

Hitachi Universal Storage Platform V Hitachi Universal Storage Platform VM

**Hitachi Performance Manager User's Guide
Performance Monitor, Volume Migration,
and Server Priority Manager**

FASTFIND LINKS

[Document Organization](#)

[Product Version](#)

[Getting Help](#)

[Contents](#)

Copyright © 2008 Hitachi Data Systems Corporation, ALL RIGHTS RESERVED

Notice: 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 Data Systems Corporation (hereinafter referred to as "Hitachi Data Systems").


Hitachi Data Systems reserves the right to make changes to this document at any time without notice and assumes no responsibility for its use. 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.

This document contains the most current information available at the time of publication. When new and/or revised information becomes available, this entire document will be updated and distributed to all registered users.

Hitachi, Hitachi logo, and Hitachi Data Systems are registered trademarks and service marks of Hitachi, Ltd. The Hitachi Data Systems logo is a trademark of Hitachi, Ltd.

Dynamic Provisioning, ShadowImage, and TrueCopy are registered trademarks or trademarks of Hitachi Data Systems.

All other brand or product names are or may be trademarks or service marks of and are used to identify products or services of their respective owners.



Contents

Preface	ix
Overview of Performance Manager	1-1
Performance Monitor	1-2
Volume Migration	1-2
Server Priority Manager	1-3
About Performance Manager Operations	2-1
Components	2-2
Overview of Performance Monitor	2-3
Understanding Statistical Storage Ranges	2-3
Parity Group Usage Statistics	2-4
Volume Usage Statistics	2-5
External Volume Group Usage Statistics	2-5
External Volume Usage Statistics	2-6
Channel Processor Usage Statistics	2-6
Disk Processor Usage Statistics	2-6
DRR Processor Usage Statistics	2-9
Write Pending Rate and Cache Memory Usage Statistics	2-9
Access Path Usage Statistics	2-10
Hard Disk Drive Workload Statistics	2-11
Port Traffic Statistics	2-12
LU Paths Traffic Statistics	2-12
Traffic between HBAs and Storage System Ports	2-13
Overview of Volume Migration	2-13
Volume Migration to Balance Workloads	2-13
Overview of Server Priority Manager	2-26
Performance of High-Priority Hosts	2-26
Upper-Limit Control	2-27
Overview of Export Tool	2-27


Interoperability with Other Products.....	2-29
Performance Monitor.....	2-29
Volume Migration.....	2-30
Server Priority Manager	2-46
Preparing for Performance Manager Operations.....	3-1
System Requirements	3-2
Storage Partition Administrators Limitations.....	3-3
Performance Monitor Limitations	3-3
Export Tool Limitations.....	3-5
Using the Performance Manager GUI	4-1
Using the Performance Monitor Windows.....	4-2
Performance Monitor Window	4-2
Performance Management Window, Physical Tab	4-3
LDEV Tab of the Performance Monitor Window.....	4-10
Port-LUN Tab of the Performance Monitor Window	4-16
WWN Tab of the Performance Monitor Window	4-23
Monitoring Options Window	4-27
Other Windows.....	4-31
Using the Volume Migration Windows.....	4-32
Manual Plan Window of the Volume Migration	4-33
Auto Plan Window of the Volume Migration.....	4-43
Attribute Window of the Volume Migration	4-48
History Window of the Volume Migration	4-56
Using the Server Priority Manager Windows.....	4-58
Port Tab of the Server Priority Manager Window.....	4-58
WWN Tab of the Server Priority Manager Window	4-62
Performance Monitor Operations	5-1
Overview of Performance Monitor Operations	5-2
Start Monitoring.....	5-2
View the Monitoring Results.....	5-4
Starting and Stopping Storage System Monitoring.....	5-6
Monitoring Resources in the Storage System.....	5-8
Viewing Usage Statistics on Parity Groups	5-8
Viewing Usage Statistics on Volumes in Parity Groups	5-10
Viewing Usage Statistics on External Volume Groups.....	5-14
Viewing Usage Statistics on External Volumes in External Volume Groups.....	5-16
Viewing Usage Statistics on Channel Processors	5-18
Viewing Usage Statistics on Disk Processors	5-19
Viewing Usage Statistics on Data Recovery and Reconstruction Processors ...	5-21

Viewing Write Pending and Cache Memory Usage Statistics.....	5-22
Viewing Usage Statistics on Access Paths.....	5-25
Monitoring Hard Disk Drives	5-27
Viewing I/O Rates for Disks	5-27
Viewing Transfer Rates for Disks.....	5-30
Monitoring Ports.....	5-34
Viewing I/O Rates for Ports	5-34
Viewing Transfer Rates for Ports	5-37
Viewing Details about the I/O and Transfer Rates	5-41
Monitoring LU Paths	5-42
Viewing LU Paths I/O Rates.....	5-42
Viewing LU Paths Transfer Rates.....	5-44
Monitoring Paths between Host Bus Adapters and Ports.....	5-46
Viewing HBA Information	5-46
Viewing I/O Rates between HBAs.....	5-46
Viewing Transfer Rates between HBAs	5-49
Volume Migration Operations.....	6-1
Starting Volume Migration.....	6-2
Reserving Migration Destinations for Volumes.....	6-3
Migrating Volumes Manually	6-5
Automating Volume Migrations.....	6-8
Setting a Fixed Parity Group	6-8
Changing the Maximum Disk Usage Rate for Each HDD Class.....	6-9
Setting Auto Migration Parameters	6-9
Remaking Auto Migration Plans.....	6-11
Viewing the Migration History Log	6-12
Server Priority Manager Operation	7-1
Overview of Server Priority Manager Operations	7-2
If One-to-One Connections Link HBAs and Ports.....	7-3
If Many-to-Many Connections Link HBAs and Ports.....	7-6
Starting Server Priority Manager	7-11
Port Tab Operations	7-13
Analyzing Traffic Statistics	7-13
Setting Priority for Ports on the Storage System	7-14
Setting Upper-Limit Values to Traffic at Non-prioritized Ports.....	7-15
Setting a Threshold.....	7-16
WWN Tab Operations	7-18
Monitoring All Traffic between HBAs and Ports.....	7-18
Analyzing Traffic Statistics	7-22
Setting Priority for Host Bus Adapters.....	7-23
Setting Upper-Limit Values for Non-Prioritized WWNs	7-24

Setting a Threshold	7-27
Changing the SPM Name of a Host Bus Adapter	7-28
Replacing a Host Bus Adapter	7-29
Grouping Host Bus Adapters	7-30
Using the Export Tool	8-1
Files to be Exported	8-2
Preparing for Using the Export Tool	8-17
Requirements for Using the Export Tool	8-17
Installing the Export Tool on a Windows Computer	8-18
Installing the Export Tool on a UNIX Computer	8-19
Using the Export Tool	8-20
Preparing a Command File	8-20
Preparing a Batch File	8-24
Running the Export Tool	8-25
Command Reference	8-30
Command Syntax	8-31
svpip Subcommand	8-33
retry Subcommand	8-34
login Subcommand	8-35
show Subcommand	8-35
group Subcommand	8-38
short-range Subcommand	8-52
long-range Subcommand	8-55
outpath Subcommand	8-59
option Subcommand	8-60
apply Subcommand	8-61
set subcommand	8-61
help Subcommand	8-63
Java Command for Exporting Data In Files	8-63
Using CCI for Manual Volume Migration	9-1
Sample Volume Migration	9-2
Volume Migration Errors	9-4
Troubleshooting	10-1
Troubleshooting Performance Monitor	10-2
Troubleshooting Volume Migration	10-2
Troubleshooting Server Priority Manager	10-3
Troubleshooting the Export Tool	10-3
Calling the Hitachi Data Systems Support Center	10-8

Acronyms and Abbreviations

Index



Preface

This document describes and provides instructions for using the following Performance Manager software for performing operations on the Hitachi Universal Storage Platform V and Hitachi Universal Storage Platform VM (USP V/VM) storage systems.

- Performance Monitor
- Volume Migration
- Server Priority Manager (henceforth, referred to as *SPM*)

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

This preface includes the following information:

- ☐ [Intended Audience](#)
- ☐ [Product Version](#)
- ☐ [Document Revision Level](#)
- ☐ [Source Document\(s\) for this Revision](#)
- ☐ [Changes in this Revision](#)
- ☐ [Document Organization](#)
- ☐ [Referenced Documents](#)
- ☐ [Document Conventions](#)
- ☐ [Convention for Storage Capacity Values](#)
- ☐ [Getting Help](#)
- ☐ [Comments](#)

Notice: The use of Performance Manager software and all other Hitachi Data Systems products is governed by the terms of your agreement(s) with Hitachi Data Systems.

Intended Audience

This document is intended for system administrators, Hitachi Data Systems representatives, and Authorized Service Providers who are involved in installing, configuring, and operating the Hitachi Universal Storage Platform V and VM storage systems.

This document assumes the following:

- The user has a background in data processing and understands RAID storage systems and their basic functions.
- The user is familiar with the Universal Storage Platform V and/or VM storage system and has read the *Universal Storage Platform V/VM User and Reference Guide*.
- The user is familiar with the Storage Navigator software for the Universal Storage Platform V/VM and has read the *Storage Navigator User's Guide*.
- The user is familiar with the operating system and Web browser software on the system hosting the Hitachi Universal Storage Platform V/VM Storage Navigator remote console software.



Note: There are different types of users for Hitachi Universal Storage Platform V/VM: storage administrators and storage partition administrators. The functions described in this manual are limited depending on the user type. For details on the limitations, see Storage Partition Administrators Limitations. For details on the user types, see the *Storage Navigator User's Guide*.

Product Version

This document revision applies to Universal Storage Platform V/VM microcode 60-02-4x and higher.

Document Revision Level

Revision	Date	Description
MK-96RD617-P	February 2007	Preliminary Release
MK-96RD617-00	April 2007	Initial Release, supersedes and replaces MK-96RD617-P
MK-96RD617-01	June 2007	Revision 1, supersedes and replaces MK-96RD617-00
MK-96RD617-02	July 2007	Revision 2, supersedes and replaces MK-96RD617-01
MK-96RD617-03	September 2007	Revision 2, supersedes and replaces MK-96RD617-02
MK-96RD617-04	November 2007	Revision 4, supersedes and replaces MK-96RD617-03
MK-96RD617-05	January 2008	Revision 5, supersedes and replaces MK-96RD617-04
MK-96RD617-06	March 2008	Revision 6, supersedes and replaces MK-96RD617-05

Source Documents for this Revision

- MK-96RD617-06c-RSD-V02

Changes in This Revision

- Added information about replacing the microprogram in [Performance Monitor](#) and [Volume Migration Errors](#).
- Added information about the Quick Format function in [Volume Migration](#).
- Added information about maintenance in [Volume Migration](#).
- Added information about IPv6 in [svpip Subcommand](#).
- Added a note in [Volume Migration Errors](#).
- Updated information about Server Priority Manager in [Using the Server Priority Manager Windows](#), [Port Tab Operations](#), and [WWN Tab Operations](#).
- Added information about the volume set as a system disk in [Manual Volume Migration](#), [Volume Migration](#), [Using the Volume Migration Windows](#), and [Attribute Window of the Volume Migration](#).
- Added notes about the Volume Migration using CCI in [Volume Migration Errors](#).
- Added notes about the source volume of Volume Migration in [Performance Monitor](#).
- Added information about the volume that can be specified as a target volume in [Volume Migration](#).

Document Organization

The following table provides an overview of the contents and organization of this document. Click the [chapter title](#) in the left column to go to that chapter. The first page of each chapter provides links to the sections in that chapter.

Chapter	Description
Overview of Performance Manager	Describes the performance management software products that allow you to monitor and tune storage system performance.
About Performance Manager Options	Provides an overview of performance manager operations.
Preparing for Performance Manager Operations	Explains the preparations for performance manager operations.
Using the Performance Manager GUI	Explains performance manager windows.
Performance Monitor Operations	Explains performance monitor operations.
Volume Migration Operations	Explains volume migration operations.

