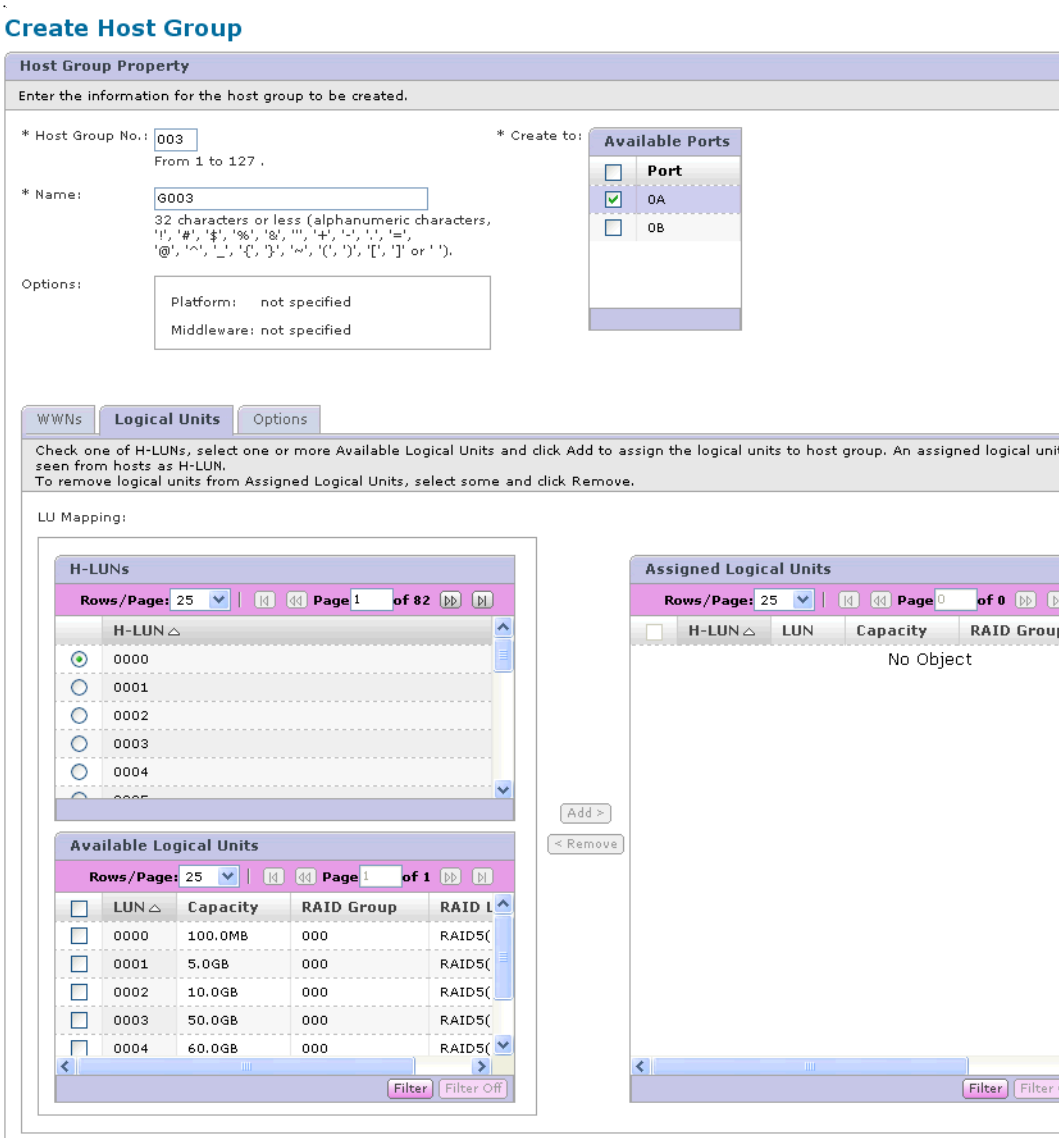

Figure 8-26: Host Group Property—Logical Units Tab

6. In the H-LUNs pane, select an available LUN. The host uses this number to identify the LUN it can connect to.
7. Click **Add**. The host LUN appears in the Assigned Logical Units list.

   To remove a host LUN, select it from the Assigned Logical Units list, and then click **Remove**.
8. Click the **Options** tab. The Create Host Group dialog box appears.
9. From the **Platform** and **Middleware** pull-down lists, select the appropriate platform and middleware, and click **OK**.
10. Follow the on-screen instructions.

Initializing Host Group 000

When you reset Host Group 000 to its default, its WWNs and logical unit settings are deleted and the host group name is reset to G000.

To initialize Host Group 0 follow these steps:

1. In the Hosts Groups window (Figure 8-23 on page 8-25), 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 logical units in Host Group 000, initialize it (see Initializing Host Group 000 on page 8-28).

1. In the Host Groups window (Figure 8-23 on page 8-25), select the appropriate host group and click **Delete Host Group**.
2. Follow the on-screen instructions.

Changing nicknames

1. In the Host Groups window (Figure 8-23 on page 8-25), click the WWNs tab. The WWNS tab appears (see Figure 8-27).

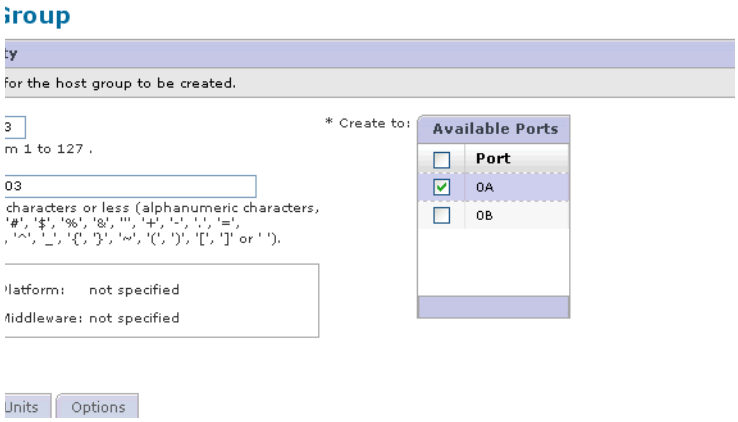

Figure 8-27: WWNs Tab

2. Select the appropriate WWN, and click **Change Nickname**.
3. Specify the nickname (up to 32 alphanumeric characters) and click **OK**.
4. Follow the on-screen instructions.

Deleting World Wide Names

1. In the Host Groups window (Figure 8-23 on page 8-25), click the WWNs tab. Figure 8-27 on page 8-29 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 logical unit of the created/edited host group
- Setting the options of the logical unit 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

The procedure for copying to the other port at the time of the host group creation is shown below.

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 8-24.
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

The procedure for copying to the other port at the time of the host group editing is shown below.

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 8-26.
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. Click **OK**.
8. Confirm the appeared message.
9. 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 8-28 on page 8-32), 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.

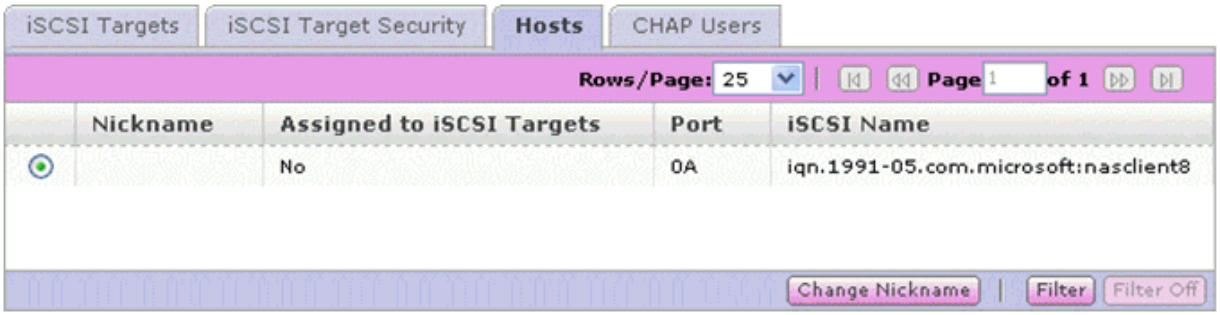

Figure 8-28: iSCSI Targets Window

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 logical unit. This setting specifies which host can access which logical unit.

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 logical units to be accessed from the Windows Host (initiator iSCSI Name A) and by the Linux Host (initiator iSCSI Name B) as shown in Figure 1-6 on page 1-9.

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 Logical Units tab enables you to assign logical units to logical unit numbers (H-LUNs) that are recognized by hosts. Figure 8-29 displays the iSCSI Target Properties - Logical Units tab.

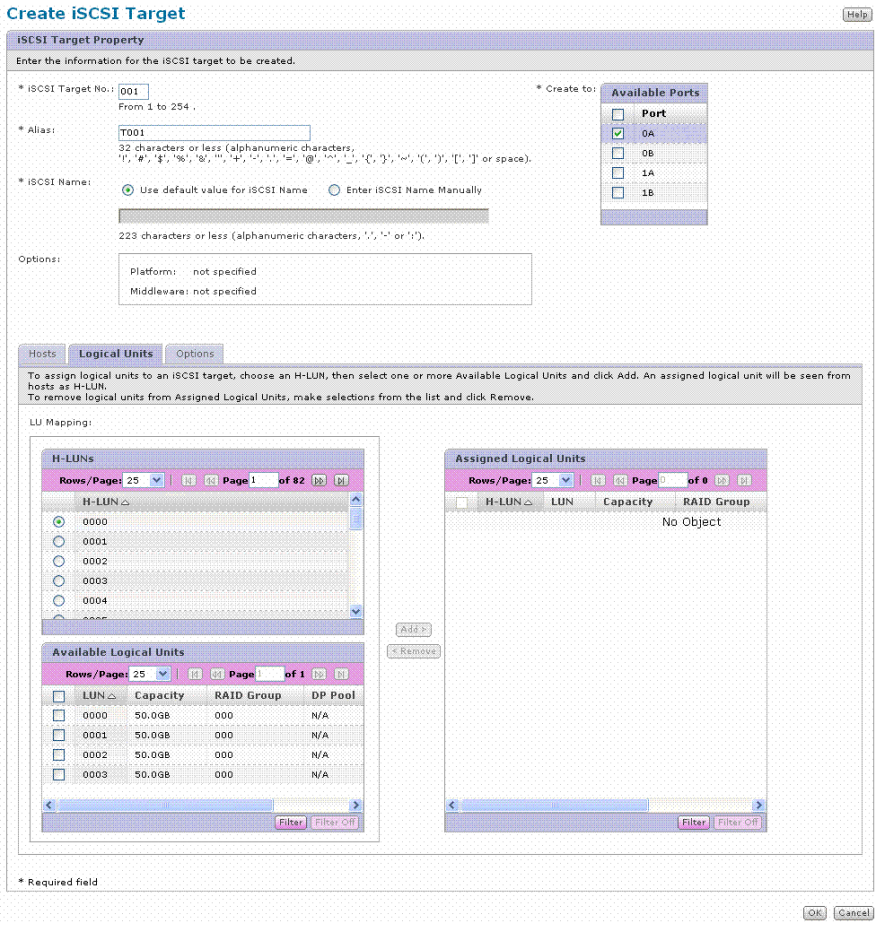

Figure 8-29: iSCSI Target Property - Logical Units tab

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 displays the iSCSI Target Property - Logical Units tab.

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.

4. Click the **iSCSI Target Security** tab, which displays the security settings for the data ports on your Hitachi Simple Modular Storage 100.

   Yes = security is enabled for the data port.

   No = security is disabled for the data port.

Figure 8-30: 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 8-31 displays the **Hosts** tab.

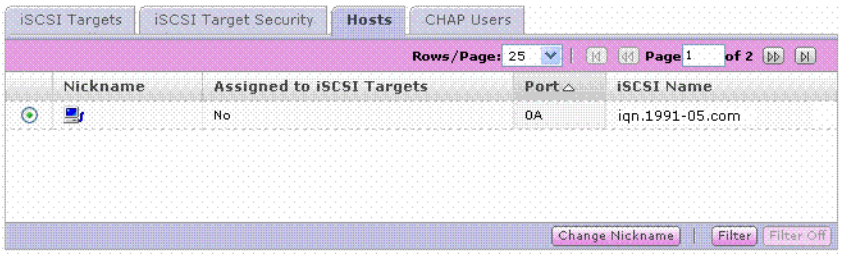

Figure 8-31: iSCSI Targets - Hosts tab

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.

Figure 8-32: iSCSI Target Security 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.

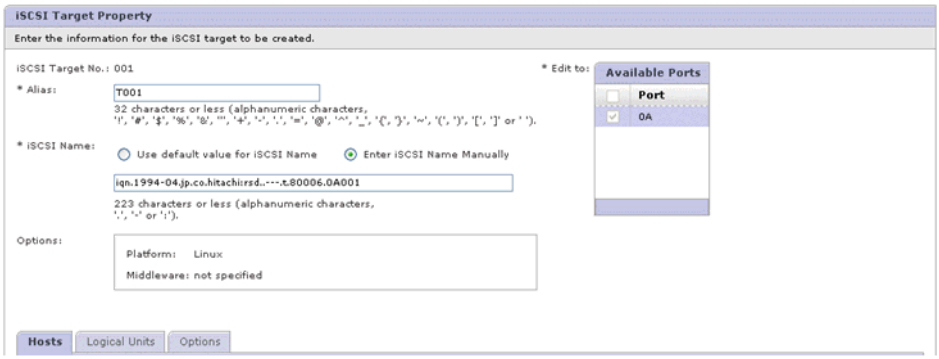

Figure 8-33: iSCSI Target Property Window

2. Enter the iSCSI Target No., Alias, or iSCSI Name.

   Note that the **Hosts** tab displays only when iSCSI Target Security is enabled.

3. If the iSCSI Target Security is enabled, set the 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 8-34, 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.

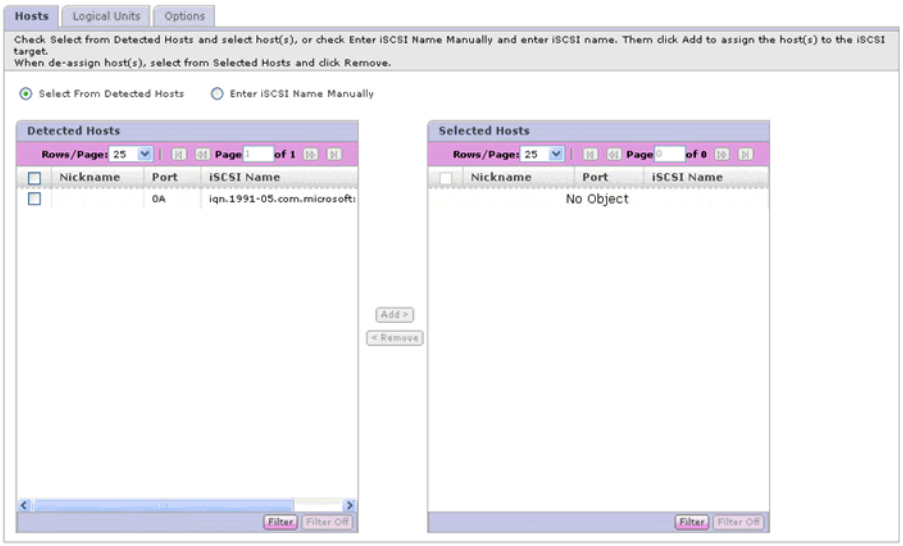

Figure 8-34: iSCSI Target Properties—Hosts Tab

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 **Logical Units** tab.
6. Select an available Host Logical Unit Number from the H-LUN list. The host uses this number to identify the LUN it can connect to and click **Add**. The added LUNs are displayed in the Selected Logical Units 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**, **Tru 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.

Deleting Targets

NOTE: Target 000 cannot be deleted. When deleting all the hosts and all the Logical Units 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.
2. Type the Alias or iSCSI Name, as required.
3. Set the host information from the **Hosts** tab.
4. Select the **Logical Units** tab.
5. Set the logical units information if necessary.
6. Select the **Options** tab.
7. Set the **Platform** and **Middleware** as required.
8. Click **OK**. The confirmation message is displayed.
9. 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 8-35 on page 8-39.

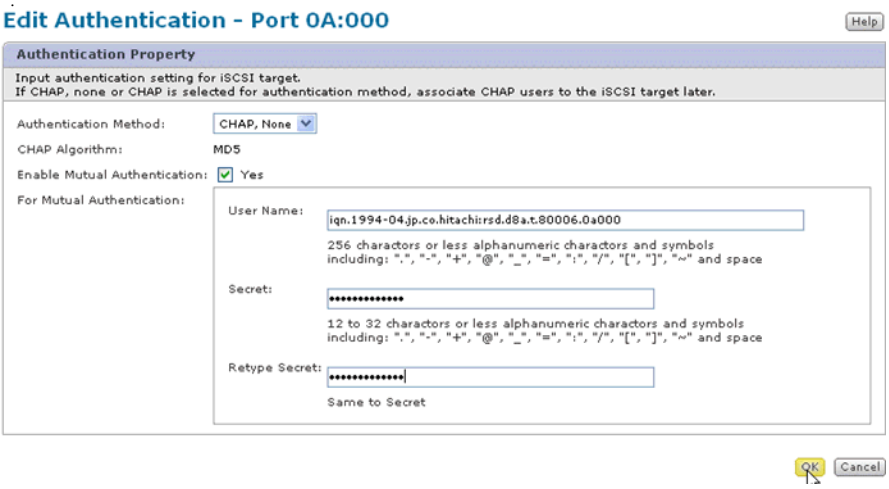

Figure 8-35: Edit Authentication Window

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 logical units 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 8-36 on page 8-40.

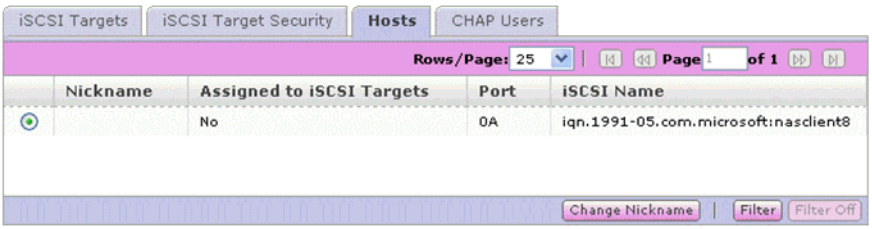

Figure 8-36: iSCSI Target Window — Hosts Tab

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 8-37 on page 8-41.

Figure 8-37: CHAP Users Window

2. Click **Create CHAP User**. The Create CHAP User window appears as shown in Figure 8-38 on page 8-41.

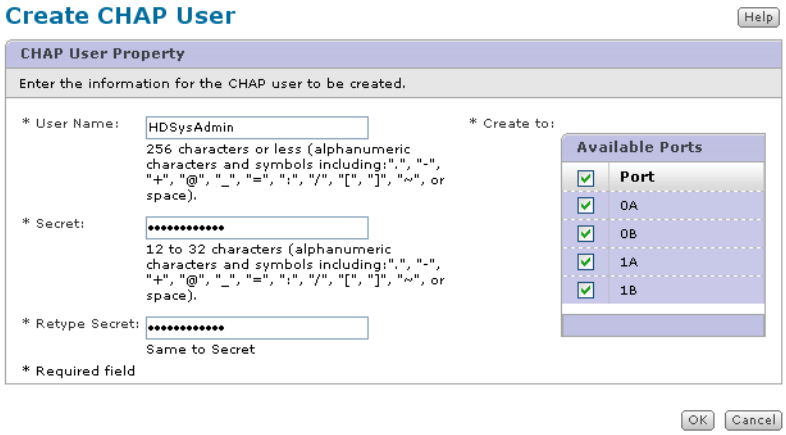

Figure 8-38: Create CHAP User Window

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 8-39 on page 8-42 shows the Edit CHAP User Window.