Chapter	Description
Server Priority Manager Operation	Explains server priority manager operations.
Using the Export Tool	Explains using the export tool.
Using CCI for Manual Volume Migration	Explains the procedure and notes for using CCI to perform manual volume migration.
Troubleshooting	Provides troubleshooting information on Performance Monitor, Volume Migration Server Priority Manager and Export Tool.
Acronyms and Abbreviations	Defines the acronyms and abbreviations used in this document.
Index	Lists the topics in this document in alphabetical order.

Referenced Documents

Hitachi Universal Storage Platform V/VM:

- *Hitachi Command Control Interface (CCI) User and Reference Guide*, MK-90RD011
- *Hitachi Compatible Mirroring for IBM FlashCopy User's Guide*, MK-96RD614
- *Hitachi Copy-on-Write Snapshot User's Guide*, MK-96RD607
- *Hitachi Dynamic Provisioning User's Guide*, MK-96RD641
- *Hitachi LUN Manager User's Guide*, MK-96RD615
- *Hitachi ShadowImage for IBM® z/OS® User's Guide*, MK-96RD619
- *Hitachi ShadowImage User's Guide*, MK-96RD618
- *Hitachi Storage Navigator Messages*, MK-96RD613
- *Hitachi Storage Navigator User's Guide*, MK-96RD621
- *Hitachi TrueCopy for IBM z/OS User's Guide*, MK-96RD623
- *Hitachi TrueCopy User's Guide*, MK-96RD622
- *Hitachi Universal Replicator for IBM z/OS User's Guide*, MK-96RD625
- *Hitachi Universal Replicator User's Guide*, MK-96RD624
- *Hitachi Virtual Partition Manager User's Guide*, MK-96RD629

Document Conventions

The terms "Universal Storage Platform V" and "Universal Storage Platform VM" refer to all models of the Hitachi Universal Storage Platform V and VM storage systems, unless otherwise noted.

This document uses the following typographic conventions:

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

This document uses the following icons to draw attention to information:

Icon	Meaning	Description
	Note	Calls attention to important and/or additional information.
	Tip	Provides helpful information, guidelines, or suggestions for performing tasks more effectively.
	Caution	Warns the user of adverse conditions and/or consequences (e.g., disruptive operations).
	WARNING	Warns the user of severe conditions and/or consequences (e.g., destructive operations).

Convention for Storage Capacity Values

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

- 1 KB = 1,000 bytes
- 1 MB = 1,000² bytes
- 1 GB = 1,000³ bytes
- 1 TB = 1,000⁴ bytes
- 1 PB = 1,000⁵ bytes

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

- 1 KB = 1,024 bytes
- 1 MB = 1,024² bytes
- 1 GB = 1,024³ bytes
- 1 TB = 1,024⁴ bytes
- 1 PB = 1,024⁵ bytes
- 1 block = 512 bytes

Getting Help

If you need to call 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 content of any error message(s) displayed on the host system(s).
- The content of any error message(s) displayed on Storage Navigator.
- The USP V/VM Storage Navigator configuration information saved on diskette using the FD Dump Tool (see the *Storage Navigator User's Guide*).
- The service information messages (SIMs), including reference codes and severity levels, displayed by Storage Navigator.

The Hitachi Data Systems customer support staff is available 24 hours/day, seven days a week. If you need technical support, please call:

- United States: (800) 446-0744
- Outside the United States: (858) 547-4526

Comments

Please send us your comments on this document. Make sure to include the document title, number, and revision. Please refer to specific section(s) and paragraph(s) whenever possible.

- **E-mail:** doc.comments@hds.com
- **Fax:** 858-695-1186
- **Mail:**
Technical Writing, M/S 35-10
Hitachi Data Systems
10277 Scripps Ranch Blvd.
San Diego, CA 92131

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

Overview of Performance Manager

The Hitachi Universal Storage Platform V and Hitachi Universal Storage Platform VM (herein after referred to as USP V/VM) includes a suite of performance management software products that allow you to monitor and tune storage system performance.

The Performance Manager suite includes the following:

- [Performance Monitor](#)
- [Volume Migration](#)
- [Server Priority Manager](#)

Restrictions: The Auto Migration function is not supported in this version.

Performance Monitor

Performance Monitor lets you obtain usage statistics about physical hard disk drives, volumes, processors or other resources in your storage system. Performance Monitor also lets you obtain statistics about workloads on disk drives and traffic between hosts and the storage system. The **Performance Management** window displays a line graph that indicates changes in the usage rates, workloads, or traffic. You can view information in the window and analyze trends in disk I/Os and detect peak I/O time. If system performance is poor, you can use information in the window to detect bottlenecks in the system.



Note: When using Performance Monitor, you must specify the volumes to be monitored in CUs. Therefore, depending on your disk subsystem configuration, the list may display performance statistics for some volumes and not display performance statistics for other volumes. This can occur, if the range of used CUs does not match the range of CUs monitored by Performance Monitor.

To correctly display performance statistics of a parity group and a LUSE volume, you must specify as follows:

- Specify all volumes that belong to the parity group as the monitoring targets.
 - Specify all volumes that make up the LUSE volume as the monitoring targets.
-

Volume Migration

Volume Migration lets you balance workload among hard disk drives, volumes (LUs) and processors to remove bottlenecks from your system. If Performance Monitor indicates that a lot of I/Os are made to particular hard disk drives, you can use Volume Migration to distribute workloads to other disk drives. For details, see Overview of Volume Migration.

Server Priority Manager

Server Priority Manager lets you tune the system to provide high-priority hosts with relatively higher throughput. Server Priority Manager can prevent production servers from suffering lowered performance. For details, see Overview of Server Priority Manager.

Figure 1-1 illustrates the performance management solution from Hitachi Data Systems.




Figure 1-1 Performance Management Solution

Figure 1-1 illustrates and simplifies the performance management process.




Figure 1-2 Performance Management Process Flow Diagram

About Performance Manager Operations

This chapter gives an overview of performance manager operations.

- [Components](#)
- [Overview of Performance Monitor](#)
- [Overview of Volume Migration](#)
- [Overview of Server Priority Manager](#)
- [Overview of Export Tool](#)
- [Interoperability with Other Products](#)

Components

To be able to use Performance Manager, you need:

- The USP V/VM storage system.
- The Performance Manager program products (At minimum, Performance Monitor is required. Volume Migration and Server Priority Manager are optional).
- A WWW client computer connected to the USP V/VM storage system via LAN.

To use Performance Manager, you must use the WWW client computer to log on to the SVP. When you are logged on, the Storage Navigator program, which is a Java™ applet, will be downloaded to the WWW client computer. You can perform Performance Manager operations in the Storage Navigator window.

For details about requirements for WWW client computers, see the *Storage Navigator User's Guide*.



Cautions:

- If Performance Monitor is not enabled, you cannot use Volume Migration or Server Priority Manager.
 - Performance management operations (Performance Monitor, Volume Migration, and Server Priority Manager) involve the collection of large amounts of monitoring data. This requires considerable Web client computer memory. It is therefore recommended that you exit the Storage Navigator program when not conducting performance management operations, to release system memory.
-

Overview of Performance Monitor

Performance Monitor tracks your storage system and lets you obtain statistics about the following:

- resources in your storage system
- workloads on disk and ports

If your system encounters some problem (for example, if server hosts suffer delayed response times), Performance Monitor can help you detect the cause of the problem.



Notes:

- An Export Tool enables you to save information on the **Performance Management** window into files, so you can use spreadsheet or database software to analyze the monitoring results. For detailed information about the Export Tool, see Overview of Export Tool.
 - Performance Monitor can also display the status of remote copies by TrueCopy, TrueCopy for IBM® z/OS®, Universal Replicator, and Universal Replicator for IBM z/OS. The displayed contents are the same as those displayed in the **Usage Monitor** windows of each remote copy function.
-

Understanding Statistical Storage Ranges

Performance Monitor has two kinds of periods (ranges) for collecting and storing statistics: *short range* and *long range*. The difference of the two ranges and the statistics at which they are targeted is as follows:

- Storing in short range

If the number of CUs to be monitored is 64 or less, Performance Monitor collects statistics at a user-specified interval that is between 1 and 15 minutes, and stores them between 1 and 15 days.

If the number of CUs to be monitored is 65 or more, Performance Monitor collects statistics at a user-specified interval that is 5, 10 or 15 minutes, and stores them between 8 hours and 1 day.

All the statistics that can be monitored by Performance Monitor are collected and stored in short range.

- Storing in long range

Performance Monitor collects statistics at fixed 15-minutes interval, and stores them for 3 months (i.e., 93 days).

The usage statistics about resources in the storage system are collected and stored also in long range, in parallel with in short range. However, some of the usage statistics about resources cannot be collected in long range.

Performance Management window can display the statistics within the range of the storing periods above. You can specify a part of the storing period to display the statistics on the lists and graphs of Performance Monitor.

All statistics, except some information related to Volume Migration, can be displayed in short range (for the storing period corresponding to the setting for the collecting interval) on **Performance Management** window. In addition, usage statistics about resources in the storage system can be displayed in both short range and long range because they are monitored in both ranges. When you display usage statistics about resources, you can select the displayed range.

For more about statistics that can be monitored in short and long ranges, see the description of **Performance Management** window in Performance Monitor Window. For more about the relationship between collection interval and the storing period of the statistics, see Monitoring Options Window.

Parity Group Usage Statistics

A *parity group* is a group of hard disk drives (HDDs) that form the basic unit of storage for the USP V/VM storage system. All HDDs in a parity group must have the same physical capacity. The USP V/VM supports three types of parity groups as follows:

- RAID-1 parity group
A RAID-1 parity group consists of two pairs of HDDs in a mirrored configuration.
- RAID-5 parity group
A RAID-5 parity group consists of four or eight HDDs. One of these HDDs is used as a parity disk.
- RAID-6 parity group
A RAID-6 parity group consists of eight HDDs. Two of these HDDs are used as a parity disk.

If the monitor data shows overall high parity group usage, you should consider installing additional HDDs and using Volume Migration to migrate the high-usage volumes to the new parity groups. If the monitor data shows that parity group usage is not balanced, you can use Volume Migration to migrate volumes from high-usage parity groups to low-usage parity groups.

For details on how to view usage statistics about parity groups, see Viewing Usage Statistics on Parity Groups.

Volume Usage Statistics

Performance Monitor displays the average and maximum usage, including sequential and random access, of each volume (LDEV) in a parity group. The volume usage is the time in use (sequential and random access) of the physical drives of each LDEV, averaged by the number of physical drives in the parity group.

If the monitor data shows overall high volume usage, you should consider installing additional hardware (e.g., HDDs, DKAs, cache). If the monitor data shows that volume usage is not balanced, you can use Volume Migration to migrate high-usage volumes to higher HDD classes and/or to lower-usage parity groups. The volume usage data can also be used to analyze the access characteristics of volumes and determine the appropriate RAID level and/or HDD type for the volumes.

For details on how to view usage statistics about volumes, see Viewing Usage Statistics on Volumes in Parity Groups.

External Volume Group Usage Statistics

If the USP V/VM storage system is connected to an external storage system by Universal Volume Manager, Performance Monitor can also monitor the usage conditions on external hard disk drives.

When you use Universal Volume Manager to map the volumes in the external storage system as volumes in the USP V/VM storage system, the mapped volumes in the external storage system are called *external volumes*. These external volumes are registered in groups by Universal Volume Manager. Performance Monitor can monitor the usage conditions for external volume groups.



Notes:

- An external volume group is just a group for managing external volumes. Unlike a parity group, it does not contain any parity information. However, some Performance Monitor windows treat external volume groups as parity groups.
 - The information that can be monitored about an external volume group differs from that of a usual parity group. For details on how to view usage conditions about external volume groups, see Viewing Usage Statistics on External Volume Groups.
-

External Volume Usage Statistics

An *external volume* is a volume existing in an external storage system that is mapped to a volume in the USP V/VM storage system using Universal Volume Manager. Performance Monitor can monitor and display the usage conditions for external volumes.

The information that can be monitored for an external volume differs from that of a typical volume. For details on how to view usage conditions about external volumes, see Viewing Usage Statistics on External Volumes in External Volume Groups.

Channel Processor Usage Statistics

A channel processor (CHP), which is contained in a channel adapter (CHA), processes host commands and controls data transfer between hosts and the cache. A channel adapter contains multiple channel processors that process host commands and control data transfer.