Figure 8-39: 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**.

9

Performance Monitor

This chapter describes Performance Monitor.

This chapter covers the following topics:

- Performance Monitor overview
- Performance Monitor operations
- Optimizing system performance
- Performance troubleshooting

Performance Monitor overview

Table 9-1 lists the Performance Monitor specifications.

Table 9-1: 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.
Graphic display	Performance Monitor displays acquired information with line graphs. It displays the graph as soon as it acquires the information or displays optional information later after making a choice from the information acquired.

Performance Monitor operations

The procedure for Performance Monitor appears below.

Initial settings

1. Verify that you have the environments and requirements for Performance Monitor (see Preinstallation information on page 2-2).
2. Collect the performance monitoring data (see Obtaining information on page 9-3).

Optional operations

1. Use the graphic displays (see Using graphic displays on page 9-3).
2. Output the performance monitor information to a file (see on page 9-28).
3. Optimize the performance (see Performance troubleshooting on page 9-6).

Optimizing system performance

This section describes how to use Performance Monitor to optimize your system.

Obtaining information

The information is obtained for each controller.

1. Start Navigator 2 and log in. The Arrays window opens
2. Click the appropriate array.
3. Expand the **Settings** list, and click **Advanced Settings**.
4. Click **Open Advanced Settings**. After some minutes, the Array Unit window opens.
5. Click **Performance**. The Performance Monitor window is displayed.
6. Click **Display Graph**.
7. Specify the interval time.
8. Select the items (up to 8) that you want to appear in the graph.
9. 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 **Display 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 **Display Graph**. The Performance Monitor Graph window appears (see Figure 9-1 on page 9-4).

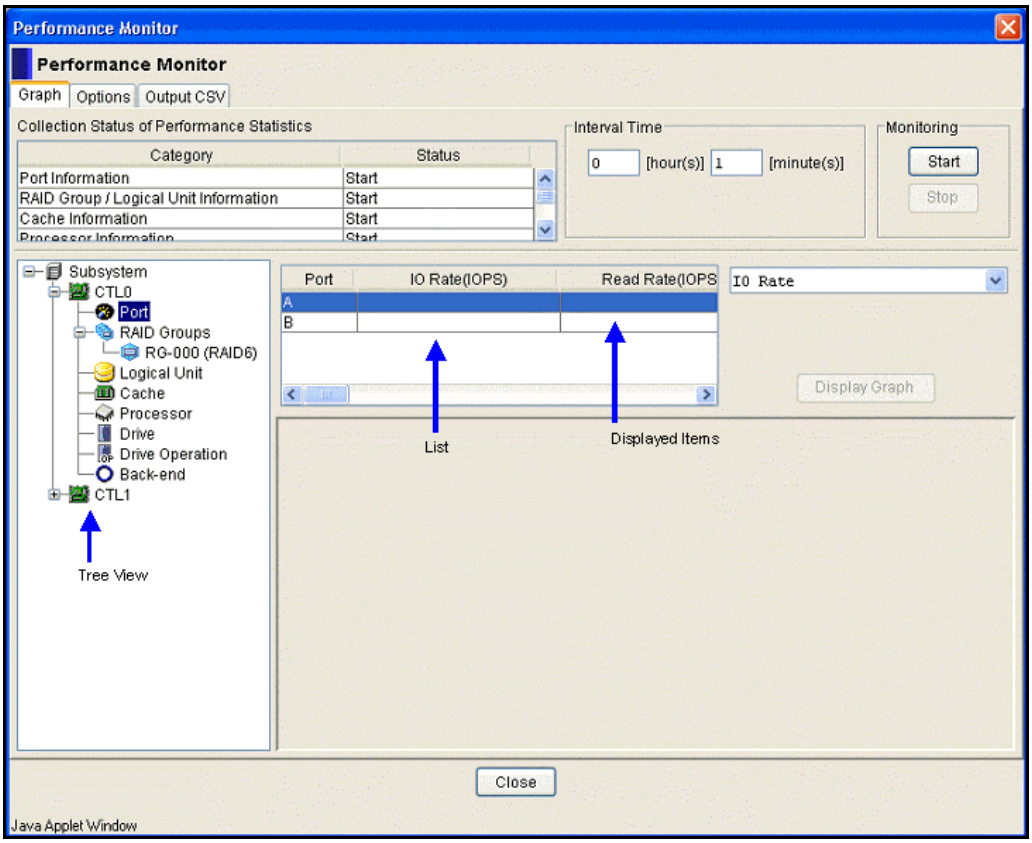

Figure 9-1: Performance Monitor Window — Graph Tab

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 on page 9-28.

An example of a Performance Monitor graph (CPU usage) is shown in Figure 9-2 on page 9-5.

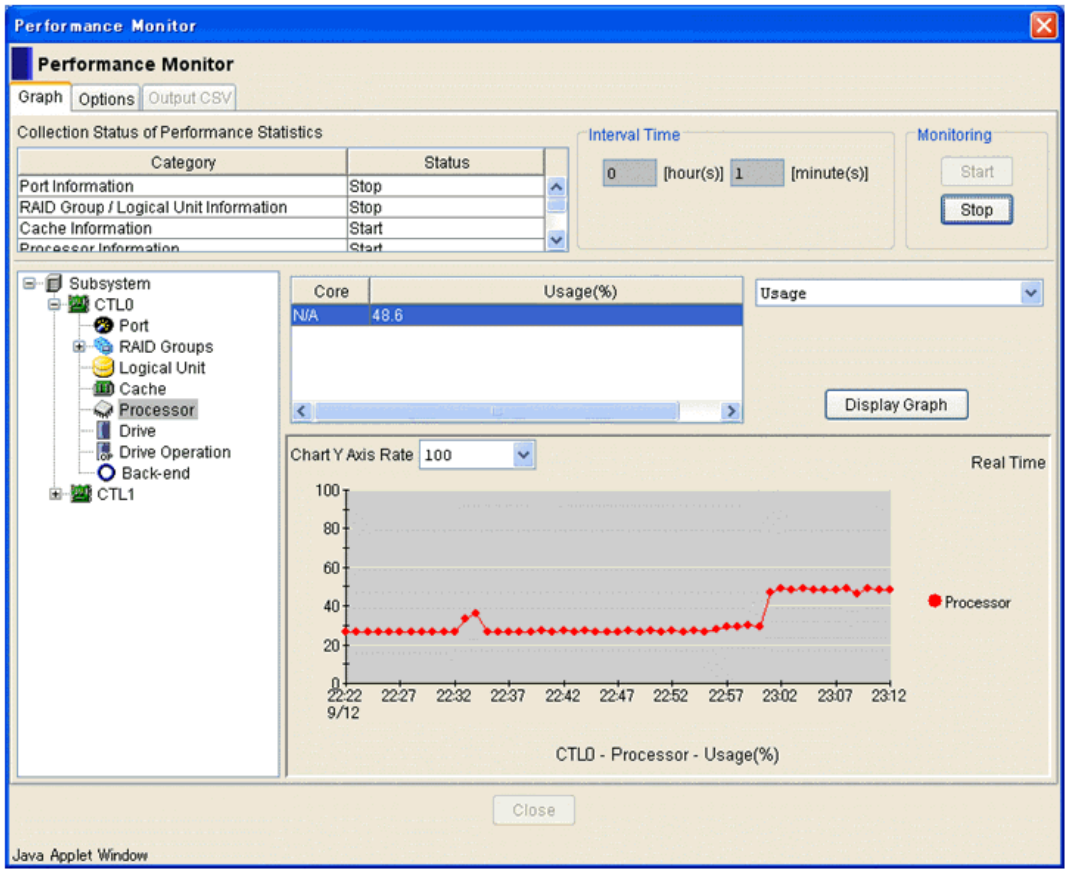

Figure 9-2: Performance Monitor — Sample Graph (CPU Usage)

Table 9-2 on page 9-6 shows the summary of each item in the Performance Monitor.

Table 9-2: 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 9-3 on page 9-6 provides descriptions of Tree View icons.

Table 9-3: Tree View Icons

Icon	Item Name	Description
	Subsystem	Represents the selected storage system. Clicking the icon displays a Tree View of icons belonging to the storage system. Information on this icon is not displayed in the list.

Icon	Item Name	Description
	CTL0/CTL1	Represents the controller on the storage system. 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 SNM2 in the case of the dual controller system, only an icon of the connected controller displays.
	Port	Represents the selected port number on the current storage system. Information on the port displays in the list.
	RAID Groups	Represents RAID groups that have been defined for the current storage system. Information on the RAID groups display in the list.
	RG-000	Represents the logical units that belong to each RAID group defined for the current storage system. Information on the logical units display in the list.
	DP Pools	Represents the Dynamic Provisioning pools that have been defined for the current storage system. Information on the DP pool displays in the list.
	Logical Unit	Represents the logical units defined for the current storage system. Information on the logical units displays in the list.
	Cache	Represents the cache resident in the current storage system. Information on the cache displays in the list.
	Processor	Represents the processor in the current storage system. Information on the processor displays in the list.
	Drive	Represents the disk drive in the current storage system. Information on the drive displays in the list.
	Drive Operation	Represents the drive operation in the current storage system. Information on the drive displays in the list.
	Back-End	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 logical unit 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 9-4 on page 9-8 details items in the Port item.

Table 9-4: 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 9-5 on page 9-9 details items in the RAID Groups DP Pools item.

Table 9-5: 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 9-6 on page 9-10 details items in the Logical Unit, Cache, and Processor items.

Table 9-6: Expanded Tree View of Logical Unit, Cache, and Processor Items

Item	Displayed Items	Description
Logical Unit DP Pool	LUN	Logical unity 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 Logical Unit)	Maximum multiplicity of commands between intervals.
	Tag Average (only Logical Unit)	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.

Item	Displayed Items	Description
	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 9-7 on page 9-11 details items in the Logical Unit, Cache, and Processor items.

Table 9-7: 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.

Item	Displayed Items	Description
	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.

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 9-3 on page 9-12 displays a sample Performance Monitor Window when Dynamic Provisioning is valid.

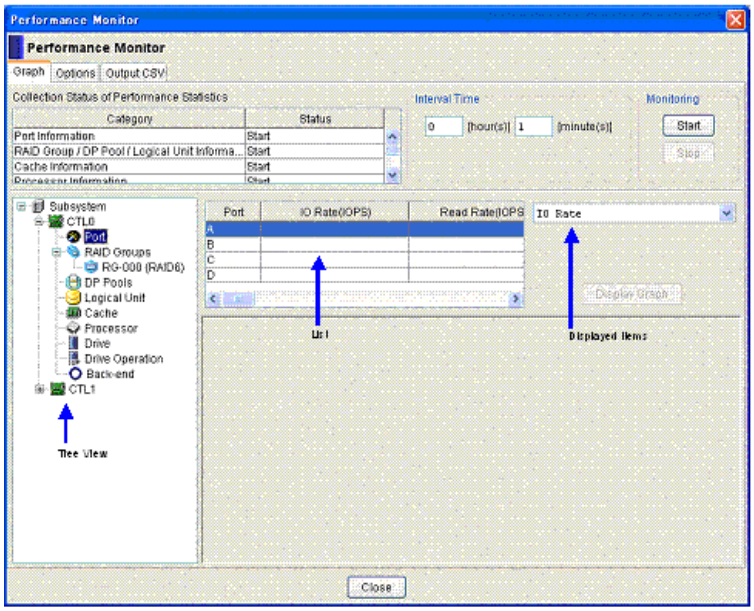

Figure 9-3: Performance Monitor: Dynamic Provisioning is Valid

Working with Graphing and Dynamic Provisioning

The Performance Monitor graph application also behaves differently when Dynamic Provisioning is valid. Figure 9-4 on page 9-13 displays a sample graph when Dynamic Provisioning is valid.

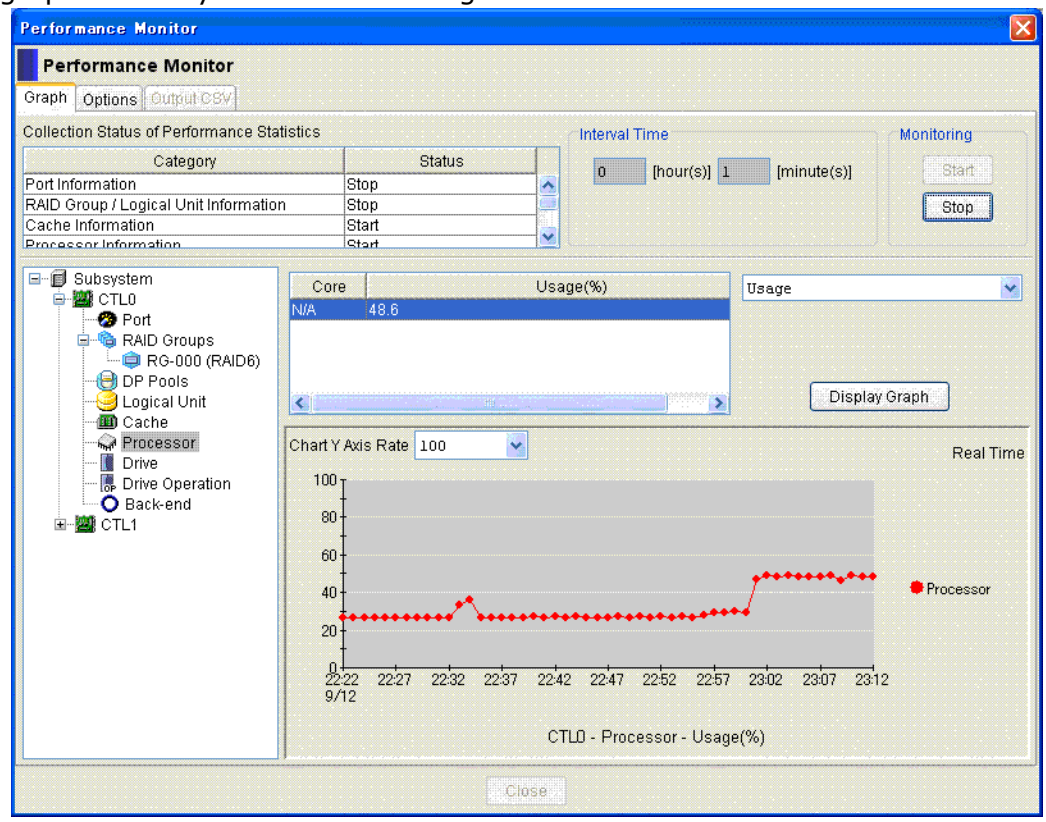

Figure 9-4: Performance Monitor Graph: Dynamic Provisioning Enabled

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

Explanation of Displayed Items

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
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
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	

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 Logical Unit 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 9-8 on page 9-17 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 Logical Units. The bolded values are default settings.

While the table is inclusive to the three object types, Note displayed items for Logical Units 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 9-8: Selectable Y Axis Values for Objects, Port Item

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
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
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	

Displayed Items	Selectable Y Axis Values
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 9-9 details Y axis values for the RAID Groups DP Pools item.

Table 9-9: 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, 5,0000, 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, 5,0000, 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, 100l, 200, 500, 1,000, 2,000
XCOPY Write Trans. Rate	

Table 9-10details Y axis values for the Logical Unit item.

Table 9-10: Selectable Y-Axis Values for Objects, Logical Unit

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
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	

Saving Monitoring Data

To save the settings you changed for Performance Monitor, perform the following steps:

1. Click the Options tab. Performance Monitor displays the Options Window that contains two sub tabs: Output Monitoring Data and Save Monitoring Data as shown in Figure 9-5 on page 9-20.

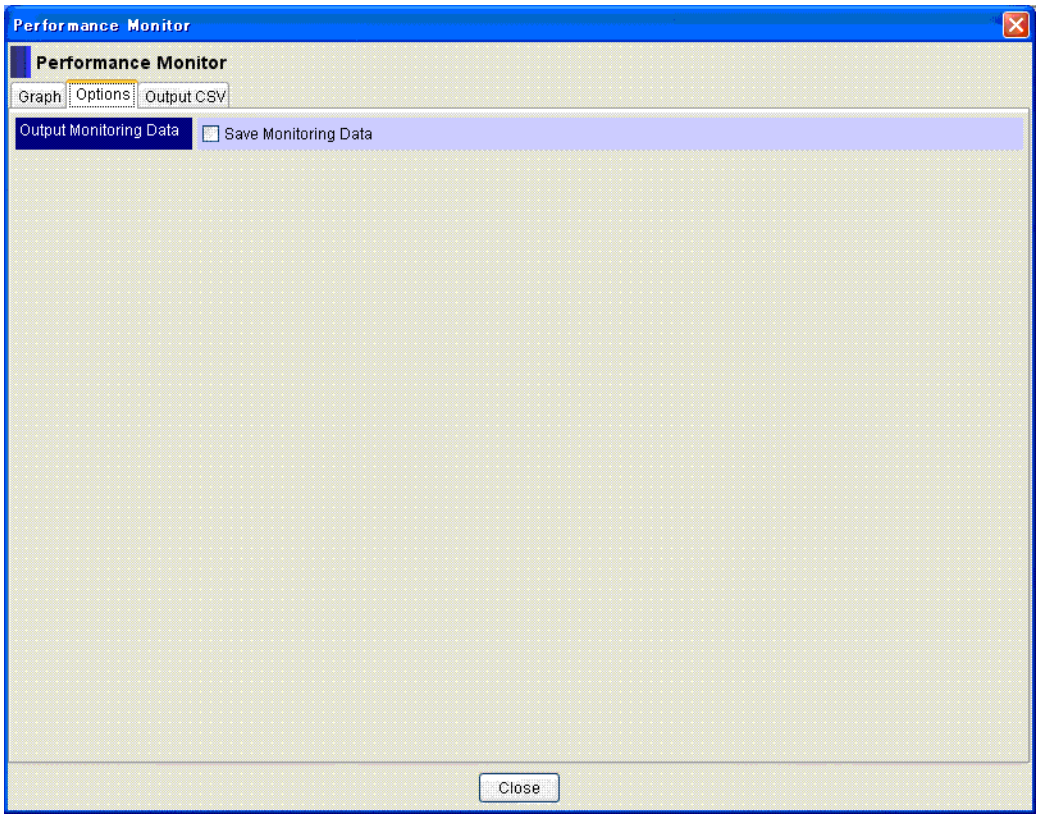

Figure 9-5: Performance Monitor - Save Monitoring Data

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, follow these steps.

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 9-6 on page 9-21.