If monitoring data shows high overall CHP usage, you should consider installing additional CHAs. If monitoring data shows that CHP usage is not balanced, you should consider moving some devices that are defined on overloaded ports to ports with lower-usage CHPs to balance front-end usage.

For details on how to view usage statistics about channel adapters and channel processors groups, see Viewing Usage Statistics on Channel Processors.



Note: A channel adapter can also be called *port controller*.

Disk Processor Usage Statistics

A disk processor (DKP), which is contained in a disk adapter (DKA), controls data transfer between the cache and the disk devices. A disk adapter contains multiple disk processors (DKPs).

If monitor data shows high DKP usage overall, you should consider installing additional HDDs and/or DKAs, and then using Volume Migration to migrate the high-write-usage volumes (especially sequential writes) to the new parity groups. If the monitor data shows that DKP usage is not balanced, you can use Volume Migration to migrate volumes from high-usage parity groups to low-usage parity groups.

When you consider migrating a volume from one parity group to another, take the following steps:

1. Refer to Table 2-1 to determine the parity groups from which you want to migrate volumes.

2. Check the usage rate of each parity group to find parity groups whose usage rate is lower than the parity groups that you want to migrate. It is recommended you migrate volumes from higher-usage parity groups to lower-usage parity groups.



Note: The information in Table 2-1 is not applicable when you are using the USP VM storage system because USP VM has only 1 pair of DKA. If you are using USP VM and you want to know usage rates of DKPs, you must check the usage rate of each parity group.

Table 2-1 How to Migrate Volumes When the Usage Rate of a Disk Processor is High (This Is Applicable When You Are Using the USP V Storage System)

Cluster	Disk Adapter	Disk Processor	How to Migrate Volumes
1	DKA1AU	DKP40-1AU DKP41-1AU DKP42-1AU DKP43-1AU	Migrate volumes from parity groups with the following IDs to another parity group: 1-X 3-X 11-X
1	DKA-AL	DKP44-1AL DKP45-1AL DKP46-1AL DKP47-1AL	Migrate volumes from parity groups with the following IDs to another parity group: 2-X 4-X 12-X
1	DKA-BU	DKP50-1BU DKP51-1BU DKP52-1BU DKP53-1BU	Migrate volumes from parity groups with the following IDs to another parity group: 5-X 13-X
1	DKA-BL	DKP54-1BL DKP55-1BL DKP56-1BL DKP57-1BL	Migrate volumes from parity groups with the following IDs to another parity group: 6-X 14-X
1	DKA-1LU	DKP60-1LU DKP61-1LU DKP62-1LU DKP63-1LU	Migrate volumes from parity groups with the following IDs to another parity group: 7-X 15-X
1	DKA-1LL	DKP64-1LL DKP65-1LL DKP66-1LL DKP67-1LL	Migrate volumes from parity groups with the following IDs to another parity group: 8-X 16-X
1	DKA-1KU	DKP70-1KU DKP71-1KU DKP72-1KU DKP73-1KU	Migrate volumes from parity groups with the following IDs to another parity group: 9-X 17-X

Cluster	Disk Adapter	Disk Processor	How to Migrate Volumes
1	DKA-1KL	DKP74-1KL DKP75-1KL DKP76-1KL DKP77-1KL	Migrate volumes from parity groups with the following IDs to another parity group: 10-X 18-X
2	DKA-2MU	DKPC0-2MU DKPC1-2MU DKPC2-2MU DKPC3-2MU	Migrate volumes from parity groups with the following IDs to another parity group: 1-X 3-X 11-X
2	DKA-2ML	DKPC4-2ML DKPC5-2ML DKPC6-2ML DKPC7-2ML	Migrate volumes from parity groups with the following IDs to another parity group: 2-X 4-X 12-X
2	DKA-2NU	DKPD0-2NU DKPD1-2NU DKPD2-2NU DKPD3-2NU	Migrate volumes from parity groups with the following IDs to another parity group: 5-X 13-X
2	DKA-2NL	DKPD4-2NL DKPD5-2NL DKPD6-2NL DKPD7-2NL	Migrate volumes from parity groups with the following IDs to another parity group: 6-X 14-X
2	DKA-2XU	DKPE0-2XU DKPE1-2XU DKPE2-2XU DKPE3-2XU	Migrate volumes from parity groups with the following IDs to another parity group: 7-X 15-X
2	DKA-2XL	DKPE4-2XL DKPE5-2XL DKPE6-2XL DKPE7-2XL	Migrate volumes from parity groups with the following IDs to another parity group: 8-X 16-X
2	DKA-2WU	DKPF0-2WU DKPF1-2WU DKPF2-2WU DKPF3-2WU	Migrate volumes from parity groups with the following IDs to another parity group: 9-X 17-X
2	DKA-2WL	DKPF4-2WL DKPF5-2WL DKPF6-2WL DKPF7-2WL	Migrate volumes from parity groups with the following IDs to another parity group: 10-X 18-X
 Note: The letter "X" is a placeholder for numerical values. For example, "parity group 1-X" indicates parity groups such as 1-1 and 1-2.			

For details on how to view usage statistics about disk adapters and disk processors, see Viewing Usage Statistics on Disk Processors.



Note: Volume Migration cannot estimate DKP usage, and may not provide any performance improvement for cases in which DKP usage values vary only slightly or for cases in which overall DRR usage values are relatively high. Volume Migration is designed for use with obvious cases of high or unbalanced DKP usage.

DRR Processor Usage Statistics

A *data recovery and reconstruction processor* (DRR) is a microprocessor (located on the DKAs) that is used to generate parity data for RAID-5 or RAID-6 parity groups. The DRR uses the formula "old data + new data + old parity" to generate new parity.

If the monitor data shows high DRR usage overall, this can indicate high write penalty condition. Please consult your Hitachi Data Systems representative about high write penalty conditions. If the monitor data shows that DRR usage is not balanced, you should consider relocating volumes using Volume Migration to balance DRR usage within the storage system.

For details on how to view usage statistics on DRRs, see Viewing Usage Statistics on Data Recovery and Reconstruction Processors.

Write Pending Rate and Cache Memory Usage Statistics

The write pending rate indicates the ratio of write-pending data to the cache memory capacity. The **Performance Management** window displays the average and the maximum write pending rate for the specified period of time. When you display monitoring results in a short range, the window also displays the average and the maximum usage statistics about the cache memory for the specified period of time. In addition, the window can display a graph that indicates how the write pending rate or the usage statistics of the cache memory changed within that period.

For details on how to view the write pending rate and the usage statistics about the cache memory, see Viewing Write Pending and Cache Memory Usage Statistics.

Access Path Usage Statistics

An *access path* is a path through which data and commands are transferred within a storage system.

In a storage system, channel adapters control data transfer between hosts and the cache memory. Disk adapters control data transfer between the cache memory and hard disk drives. Data transfer does not occur between channel adapters and disk adapters. Data is transferred via the cache switch (CSW) to the cache memory.

When hosts issue commands, the commands are transferred via channel adapters to the shared memory (SM). The content of the shared memory is checked by disk adapters.




Figure 2-1 Access Paths

Performance Monitor tracks and displays the usage rate for the following access paths.

- Paths between channel adapters and the cache switch
- Paths between disk adapters and the cache switch
- Paths between the cache switch and the cache memory
- Paths between channel adapters and the shared memory
- Paths between disk adapters and the shared memory

For details on how to view usage statistics about access paths, see Viewing Usage Statistics on Access Paths.

Hard Disk Drive Workload Statistics

If particular hard disk drives or data are heavily accessed, system performance might deteriorate. Performance Monitor lets you view statistics about parity groups and logical devices to help you detect bottlenecks in your system. If you mapped volumes in an external storage system, Performance Monitor can also monitor the access workloads of the external volume groups and the external volumes. Performance Monitor displays a line graph that indicates changes in access workloads, so that you can detect the peak I/O access times.



Note: You will be unable to view workload statistics that expired a specific storing period because such statistics are erased from the storage system. The storing period of statistics is only short range (between 8 hours to 15 days) and that changes depending on the collecting interval and the number of CUs to be monitored specified by the user. For details on the storing period of statistics, see Understanding Statistical Storage Ranges.

Workload information mainly displayed by Performance Monitor is as follows:

- I/O rate

The *I/O rate* indicates how many I/Os are made to the hard disk drive in one second. If the I/O rate is high, the hosts might consume a lot of time for accessing disks and the response time might be long.

- Transfer rate

The *transfer rate* indicates the size of data transferred to the hard disk drive in one second. If the transfer rate is high, the hosts might consume a lot of time for accessing disks and the response time might be long.

- The read hit ratio

For a read I/O, when the requested data is already in cache, the operation is classified as a read hit. For example, if ten read requests have been made from hosts to devices in a given time period and the read data was already on the cache memory three times out of ten, the read hit ratio for that time period is 30 percent. A higher read hit ratio implies higher processing speed because fewer data transfers are made between devices and the cache memory.

- The write hit ratio

For a write I/O, when the requested data is already in cache, the operation is classified as a write hit. For example, if ten write requests were made from hosts to devices in a given time period and the write data was already on the cache memory three cases out of ten, the write hit ratio for that time period is 30 percent. A higher write hit ratio implies higher processing speed because fewer data transfers are made between devices and the cache memory.

Apart from the items listed above, Performance Monitor also displays additional information about hard disk drive workloads. For details on how to view workload statistics about hard disk drives, see [Monitoring Hard Disk Drives](#).

Port Traffic Statistics

Performance Monitor tracks host ports and storage system ports to obtain statistics about I/O rates and transfer rates at these ports. If you analyze these I/O rates and transfer rates, you can determine which hosts issue a lot of I/O requests to the disk and which hosts transfer a lot of data to the disk. For details on how to view statistics about traffic at ports, see [Monitoring Ports](#).



Note: You will be unable to view workload statistics that expired a specific storing period because such statistics are erased from the storage system. The storing period of statistics is only short range (between 8 hours to 15 days) and that changes depending on the collecting interval specified by the user. For details on the storing period of statistics, see [Understanding Statistical Storage Ranges](#).

Important: Performance Monitor can obtain statistics about traffics of ports connected to open-system host groups only. The statistics about traffics of ports connected to mainframe host groups cannot be obtained.

LU Paths Traffic Statistics

Performance Monitor tracks LU paths to obtain statistics about I/O rates and transfer rates at these LU paths. If you analyze these I/O rates and transfer rates, you can detect LU paths through which a lot of I/O requests are made to the disk. You can also determine the LU paths through which a lot of data are transferred to the disk. For details on how to view workload statistics about LU paths, see [Monitoring LU Paths](#).



Note: You will be unable to view workload statistics that expired a specific storing period because such statistics are erased from the storage system. The storing period of statistics is only short range (between 8 hours to 15 days) and that changes depending on the collecting interval specified by the user. For details on the storing period of statistics, see [Understanding Statistical Storage Ranges](#).



Note: The traffic statistics reported for an LU is aggregated across all LU paths defined for an LU.

- I/O rate is the sum of I/Os across all LU paths defined for an LU.
 - Transfer rate is the total transfer rate across all LU paths defined for an LU.
 - **Response Time** is the average response time across all LU paths defined for an LU.
-

Traffic between HBAs and Storage System Ports

Host bus adapters (HBAs) are adapters contained in hosts. HBAs, which serve as ports on hosts, are connected to ports on the storage system.

If Server Priority Manager is enabled, Performance Monitor lets you view statistics about traffic between HBAs and storage system ports. The traffic statistics reveals the number of I/O requests that have been made from hosts and also reveals the size of data transferred between hosts and storage system ports. For details on how to view traffic statistics about HBAs, see Viewing HBA Information.



Note: You will be unable to view workload statistics that expired a specific storing period because such statistics are erased from the storage system. The storing period of statistics is only short range (between 8 hours to 15 days) and that changes depending on the collecting interval specified by the user. For details on the storing period of statistics, see Understanding Statistical Storage Ranges.

Overview of Volume Migration

Volume Migration enables you to optimize your data storage and migrate volumes on the storage system. Volume Migration tuning operations can be used to balance resource utilization and resolve bottlenecks of system activity.

Volume Migration to Balance Workloads

If I/O accesses from hosts converge on particular hard disk drives, the hosts might suffer from slower response times. To overcome this situation, Volume Migration lets you migrate high-usage volumes to a low-usage hard disk drive or a fast hard disk drive. Volume migration can balance workloads among hard disk drives and thus remove bottlenecks from the system. Volume migration can also balance workloads among disk processors, which control disk I/Os.

Volume Migration operations are either manual migrations or automated (auto) migrations.

- In manual migrations, the system administrator specifies the volumes to be migrated and the migration destinations, and then executes migrations.
- In auto migrations, Volume Migration automatically determines the volumes to be migrated and the migration destinations, and also executes migrations automatically.

**Notes:**

- Manual migration supports external volumes, but auto migration does not. Even in manual migration, you cannot examine the estimated usage rate after migration because the usage rates of external volumes cannot be collected. For details on external volumes, see External Volume Group Usage Statistics.
- Manual migration supports Dynamic Provisioning virtual volumes, but auto migration does not. If source volumes for manual migration are Dynamic Provisioning virtual volumes, the target volumes can be internal volumes, external volumes, or Dynamic Provisioning virtual volumes.
- Volume Migration operations are completely non-disruptive - the data being migrated can remain online to all hosts for read and write I/O operations throughout the entire volume migration process.
- Keep in mind that storage system-tuning operations can improve performance in one area while at the same time decreasing performance in another. Suppose that parity groups A and B have average usage values of 20% and 90%, respectively. Volume Migration estimates that if one volume is migrated from parity group B to parity group A, the usage values will become 55% and 55%. If you perform this migration operation, the I/O response time for parity group B will probably decrease, and the I/O response time for parity group A may increase, while the overall throughput may increase or decrease.
- Volume Migration should be performed only when you can expect a large improvement in storage system performance. Volume Migration may not provide significant improvement for cases in which parity group or volume usage varies only slightly, or for cases in which overall DKP or DRR usage is relatively high.
- From an open-system host, users can also use Command Control Interface (CCI) to perform manual volume migration by using commands. To perform volume migration by CCI, Volume Migration should be installed on the storage system. For details on manual volume migration by using CCI, see [Using CCI for Manual Volume Migration](#). For details on CCI, see the *Command Control Interface (CCI) User and Reference Guide*.



Caution: When an error condition exists in the USP V/VM storage system, resource usage can increase or become unbalanced. Do not use data collected during an error condition as the basis for planning Volume Migration operations.

Manual Volume Migration

This section explains how to migrate volumes in parity groups manually, citing usage rates before and after migration.

Volume Migration enables you to balance workloads among parity groups to improve system performance. To balance workloads, you select high-usage volumes from a high-usage parity group and then move the volumes to a different parity group.



Notes:

- You can also migrate a Dynamic Provisioning virtual volume or an external volume manually. However, you cannot examine the estimated usage rate after migration because the usage rate of an external volume cannot be collected.
- In some cases, the volumes of OPEN-V emulation type might not be displayed as target volumes. If Inquiry commands are executed to volumes displayed as OPEN-V in the Volume Migration window, the replies from some volumes could be OPEN-V and from other volumes OPEN-0V. You cannot distinguish OPEN-V volumes from OPEN-0V volumes in the LDEV's list of the Volume Migration window. If you want to distinguish between OPEN-V and OPEN-0V volumes, open the Basic Information Display dialog box in Storage Navigator.
- The volume migration between OPEN-V and OPEN-V is permitted, as is between OPEN-0V and OPEN-0V. But the volume migration between OPEN-V and OPEN-0V is not currently permissible.
- The system disk cannot be migrated.

Before migrating a volume from one parity group to another, you must create a migration plan. When you create a migration plan, you specify the following:

- a volume for migrating
- the migration destination for the specified volume

To migrate volumes, you execute your migration plan.

Figure 2-2 gives an example of a migration plan as it would be displayed in the **Manual Plan** window of Volume Migration. In this figure, **Source LDEV** indicates the volume that should be migrated. **Target LDEV** indicates the migration destination for the volume in **Source LDEV**.

Stat	DEL	%	Source LDEV							Target LDEV					
			LDEV	Emulation	Capacity	RAID	PG	HDD	CLPR	LDEV	RAID	PG	HDD	CLPR	Owner
-	-	-	 00:03:00	OPEN-V	2.28 GB	N/A	2-1	DKR1B-J047FC	00:CLPR0	 00:01:00	N/A	1-2	DKR1B-J047FC	01:CLPR1	USP

Figure 2-2 An Example of Migration Plans Displayed in the Volume Migration Window

Before you create a migration plan, you must determine which volume should be migrated to which parity group. When you determine a volume to be migrated, do the following:

1. Check usage rates for volumes.
2. Choose a volume to be migrated from among high-usage volumes.

Manual Migration Plans		LDEVs				
		LDEV	Emulation	Capacity	Ave.(%)(Total)	Max.(%)
Subsystem Parity Group 1-1[5-5](-/- % 00:CLPR0) 1-2[5-1:5-2](-/- % 01:CLPR 1-3[5-3:5-4](-/- % 02:CLPR 1-4(-/- % 03:CLPR3) 2-1(-/- % 00:CLPR0) 2-2(-/- % 01:CLPR1) 3-1(-/- % 02:CLPR2) 3-2(-/- % 03:CLPR3)		00:00:00	OPEN-V	2.28 GB	-	-
		00:00:01	OPEN-V	2.28 GB	-	-
		00:00:02	OPEN-V	2.28 GB	-	-
		00:00:03	OPEN-V	2.28 GB	-	-
		01:00:04	OPEN-V	2.28 GB	-	-
		01:00:05	OPEN-V	2.28 GB	-	-
		01:00:06	OPEN-V	2.28 GB	-	-
		01:00:07	OPEN-V	2.28 GB	-	-

Figure 2-3 Usage Rate for a Volume Displayed in the Window

3. After determining a migration destination, find fast parity groups (that is, fast hard disk drives). Next, you should choose a low-usage parity group from the fast parity groups.

To find fast hard disk drives, display the **Volume Migration** window and check *classes*. *Classes* indicate disk access speed for parity groups. Fastest parity groups (or groups of fastest disks) belong to class A. Parity groups in class B are not as fast as parity groups in class A, but are faster than parity groups in class C. Parity groups in higher classes (such as class A) provide faster disk access than parity groups in lower classes (such as class C).



Note: In the USP V/VM documentation, the terms *class* and *HDD class* mean the same thing.




Figure 2-4 Classes and Parity Groups that Appear in the Volume Migration Window

Next, you should choose a low-usage parity group from the fast parity groups. You can find low-usage volumes in the **Volume Migration** Window.




Figure 2-5 Parity Group Usage Displayed in the Volume Migration Window



Note: When you use Volume Migration for the first time, you must reserve volumes for migration destinations in advance. These reserved volumes are marked with the  icon. When you specify a volume to be migrated, icons of some reserved volumes () that can be used as a migration destination will change to the  icons. For details on reserving volumes for migration destinations, see Reserving Migration Destinations for Volumes.

If you specify the volume you wish to migrate in the window and click **Target**, Volume Migration displays the list of disks of the parity groups that can be used as destinations. If you select the destination disk area and click **Graph**, Volume Migration calculates the estimated volume usage rate after migration, and displays the estimated usage rate in the window. However, Volume Migration cannot estimate the usage rate of any external volumes. In case of migrating an external volume, the transition of the external volume usage rate is not displayed even if you click **Plan**. If the volume usage rate of the migration source cannot be calculated or if you do not need to know the transition of the usage rate, you can migrate the volume without displaying the graph.

Volume Migration uses past resource usage rates, hard disk drive type information, and the RAID level to calculate the estimated usage rate. In the example in Figure 2-6, if the volume 00:03:06 in the parity group 2-1 is migrated to the reserved volume 01:01 in the parity group 1-2, the parity group 2-1 can expect 10 percent decrease in the average usage rate (from 21 percent to 11 percent). The destination parity group 1-2 can expect a 13 percent increase in the average usage rate (from 12 percent to 25 percent). If the estimated result is not satisfactory, you can change the migration destination.




Figure 2-6 Estimated Usage Rate for Parity Groups

You can create multiple migration plans. If there are any other high-usage volumes, you can continue to create migration plans. If you finish creating migration plans, you can execute the migration plans to start volume migrations.

For more detailed information about volume migrations, see *Migrating Volumes Manually*.

Hosts can perform I/O operations on volumes that are being migrated. This section explains why.



Note: The USP V/VM documents (including this manual) use the term *source volumes* to refer to volumes that should be migrated. The term *target volumes* is used for volumes that should become migration destinations.

The Volume Migration operation consists of two steps:

- Copy the data on the Volume Migration source volume to the Volume Migration target volume.
- Transfer host access to the target volume to complete the migration.

The Volume Migration source volume can be online to hosts during migration. The target volume is reserved prior to migration to prevent host access during migration. The source and target volumes can be located anywhere within the storage system.

The Volume Migration copy operation copies the entire contents of the source volume to the target volume cylinder by cylinder (24 tracks at a time, not including the diagnostic and unassigned alternate tracks). If the source volume is updated by write I/Os during the copy operation, the updates will be recorded on the cylinder map of a volume. If there are differential data formed by the update, the differential data will be copied from the source volume to the target volume and this will be repeated until all differential data on the source volume is copied. However, there are upper limits to the number of times the copy is executed and the limit depends on the capacity of the migrated volume. The limit increases as the capacity of the migrated volume increases. When the differential data still exists after copy is repeated to the upper limit, the volume migration fails. Execute the migration by reducing the workloads between hosts and the storage system. The indication of the workloads between hosts and the storage system shall be 50 IOPS or below for updating I/O.

When the volumes are fully synchronized (i.e., when there is no differential data on source volume), the USP V/VM completes the migration by swapping the reserve attribute (target is normal, source is reserved) and redirecting host access to the target volume. Volume Migration performs all migration operations sequentially (i.e., one volume migration at a time).



Note: Approximate copy times with no other I/O activity are OPEN-3 = 1.5 min, OPEN-9 = 4.5 min, OPEN-E = 9 min.



Note: For auto and manual migration operations, Volume Migration checks the current storage system write pending rate and cancels a migration operation if the write pending rate is higher than 60%. For auto migration, Volume Migration also checks the current disk utilization of the source volumes and source and target parity groups, and cancels the auto migration operation if current disk utilization is higher than the user-specified maximum disk utilization for auto migration.




Figure 2-7 Data Flow during a Volume Migration Operation




Figure 2-8 Data Flow after a Volume Migration Operation



Note: Immediately after a volume to be migrated and a volume of a migration destination replace each other, RAID level before the replacement might be displayed or an internal error might occur in the storage system. In this case, select the **Refresh** button to update the display of the window.

Automatic Volume Migration

Volume Migration allows you to create and execute migration plans. Volume Migration also allows you to automate creation and execution of volume migrations plans. If you automate creation and execution of migration plans, you do not need to analyze usage rates of storage system resources and speed of hard disk drives. Volume Migration performs automatically after determining the volumes to be migrated and the migration destinations.

If you want to analyze processor and access path usage when creating migration plans, use the manual migration function rather than the auto migration function.



Caution: Always turn off the auto migration function and cancel any existing auto migration plan before performing manual migration operations. Do not perform manual migration operations while the auto migration function is active.



Notes:

- If you use the manual migration function, you can analyze disk usage, processor usage, and access path usage when you create migration plans. If you use the auto migration function, Volume Migration only uses disk usage and does not analyze processor usage and access path usage when creating migration plans.
 - Auto migration does not support external volumes.
-

To automate the creation and execution of migration plans, do the following:

1. **Use Performance Monitor to monitor disk usage.** Volume Migration automatically analyzes past disk usage and then creates migration plans. For a description of how to use Performance Monitor to collect disk usage statistics so that they can be analyzed by Volume Migration, see *Migrating Volumes Manually*.
2. **If necessary, specify fixed parity groups.** If you automate volume migrations, some volumes and parity groups might suffer degraded I/O speeds. For example, if a volume 00:00:01 is migrated from the parity group Grp1-1 of the class A to the parity group Grp-2-1 of the class B, the volume 00:00:01 is migrated to a lower-speed hard disk drive and thus might suffer from delay in I/O speed.