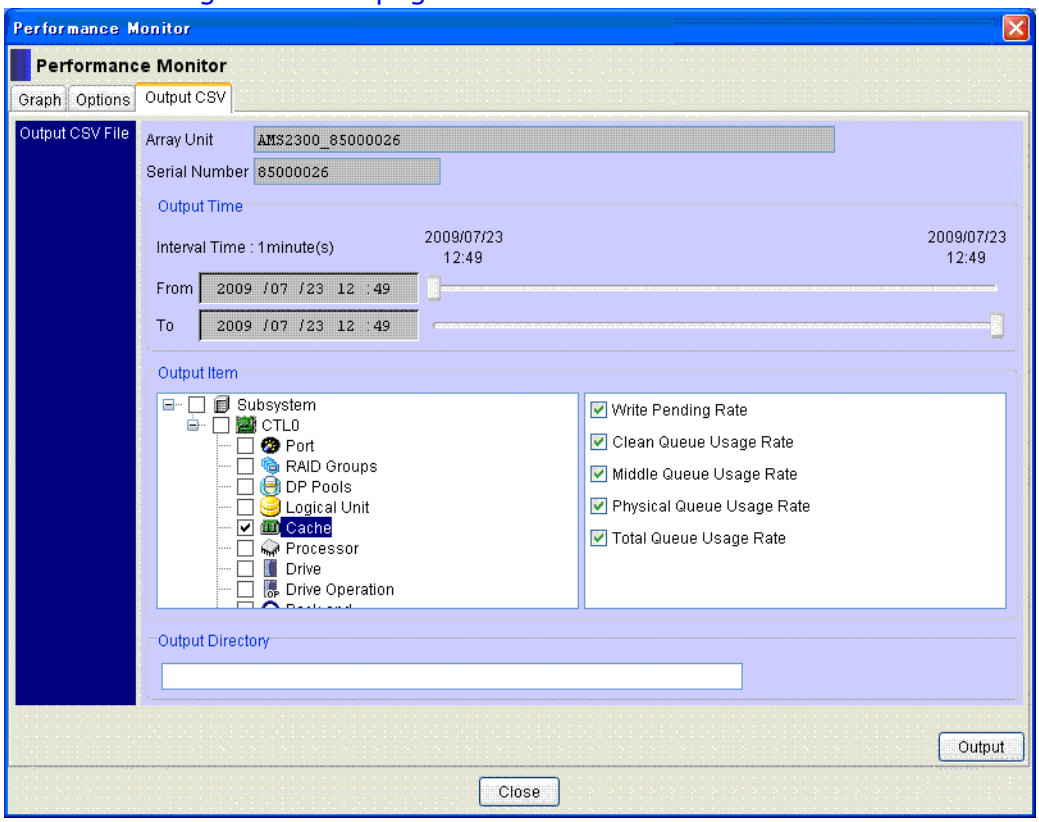

Figure 9-6: Output CSV Tab: Dynamic Provisioning Valid

Figure 9-11 on page 9-21 provides descriptions of objects displayed in the Output CSV Window.

Table 9-11: 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 9-12 on page 9-22 lists filenames for the Port object.

Table 9-12: 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_XcpyTime.csv
XCOPY Max Time	CTL0_Port_XcopyMaxTime.csv
XCOPY Read Rate	CTL0_Port_XcopyReadRate.csv
XCOPY Read Trans. Rate	CTL0_Port_XcopyReadTransRate.csv

Table 9-13 on page 9-22 details CSV filenames for list items for RAID Groups and DP Pool objects.

Table 9-13: CSV Filenames: RAID Groups and DP Pool Objects

Object	List Items	CSV Filename
RAID Groups	IO Rate	CTL0_Rg_IORate*nn*.csv
	Read Rate	CTL0_Rg_ReadRate*nn*.csv
	Write Rate	CTL0_Rg_WriteRate*nn*.csv
	Read Hit	CTL0_Rg_ReadHit*nn*.csv
	Write Hit	CTL0_Rg_WriteHit*nn*.csv
	Trans. Rate	CTL0_Rg_TransRate*nn*.csv
	Read Trans. Rate	CTL0_Rg_ReadTransRate*nn*.csv
	Write Trans. Rate	CTL0_Rg_WriteTransRate*nn*.csv
DP Pools	IO Rate	CTL0_DPPool_IORate*nn*.csv
	Read Rate	CTL0_DPPool_ReadRate*nn*.csv
	Write Rate	CTL0_DPPool_WriteRate*nn*.csv
	Read Hit	CTL0_DPPool_ReadHit*nn*.csv

Object	List Items	CSV Filename
	Write Hit	CTL0_DPPool_WriteHit*nn*.csv
	Trans. Rate	CTL0_DPPool_TransRate*nn*.csv
	Read Trans. Rate	CTL0_DPPool_ReadTransRate*nn*.csv
	Write Trans. Rate	CTL0_DPPool_WriteTransRate*nn*.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 9-14 on page 9-23details CSV filenames for list items associated with Logical Units and Processor objects.

Table 9-14: CSV Filenames: Logical Units and Processor Objects

Object	List Items	CSV Filename
Logical Unit	IO Rate	CTL0_Lu_IORate*nn*.csv
	Read Rate	CTL0_Lu_ReadRate*nn*.csv
	Write Rate	CTL0_Lu_WriteRate*nn*.csv
	Read Hit	CTL0_Lu_ReadHit*nn*.csv
	Write Hit	CTL0_Lu_WriteHit*nn*.csv
	Trans. Rate	CTL0_Lu_TransRate*nn*.csv
	Read Trans. Rate	CTL0_Lu_ReadTransRate*nn*.csv
	Write Trans. Rate	CTL0_Lu_WriteTransRate*nn*.csv
	CTL CMD IO Rate	CTL0_Lu_CTL_CMD_IORate*nn*.csv
	Data CMD IO Rate	CTL0_Lu_CMD_TransRate*nn*.csv
	CTL CMD Trans. Rate	CTL0_Lu_CTL_CMD_TransRate*nn*.csv
	Data CMD Trans. Rate	CTL0_Lu_data_CMD_Trans_Time*nn*.csv
	XCOPY Rate	CTL0_Lu_XcopyRate*nn*.csv
	XCOPY Time	CTL0_Lu_XcopyTime*nn*.csv
	XCOPY Max Time	CTL0_Lu_XcopyMaxTime*nn*.csv
	XCOPY Read Rate	CTL0_Lu_XcopyReadRate*nn*.csv
	XCOPY Read Trans. Rate	CTL0_Lu_XcopyReadTransRate*nn*.csv
	XCOPY Write Rate	CTL0_LuXcopyWriteRate*nn*.csv
	XCOPY Write Trans. Rate	CTL0_Lu_XcopyWriteTransRate*nn*.csv
Processor	Usage	CTL0_Processor_Usage.csv

Table 9-15 on page 9-24 details CSV filenames for list items associated with Cache, Drive, and Drive Operation objects.

Table 9-15: 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_IORate*nn*.csv
	Read Rate	CTL0_Drive_ReadRate*nn*.csv
	Write Rate	CTL0_Drive_WriteRate*nn*.csv
	Trans. Rate	CTL0_Drive_TransRate*nn*.csv
	Read Trans. Rate	CTL0_Drive_ReadTransRate*nn*.csv
	Write Trans. Rate	CTL0_Drive_WriteTransRate*nn*.csv
	Online Verify Rate	CTL0_Drive_OnlineVerifyRate*nn*.csv
Drive Operation	Operating Rate	CTL0_DriveOpe_OperatingRate*nn*.csv
	Max Tag Count	CTL0_DriveOpe_MaxtagCount*nn*.csv

Enabling Performance Measuring Items

The Performance Measuring tool enables you to enable specific types of performance monitoring. To access the Performance Measuring tool, perform the following steps.

1. Start Navigator 2 and log in. The Arrays window opens
2. Click the appropriate array.
3. Expand the **Settings** list, and click **Advanced Settings**.
4. Click **Open Advanced Settings**. After some minutes, the Array Unit window opens.
5. Click **Performance**. The Performance Monitor window is displayed.
6. Click **Monitoring**. The Monitoring - Performance Measurement Items window displays as shown in Figure 9-7 on page 9-25.

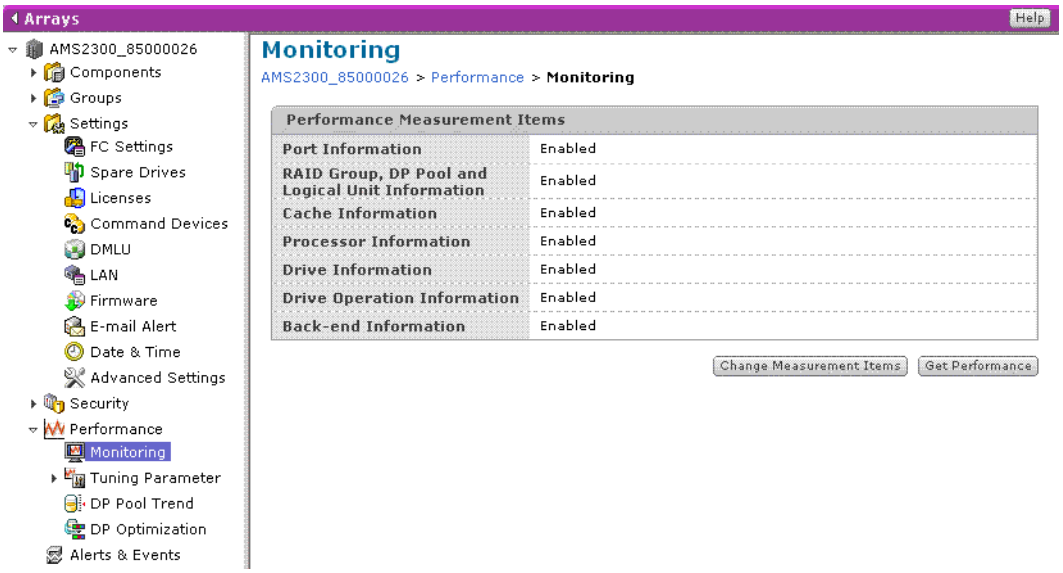

Figure 9-7: Monitoring - Performance Measurement Items

7. Click on the Change Measurement Name Button. The Change Measurement Items dialog box displays with six performance statistics as shown in Figure 9-8 on page 9-25.

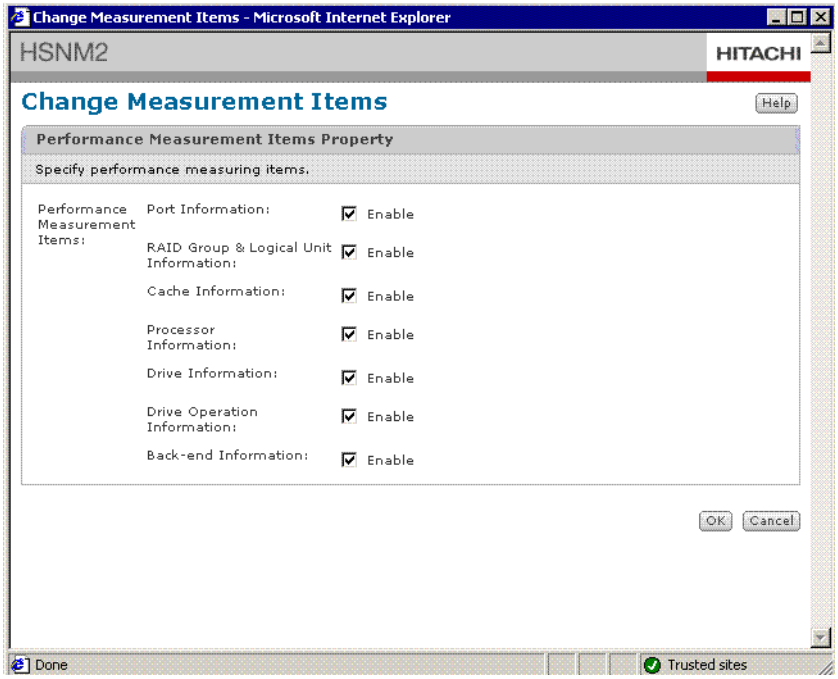

Figure 9-8: Change Measurement Items dialog box

describes each of the performance statistics.

Table 9-16: Performance Statistics

Item	Description
Port Information	Displays information about the port.
RAID Group, DP Pool and Logical Unit Information	Displays information about RAID groups, Dynamic provisioning pools and logical units.
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.

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.

8. 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 Logical Unit Information

The storage system acquires all array RAID group/DP pool information of logical units. 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 ache-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 4-42 for information on how to analyze the problem.

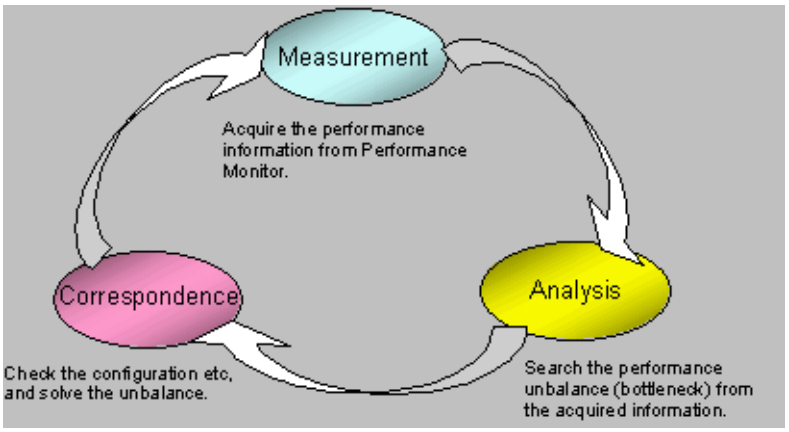

Figure 9-9: Performance Optimization Analysis

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 logical unit load can be obtained from the I/O and transfer rate of each logical unit.

When the loads between controllers differ considerably, the array disperses the loads (load balancing). However, when this does not work, change the logical unit 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 logical unit 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 logical unit 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 logical unit 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 logical unit; however controller imbalance can occur.

10

SNMP Agent Support

This chapter describes Simple Network Management Protocol (SNMP) Agent support.

This chapter covers the following topics:

- SNMP Agent Support overview
- SNMP functions
- SNMP Agent Support operations
- Managing SNMP Agent Support

SNMP Agent Support overview

A workstation with SNMP agent support is required on a LAN.

Figure 10-1 shows a private LAN connection.

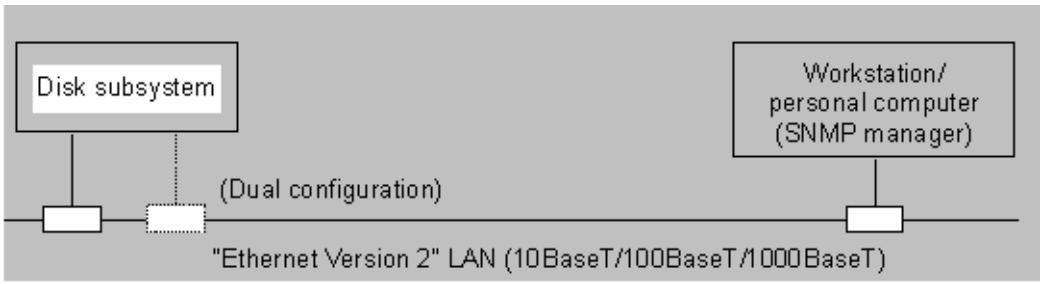

Figure 10-1: Private LAN Connection

Figure 10-2 shows a public LAN connection. One Gateway address (default Gateway address) can be set for each controller.

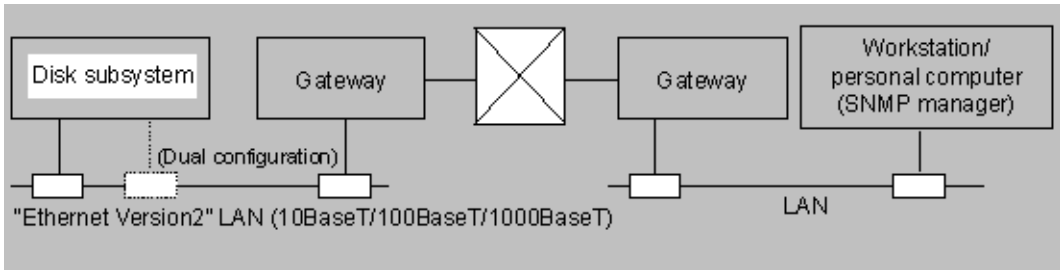

Figure 10-2: Public LAN Connection

Table 10-1: Operations

Item	Description
GET	Obtains a MIB object value. Normal operation is assumed when GET REQUEST (request from the SNMP manager) and GET RESPONSE (response from the agent) are completed.
GETNEXT	Continuously searches MIB objects. Normal operation is assumed when GETNEXT REQUEST (request from the SNMP manager) and GET RESPONSE (response from the agent) are completed.
TRAP	Reports an event (error or status change) to the SNMP manager. When an event occurs, the agent sends a TRAP to the manager.

Figure 10-3 on page 10-3 shows communications between the SNMP manager and the SNMP agent for a supported SNMP operation.

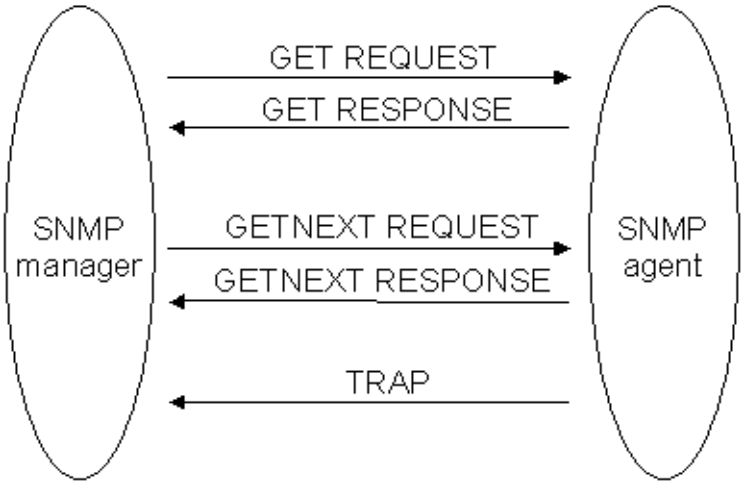

Figure 10-3: SNMP Communication

Error status

When an error is detected, the array sends an SNMP message (GET RESPONSE) to the manager, with the error status, as shown in Table 10-2 on page 10-3.

If any of these errors are detected in the SNMP manager's request, the array does not respond.

- The community name does not match the setting. The array does not respond; however, it sends a 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 longer than 484 bytes.

Table 10-2: Error Status

Error Status	Description
noError (0)	No error detected. The requested MIB object value is placed in the SNMP message to be sent.
tooBig (1)	The SNMP message is too large (more than 484 bytes) to contain the operation result.
noSuchName (2)	The requested MIB object could not be found. The GETNEXT REQUEST was received. The requested MIB object value is not set in the SNMP message. The requested process (SET REQUEST) is not executed.
badValue (3)	Does not occur.
readOnly (4)	Does not occur.
genErr (5)	The operation cannot be executed.

Dual controller GET/TRAP specifications

The GET/TRAP specifications for dual system configuration are shown in Table 10-3.