Volume Migration lets you prevent a parity group from experiencing degraded I/O speeds. If you specify a parity group as fixed, it will not be affected by the auto migration function. For example, if you specify Grp1-1 as a fixed parity group, volumes in Grp1-1 will not be migrated to any other parity group. Also, no volumes will be migrated to Grp1-1. For details on how to specify fixed parity groups, see *Setting a Fixed Parity Group*.



Note: In manual migration, a fixed parity group operates same as a normal parity group. In this case, you can specify a volume in a fixed parity group as a target or source volume.

3. **Review the disk usage limit for each HDD class.** Each HDD class is assigned a disk usage limit by default. Disk usage limits affect system performance after volume migrations. If the disk usage limit for one class is high, many volumes might be migrated to parity groups in that class. If the disk usage limit for one class is low, fewer volumes might be migrated to parity groups in that class.

Volume Migration lets you change the disk usage limit. The default usage limits are rough estimates that may not provide the best results. For details on how to change the disk usage limit, see *Changing the Maximum Disk Usage Rate for Each HDD Class*.

4. **Specify the time period for which you want to improve disk access performance.** As explained earlier, Volume Migration automatically analyzes past disk usage and creates migration plans. To enable Volume Migration to analyze disk usage, you must specify a time period to narrow disk usage statistics to be analyzed. It is recommended that you specify the time period on which you want to improve disk access performance. For example, if you want to improve disk access performance every day from 9:00 a.m. to 5:00 p.m., specify *every day from 9:00 a.m. to 5:00 p.m.* Volume Migration analyzes disk usage statistics for the specified time period and creates migration plans based on the analysis. For how to specify a time period, see *Setting Auto Migration Parameters*.
5. **Specify the time when Volume Migration executes migration plans.** It is recommended that you specify a time when disk usage rate is low because volume migrations typically produce heavy workloads on hard disk drives.

Volume Migration executes migration plans every day at the specified time. For details on how to specify an execution time, see *Setting Auto Migration Parameters*.
6. **Specify parameters that affect the system performance.** The **Volume Migration** windows let you specify parameters that affect the system performance after volumes are migrated. You can specify these parameters, if necessary.
7. **Set the auto migration function to Enable.** To automate creation and execution of migration plans, you must set the Auto Migration Function option to **Enable**.

Volume Migration can now create and execute migration plans. Volume Migration creates plans 30 minutes later than the time period specified in step 4 (For example, if 13:00 - 15:00 is specified, plan creation will be performed at 15:30). The migration plans will be executed every day at the time specified in step 5.



Note: You can change or cancel automatically created migration plans before they are executed. For details, see *Setting Auto Migration Parameters* and *Remaking Auto Migration Plans*.

After auto migration finishes, you can use Performance Monitor to check whether the workloads on hard disk drives are balanced and the system performance is improved. If the migration results are not satisfactory, you can change settings for fixed parity groups and disk usage limits or change other parameters and then retry auto migrations.

The auto migration function allows you to specify the maximum usage rate for each HDD class in the USP V/VM storage system. The maximum usage rate applies to all the parity groups in the HDD class. This section described how the specified usage limit affects auto migration behavior.

The auto migration function allows you to specify the maximum usage rate for each HDD class in the USP V/VM storage system. When a parity group exceeds this limit, the auto migration function makes a plan to migrate one or more volumes in this parity group to another parity group in a higher HDD class or to a parity group in the same HDD class with lower utilization. This storage tuning method addresses and can eliminate bottlenecks of hard disk drive activity and also provides load balancing of hard disk drive usage. Volume Migration uses the estimated usage rates to verify each proposed auto migration, and will not perform a migration operation which might result in a target parity group exceeding the user-specified maximum disk usage rate.



Note: The default maximum usage values are rough estimates only and may not provide the best results.

Figure 2-9 illustrates how the auto migration function identifies parity groups that exceed the user-specified usage limit, and then selects the high-usage volumes as the source volumes to be migrated to parity groups in higher HDD classes or to other parity groups in the same HDD class with lower utilization.




Figure 2-9 Migrating to Improve Disk Usage

When the parity groups in the highest HDD classes have run out of reserved volumes (are empty), Volume Migration automatically migrates low-usage volumes from higher HDD class groups to lower HDD class groups to maintain available reserve volumes.




Figure 2-10 Keeping Reserve Volumes in the High HDD Class Available

As explained above, the auto migration function can move a high-usage volume to a higher HDD class group, forcing a low-usage volume out of that HDD class group. To be able to move one volume into a different HDD class group and then force another volume out of that class group, the auto migration function requires a minimum of five (5) percent difference in estimated disk usage between the two volumes. If the difference in disk usage is less than 5%, this migration is considered an ineffective strategy, and the volume is not moved.




Figure 2-11 Migrating a Volume Automatically to a Lower HDD Class

The auto migration function can also move a volume from one parity group to another group of the same HDD class, forcing another volume out of the destination parity group. To be able to move one volume into a parity group of the same HDD class and force another volume out of that parity group, the auto migration function requires a minimum of twenty (20) percent difference in estimated disk usage between the two volumes. If the difference in disk usage is less than 20%, this migration is considered ineffective, and the volume is not moved.




Figure 2-12 Migrating a Volume Automatically Within the Same HDD Class

Overview of Server Priority Manager

When Server Priority Manager is used, I/O operations from hosts requiring high performance are given higher priority over I/O operations from other hosts.

Performance of High-Priority Hosts

In an SAN (storage area network) environment, the storage system is usually connected with a lot of host servers. Some types of the host servers often require high performance but others might not require as high performance.

For example, *production servers* usually require high performance. Production servers, which include database servers and application servers, are used to perform daily tasks of business organizations. If production servers suffer lowered performance, productivity in business activities is likely to be damaged. For this reason, the system administrator needs to maintain performance of production servers at a higher level.

Computer systems in business organizations often include *development servers* as well as production servers. Development servers are used for developing, testing and debugging business applications. If development servers suffer lowered performance, it would bring undesirable results to developers. However, a decline in development server performance would not bring as much negative impact to the entire organization as a decline in production server performance. In this sense, production servers should be given higher priority over development servers.

Server Priority Manager allows you to limit the number of I/Os requests from development servers to the storage system. Server Priority Manager also allows you to limit the size of data that should be transferred between the development servers and the storage system. Production servers can expect reduced response time. Production server performance can be maintained at a higher level.

Throughout this document, the term *upper limit control* is used to refer to an act of limiting performance of low-priority host servers in order to maintain high-priority host servers at a higher level.

Upper-Limit Control

Upper-limit control can help production servers to perform at higher levels, but it is not necessarily useful when production servers are not busy.

For example, if the number of I/Os from production servers greatly increases from 9:00 a.m. to 3:00 p.m. and decreases significantly after 3:00 p.m., upper-limit control would suppress performance of development servers even after 3:00 p.m. Development servers should be free from upper-limit control when production servers are not busy.

Server Priority Manager provides a function called *threshold control*. If threshold control is used, upper limit control is automatically disabled when traffic between production servers and the storage system decreases to a certain level. A *threshold* is a value that indicates the timing at which upper limit control is disabled. For example, if a threshold of 500 IO/s (500 I/Os per second) is applied to the entire storage system, development servers are free from the limit on the I/O rate (i.e. the number of I/Os per second) when the number of I/Os from all the production servers is below 500 IO/s. If the number of I/Os from the production servers increases and exceeds 500 IO/s, upper limit control is restored to limit the number of I/Os from the development servers again.

The threshold can be used to control the I/O rate (the number of I/Os per second) or the transfer rate (the size of data transferred per second). For example, if a threshold of 20 MB/s (20 megabytes per second) is set to a storage system, the I/O rate limit for development servers is disabled when the amount of data transferred between the storage system and all the production servers is below 20 MB/s.

Overview of Export Tool

The Export Tool enables you to export monitoring data (i.e., statistics) that can display in the **Performance Management** window to text files. The Export Tool also enables you to export monitoring data on remote copy operations performed by TrueCopy, TrueCopy for IBM z/OS, Universal Replicator, and Universal Replicator for IBM z/OS. If you export monitoring data to text files, you can import the monitoring data into word processor documents. You can also import the monitoring data into spreadsheet or database software to analyze the monitoring data.

Figure 2-13 is an example of a text file imported into spreadsheet software:

LU_IOPS.csv				
Serial number : 60001(RAID600)		Indicates that the subsystem serial number is 60001. The word in parentheses is a code that indicates the subsystem type.		
From : 2007/03/28 18:57		Indicates that the data was obtained from 18:57 to 19:01 on March 28, 2007. Sampling rate: 1 indicates that monitoring was performed every minute (at a one-minute interval).		
To : 2007/03/28 19:01				
sampling rate : 1				
No.	time	CL1-A.00(1A-G00).0000	CL1-A.00(1A-G00).0001	CL1-A.00(1A-G00).0002
1	2006/3/28 18:58	10	20	0
2	2006/3/28 18:59	12	20	1
3	2006/3/28 19:00	11	21	2
4	2006/3/28 19:01	13	20	2
		Indicates monitoring data. If a value is -1, the value indicates that Performance Monitor failed to obtain the data for some reason. For detailed information, refer to the troubleshooting information later in this appendix.		

Notes:

- In this LU_IOPS.csv file, the last four digits of a table column heading (such as **0001** and **0002**) indicate an LUN. For example, the heading **CL1-A.00(1A-G00).0001** indicates the port **CL1-A**, the host group ID **00**, the host group name **1A-G00**, and the LUN **0001**.
- If you export monitoring data about concatenated parity groups, the resulting CSV file do not contain column headings for the concatenated parity groups. For example, if you export monitoring data about a concatenated parity group named **1-3[1-4]**, you will be unable to find **1-3[1-4]** in column headings. To locate monitoring data about **1-3[1-4]**, find the **1-3** column or the **1-4** column. Either of these columns contains monitoring data about **1-3[1-4]**.

Figure 2-13 Example of a Text File**Notes:**

- When you run the Export Tool, text files are usually compressed in a ZIP-format archive file. To be able to open a text file, you must use decompress the ZIP file to extract the text files.
- Text files are in CSV (comma-separated value) format, in which values are delimited by commas. Many spreadsheet applications can be used to open CSV files.
- Do not run multiple instances of the Export Tool simultaneously. If you run multiple instances, the SVP may be overloaded and a timeout error may occur.

Interoperability with Other Products

Performance Monitor

Keep the following in mind while viewing **Performance Management** windows.

- User types

If the user type of your user ID is *storage partition administrator*, the functions you can use are limited. For details, see Storage Partition Administrators Limitations.

- Powering off the storage system

If you power off the storage system during monitoring, monitoring stops only while the storage system is powered off. When you power on the storage system again, Performance Monitor cannot display the monitoring data while the storage system is powered off. The monitoring data immediately after powering on again might contain extremely large values.

- Viewing the **Physical** tab

- You will be able to view usage statistics that have been obtained for the last three months (i.e., 93 days) in long-range monitoring, and for the last 15 days in short-range monitoring. You will not be able to view usage statistics that expired after these storing periods because such statistics are erased from the storage system.
- In short range, if I/O workloads between hosts and the storage system become heavy, the storage system gives higher priority to I/O processing than monitoring processing, therefore, a part of monitoring data might be missing. In case that monitoring data are missing frequently, use the Gathering Interval option in the Monitoring Options window to change the collection interval longer. For details, see Start Monitoring and Monitoring Options Window.
- Short-range monitoring data and long-range monitoring data may have some margin of error.

- Viewing the **LDEV**, **Port-LUN**, and **WWN** tab

- Monitoring results are stored for the last 8 hours to 15 days depending on the specified gathering interval. If the storing period has passed since a monitoring result was obtained, the result is erased from the storage system and you will not be able to view that monitoring result.
- If I/O workloads between hosts and the storage system become heavy, the storage system gives higher priority to I/O processing than monitoring processing, therefore, a part of monitoring data might be missing. In case that monitoring data are missing frequently, use the **Gathering Interval** option in the **Monitoring Options** window to change the collection interval longer. For details, see Start Monitoring and Monitoring Options Window.