Table 10-3: GET/TRAP Specifications—AMS 2000 Family

Connection status	Controller status	GET/TRAP specification				Remarks
		Controller 0		Controller 1		
Both controller	① Both controllers are normal	GET	○	GET	○	Master controller: 0
		TRAP	○	TRAP	Δ	
	②Controller 1 is blockaded	GET	○	GET	×	Master controller: 0
		TRAP	○	TRAP	×	If controller 1 is recovered, the system goes to ①.
	③110 Controller 0 is blockaded	GET	×	GET	○	Master controller: 1
		TRAP	×	TRAP	○	
	④ Controller 0 is recovered (the board was replaced while the power is on)	GET	○	GET	○	Master controller: 1
		TRAP	Δ	TRAP	○	The system goes to ① when restarted (P/S ON).
Controller 0 only	⑤Both controllers are normal	GET	○	GET	×	Master controller: 0
		TRAP	○	TRAP	×	
	⑥ Controller 1 is blockaded	GET	○	GET	×	
		TRAP	○	TRAP	×	
	⑦ Controller 0 is blockaded	GET	×	GET	×	Master controller: 1
		TRAP	×	TRAP	×	
	⑧ Controller 0 is recovered (the board was replaced while the power is on)	GET	○	GET	×	Master controller: 1
		TRAP	Δ	TRAP	×	The system goes to ⑤ when restarted (P/S ON).

○: GET and TRAP are possible. (The drive blockade and the occurrence detected by the other controller is excluded.)
×: GET and TRAP are impossible.
Δ: A trap is reported only for an own controller blockade, and a drive blockade (drive extraction is not included) detected by the own controller.

Note: A trap is reported for an error that has been detected when a controller board is replaced while the power is on or the power is turned on. Therefore, traps other than the above are also reported.

For a dual system configuration, SNMP managers should be divided as shown in Figure 10-4 on page 10-5.

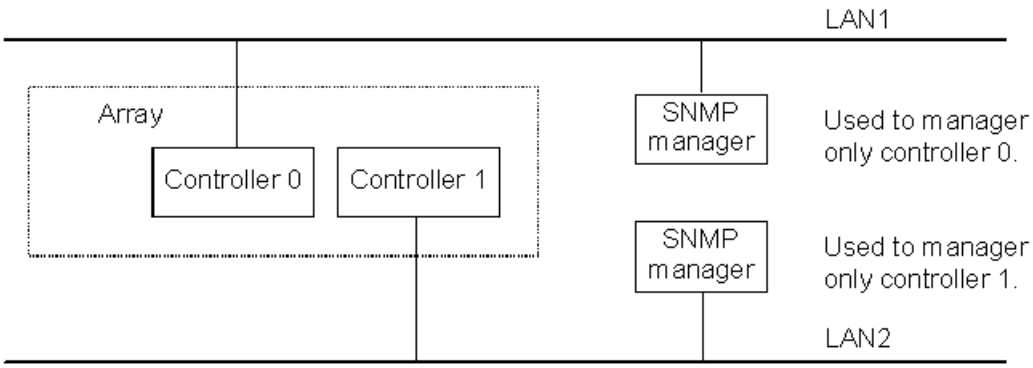

Figure 10-4: Divided SNMP Managers

Only the master controller reports TRAPs for fan, power supply, and battery failures. Even though you can specify more than one SNMP manager, each SNMP manager should be set so that it controls both controllers.

- A device that executes broadcast, etc. should not be connected to the LAN where the array is connected, or host command processing deteriorates.
- If the IP address of the SNMP manager is changed when the DHCP function is used, the TRAP cannot be reported.
- If you change the array IP Address of the disk subsystem is changed, restart the array.
- Contact service personnel if a failure occurs.

SNMP functions

The following sections provide information on the functions that report disk array failures to the SNMP manager.

TRAP reporting

The following events are detected through TRAPs.

Standard TRAPs

- Power supply power on
- SNMP access error (incorrect community name)
- Startup of the SNMP Agent Support Function (occurs when installing or enabling SNMP Agent Support Function)
- Changing the settings of the SNMP Agent Support Function
- Incorrect community string name given when acquiring MIB information

Extended TRAPs

The following list identifies extended TRAPs.

- Unrecoverable data (multiple failures of drives) (see Note 2)
- Blocked path, which is reported only if the TrueCopy remote replication feature is enabled.
- Additional battery failure
- Battery failure
- Cache backup circuit failure
- Cache memory failure
- Cycle time threshold over
- Data pool no free
- Data pool threshold over
- Drive blocking (data drive)
- Enclosure controller failure
- Failure (Modular Volume Migration)
- Failure (ShadowImage)
- Failure (SnapShot)
- Failure (TCE)
- Failure (TrueCopy)
- Fan failure
- Firmware replacement executed
- Host connector failure
- Interface board failure
- Own controller failure (see Note 1 and Note 2)
- Path blockade (see Note 4)
- Port error threshold over
- Power supply failure
- Slave controller failure (see Note 2)
- Spare drive failure
- UPS failure
- Warning disk array (see Note 3)

NOTE: When a controller blockade occurs, the array issues TRAPs that show the blockade. The controller can automatically recover, depending on the failure.

NOTE: The TRAP that shows the array warning status can appear through preventive maintenance, periodic part replacement, or service personnel fieldwork.

NOTE: 3. Path blockade is reported only when the TrueCopy or TCE feature is enabled.

NOTE: A TRAP is issued if multiple hard disk failures occur in drives and the data in the RAID group and logical units are not recoverable. For example, a TRAP is issued if failures occur in three hard disk units under RAID 6.

Table 10-4 details specific trap codes for extended traps.

Table 10-4: Supported Extended Traps

No.	Specific Trap Code	Meaning
1	11	warning (see note below).
2	252	repairDriveRequest
3	253	saSubsystemReplacement
4	308	luFailure
5	316	Port Error Threshold

NOTE: The warning status of the array can be automatically set in the warning information by preventive maintenance, periodic part replacement or a fieldwork of the service.

Request processing

Request processing enables the SNMP manager to refer to MIB objects (the function to set MIB objects is not provided).

The information that appears is as follows:

- Device information (product name and firmware revision)
- Warnings (see below)
- Command execution condition information

The warning information that can be acquired by the array is shown below.

- Additional battery failure
- Array warning (see Note 2)
- Battery failure
- Cache backup circuit failure
- Cache memory failure
- Data pool is over threshold
- Data pool is unavailable
- Drive blockade (data or spare drive)

- Enclosure controller failure
- Fan failure
- Host connector failure
- Interface board failure
- Power supply failure
- PSUE (Modular Volume Migration)
- PSUE (ShadowImage)
- PSUE (SnapShot)
- Slave controller failure (see Note 1)
- UPS failure

NOTE: 1. When the other controller is blocked, the blockade is set in the warning information. However, the controller blockade may recover automatically depending on the cause of the failure.

NOTE: 2. The array warning status can appear through preventive maintenance, periodic part replacement, or service personnel fieldwork.

Additional SNMP environment requirements

- Firmware version 0862 or newer is required if using IPv6 operating under an IPv6 environment for AMS 2000 Family storage systems
- Hitachi Storage Navigator Modular 2 requires v6.20 or newer if operating under an IPv6 environment
- Warnings (see below)
- The warning information that can be acquired by the array is shown below.

SNMP Agent Support operations

The procedure for SNMP Agent Support appears below.

1. Verify that you have the environments and requirements for SNMP Agent Support (see Preinstallation information on page 2-2).
2. Set the `config.txt` and `name.txt` files (see Creating environmental information files on page 10-10).
3. Set the MIB information (see Settings on page 10-13).
4. Confirm the TRAP and REQUEST connection (see Verifying SNMP connections on page 10-16).

Managing SNMP Agent Support

This section describes how to manage SNMP Agent functions, including:

- SNMP setup
- Creating SNMP Environmental file information
- Registering SNMP Environment information
- Referencing SNMP Environment information
- Verifying the SNMP connection
- Detecting failures

SNMP setup

The setup procedure includes setting up the array, the SNMP manager, and verifying the connection. This enables communication between the array and SNMP manager.

Disk array-side setup

1. Specify the LAN information (IP Address, Sub Net Mask, and Default Gateway Address). For more information, see the SNM2 Online Help.
2. Make sure the SNMP Agent Support license is installed and enabled. For more information, see Requirements for installing and enabling features on page 1-17.
3. Create the SNMP environment information file. The SNMP environment file consists of the following.
   - The operating environment setting file (`Config.txt`). This sets the IP address and community of the SNMP manager to send TRAPs, and so forth. (The Community name described in the `config.txt` file in the provided CD.)
   - The array name setting file (`Name.txt`). (Sets the array names.)

   For more information, see Creating environmental information files on page 10-10.
4. Register the SNMP environment information file in the array. For more information, see Registering SNMP environmental information on page 10-14.

SNMP Manager-side setup

1. Transfer the MIB definition file into the SNMP manager.
2. Register the array in the SNMP manager. Refer to the manuals of the SNMP manager for operating procedures.

Checking the connection

Verify the connection between the array and SNMP manager (see Verifying SNMP connections on page 10-16).

Creating environmental information files

To use the SNMP agent, the SNMP environment information file is created and registered in the array. The following files are created as the SNMP environment information file:

- Operation environment setting file (`Config.txt`).
- Array name setting file (`Name.txt`).

The SNMP environment information file is created and registered at the SNMP initial setting and when an operating environment is changed.

Only one set (two files) per array is created in dual controller configurations. Therefore, you cannot set different information for each controller.

Environment setting file

This section describes how to create environment setting files.

Format file

This file is in text form on a CD. The file name is `Config.txt`.

Settings

There are four basic settings that can be made to the environment setting file:

NOTE: TRAP sending is the only setting that is required.

1. sysContact (MIB information).

   Manager information for the contact (name, department, extension No., etc.). This is an internal object value of the MIB-II system group in ASCII form, up to 255 characters.

2. sysLocation (MIB information).

   Where where the device is installed. This is an internal object value of the MIB-II system group in ASCII form, up to 255 characters.

3. Community information setting (MIB information).

   Name of the community permitted access. Multiple community names can be set.

4. TRAP sending (TRAP report).

   Required settings used to send a TRAP:

   - Destination manager IP address
   - Destination port number
   - Community name for the TRAP. Multiple information combinations can be set.

Creating files

Use the following procedure to set each numbered item from the previous section (**sysContact**, **sysLocation**, and so forth):

❒ Setting sysContact (manager's name/items for contact)

- Add a line beginning with "INITIAL" to the file to set the sysContact value:

  ```
  INITIAL sysContact <user set information>
  ```

- User set information cannot exceed 255 alphanumeric characters.
- For any characters (space, tab, "-", """, etc.) other than the letters a to z (uppercase and lowercase) and the numerals are used by the user to set information, the characters must be enclosed with double quotation marks (").
- Do no include line-feed symbols.
- If you are not setting the sysContact value, you can leave the value with closed quote as follows:

  ```
  INITIAL <sysContact "">
  ```

  Or...you can delete this line from your configuration file.

❒ Setting sysLocation (array installation location):

- Add a line beginning with "INITIAL" to the file to set the sysLocation value:

  ```
  INITIAL <sysLocation user set information>
  ```

- User set information cannot exceed 255 alphanumeric characters.
- For any characters (space, tab, "-", """, etc.) other than the letters a to z (uppercase and lowercase) and the numerals are used by the user to set information, the characters must be enclosed with double quotation marks (").
- Do no include line-feed symbols.
- If you are not setting the sysLocation value, you can leave the value with closed quote as follows:

  ```
  INITIAL <sysLocation "">
  ```

  Or...you can delete this line from your configuration file.

❒ Setting community information:

- Add a line that begins with **COMMUNITY** in the file to specify the community string that allows the array to receive requests as shown:

  ```
  COMMUNITY <community string>
  ALLOW ALL OPERATIONS
  ```

  NOTE: The array will accept all community string (names) if you do not define this parameter.

- The community string must be described in alphanumeric characters only.

- If any characters (space, tab, "-", "\", etc.) other than the letters a to z (uppercase and lowercase) and the numerals are used in the community name, the characters must be enclosed with double quotation marks ("). The community name cannot contain line-feed codes.
- To accept all community string (names), delete the above 2 lines including the line that begins with **COMMUNITY**.

❒ Setting Addresses for Receiving TRAPs:

NOTE: Multiple addresses may be configured and up to 3 SNMP manager may be configured.

- Add a line that begins with **MANAGER** in the file to specify the SNMP manager that receives TRAP requests from the array as shown:

```
MANAGER <SNMP manager IP address>
SEND ALL TRAPS TO PORT <port number>
WITH COMMUNITY <community string name>
```

- Enter the IP address to select the object SNMP manager. Do not specify a host name.
- Enter IP addresses with the leading 0's in each dotted quad suppressed (for example, specify 111.22.3.55 for 111.022.003.055).
- Enter the UDP destination port number to be set when sending a TRAP to the SNMP manager. The port number 162 is the usual port number used by the SNMP manager to receive TRAPs.
- For the Community string name, a community name, which is set in an SNMP message when sending a TRAP, is specified with alpha numerics. If any characters (space, tab, "-", "\", etc.) other than the letters a to z (uppercase and lowercase) and the numerals are used in the community name, enclose them with double quotation marks (").
- Do no include line-feed symbols. If the community name does not contain a close (the line that begins **WITH COMMUNITY**), add **public** to the Community string name.

NOTES:
This file cannot exceed 1,140 bytes. Also, the total length of "sysContact", "sysLocation", and "sysName" (to be explained later) should not exceed 280 characters (when the community that has access rights does not exceed 10 characters) so that all the objects in the MIB-II system group can be obtained with the one `GET` request. This prevents a "tooBig" error message. Refer to the following example.

Operation Environment Setting File example for IPv4:

```
INITIAL sysContact "System Admin"


INITIAL sysLocation "Computer Room A on Hitachi 10F north"


COMMUNITY tagmastore
ALLOW ALL OPERATIONS


MANAGER 123.45.67.89
SEND ALL TRAPS TO PORT 162
WITH COMMUNITY "HITACHI"
```

Operation Environment Setting File example for IPv6:

```
INITIAL sysContact "System Admin"


INITIAL sysLocation "Computer Room A on Hitachi 10F north"


COMMUNITY tagmastore
ALLOW ALL OPERATIONS


MANAGER 2001::1::20a:87ff:fec6:1928
SEND ALL TRAPS TO PORT 162
WITH COMMUNITY "HITACHI"
```

Array name setting file

This section contains the following:

- File format
- Settings
- Creating files

File format

1. This file should be in text format on a DOS-formatted disk. The file name is `Name.txt`.

Settings

Set the following for the disk array name:

- **sysName**

  This is the name of the array that is being managed.

NOTE: The internal object value of MIB-II system group in ASCII character string cannot exceed 255 characters.

Creating files

To set the value of `sysName`, register the information continuously. Since the entire contents of this file are regarded as the `sysName` value, the file should not exceed 255 characters.

- Do no include line-feed symbols.
- Use only alphanumeric characters. For example:

```
Hitachi Disk Array
```

NOTE: The total length of "sysContact", "sysLocation", and "sysName" (to be explained later) should not exceed 280 characters (when the community that has access rights does not exceed 10 characters) so that all the objects in the MIB-II system group can be obtained with the one `GET` request. This prevents a "tooBig" error message. Refer to the following example.

Registering SNMP environmental information

Follow these steps to register the SNMP environment information file:

1. Start Navigator 2 and log in. The Arrays window appears
2. Click the appropriate array.
3. Select **SNMP Agent** in the Settings tree view. The SNMP Agent window appears.
4. Click **Edit SNMP Settings**. The Edit SNMP Settings window is displayed (Figure 10-5).

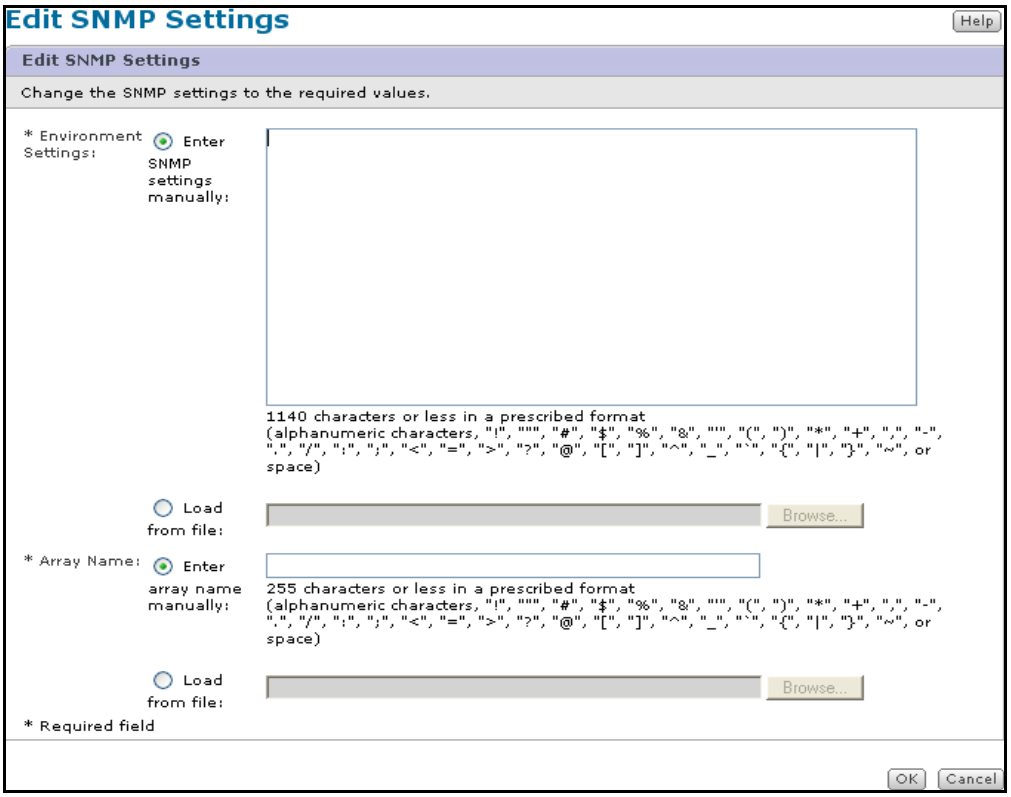

Figure 10-5: Edit SNMP Settings Files

5. Select either **Enter SNMP settings manually** or load the setting values by clicking the **Load from file** option and browsing to the file location.
   - When you select the Enter SNMP settings manually, enter it directly in the screen according to Environment setting file on page 10-10.
   - When you select **Load from file:**, specify a path to the SNMP environmental information file (`config.txt`) for the Load from. You can also specify the path in which the SNMP environmental information file is stored using the Browse button.
   - Select **Enter Array Name manually** or **Load from file:** to set the array name. When you select Enter array name manually, enter it directly in the screen according to Array name setting file on page 10-13.
   - When you select **Load from file:**, specify a path to the array name setting file (`name.txt`) for the Load from. You can also specify the path in which the array name setting file is stored using the Browse button.
   - If only one file is set, specify only a file to set.
6. Click **OK**. A confirmation message is appears.
7. Click **Close**.

Referencing the SNMP environment information file

You can view the current SNMP environment information for the SNMP agent.

Follow these steps to register the SNMP environment information file.

1. Start Navigator 2 and log in. The Arrays window appears
2. Click the appropriate array.
3. Select **SNMP Agent** in the Settings tree view. The SNMP Agent window is displayed. See Figure 10-6 on page 10-16.

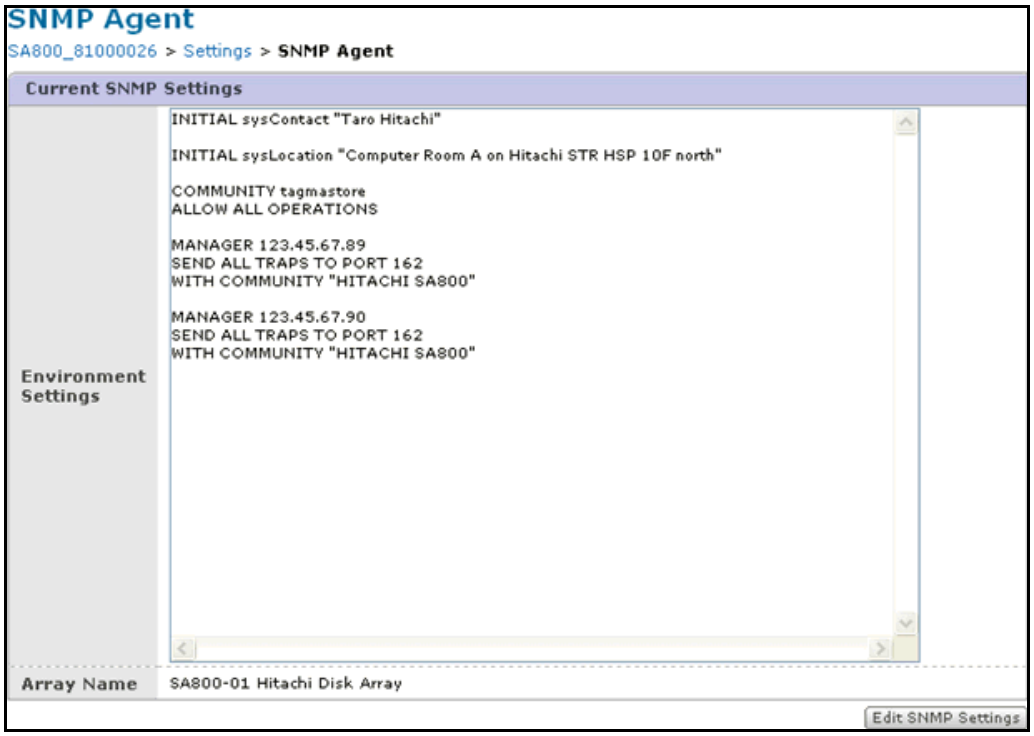

Figure 10-6: SNMP Agent Window — Current SNMP Settings

4. You may copy the information and store it in a file for saving on the local host.

Verifying SNMP connections

This section describes how to verify SNMP connections between the array and the SNMP manager.

Checking TRAP connections

Set the SNMP agent to invalid, and then to valid. Check that a standard TRAP, "warmStart," has been received by the SNMP managers that are set as TRAP receivers in the SNMP environment information file (`Config.txt`).

Checking REQUEST connections

Send a **MIB GET** request to the array from all the SNMP managers to be connected to the array. Verify that the array responds.

Detecting failures

This section describes how to detect array failures.

Obtaining information

Obtain MIB information (`dfRegressionStatus)` periodically. This MIB value is set to **0** when there are no failures.

Reporting errors

If an error results in a TRAP, the array reports the error to the SNMP manager. This TRAP allows you to detect array failures when they occur. However the UDP protocol can incorrectly report the TRAP to the SNMP manager. If a controller goes down, the `systemDown` TRAP may not be issued.

Detecting errors

Errors are detected with the MIB information. Even if the TRAP is not reported, you can detect errors if the MIB value (`dfRegressionStatus`) is not **0**.

For example, if a drive is blocked the `dfRegressionStatus = 69`.

When continuous requests are not responding, it is likely because of a controller blockade.

11

Modular Volume Migration

This chapter describes Modular Volume Migration.

This chapter covers the following topics:

- Modular Volume Migration overview
- Modular Volume Migration operations
- Managing Modular Volume Migration

Modular Volume Migration overview

Figure 11-2 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 (AMS2100) • 2,047 (AMS2300 and AMS 2500) **Note:** The maximum number of the pairs is limited when using ShadowImage. For more information, see Using with ShadowImage on page 11-12.
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 11-12.
Number of reserved logical units	• 1,023 (AMS2100) • 2,047 (AMS2300 and AMS2500)
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. RAID 0 cannot be set for the SATA disk drive.
RAID level combinations	All combinations are supported.
Types of P-VOL/S-VOL drives	Logical units consisting of SAS and SATA drives can be assigned to any P-VOLs and S-VOLs. You can specify a logical unit consisting of SAS drives, and a logical unit consisting of SATA 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 logical units	You cannot delete logical units or RAID groups while they are being migrated.
Handling of logical units	You cannot format, delete, expand, or reduce logical units 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 logical unit as a P-VOL or an S-VOL while it is being formatted. Execute the migration after the formatting is completed.
Logical unit restrictions	Data pool LU, DMLU, and command devices (CCI) cannot be specified as a P-VOL or an S-VOL.
Concurrent use of unified logical units	The unified logical units migrate after the unification. Using unified logical units on page 11-9.

Table 11-1: Volume Migration Specifications (Continued)

Item	Description
Concurrent use of Data Retention	When the access attribute is not Read/Write, the logical unit cannot be specified as an S-VOL. The logical unit 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 11-11.
Concurrent use of SNMP Agent	Available
Concurrent use of LUN Manager	Available
Concurrent use of Cache Residency Manager	The Cache Residency logical unit cannot be set to P-VOL or S-VOL.
Concurrent use of Cache Partition Manager	Available. Note that a logical unit 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	When a P-VOL or an S-VOL is included in a RAID group for which the Power Saving has been specified, you cannot use Volume Migration.
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 logical unit (V-VOL) is defined.
Concurrent use of TrueCopy	A P-VOL and an S-VOL of TrueCopy cannot be specified as a P-VOL or an S-VOL of Volume Migration unless their pair status is Simplex.
Concurrent use of TCE	A P-VOL and an S-VOL of TrueCopy 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. However, the DP-VOLs created by Dynamic Provisioning cannot be used. The normal LU can be set as a P-VOL, an S-VOL, or reserve LU.
Failures	The migration fails if the copying from the P-VOL to the S-VOL stops. The migration also fails when a logical unit 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: Reserved Logical Unit 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	P-VOL or S-VOL of TrueCopy
Concurrent use of TCE	P-VOL or S-VOL of TCE
Concurrent use of Data Retention	Data Retention logical unit.
Concurrent use of Dynamic Provisioning	The DP-VOLs created by Dynamic Provisioning
Logical unit restrictions for special uses	Data pool LU, DMLU, command device (CCI).
Other	Unformatted logical unit. However, a logical unit being formatted can be set as reserved even though the formatting is not completed.

Environments and Requirements

Table 11-3 shows environments and requirements.

Table 11-3: Environments and Requirements

Item	Description
Environments	Firmware Version 0832/B or later is required for AMS2100 or AMS2300 array of the hardware revision 0100. Version 0840/A or later is required for AMS2500 array of the hardware revision 0100. Version 0890/A or later is required for AMS2100/AMS2300/AMS2500 of the H/W Rev. 0200. Navigator 2l Version 3.21 or later is required for the management PC for AMS2100 or AMS2300 array of the hardware revisoin 0100. Version 4.00 or more is required for the management PC for AMS2500 array of the hardware revision 0100. Version 9.00 or more is required for management PC for AMS200/AMS2300/AMS2500 of the hardware revision 0200. CCI Version 01-21-0306 or more is required for host only when CCI is used for the operation. License key for Volume Migration.

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 LU size must be greater thanh or equal to 33 MB.) DMLU: Max 2 (The DMLU size must be greater than or equal to 10 GB. It is recommended that two DMLUs are set according to be created in different RAID groups. Size of LU: The P-VOL size must equal the S-VOL LU size.

The hardware revision can be displayed when you select an individual array from the Arrays list using version 9.0 or later of Navigator 2. Figure 11-1 displays the Arrays List window with the hardware revision noted.

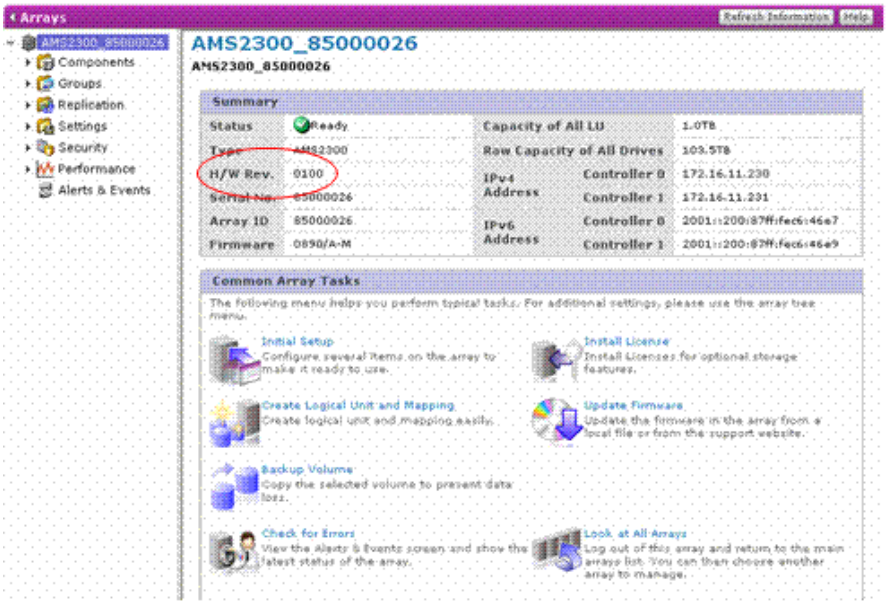

Figure 11-1: Arrays List window with hardware revision noted

Setting up Volume Migration

This section explains guidelines to observe when setting up Volume Migration.

Setting Logical Units 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 11-2 on page 11-7). 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.

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.

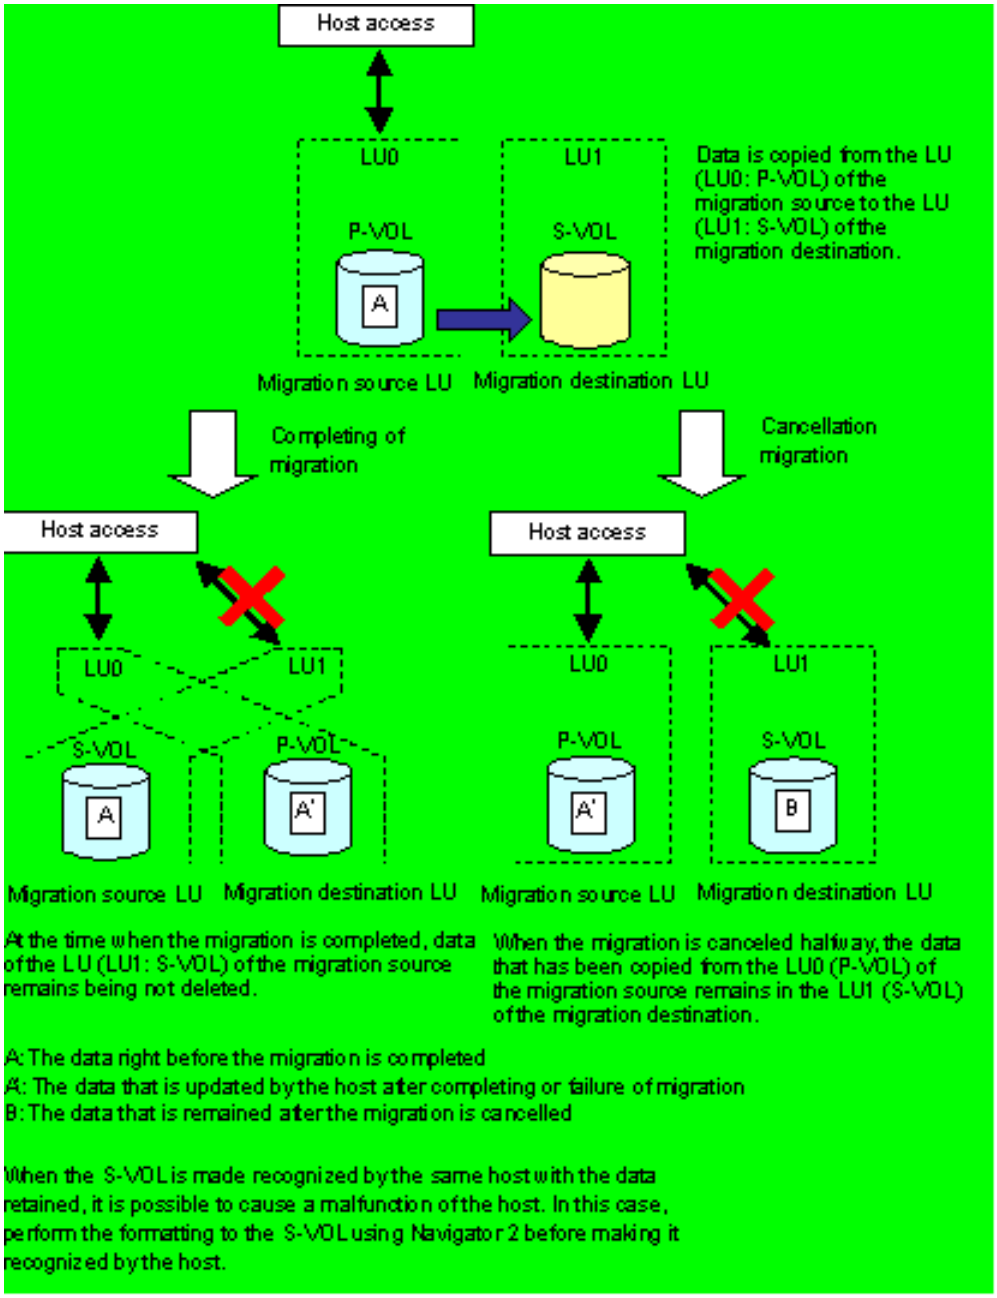

Figure 11-2: Volume Migration Host Access

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 2003/Windows Server 2008

- 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 logical unit is used with a drive character assigned, the drive character is taken to the migration logical unit. However, when both logical units 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 2003/Windows Server 2008 and Dynamic Disk

- Do not allow the P-VOL and S-VOL to be recognized by the host at the same time.

Performance

- Migration affects the performance of the host I/O to P-VOL and other logical units. 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.
- Do not concurrently migrate logical volumes that are in the same RAID group.
- Do not run Volume Migration from/to LUs 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 LUs 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 logical units

A unified logical volume (Figure 11-3) 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 logical units).

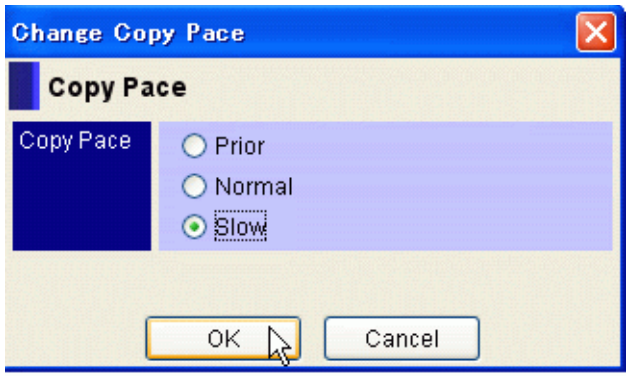

Figure 11-3: Unified Logical Units Assigned to P-VOL or S-Vol (Capacity)

The number of logical units that can be unified as components of a P-VOL or S-VOL is 128 (Figure 11-4).

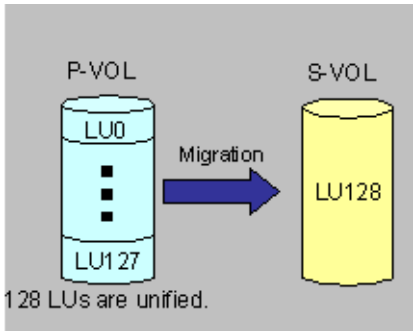

Figure 11-4: Unified Logical Units Assigned to P-VOL or S-Vol (Unification)

The logical units, including the unified logical units assigned to the P-VOL and S-VOL, cannot be on the same RAID level, or have the same number of disks (Figure 11-5 on page 11-10 and Figure 11-6 on page 11-10).

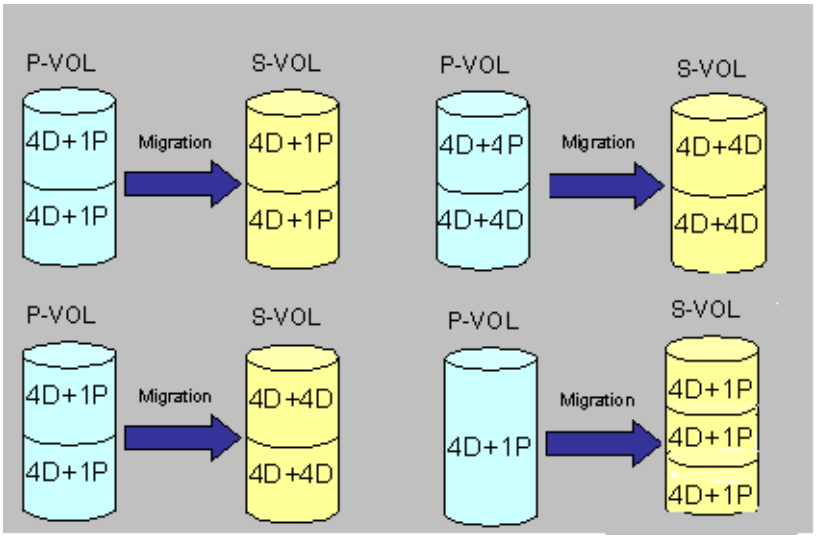

Figure 11-5: RAID Level Combination

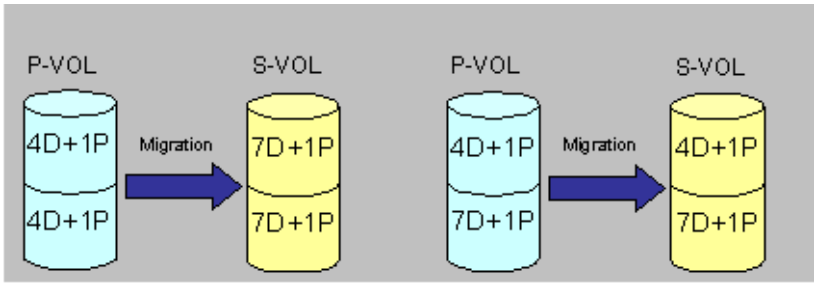

Figure 11-6: Disk Number Combination

Do not migrate when the P-VOL and the S-VOL logical units belong to the same RAID group.

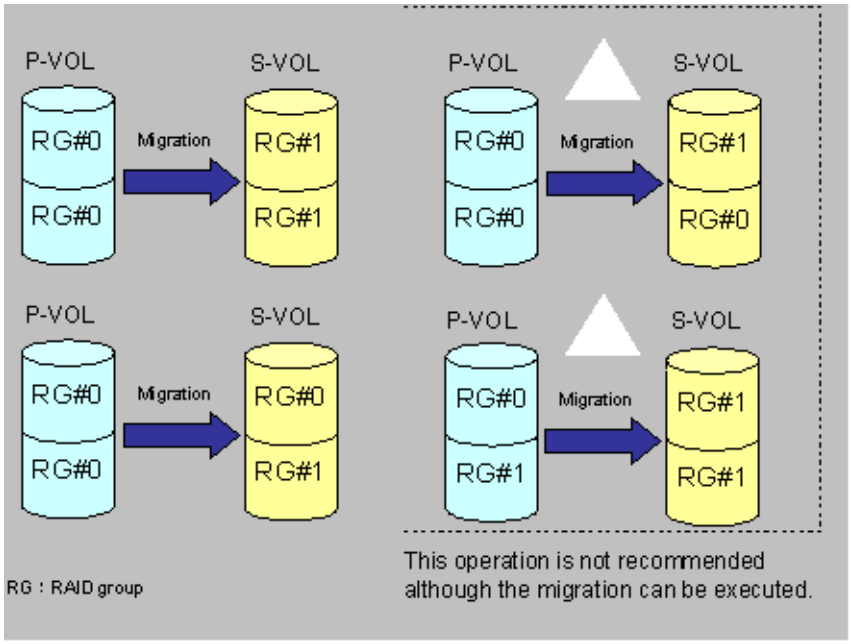

Figure 11-7: Logical Unit RAID Group Combinations

Using with the Data Retention Utility

The logical unit that executed the migration carries the access attribute and the retention term set by Data Retention, to the destination logical unit. If the access attribute is not Read/Write, the logical unit cannot be specified as an S-VOL.