- The statistics of monitoring data of pool volumes is included in the statistics of monitoring data of V-VOLs. For this reason, the pool volumes are not displayed in **LDEV** tab.
- Viewing the **WWN** tab

To start monitoring traffic between host bus adapters and storage system ports, you must make settings before starting monitoring. For details, see Monitoring All Traffic between HBAs and Ports and Setting Priority for Host Bus Adapters.
- Displaying monitoring data
 - In lists of the **Performance Management** window contents, a hyphen (-) might be displayed in monitoring data columns. It means that the statistics of that monitoring item cannot be collected.
 - If the SVP is overloaded, more time than the gathering interval allots might be required for updating the display of monitoring data. In this case, some portion of monitoring data will not be displayed in the window. For example, suppose that the gathering interval is 1 minute. In this case, if the display in the **Performance Management** window is updated at 9:00 and the next update occurs at 9:02, the window (including the graph) does not display the monitoring result for the period of 9:00 to 9:01. This situation occurs when you use a Storage Navigator computer, as well as when the SVP is used to perform maintenance operations for the DKC.
 - After you set **Monitoring Switch** to **Enable**, the SVP might be overloaded for up to 15 minutes if Performance Monitor receives a couple of data items.
 - After LDEVs are installed or CUs to be monitored are added, the SVP might be overloaded for up to 15 minutes if Performance Monitor receives a couple of data items.
- Replacing the microprogram

After the microprogram was replaced, the monitoring data is not stored until a service engineer releases the SVP from **Modify** mode. Therefore, inaccurate monitoring data may be displayed temporarily.

Volume Migration

The requirements and restrictions for performing Volume Migration operations are:

- **User types.** If the user type of your user ID is *storage partition administrator*, you cannot use Volume Migration. For details on the limitations when using Performance Manager logged in as a storage partition administrator, see Storage Partition Administrators Limitations.
- **Combinations of volumes.** The source and target volumes must be managed by the same USP V/VM storage system. (An external volume can be used for source or target volume.)

Also, a combination of source and target volumes should satisfy the following conditions:

- The **LDKC:CU:LDEV** numbers of volumes take the value between 00:00:00 and 00:FE:FF or between 01:00:00 and 01:FE:FF.
- The source and target volumes have the same emulation type and capacity.
- If the emulation type is not OPEN-V, both the source and target volumes are *custom-sized volume* (CV), or both the volumes are normal volumes.

Specify the volumes that satisfy the conditions above as the source and target volumes. The source and target volumes must be specified by LDEV ID, not VOLSER or port/TID/LUN.

In some cases, the volumes of OPEN-V emulation type might not be displayed as target volumes. If Inquiry commands are executed to volumes displayed as OPEN-V in the Volume Migration window, the replies from some volumes could be OPEN-V and from other volumes OPEN-0V. You cannot distinguish OPEN-V volumes from OPEN-0V volumes in the LDEV's list of the Volume Migration window. If you want to distinguish between OPEN-V and OPEN-0V volumes, open the Basic Information Display dialog box in Storage Navigator.

The volume migration between OPEN-V and OPEN-V is permitted, as is between OPEN-0V and OPEN-0V. But the volume migration between OPEN-V and OPEN-0V is not currently permissible.

- **Source volumes.** The following volumes cannot be used as source volumes:
 - Volumes set as command devices (reserved for use by the host).
 - Volumes which are in an abnormal or inaccessible condition (e.g., pinned track, fenced).
 - Volumes used by the Compatible XRC host software function.



Note: If the status of an XRC volume is suspended because the session was canceled, the volume can be used as a source volume.

- Volumes used by the IBM® 3990 Concurrent Copy (CC) host software function.
- Volumes which have Cache Residency Manager data or Cache Residency Manager for IBM z/OS® data stored in cache.
- Volumes used by Universal Replicator or Universal Replicator for IBM z/OS as a journal volume.
- Volumes which are reserved by a migration program other than Volume Migration.
- Volumes on which CCI is performing volume migration operation.
- Volumes that form a Copy-on-Write Snapshot pair.

- Virtual volumes and pool volumes that are used by Copy-on-Write Snapshot.
- Pool volumes used by Dynamic Provisioning.
- Volumes being shredded.
- Volumes on which Quick Format is being performed.
- System disks.

TrueCopy for IBM z/OS: If the status of a TrueCopy for IBM z/OS volume is *suspended* or *simplex*, the volume can be used as a source volume. If not, the volume cannot be used as a source volume.

- When a TrueCopy for IBM z/OS volume is deleted from the MCU, the status of both volumes changes to *simplex*, and both volumes can be used as source volumes.
- When a TrueCopy for IBM z/OS volume is deleted from the RCU, the status of the M-VOL changes to *suspended* and the status of the R-VOL changes to *simplex*, and both volumes can be used as source volumes.

If you specify a TrueCopy for IBM z/OS volume as a source volume, you cannot specify an external volume as the target volume. If you want to migrate the TrueCopy Asynchronous for IBM z/OS volumes, you must specify the target volumes that belong to the same CLPR.

TrueCopy: If the status of a TrueCopy volume is PSUS, PSUE, or SMPL, the volume can be used as a source volume. If not, the volume cannot be used as a source volume.

- When a TrueCopy pair is deleted from the MCU, the status of both volumes changes to SMPL, and both volumes can be used as source volumes.
- When a TrueCopy pair is deleted from the RCU, the status of the P-VOL changes to PSUS and the status of the S-VOL changes to SMPL, and both volumes can be used as source volumes.

If you specify a TrueCopy volume as a source volume, you cannot specify an external volume as the target volume. If you specify a TrueCopy volume that is a normal volume as a source volume, you must specify a normal volume as the target volume. If you specify a TrueCopy volume that is a Dynamic Provisioning virtual volume as a source volume, you must specify a Dynamic Provisioning virtual volume as the target volume. If you want to migrate the TrueCopy Asynchronous volumes, you must specify the target volumes that belong to the same CLPR.

Universal Replicator for IBM z/OS: If the status of a Universal Replicator for IBM z/OS volume is Pending duplex or Duplex, the volume cannot be used as a source volume. If you specify a Universal Replicator for IBM z/OS volume as a source volume, you cannot specify an external volume as the target volume. S-VOL of a URz delta resync pair cannot be used as a source volume. If you specify a Universal Replicator volume that is a normal volume as a source volume, you must specify a normal volume as the target volume. If you specify a Universal Replicator volume that is a Dynamic Provisioning virtual volume as a source volume, you must specify a Dynamic Provisioning virtual volume as the target volume. If you want to migrate the Universal Replicator for IBM z/OS volumes, you must specify the target volumes that belong to the same CLPR.

Universal Replicator: If the status of a Universal Replicator volume is COPY or PAIR, the volume cannot be used as a source volume. If you specify a Universal Replicator volume as a source volume, you cannot specify an external volume as the target volume. S-VOL of a UR delta resync pair cannot be used as a source volume. If you want to migrate the Universal Replicator volumes, you must specify the target volumes that belong to the same CLPR.

ShadowImage for IBM z/OS® and ShadowImage: Usage of ShadowImage for IBM z/OS and ShadowImage volumes as a source volume depends on the status and cascade configuration of the volume as described below.

- The following volumes cannot be used as a source volume:
 - ShadowImage for IBM z/OS volumes whose status is SP-Pend or V-Split.
 - ShadowImage volumes whose status is COPY(SP) or PSUS(SP).
- Table 2-2 specifies which non-cascade volumes can be used as source volumes.

Table 2-2 Non-Cascade Volumes and Volume Migration Source Volumes

Pair Configuration [Status ≠ COPY(SP) or PSUS(SP)]	Can P-VOL be used as a source volume?	Can S-VOLs be used as source volumes?
Ratio of P-VOL to S-VOLs is 1:1	Yes	Yes
Ratio of P-VOL to S-VOLs is 1:2	Yes	Yes
Ratio of P-VOL to S-VOLs is 1:3	No	Yes

- Table 2-3 specifies which cascade volumes can be used as source volumes.

Table 2-3 Cascade Volumes and Volume Migration Source Volumes

Pair Configuration [Status ≠ COPY(SP) or PSUS(SP)]	Can P-VOL be used as a source volume?	Can S-VOLs be used as source volumes?
L1 pair, ratio of P-VOL to S-VOLs is 1:1	Yes	Yes
L1 pair, ratio of P-VOL to S-VOLs is 1:2	Yes	Yes

L1 pair, ratio of P-VOL to S-VOLs is 1:3	No	Yes
L2 pair, ratio of P-VOL to S-VOLs is 1:1	Yes	No
L2 pair, ratio of P-VOL to S-VOLs is 1:2	No	No



Caution: If any of the following operations is performed on a Volume Migration source volume during migration, the volume migration process stops:

- Compatible XRC operation.
- CC operation.
- TrueCopy for IBM z/OS operation that changes the volume status to something other than suspended or simplex.
- TrueCopy operation that changes the volume status to something other than PSUS, PSUE, or SMPL.
- Universal Replicator operation or Universal Replicator for IBM z/OS operation that changes the volume status to COPY.
- ShadowImage for IBM z/OS operation that changes the volume status to SP-Pend or V-Split.
- ShadowImage operation that changes the volume status to COPY(SP) or PSUS(SP).
- Universal Replicator operation or Universal Replicator for IBM z/OS operation.

- **Dynamic Provisioning:** If some Dynamic Provisioning virtual volumes related to one pool are used as a source volume, the other Dynamic Provisioning virtual volume related to the same pool cannot be specified as the target volume.
- **Expanded LUs as source volumes.** If you want to specify an expanded LU as a source volume for migration, you can specify individual LDEVs within the expanded LU (e.g., LDEVs with high usage), and Volume Migration will migrate only the specified LDEVs. If desired, you can also specify all LDEVs of the expanded LU to relocate the entire expanded LU. In this case, you must make sure that the required reserved target LDEVs are available.
- **Target volumes.** Target volumes must be reserved prior to migration using the Volume Migration software. Hosts cannot access reserved volumes.

The following volumes cannot be reserved as target volumes:

- Expanded LUs
- Volumes which are set as command devices (devices reserved for use by the host).
- Volumes which are assigned to ShadowImage or TrueCopy pairs.

- Volumes that are used by the Compatible XRC host software function.



Note: If the status of a Compatible XRC volume is suspended because the session was canceled, the volume can be used as a source volume.

- Volumes which are used by the IBM 3990 Concurrent Copy (CC) host software function.
- Volumes which are reserved for ShadowImage operations.
- Volumes which have Cache Residency Manager data or Cache Residency Manager for IBM z/OS data stored in cache.
- Volumes used by Universal Replicator or Universal Replicator for IBM z/OS as a data volume or journal volume.
- Volumes in which Volume Retention Manager specified the Read Only or Protect attribute.
- Volumes which Volume Security disabled to be used as secondary volumes.
- Volumes which Data Retention Utility specified the Read Only or Protect attribute, or disabled to be used as secondary volumes.
- Volumes which are reserved by a migration program other than Volume Migration.
- Volumes on which CCI is performing volume migration operation.
- Volumes that form a Copy-on-Write Snapshot pair.
- Pool volumes used by Dynamic Provisioning.
- Volumes which are in an abnormal or inaccessible condition (e.g., pinned track, fenced).
- Volumes on which Quick Format is being performed.
- System disks.

You can also reserve external volumes as target volumes. However, if you specify a TrueCopy for IBM z/OS volume or TrueCopy volume as a source volume, you cannot specify an external volume as the target volume.

You can also reserve Dynamic Provisioning virtual volumes as target volumes. However, a Dynamic Provisioning virtual volume which is related to the same pool as that related to the source volume cannot be specified as the target volume.

- **Number of migration plans that can be executed concurrently:** In manual migration, the number of migration plans that can be executed concurrently might be restricted depending on the usage of other USP V/VM programs, and also depending on the emulation types and sizes of the migrated volumes.

Volume Migration uses the *differential tables* to manage differentials occurred by write operation while migrating volumes. Differential tables are the resources shared by Copy-on-Write Snapshot, Volume Migration, ShadowImage, ShadowImage for IBM z/OS, Compatible FlashCopy®, and Compatible FlashCopy V2. Therefore, the maximum number of differential tables that Volume Migration can use is the number of total differential tables in the storage system minus the number of that used by other programs.