The status of the migration for a Read Only logical unit appears in Figure 11-8 on page 11-12. When the migration of the Read Only LU0 to the LU1 is executed, the Read Only attribute is carried to the destination logical unit. Therefore, LU0 is Read Only. When the migration pair is released and LU1 is deleted from the reserved logical unit, a host can Read/Write to the LU1

.

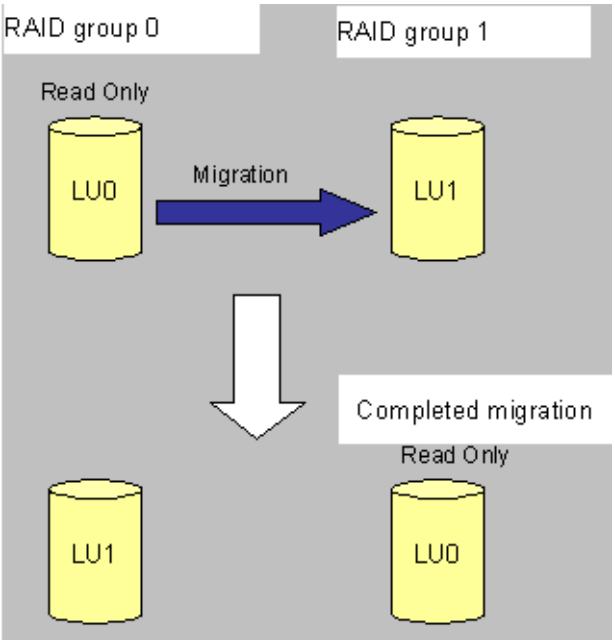

Figure 11-8: Read Only

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 11-9 on page 11-13 and Figure 11-10 on page 11-13).

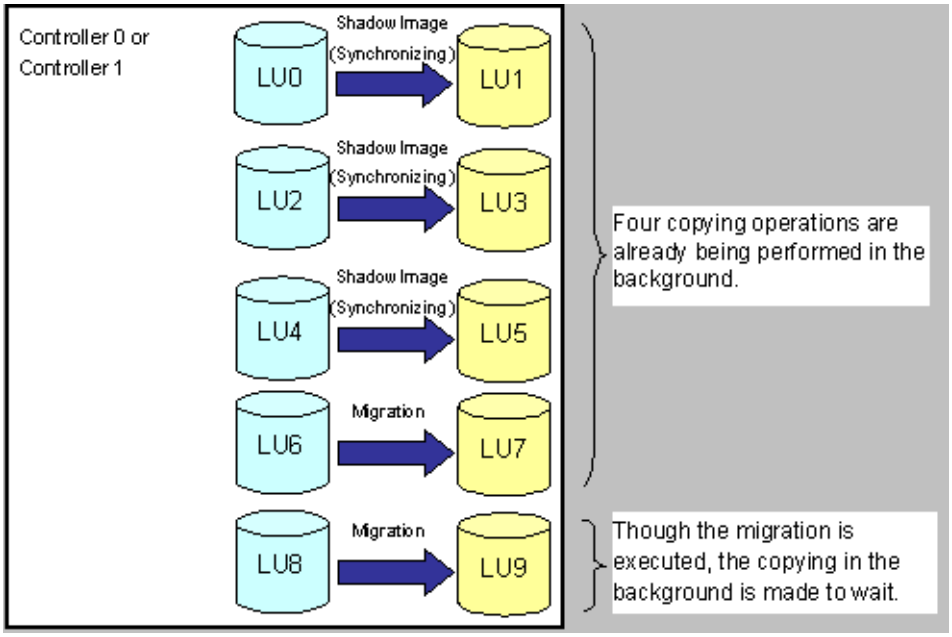

Figure 11-9: Copy Operation where Volume Migration Pauses

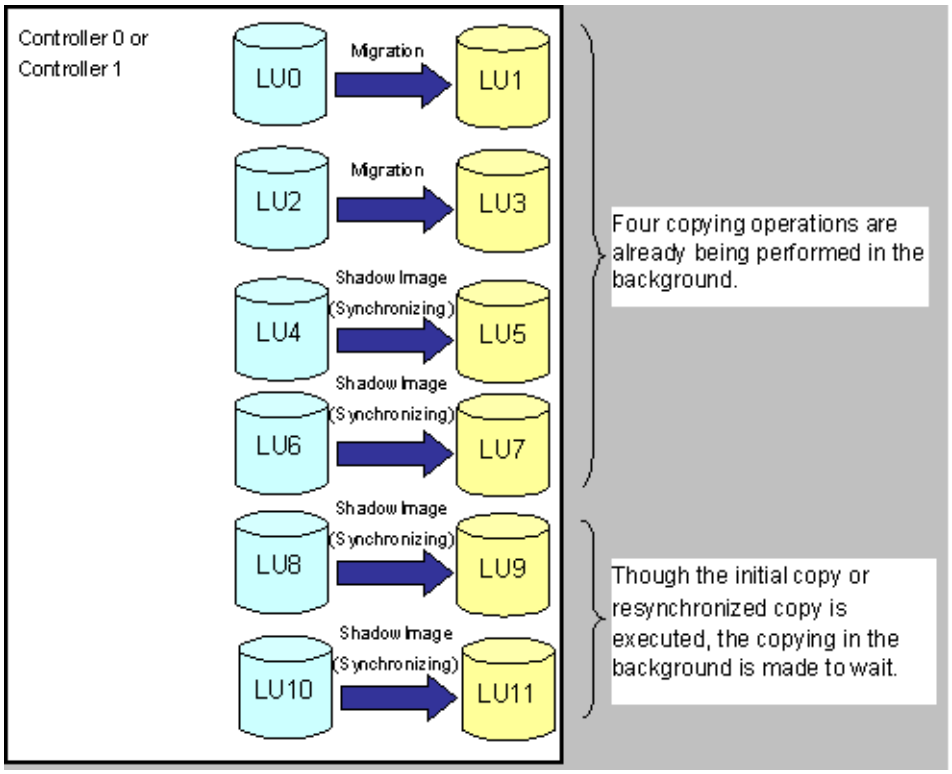

Figure 11-10: Copy Operation where ShadowImage Operation Pauses

Using with Cache Partition Manager

It is possible to use Volume Migration with Cache Partition Manager. Note that an LU that belongs to a partition cannot carry over. When a migration process completes, an LU belonging to a partition is changed to destination partition.

Modular Volume Migration operations

The procedure for Volume Migration (transferring a logical volume) appears below.

1. Verify that you have the environments and requirements for Volume Migration (see Preinstallation information on page 2-2).
2. Set the DMLU (see Adding reserved logical units on page 11-16).
3. Create a logical unit in RAID group 1 and format it. The size of the logical unit must be same as the one you are migrating. When the logical unit that has already been formatted is to be the logical unit of the migration destination, it is not necessary to format it again.
4. Set LU X as a reserved logical unit (see Adding reserved logical units on page 11-16).
5. Migrate. Specify the LU0 and the LU1 for the P-VOL and the S-VOL, respectively.

NOTE: You cannot migrate while the reserved logical unit 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 11-23).
7. When the migration pair status is Completed, release the migration pair. The relation between the P-VOL/S-VOL of LU0/LU1 is released and the two logical units 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, LU0 has been migrated to the RAID group 1 where LU1 was created, and LU1 has been migrated to the RAID group 0 where LU0 was. If the migration fails, LU0 is not migrated from the original RAID group 0 (see In the resulting message boxes, click Format LU if you want to format the removed Reserve LUs. Otherwise click Close.Migrating volumes on page 11-18).
9. The LU1 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 the LU1 from the reserved logical unit. The LU1 deleted from the reserved LU can be used for the usual system operation

as a formatted logical unit (see In the resulting message boxes, click Format LU if you want to format the removed Reserve LUs. Otherwise click Close.Migrating volumes on page 11-18).

Managing Modular Volume Migration

This section describes how to migrate volumes using the Modular Volume Migration tool.

Volume Migration runs under the Java applet used for some of the storage features. Please see Advanced Settings Java Applet on page 1-20 for more information on JRE and Java console settings.

Adding reserved logical units

When mapping mode is enabled, the host cannot access the logical unit if it has been allocated to the reserved logical unit.

NOTE: When the mapping mode displays, the host cannot access the logical unit if it has been allocated to the reserved logical unit. Also when the mapping mode is enabled, the host cannot access the logical unit if the mapped logical unit has been allocated to the reserved logical unit.

WARNING! Stop host access to the logical unit before adding reserved logical units for migration.

Follow these steps to add reserved logical units for migration.

1. Start Navigator 2 and log in. The Arrays window appears
2. Click the appropriate array.
3. Click **Show & Configure Array**.
4. Select the **Reserve LUs** icon in the Volume Migration tree view as shown in Figure 11-11.

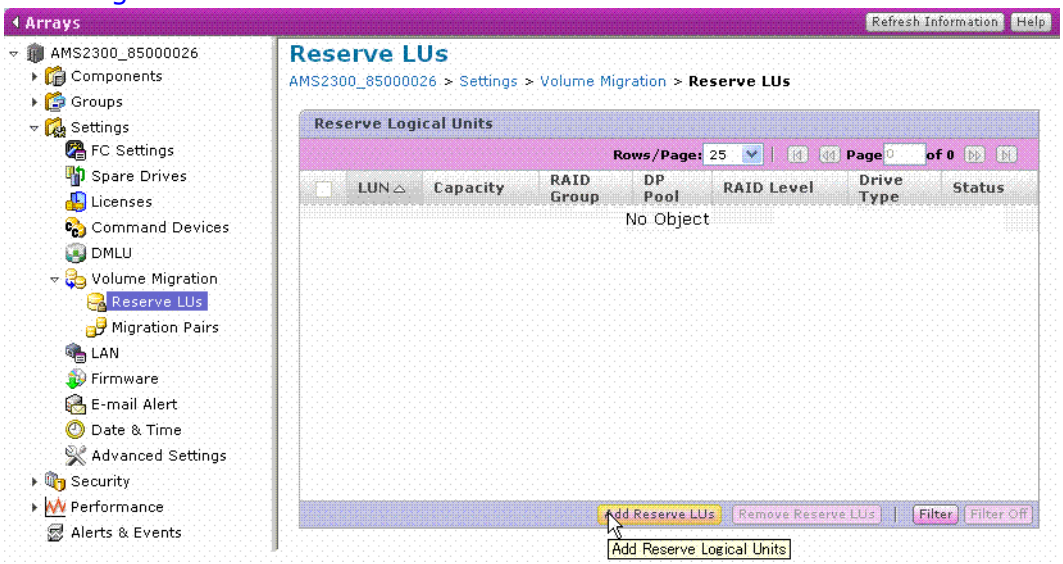

Figure 11-11: Reserve LUs dialog box

5. Click **Add Reserve LUs**. The **Add Reserve Logical Units** panel displays as shown in Figure 11-12.

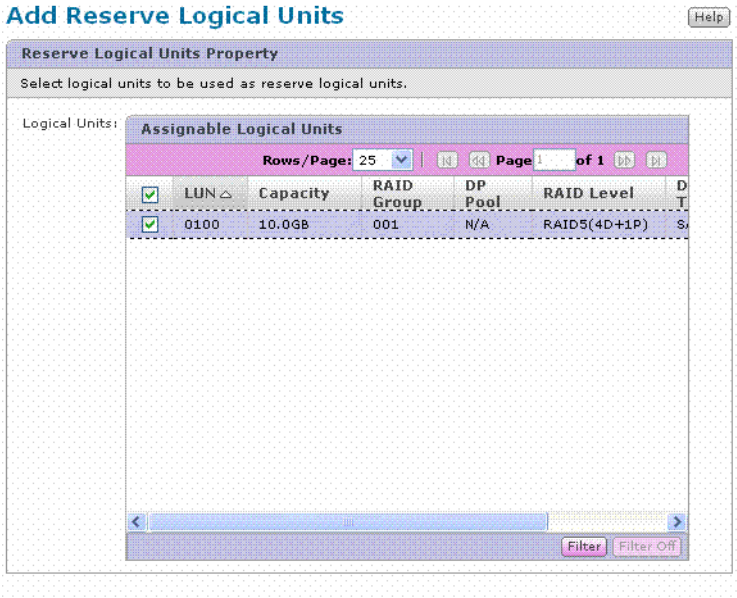

Figure 11-12: Add Reserve Logical Units panel

6. Select the LUN for the reserved logical unit and click **OK**.
7. In the resulting message boxes, click **Confirm**.
8. In the resulting message boxes, click **Close**.

Deleting reserved logical units

When canceling or releasing the volume migration pair, delete the reserve logical unit, or change the mapping. For more information, see Table 11-1 on page 11-2 and Setting up Volume Migration on page 11-5.

NOTE: Be careful when the host recognizes the logical unit that has been used by Volume Migration. After releasing the Volume Migration pair or canceling Volume Migration, delete the reserved logical unit or change the logical unit mapping.

Follow these steps to delete reserved logical units.

1. From the Reserve LUs dialog box, select the LUN to be deleted as shown in Figure 11-13.

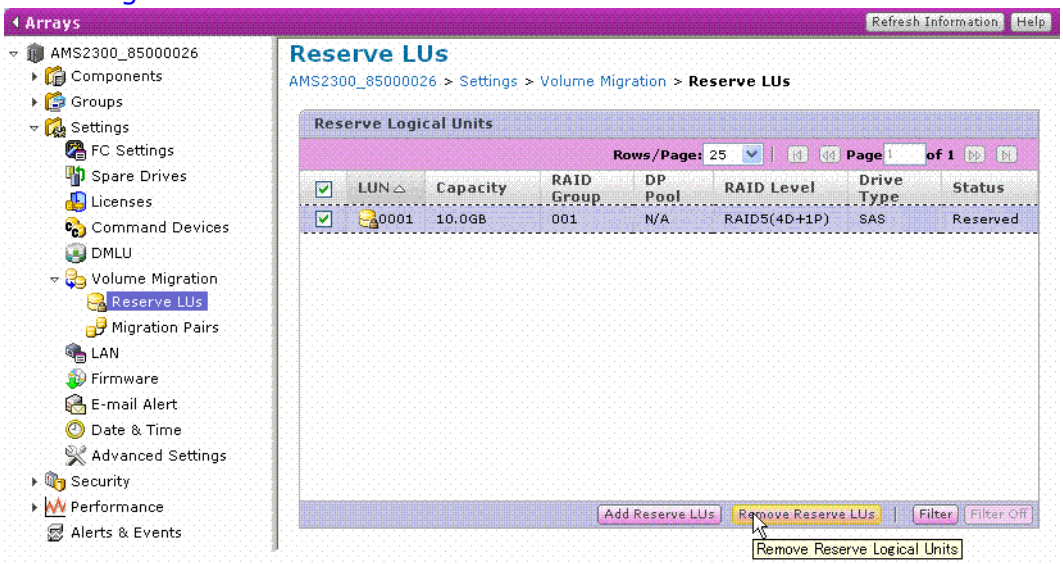

Figure 11-13: Reserve LUs dialog box - LUN Selected for Deletion

2. In the resulting message boxes, click **Confirm**.
3. In the resulting message boxes, click **Format LU** if you want to format the removed Reserve LUs. 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. Expand the **Settings** list, and click **Advanced Settings**.
4. Click **Open Advanced Settings**. After some minutes, the Array Unit window appears, as shown in Figure 11-14.

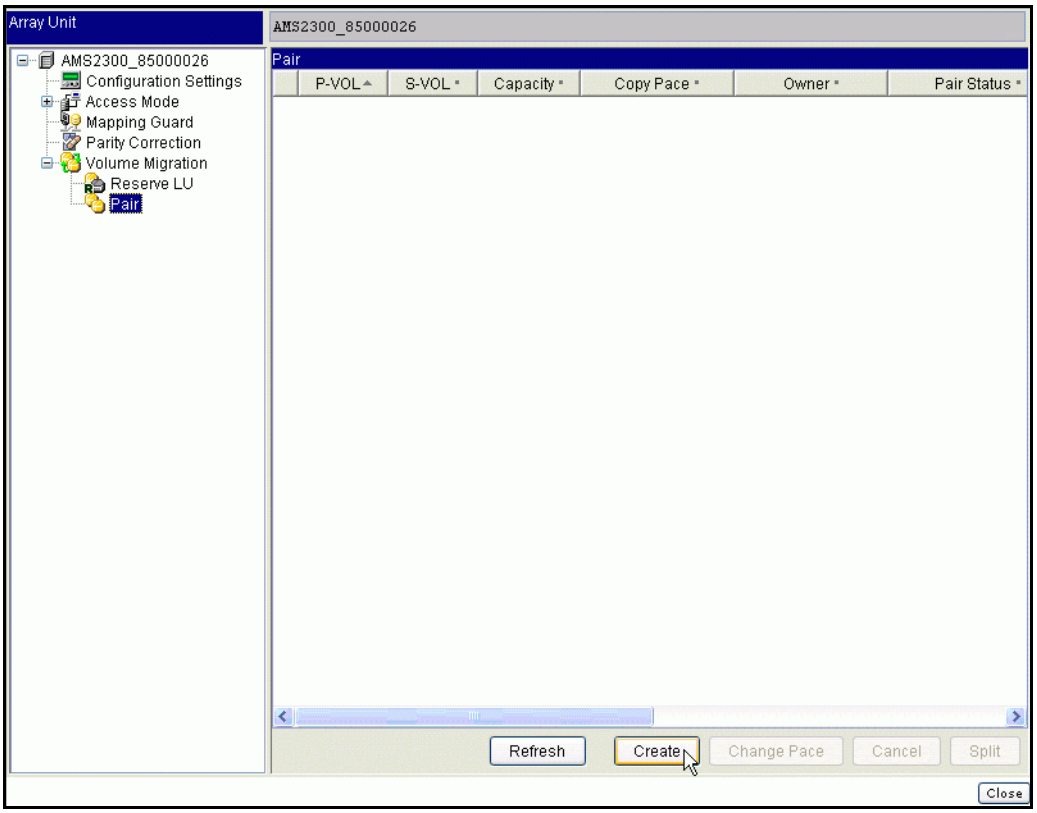

Figure 11-14: Array Unit Window-Pair

5. Expand the **Volume Migration** list, and click **Pair**.
6. Click **Create**.
7. Click **OK** in the Confirmation dialog box.
8. Click **Select** to specify a P-VOL. The Create Volume Migration Pair window appears, as shown in Figure 11-15.