Moreover, the number of differential tables needed to migrate one volume differs depending on the emulation type and size of the volume. Thus, the maximum number of migration plans that can be executed concurrently changes depending on the number of differential tables that Volume Migration can use, and depending on the emulation types and sizes of volumes to be migrated.

For details on how to calculate the number of migration plans that can be executed concurrently, see How to Calculate the Number of Concurrent Migration Plans. For details on the number of differential table required for migrating a volume, see Calculating Differential Tables for Mainframe Volume Migration and Calculating Differential Tables for Open-System Volume Migration.

- **Auto migration planning.** Volume Migration will **not** plan an auto migration operation when:
 - There is no reserve volume with the same emulation and capacity as the source volume in the same or next higher/lower HDD class (auto migration is performed only between consecutive classes or within the same class).
 - The usage rate for the target parity group cannot be estimated, because some condition is preventing Volume Migration from estimating expected usage (e.g., invalid monitor data).
 - The estimated usage of the target parity group or volume is over the user-specified maximum disk utilization value for that HDD class.
 - The estimated performance improvement is not large enough.
- **Auto migration execution.** Volume Migration will **not** execute an auto migration operation when:
 - The current usage (last 15 min) of the source or target parity group or volume is over the maximum disk utilization rate for auto migration.
 - The current write pending rate for the storage system is 60% or higher.
- **Manual migration execution.** Volume Migration will **not** execute a manual migration operation when the current write pending rate for the storage system is 60% or higher.
- **If heavy workloads are imposed on your storage system:** Avoid volume migration operations when heavy workloads are imposed onto your storage system due to I/O requests from hosts. If you should perform volume migration, the migration could fail. If migration fails, please reduce the workloads on your storage system and then retry volume migration.

Copy Threshold Option: If the load of storage system increases, host server I/O performance (response) may be degraded. If Volume Migration performs migration operations when the load of storage system is heavy, it is more likely that host server I/O performance (response) may be degraded. The Copy Threshold option temporarily stops Volume Migration operation when the load of the storage system is heavy. You can configure this option to minimize the degradation of host I/O performance (response) by temporarily stopping Volume Migration operations when the load of the storage system is heavy.

Migration of this volume is likely to fail when hosts frequently update volumes being migrated. You can configure the Copy Threshold option to temporarily stop the copy processing of Volume Migration when the load of the storage system is at capacity. But the use of this option is not recommended for hosts that frequently update volumes which are likely to be migrated, because of the increased possibility of failure during volume migration.

For information about the setting of the Copy Threshold option, call the Hitachi Data Systems Support Center (see Calling the Hitachi Data Systems Support Center).



Note: Copy operations that are stopped by the Copy Threshold option will resume once the load on the storage system becomes light. But if this option is turned on, not only will Volume Migration operations stop if the load on the storage system becomes heavy, but the copy operations of the following program products will also be interrupted:

- ShadowImage
 - ShadowImage for IBM z/OS
 - Compatible FlashCopy
 - Compatible FlashCopy V2
 - Copy-on-Write Snapshot
- **Progress Value:** Even though a Volume Migration display may show that migration progress is 100 percent complete, it is possible that the migration is unfinished. This situation can arise when volumes with heavy workloads (update I/Os between hosts and the storage system) are migrated together with volumes with light workloads. In this situation, the migration of the volumes may not have completed even though the progress value shows 100%. If this problem occurs, execute either of the following to migrate the volumes:
 - Reduce the workloads for update I/Os between hosts and the storage system and complete the migration.
 - Stop the volume migration for the volume with heavy workloads for update I/Os between hosts and the storage system, and execute the migration of another volume and then retry the migration which has been stopped.

- **Effects on Other Program Products:** When the workloads updating I/Os between hosts and the storage system are heavy, more time will be required to complete volume migration because the copying of differential data is repeatedly executed. When differential data is being copied, longer copying time may also be required for other program products. Table 2-4 shows the amount of additional copy time expected for affected program products. The maximum increase may be as much as doubled. The amount of increase in the copy time depends on the number of pairs set to the program product.

The term "other program products" refers to the following program products.

- Volume Migration
- ShadowImage
- ShadowImage for IBM z/OS
- Compatible FlashCopy

Table 2-4 Capacity of Migrated Volumes and Estimated Delay in Copy Time

Capacity of Migrated Volume (MB)	Estimated delay in the copying speed (in minutes)
0 - 1,000	4
1,001 - 5,000	18
5,001 - 10,000	37
10,001 - 50,000	186
50,001 - 100,000	372
100,001 - 500,000	1,860
500,001 - 1,000,000	3,720
1,000,001 - 2,150,400	9,667

The above estimates are calculated based on the assumption that the workload for update I/Os for the migrated volume is 50 IOPS for one volume.

- **Maintenance.** Do not perform Volume Migration operations during storage system maintenance activities (e.g., installation, replacement, and un-installation of cache or drives, or replacement of the microprogram). The Volume Migration operation might fail if the volume migration operation is executed in above situation.
- **Auto migration parameters to be initialized when the storage system configuration changes.** Auto migration parameters are initialized when a change is made to the storage system configuration. For example, the parameters are initialized when new drives or LDEVs are installed. The parameters include the disk usage limit and the migration time. After a change is made to the storage system configuration, you must specify auto migration parameters again.

- **Notes on powering off the storage system:** Before turning the power off at the storage system, confirm that volume migration has finished. *Do not turn off the power until the migration is finished.*

If you turn the power off during migration some of the data are not migrated. If later the power is restored, Volume Migration resumes. Volume Migration first attempts to copy un-migrated data that was left in the shared memory (which is a volatile memory) when the power failed. However, if that data was lost from the shared memory, Volume Migration begins recopying all the data - both the un-copied data not previously migrated as well as the data previously copied to the migration destination before the power interruption.



Note: The USP V/VM documentation sometimes uses the term "**PS ON**", which refers to an operation for turning on the power supply to the storage system. Also, the 9900 documentation sometimes uses the term "**PS OFF**", which refers to an operation for turning off the power supply to the storage system.

- **Note on setting Cache Residency:** Cache Residency must not be set on the volume being migrated.

How to Calculate the Number of Concurrent Migration Plans

In using Volume Migration for manual volume migration, the number of migration plans that can be executed concurrently might be restricted. The number of migration plans that can be executed concurrently depends on the following conditions.

- How much shared memory is available for differential tables:

You can install additional shared memory for differential tables.

If additional shared memory for differential tables is installed, you may use 26,176, 57,600, or 104,768 differential tables. If USP V is being used, you may also use 146,688, or 209,600 differential tables.

To install additional shared memory for differential tables, call the Hitachi Data Systems Support Center (see Calling the Hitachi Data Systems Support Center).

- How much shared memory is available for pair tables:

You can install additional shared memory for differential tables.

You may use 8,192 or 16,384 pair tables if additional shared memory for pair tables is installed.

To install additional shared memory for pair tables, call the Hitachi Data Systems Support Center.

- The emulation type and capacity of each volume to be migrated:
The number of differential tables and pair tables needed to migrate one volume differs depending on the emulation type and size of the volume. For the number of differential tables and pair tables needed for migrating a mainframe volume, see *Calculating Differential Tables for Mainframe Volume Migration*. For the number of differential tables needed for migrating an open-system volume, see *Calculating Differential Tables for Open-System Volume Migration*.

You can estimate the maximum number of migration plans that can be executed concurrently by applying the above conditions into the following equation:

- $\Sigma(\alpha) \leq (\beta)$ and $\Sigma(\gamma) \leq (\delta)$
 $\Sigma(\alpha)$ stands for the total number of differential tables needed for migrating all volumes.
 (β) stands for the number of differential tables available in the storage system.
 $\Sigma(\gamma)$ stands for the total number of pair tables needed for migrating all volumes.
 (δ) stands for the number of pair tables available in the storage system.
For example, if you want to create 20 migration plans of OPEN-3 volumes (size of a volume is 2,403,360 kilobytes), the number of required differential tables is 3, and the number of pair tables is 1, as calculated by the equation in section *Calculating Differential Tables for Open-System Volume Migration*.
 $[(3 \times 20) = 60] \leq 26,176$ and $[(1 \times 20) = 20] \leq 8,192$
Since this equation is true, you can create all the migration plans that you wish to create.

In this section, we mentioned the calculation of the maximum number of migration plans when only Volume Migration is running. However, in fact, the total number of differential tables used by Copy-on-Write Snapshot, ShadowImage, ShadowImage for IBM z/OS, Compatible FlashCopy, Compatible FlashCopy V2, the total number of pair tables used by ShadowImage, ShadowImage for IBM z/OS, Compatible FlashCopy and Volume Migration should be within the value of (β) and (δ) . For details on how to calculate the number of differential tables and pair tables used by the programs other than Volume Migration, see the following manuals:

- For ShadowImage, see *ShadowImage User's Guide*.
- For ShadowImage for IBM z/OS, see *ShadowImage for IBM z/OS User's Guide*.
- For Compatible FlashCopy and Compatible FlashCopy V2, see *Compatible Mirroring for IBM FlashCopy User's Guide*.
- For Copy-on-Write Snapshot, see *Copy-on-Write Snapshot User's Guide*.

Calculating Differential Tables for Mainframe Volume Migration

When you migrate mainframe volumes, use the following expression to calculate the total number of the required differential tables and pair tables per migration plan.

$$\text{Total number of the differential tables per migration plan} = ((X) + (Y)) \times 15 \div (Z)$$

(X): The number of the cylinders of the volume to be migrated.*¹

(Y): The number of the control cylinders.

(Z): The number of the slots that can be managed by a differential table.

(20,448)

*1: If the volume is divided by the VLL function, this value means the number of the cylinders of the divided volume.



Note: Round up the number to the nearest whole number.

For example, in case of a volume which emulation type is 3390-3, and when provided that the number of the cylinders of the volume is 3,390 ((X) in the expression above), the calculation of the total number of the differential table is as follows.

$$(3,339 + 6) \times 15 \div (20,448) = 2.453785211$$

When you round up 2.453785211 to the nearest whole number, it becomes 3. Therefore, the total number of the differential table for one migration plan is 3 when emulation type is 3390-3.

One pair table can be used for 36 differential tables. Therefore, the number of pair tables for one migration plan is 1 when emulation type is 3390-3. Two pair tables are used if the emulation type is 3390-M and the number of volume cylinders is default number. The emulation type is 3390-M if multiple pair tables are used on a mainframe system.

The following table shows the number of the control cylinders according to the emulation types.

Table 2-5 Number of the Control Cylinders According to Emulation Types

Emulation Type	Number of the Control Cylinders
3380-3	7
3380-3A	7
3380-3B	7
3380-3C	7
3380-F	22
3380-K	7
3380-KA	7
3380-KB	7
3380-KC	7
3390-3	6
3390-3A	6
3390-3B	6
3390-3C	6
3390-3R	6
3390-9	25
3390-9A	25
3390-9B	25
3390-9C	25
3390-L	23
3390-LA	23
3390-LB	23
3390-M	53
3390-MA	53
3390-MB	53
3390-MC	53
3390-LC	23
NF80-F	22
NF80-K	7
NF80-KA	7
NF80-KB	7
NF80-KC	7

Calculating Differential Tables for Open-System Volume Migration

When you migrate open-system volumes, use the expression in Table 2-6 to calculate the total number of the required differential and pair tables per migration plan.