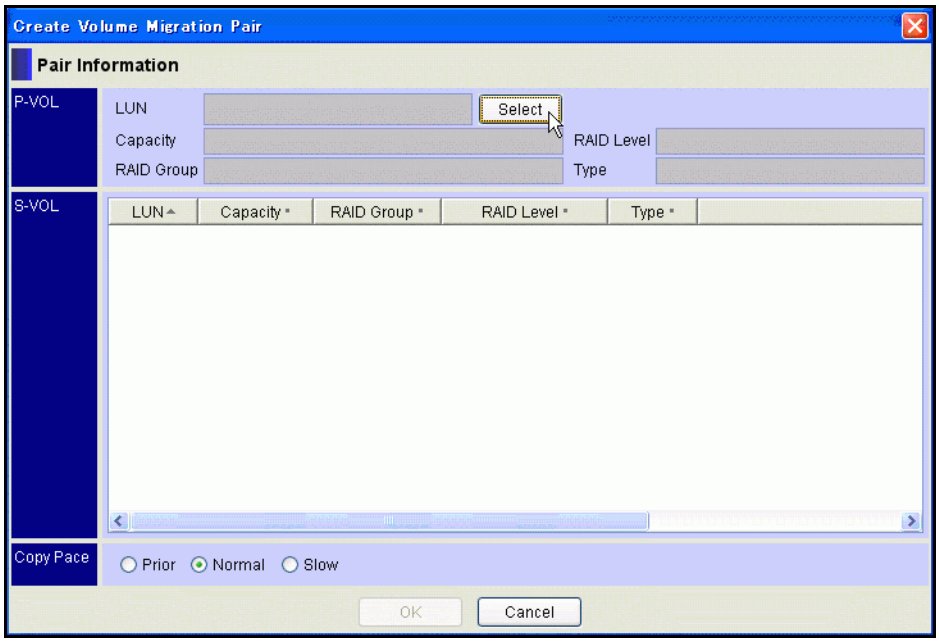

Figure 11-15: Create Volume Migration Pair

9. Select the LUN for the P-VOL, and click **OK**. The Select P-VOL dialog appears, as shown in Figure 11-16.

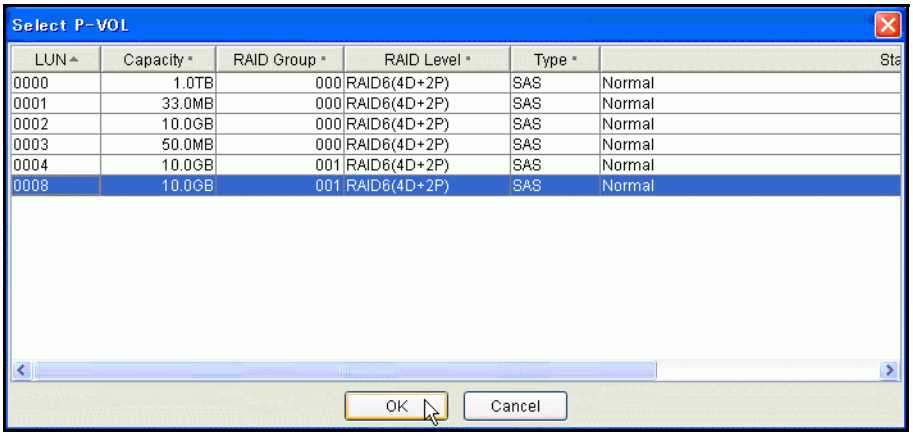

Figure 11-16: Select P-VOL Dialog

10. Select the LUN for the S-VOL and Copy Pace, click **OK** as shown in Figure 11-17.

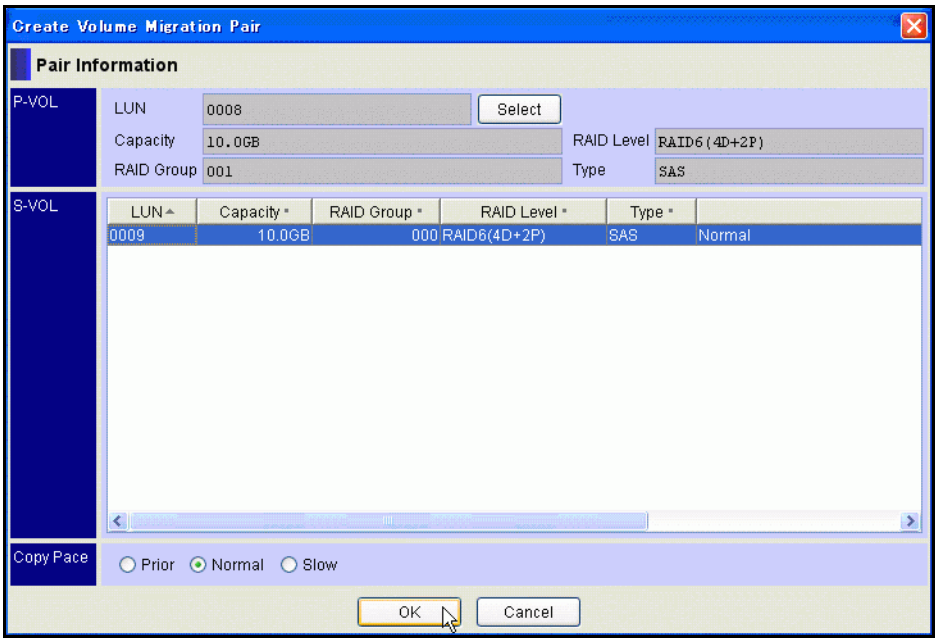

Figure 11-17: Create Volume Migration Pair Dialog

11. 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
- Normal
- Slow

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. Expand the **Settings** list, and click **Advanced Settings**.
4. Click **Open Advanced Settings**. After some minutes, the Array Unit window appears, as shown in Figure 11-14 on page 11-19.
5. Expand the **Volume Migration** list, and click **Pair**.
6. Select the pair whose copy pace you are modifying, and click **Change Pace**. The Change Copy Pace dialog box appears, as shown in Figure 11-18.

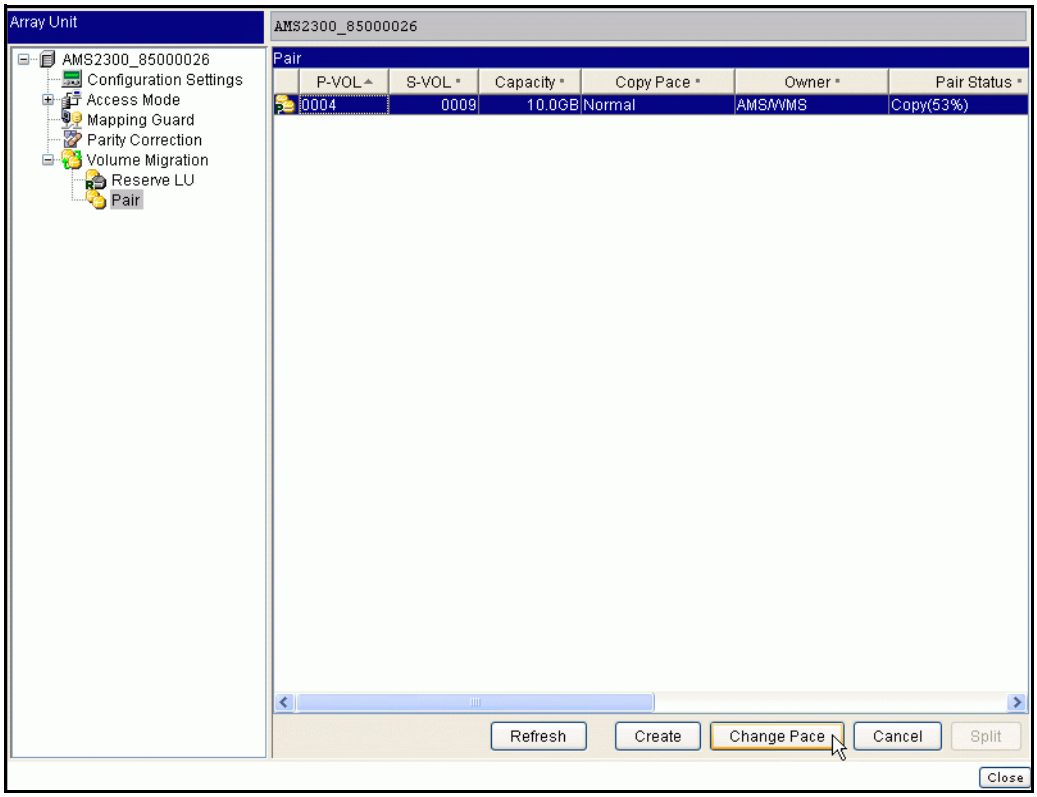

Figure 11-18: Change Copy Pace Dialog Box

7. Select the copy pace and click **OK**. The Change Copy Pace panel appears, as shown in Figure 11-19.

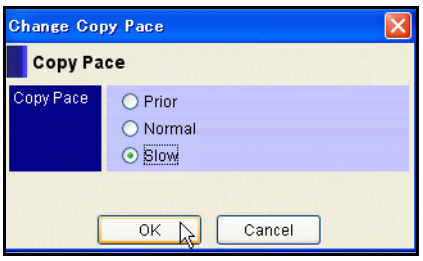

Figure 11-19: Change Copy Pace Panel

8. In the resulting message box, click **OK**, as shown in.
9. Follow the on-screen instructions.

Confirming Volume Migration Pairs

Figure 11-20 shows the pair migration status.

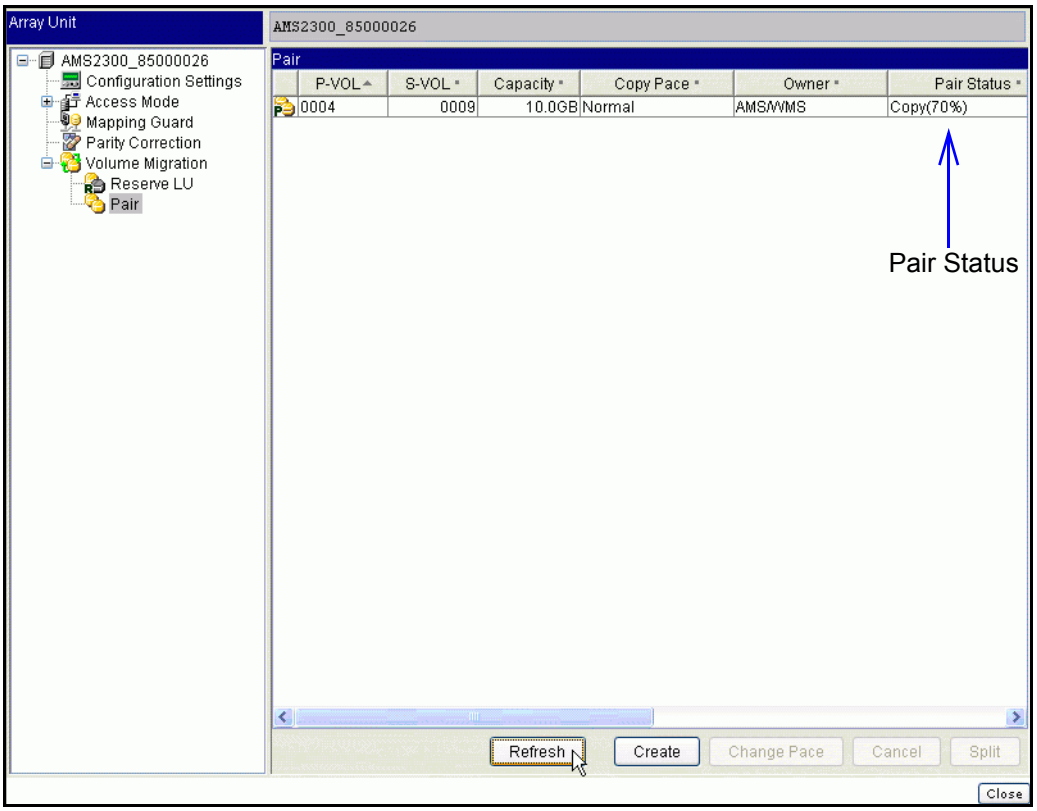

Figure 11-20: Array Unit Window-P-VOL and S-VOL Migration

- **P-VOL** - The logical unit number appears for the P-VOL.
- **S-VOL** - The logical unit 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.

Splitting Volume Migration pairs

A pair can only be split if it is in Completed or Error status.

1. Start Navigator 2 and log in. The Arrays window appears
2. Click the appropriate array.
3. Expand the **Settings** list, and click **Advanced Settings**.
4. Click **Open Advanced Settings**. After some minutes, the Array Unit window appears.
5. Expand the **Volume Migration** list, and click **Pair**.
6. Select the migration pair to split, and click **Split** as shown in Figure 11-21.

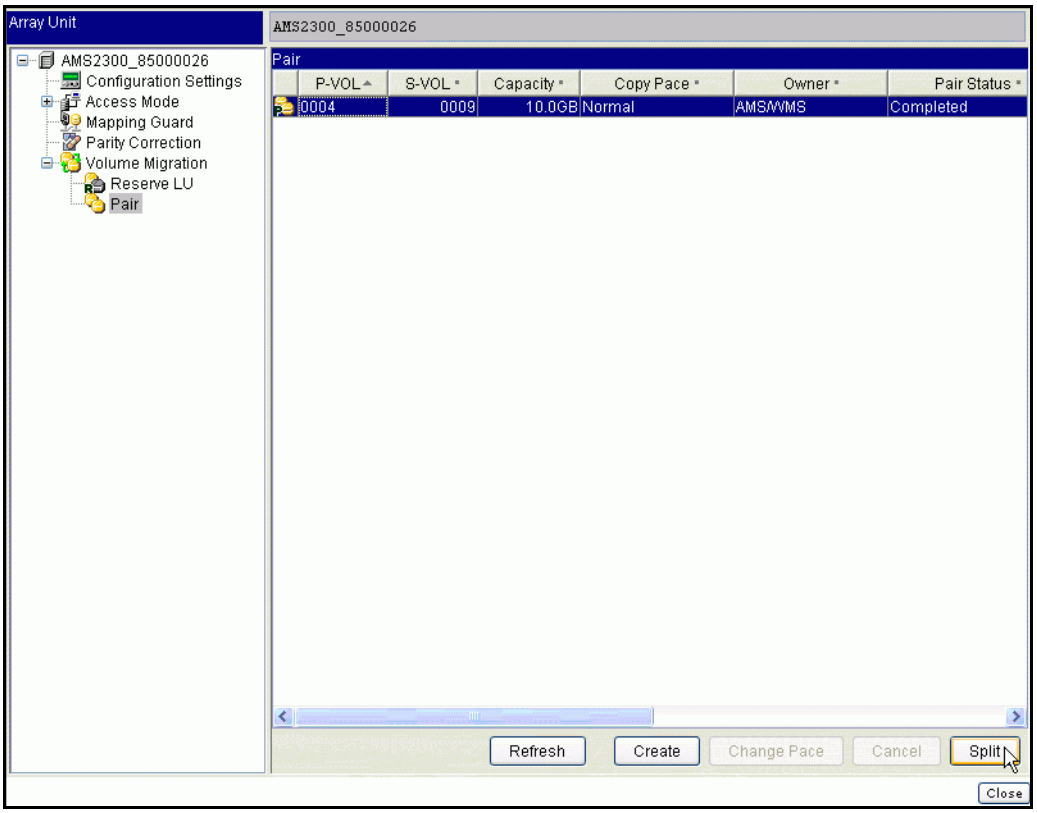

Figure 11-21: Pair Splitting

7. Follow the on-screen instructions.

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, follow these steps.

1. Start Navigator 2 and log in. The Arrays window appears
2. Click the appropriate array.
3. Expand the **Settings** list, and click **Advanced Settings**.
4. Click **Open Advanced Settings**. After some minutes, the Array Unit window appears.
5. Expand the **Volume Migration** list, and click **Pair**.
6. Select the P-VOL you are canceling, and click **Cancel** as shown in Figure 11-22.

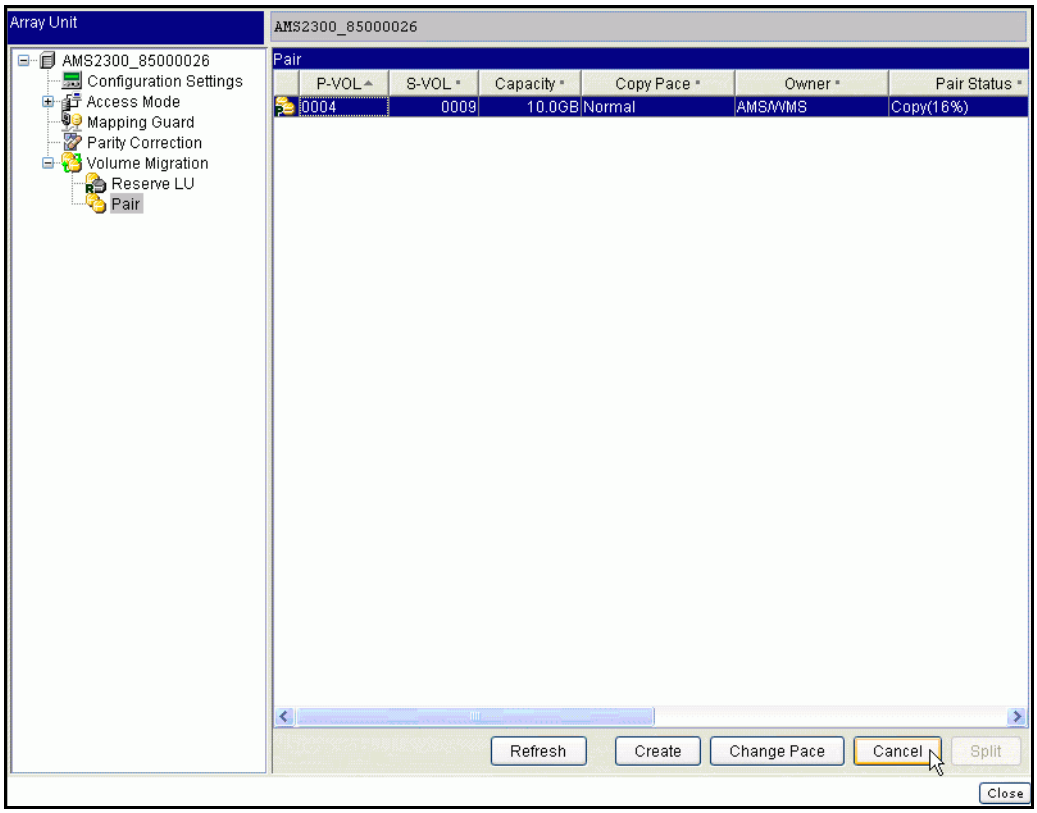

Figure 11-22: Pair Cancellation

7. Follow the on-screen instructions.

Appendix A — Logical unit expansion/reduction

This Appendix describes logical unit expansion and reduction operations and includes the following:

- Overview
- Expanding capacity using LUN "Grow/Shrink"
- Changing logical unit capacity
- Displaying unified logical unit properties
- Separating unified logical units

Overview

There are two ways to expand or reduce the size of logical units under Navigator 2:

- Logical unit expansion and reduction (grow/shrink) — Increases the size of a logical unit by incorporating un-used disk space into a selected logical unit. This feature can also be used to reduce the amount of space assigned to a logical unit.
- Logical Unit Size Expansion (LUSE) — Increases the size of a logical unit by concatenating selected logical units together. The result is referred to as a unified logical unit.

In most cases, it is probably desirable to take un-used disk space and "grow" the capacity of an existing logical unit. However, you cannot take un-used space from another RAID group and add it to a logical unit. You must use the LUSE feature in the case where your available or additional capacity must come from another RAID group.

Expanding capacity using LUN "Grow/Shrink"

Under Navigator 2, from the Change Logical Unit Capacity window, you can expand (grow) or reduce (shrink) a logical unit.

The Change Logical Unit Capacity window may be accessed from the following two windows:

- Logical Units tab from the RAID Groups window
- Property window for RAID Group

Requirements and conditions for expanding logical units

The follow describe the requirements and conditions required to expand the size of a logical unit.

1. You cannot delete an LU where the status of the forced parity correction is any of the following status messages observed in the RAID group of the expansion target:
   - "Correcting"
   - "Waiting"
   - "Waiting Drive Reconstruction"
   - "Unexecuted", "Unexecuted 1", or "Unexecuted 2"

   If any of the above messages are displayed, you need to execute a forced parity correction for the LU and change its LU status to "Correction Completed" or "Skip," and then delete this LU.
2. You cannot execute an LU expansion when they are being formatted.
3. You cannot execute an LU expansion for a unified LU that is configured using LUs of two or more RAID groups.

4. You cannot execute an LU expansion when a drive restoration is in progress on the target LU. Execute after completing the drive restoration.
5. You cannot expand the following LU types:
   - LUs used in a ShadowImage pair
   - LUs used as a Copy-on-write SnapShot pair
   - LUs used as a TrueCopy or TrueCopy Extended Distance pair
   - LUs or reserve LUs for Modular Volume Migration
   - LUs in which Cache Residency Manager is set
   - LUs that are being formatted
   - LUs used as a command device
   - Differential Management Logical Units (DMLUs)
   - LUs that are registered in the data pool
   - LUs that are in the RAID group during a RAID group expansion
6. You cannot expand LUs of the RAID group in which the Power Saving function is set. Change the status of the Power Saving feature to “Normal (spin-up)” and then expand the LUs.

Additionally, please note the following:

1. You cannot reduce an LU where its properties are set to “Read Only,” “Protect,” or “Can't Guard.”
2. You cannot reduce an LU where the S-VOL setting is Setting Impossible (disabled) and the mode is any of the following:
   - “Read Capacity 0 (Zer)”
   - “Inquiry Command Shielding (Zer/Inv)”.
3. You cannot reduce an LU if the Data Retention Utility is enabled and its properties are set to “Read/Write”, or if the S-VOL is set to “Setting Possible (Enabled)”, and mode to “Unset.”
4. You must wait until after the drive is restored to reduce the LUs if the dynamic sparing/correction copy/copy back is in progress.

Changing logical unit capacity

This section assumes that you are logged on to Navigator 2 and know how to use its features and functions.

CAUTION! Before beginning any LUN expansion procedure please read and follow all of the following instructions.

- 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 logical unit will not be created.
- Back up the unified logical units before modifying them.
- Format the unified logical units to delete the volume label which the operating system adds to logical units.
- Create unified logical units only from logical units within the same array.
- You must format a logical unit that is undefined before you can use it.

- You can increase the size of a logical unit using the available free space within the RAID group to which it belongs.
- You can change the capacity of a logical unit adding existing logical units to a selected logical unit.

To change the capacity of a logical unit:

1. From the array tree, click **Groups > Logical Units**. The Logical Units window appears.
2. Select the LUN on which you want to change capacity.
3. Click **Change LU Capacity**. The Change Logical Unit Capacity is displayed. The current properties of the selected LUN are displayed, including LUN, Current Capacity, and Free space (except for logical units in a DP pool).
4. From the **Basic** tab, choose the desired operation to expand or shrink capacity:
   - To expand (“grow”) or reduce (“shrink”) the LUN capacity:
   a. Click **Input** and enter the capacity value you want. ALL is the default option from the pull-down list.
   b. To set available capacity automatically, go to the step 5.
   c. To set the capacity manually, go to step 6.

   - To unify or concatenate logical units to LUN expansion (LUSE):
   a. Click **Add** logical units and select a LUN or LUNs from the **Available Logical Units** table. You can select logical units from across available RAID groups.
   b. Go to step 6.

- To separate LUNs that have been joined together:

  a. Select either Separate last logical units or Separate all logical units and check the desired box on the Available Logical Units table.

  b. Go to step 6.

5. Click the **Input Capacity Options** tab. Select **Set Manually** and check the desired box from the Free Space table. If you set the LUN of an existing logical unit, add the LUN to this field. The LUN assignment numbers are in descending order.
6. Click **OK** and the result dialog is displayed.
7. Click **Close**. Confirm your changes.

Displaying unified logical unit properties

You can view the properties of logical units that you unified using the LUSE feature.

To display the list of unified logical units:

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, click **Groups > Logical Units**. The Logical Units window and table is displayed.
3. From the LUN column, click the LUN number. The LUN screen for that logical unit appears. The properties for the logical unit are displayed.
4. Click the **Sub Logical Units** tab to view the logical units that have been unified into the current LUN.

Separating unified logical units

The process of separating logical units that have been unified under the LUSE feature may be done in two ways:

- Separate the last logical unit that was added to the unified logical unit
- Separate all of the logical units that make up the unified logical unit

Separating the last logical unit

This process is the reverse of adding a logical unit to a unified logical unit.

To remove the last logical unit that was added to a unified logical unit:

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 **LUN Expansion**. The LUN Expansion window is displays the list of unified logical units in the array and a set of parameters for each listed unit.
4. In the LUN Expansion window, click the LUN that you want to separate.
5. In the unified logical unit property window, click **Separate Last LUs**. A confirmation dialog box is displayed.
6. Review the warning message, and then click **Confirm**. A message box stating that the logical unit has been successfully separated is displayed.
7. Click **Close** to exit the message box and return to the unified logical unit properties window.
8. Review the contents of the window and verify that the logical unit was separated from the unified logical unit.

Separating all logical units

This process will separate all unified logical units.

To remove the all of the logical units added to a unified logical unit:

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 **LUN Expansion**. The LUN Expansion window is displays the list of unified logical units in the array and a set of parameters for each listed unit.
4. In the LUN Expansion window, click the LUN that you want to separate.
5. In the unified logical unit property window, click **Separate All LUs**. A confirmation dialog box is displayed.
6. Review the warning message, and then click **Confirm**. A message box stating that the logical unit has been successfully separated is displayed.
7. Click **Close** to exit the message box and return to the unified logical unit properties window.
8. Review the contents of the window and verify that the logical unit was separated from the unified logical unit.

Available unified LUN information

The following information appears in the LUN window for unified logical units.

Table A-1: Unified LUN Information

Item	Description
LUN	This shows the name of the unified logical unit.
Capacity	This shows the size in GB of the unified logical unit.
RAID Group	This shows the name of the RAID group that the unified logical unit is part of.
RAID Level	This shows the RAID level of the unified logical unit. This is RAID6.
TYPE	This shows the type of hard drive the logical unit resides on: SATA or SAS.
Status	This shows the status of the unified logical unit: Normal or Alarm.
Default Cache Partition No.	The is the cache partition number assigned by default.
Mapped to Host Group/iSCSI Target	Provides Host Group or iSCSI target mapping information.
Default and Pair Cache partition No.	Provides information on the default or pair cache partitions.
Current Cache Partition No.	The current cache partition to which the logical unit is assigned.
LUN Expansion	Defines the role of the logical unit in a unified LUN.

Glossary

This glossary provides definitions of general storage networking terms as well as specific terms related to the technology that supports Hitachi Data Systems products. Click the letter of the glossary section to display that 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

1000BASE-T

A specification for Gigabit Ethernet over copper wire. The standard defines 1 Gbps data transfer over distances of up to 100 meters using four pairs of Category 5 balanced copper cabling and a 5-level coding scheme.

Array

A set of hard disks grouped logically together to function as one contiguous storage space.

ATA

Advanced Technology Attachment, a disk drive implementation that integrates the controller on the disk drive.

BIOS

Basic Input Output System, built-in software code that determines the functions that a computing device can perform without accessing programs from a disk.

Bps

Bits per second, the standard measure of data transmission speeds.

BSD syslog protocol

This protocol has been used for the transmission of event notification messages across networks for many years. While this protocol was originally developed on the University of California Berkeley Software Distribution (BSD) TCP/IP system implementations, its value to operations and management has led it to be ported to many other operating systems as well as being embedded into many other networked devices.

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; the volume it can contain or hold. In communications,

capacity refers to the maximum possible data transfer rate of a communications channel under ideal conditions.

Challenge Handshake Authentication Protocol

A security protocol that requires users to enter a secret for access.

CHAP

See Challenge Handshake Authentication Protocol.

command control interface (CCI)

Hitachi's Command Control Interface software provides command line control of Hitachi array and software operations through the use of commands issued from a system host. Hitachi's CCI also provides a scripting function for defining multiple operations.

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.

DHCP

Dynamic Host Configuration Protocol, allows a computer to join an IP-based network without having a pre-configured IP address. DHCP is a protocol that assigns unique IP addresses to devices, then releases and renews these addresses as devices leave and re-join the network.

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.

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.

Fabric

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

FC

Fibre channel.

Firmware

Software embedded into a storage device. It may also be referred to as Microcode.

Full-duplex

The concurrent transmission and the reception of data on a single link.

Gbps

Gigabit per second.

GUI

Graphical user interface.

HBA

Host bus adapter, a circuit board and/or integrated circuit adapter installed in a workstation or server that provides input/output processing and physical connectivity between a server and a storage device. An iSCSI HBA implements the iSCSI and TCP/IP protocols in a combination of a software storage driver and hardware.

HDD

Hard disk drive.

Initiator

A system component that originates an I/O command over an I/O bus or network, such as an I/O adapters or network interface cards.

I/O

Input/output.

IP

Internet Protocol, specifies the format of packets and addressing scheme. Most networks combine IP with a higher-level protocol called Transmission Control Protocol (TCP), which establishes a virtual connection between a destination and a source.

IP address

An identifier for a computer or device on a TCP/IP network. Networks using the TCP/IP protocol route messages based on the IP address of the destination. The format of an IP address is a 32-bit numeric address written as four numbers separated by periods. Each number can be zero to 255 (for example, 192.168.0.200).

IP-SAN

Block-level Storage Area Networks over TCP/IP using the iSCSI protocol.

iSCSI

Internet SCSI, an IP-based standard for connecting data storage devices over a network and transferring data using SCSI commands over IP networks. iSCSI enables a Storage Area Network to be deployed in a Local Area Network.

iSNS

Internet Storage Name Service, a protocol that allows automated discovery, management and configuration of iSCSI devices on a TCP/IP network.

L

LAN

Local Area Network, a computer network that spans a relatively small area, such as a single building or group of buildings.

LU

Logical unit.

LUN

Logical unit number.

Middleware

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

MIB

Message Information Block.

NIC

Network Interface Card, an expansion board in a computer that allows the computer to connect to a network.

NTP

Network Time Protocol, a protocol for synchronizing the clocks of computer systems over packet-switched, variable-latency data networks. NTP uses UDP port 123 as its transport layer. It is designed particularly to resist the effects of variable latency (jitter).

Pool volume

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

primary volume (P-VOL)

The storage volume in a volume pair. It is used as the source of a copy operation. In copy operations a copy source volume is called the P-VOL while the copy destination volume is called S-VOL (secondary volume).

RAID

Redundant Array of Independent Disks, 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 of user data in the event that one of the array's member disks or the access path to it fails. SNIA.

RAID 6

An extension of the RAID 5 array, that allows for two simultaneous drive failures without downtime or data loss.recovery point objective (RPO).

After a recovery operation, the recovery point objective (RPO) is the maximum desired time period, prior to a disaster, in which changes to data may be lost. This measure determines up to what point in time data should be recovered. Data changes preceding the disaster are preserved by recovery.

SAN

Storage Area Network, a network of shared storage devices that contain disks for storing data.

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.

SATA

Serial ATA is a computer bus technology primarily designed for the transfer of data to and from hard disks and optical drives. SATA is the evolution of the legacy Advanced Technology Attachment (ATA) interface from a parallel bus to serial connection architecture.

SCSI

Small Computer System Interface, a parallel interface standard that provides faster data transmission rates than standard serial and parallel ports.

Session

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

secondary volume (S-VOL)

A replica of the primary volume (P-VOL) at the time of a backup and is kept on a standby storage system. Recurring differential data updates are performed to keep the data in the S-VOL consistent with data in the P-VOL.

SMTP

Simple Mail Transfer Protocol, a protocol used to receive and store email data directly from email servers.

Software initiator

A software application initiator communicates with a target device. A software initiator does not require specialized hardware because all processing is done in software, using standard network adapters.

Storage Navigator Modular 2

A multi-featured scalable storage management application that is used to configure and manage the storage functions of Hitachi arrays. Also referred to as Navigator 2.

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 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.

Target

Devices that receive iSCSI requests that originate from an iSCSI initiator.

TOE

A dedicated chip or adapter that handles much of the TCP/IP processing directly in hardware. TCP/IP transmission is inherently a CPU-intensive operation. Therefore, using dedicated hardware that can operate in parallel with the main processor allows for superior system performance. Although all iSCSI HBAs have a TOE, a generic TOE only implements TCP/IP, while an iSCSI HBA implements the iSCSI protocol in addition to TCP/IP.

User Datagram Protocol (UDP)

UDP is one of the core protocols of the Internet protocol suite. Using UDP, programs on networked computers can send short messages sometimes known as datagrams (using Datagram Sockets) to one another.

UDP does not guarantee reliability or ordering in the way that TCP does. Datagrams may arrive out of order, appear duplicated, or go missing without notice. Avoiding the overhead of checking whether every packet actually arrived makes UDP faster and more efficient, at least for applications that do not need guaranteed delivery. Time-sensitive applications often use UDP because dropped packets are preferable to delayed packets. UDP's stateless nature is also useful for servers that answer small queries from huge numbers of clients. Unlike TCP, UDP is compatible with packet broadcast (sending to all on local network) and multicasting (send to all subscribers).

World Wide Name (WWN)

A unique identifier for an open systems host. It consists of a 64-bit physical address (the IEEE 48-bit format with a 12-bit extension and a 4-bit prefix). The WWN is essential for defining the SANtinel™ parameters because it determines whether the open systems host is to be allowed or denied access to a specified logical unit or a group of logical units.

Index

A

access control. *See* Account Authentication
Account Authentication
- account types 3-4
- adding accounts 3-11
- default account 3-4
- deleting accounts 3-15
- modifying accounts 3-14
- overview 1-2
- permissions and roles 3-5–3-7
- session timeout settings 3-16
- setup guidelines 4-2
- viewing accounts 3-10

account types 3-4
Advanced Settings 1-20
audit logging
- external syslog servers 4-6
- initializing logs 4-6
- protocol compliance 1-18, 2-3
- setup guidelines 4-3
- syslog server 1-18, 2-3
- transferring log data 4-3
- viewing log data 4-5–4-6

Audit Logging. *See* audit logging

C

Cache Partition Manager
- adding cache partitions 5-14
- adding or reducing cache 5-12
- assigning partitions 5-16
- changing owner controllers 5-17
- changing partitions 5-17
- deleting partitions 5-15
- load balancing 5-11
- setting a pair cache partition 5-16
- setup guidelines 5-12–5-13
- SnapShot and TCE installation 5-18–5-19

Cache Residency Manager
- setting residency LUs ??–6-13, 6-14–??
- setup guidelines 6-14

CHAP network security 1-9
copy speed (pace). *See* Modular Volume Migration
creating
- Host Groups (FC) 8-26
- iSCSI targets 8-32

D

Data Retention Utility
- Expiration Lock configuration 7-8
- setting attributes 7-8
- setup guidelines 7-6, 7-6
- S-VOL configuration 7-8

deleting accounts. *See* Account Authentication
Dynamic Provisioning
- disabled, partition capacity 5-4
- logical unit capacity 6-4
- partition capacity 5-6
- small segments 5-4

F

features
- Account Authentication 3-2
- Audit Logging 3-8
- Cache Partition Manager 5-2
- Cache Residency Manager 6-2
- Data Retention Utility 7-2
- Volume Migration 11-2

fibre channel
- adding host groups 8-24
- deleting host groups 8-29
- initializing Host Group 000 8-28

fibre channel setup workflow. *See* LUN Manager

H

hosts, mapping to LUs 1-9

I

iSCSI
- adding targets 8-36
- configuration 1-10
- creating a target 8-32
- creating iSCSI targets 8-32
- deleting targets 8-38
- description 1-10
- editing authentication properties 8-38
- editing target information 8-38
- host platform options 8-37
- initializing Target 000 8-39
- nicknames, changing 8-40
- system configuration 8-10
- Target 000 8-38
- using CHAP 8-31, 8-40, ??–9-5

iSCSI setup workflow. *See* LUN Manager

J

Java applet, timeout period 1-20
Java applet. *See also* Advanced Settings
Java runtime requirements 1-20

L

logical units
- expanding A-1

LUN expansion. *See* logical units, expanding
LUN Manager
- adding host groups 8-24–8-28
- connecting hosts to ports 1-9
- creating iSCSI targets 8-32
- fibre channel features 1-8
- fibre channel setup workflow 8-23
- Host Group 000 8-28
- host group security, fibre channel 8-24
- iSCSI features 1-9
- iSCSI setup workflow 8-23

LUSE. *See* logical units, expanding

M

Management Information Base (MIB). *See* SNMP
migrating volumes. *See* Modular Volume Migration
Modular Volume Migration
- copy pace, changing 11-22
- migration pairs, canceling. 11-25
- migration pairs, confirming 11-23
- migration pairs, splitting 11-24
- Reserved LUs, adding 11-16
- Reserved LUs, deleting 11-17
- setup guidelines 11-14–11-15

N

NTP, using SNMP 1-20, 2-5

P

password, default. *See* account types
Performance Monitor
- exporting information 9-22
- obtaining system information 9-3
- performance imbalance 9-29–9-30
- troubleshooting performance issues 9-29
- using graphs 9-3–9-5

permissions. *See* Account Authentication

S

security, setting iSCSI target 8-33, 8-34
SNMP
- agent setup workflow 10-8
- disk array-side configuration 10-9
- failure detection 10-18
- Get/Trap specifications 10-4
- IPv6 requirements 10-8
- message limitations 1-13
- MIB information 1-19, 2-4, 10-18
- REQUEST connections 10-17
- request processing 1-13
- SNMP manager-side configuration 10-9
- trap connections, verifying 10-16
- trap issuing 1-12

SNMP agent support
- LAN/workstation requirements 1-11
- overview 1-11

SNMP manager, dual-controller environment 1-20, 2-4
syslog server. *See* audit logging
system configuration 8-10

T

timeout length, changing 3-16
timeout, Java applet 1-20

Hitachi Data Systems

Corporate Headquarters
750 Central Expressway
Santa Clara, California 95050-2627
U.S.A.
Phone: 1 408 970 1000
www.hds.com
info@hds.com

Asia Pacific and Americas
750 Central Expressway
Santa Clara, California 95050-2627
U.S.A.
Phone: 1 408 970 1000
info@hds.com

Europe Headquarters
Sefton Park
Stoke Poges
Buckinghamshire SL2 4HD
United Kingdom
Phone: + 44 (0)1753 618000
info.eu@hds.com

Hitachi Data Systems

MK-97DF8148-07