Table 2-6 Total Number of the Differential Tables Per Migration Plan

Emulation Type	Expression
OPEN-3	Total number of the differential tables per migration plan = $((X) \div 48 + (Y) \times 15) \div (Z)$ (X): The capacity of the volume to be migrated. (kilobyte)*¹ (Y): The number of the control cylinders. (Z): The number of the slots that can be managed by a differential table. (20,448) Note: Round up the number to the nearest whole number. For example, if the emulation type of a volume is OPEN-3, and when provided that the number of the cylinders of the volume is 2,403,360 kilobytes ((X) in the expression above), the calculation of the total number of the differential tables is as follows. $(2,403,360 \div 48 + 8 \times 15) \div 20,448 = 2.454518779$ When you round up 2.454518779 to the nearest whole number, it becomes 3. Therefore, the total number of the differential tables for one migration plan is 3 when emulation type is OPEN-3. One pair table can be used for 36 differential tables. Therefore the total number of the pair tables for one migration plan is 1 when emulation type is OPEN-3.
OPEN-8	
OPEN-9	
OPEN-E	
OPEN-L	
OPEN-V	Total number of the differential tables per migration plan = $((X) \div 256) \div (Z)$ (X): The capacity of the volume to be migrated. (kilobyte)*¹ (Z): The number of the slots that can be managed by a differential table. (20,448) Note: Round up the number to the nearest whole number. For example, if the emulation type of a volume is OPEN-V, and when provided that the number of the cylinders of the volume is 3,019,898,880 kilobytes ((X) in the expression above), the calculation of the total number of the differential tables is as follows. $(3,019,898,880 \div 256) \div 20,448 = 576.9014085$ When you round up 576.9014085 to the nearest whole number, it becomes 577. Therefore, the total number of the differential tables for one migration plan is 577 when emulation type is OPEN-V. One pair table can be used for 36 differential tables. Therefore the total number of the pair tables for one migration plan is 17 when emulation type is OPEN-V. The emulation type OPEN-V can only be used, if you use multiple pair tables on the open system.

*1: If the volume is divided by the VLL function, this value means the capacity of the divided volume. If the emulation type is OPEN-L, the VLL function is unavailable.

The following table shows the number of the control cylinders according to the emulation types.

Table 2-7 Number of the Control Cylinders According to the Emulation Types

Emulation Type	Number of the Control Cylinders
OPEN-3	8 (5,760KB)
OPEN-8	27 (19,440KB)
OPEN-9	
OPEN-E	19 (13,680KB)
OPEN-L	7 (5,040KB)
OPEN-V	0 (0KB)

Notes on Migrating Dynamic Provisioning Virtual Volumes

This section explains the appropriate combination of source volumes and target volumes that are also Dynamic Provisioning virtual volumes. Depending on the state of the pool related to the target volume that is also a Dynamic Provisioning virtual volume, a warning message may appear or manual migration plan may not be created. The following table shows the combination of a source volume and a target volume. In the table, the Dynamic Provisioning virtual volume is called DP-VOL.

Table 2-8 Combination of a Source Volume and a Target Volume

The Source Volume		The internal volume as the target volume	DP-VOL as the target volume			
			Normal state			Blocked state due to full capacity
			Normal state	The case that the data size will exceed the pool threshold after completion of the migration	The case that the volume will be full after completion of the migration	
The internal volume as the source volume		a migration plan can be created	a migration plan can be created*1	a migration plan can be created*1, and the warning message appears	a migration plan cannot be created	a migration plan cannot be created
DP-VOL as the source volume	Normal state	a migration plan can be created*1				
	Blocked state	a migration plan cannot be created	a migration plan cannot be created	a migration plan cannot be created		
	Blocked state due to full capacity	a migration plan can be created*1	a migration plan can be created*1	a migration plan can be created*1, and the warning message appears		

*1: If the migrated volume is going to be migrated again, the migration must be executed at the interval which is calculated on the following formula.

(Pool Capacity(TB) × 3 (seconds)) + 40 (minutes)

If the operations that change configuration of the storage system are executed after the operations of Volume Migration, re-migration of the volume must be executed at the interval which is calculated on the above formula after the operations that change configuration of the storage system. However, the migration might not be executed after the time which is calculated on above formula depending on workloads on your storage system.

Consider the following when moving DP-VOLs using Volume Migration:


- DP-VOLs have two emulation types, OPEN-0V and OPEN-V. You can check the emulation type of DP-VOLs in the LDEV window of the Basic Information window of Storage Navigator. For more information about the LDEV window, see the *Storage Navigator User's Guide*.
- You cannot move OPEN-0V DP-VOLs to OPEN-V DP-VOLs or OPEN-V normal volumes. And reversely, You cannot move OPEN-V DP-VOLs or OPEN-V normal volumes to OPEN-0V DP-VOLs. If you need to move volumes in this way, change the emulation type from OPEN-0V to OPEN-V. For more information about changing an emulation type and the V-VOL window, see the *Dynamic Provisioning User's Guide*.

Server Priority Manager

- **User types.** If the user type of your user ID is *storage partition administrator*, you cannot use Server Priority Manager. For details on the limitations when using Performance Manager logged in as a storage partition administrator, see Storage Partition Administrators Limitations.
- **I/O rates and transfer rates.** Server Priority Manager runs based on I/O rates and transfer rates measured by Performance Monitor. Performance Monitor measures I/O rates and transfer rates every second, and calculates the average I/O rate and the average transfer rate for every gathering interval (specified between 1 and 15 minutes) regularly.

Suppose that 1 minute is specified as the gathering interval and the I/O rate at the port 1-A changes as illustrated in Graph 1 in Figure 2-14. When you use Performance Monitor to display the I/O rate graph for 1A, the line in the graph indicates changes in the average I/O rate calculated every minute (refer to Graph 2). If you select the **Detail** check box in the **Performance Management** windows, the graph displays changes in the maximum, average, and minimum I/O rates in one minute.

Server Priority Manager applies upper limits and thresholds to the average I/O rate or the average transfer rate calculated every gathering interval. For example, in Figure 2-14 in which the gathering interval is 1 minute, if you set an upper limit of 150 IO/s to the port 1A, the highest data point in the line *CL1-A* in Graph 2 and the line *Ave.(1 min)* in Graph 3 is somewhere around 150 IO/s. It is possible that the lines *Max (1 min.)* and *Min (1 min.)* in Graph 3 might exceed the upper limit.



Graph 1: Actual I/O rate (measured every second)




Figure 2-14 Line Graphs Indicating Changes in Port Traffic

- **Note on using TrueCopy:** Server Priority Manager monitors write I/O requests issued from initiator ports of your storage system.
- **Note on using the remote copy functions:** When the remote copy functions (TrueCopy, TrueCopy for IBM z/OS, Universal Replicator, and Universal Replicator for IBM z/OS) are used in your environment, Server Priority Manager monitors write I/O requests issued from initiator ports of your storage system.

If you specify an RCU target port as a prioritized port, I/O requests from the initiator port will **not** be a target of threshold control.

If you specify an RCU target port as a non-prioritized port, I/O requests from the initiator port will **not** be a target of upper limit control.

- **Note on the statistics of Initiator/External ports:** The initiator ports and external ports of your storage subsystem are not controlled by Server Priority Manager. Although you can set **Prioritize** or **Non-Prioritize** to initiator ports and external ports by using Server Priority Manager, the initiator ports and the external ports become the prioritized ports that are not under threshold control, regardless of whether the setting of the ports are **Prioritize** or **Non-Prioritize**. If the port attributes are changed from Initiator/External into Target/RCU Target, the settings by Server Priority Manager take effect instantly and the ports are subject to threshold or upper limit control.

The statistics of **All Prio.** and **All Non-Prio.** that is indicated in the **Port-LUN** tab of **Performance Management** windows are sum total of statistics on Target/RCU Target ports that are controlled by Server Priority Manager. The statistics of **All Prio.** and **All Non-Prio.** does not include the statistics of Initiator/External ports. Because the statistics of Initiator/External ports and Target/RCU Target ports are based on different calculation methods, it is impossible to sum up the statistics of Initiator/External ports and Target/RCU Target ports.

Preparing for Performance Manager Operations

This chapter explains the preparations for performance manager operations.

- [System Requirements](#)
- [Storage Partition Administrators Limitations](#)

System Requirements

To use Performance Manager, you need:

- USP V/VM storage system
- Performance Manager software (Volume Migration and Server Priority Manager are optional, but Performance Monitor is required)
- A Web client computer (intended for use as a Storage Navigator) connected to USP V/VM via LAN.

To use Performance Manager, you use the Web client computer to log on to the SVP (Web server). When you are logged on, the Storage Navigator program, which is a Java application program, will be downloaded to the Web client computer. You can then perform Performance Monitor operations in the **Storage Navigator** window.

For a summary of Web client computer requirements, see the *Storage Navigator User's Guide*.

Performance Manager operations require the Storage Navigator program, which is downloaded to your WWW client computer. Your WWW client computer must be connected to the USP V/VM storage system via LAN. Browser settings are also required on your WWW client computer. For details, see the *Storage Navigator User's Guide*.



Caution: Performance Manager operations (Performance Monitor, Volume Migration, and Server Priority Manager) involve the collection of large amounts of monitoring data. This requires considerable Web client computer memory. It is therefore recommended that you exit the Storage Navigator program when not conducting Performance Manager operations to release system memory.

For details on how to install Performance Monitor, Volume Migration and Server Priority Manager, see the *Storage Navigator User's Guide*.

Storage Partition Administrators Limitations

If your user ID is of the *storage partition administrator* type, you can use only Performance Monitor and the Export Tool among Performance Manager programs. Volume Migration and Server Priority Manager are not available to storage partition administrators.

This section explains the permission-driven limitations of Performance Monitor and the Export Tool.


Performance Monitor Limitations

The Performance Monitor functions, which are limited when you logged in as a storage partition administrator, are shown in Table 3-1. For the window of Performance Monitor displayed when you logged in as a storage partition administrator, see Figure 3-1.

Table 3-1 Limitations for Storage Partition Administrators (Performance Monitor)

Window	Limited function
Physical tab in the Performance Management window	The tree and list display only the information in the SLPR (storage management logical partition) allocated to the user ID. The Volume Migration button is not displayed. Therefore, the user cannot start Volume Migration.
LDEV tab in the Performance Management window	The tree and list display only the information in the SLPR allocated to the user ID.
Port-LUN tab in the Performance Management window	The tree and list display only the information in the SLPR allocated to the user ID. The SPM button is not displayed. Therefore, the user cannot start Server Priority Manager.
WWN tab in the Performance Management window	The WWN tab is not displayed. The user cannot view the traffics between host bus adapters and ports.
TC Monitor window TCz Monitor window UR Monitor window URz Monitor window	These windows are not displayed. The user cannot view the information about remote copy operations performed by TrueCopy, TrueCopy for IBM z/OS, Universal Replicator, and Universal Replicator for IBM z/OS.
Monitoring Options window	The Monitoring Options windows are not displayed. The user cannot start or stop monitoring, or change the gathering interval.

The **TC Monitor** window, **TCz Monitor** window, **UR Monitor** window, **URz Monitor** window, and the **Monitoring Options** window are not displayed.



The **WWN** tab is not displayed.

The tree and list display only the information in the SLPR allocated to the user ID.

The **SPM** button and the **Volume Migration** button are not displayed.

Figure 3-1 Performance Management Window displayed When You Logged in as a Storage Partition Administrator

Export Tool Limitations

The Export Tool functions which are limited when you logged in as a storage partition administrator are as follows:

- Only the monitoring data about SLPR allocated to the user ID can be exported into files.
- When a storage partition administrator use the **group** subcommand with specifying the **PPCG** or **PPCGWWN** operand to export the monitoring data about SPM groups or the host bus adapters belonging to these SPM groups, an error will occur in the following conditions:
 - One SPM group contains multiple host bus adapters which are allocated to different SLPRs.
 - One host bus adapter is connected to multiple ports which exist in different SLPRs.
- The monitoring data about remote copy operations performed by TrueCopy, TrueCopy for IBM z/OS, Universal Replicator, and Universal Replicator for IBM z/OS cannot be exported.
- A storage partition administrator cannot use the **set** subcommand to start or stop monitoring, or change the gathering interval.

Using the Performance Manager GUI

This chapter explains performance manager windows.

- [Using the Performance Monitor Windows](#)
- [Using the Volume Migration Windows](#)
- [Using the Server Priority Manager Windows](#)

Using the Performance Monitor Windows

This chapter describes **Performance Management** windows and operations. This chapter explains in the following order: each window of Performance Monitor, the procedure for starting and stopping monitoring, and various operations about obtaining and viewing statistics such as resource usage.



Caution: If the user type of your user ID is *storage partition administrator*, the functions you can use are limited. For details, see Storage Partition Administrators Limitations.

Performance Monitor Window

Performance Monitor has the following windows:

- Performance Management window

This window displays the monitoring results about the storage system performance collected by Performance Monitor. You can change the information to be viewed by selecting each tab in the tree. The information displayed by selecting each tab is as follows:

- **Physical** tab
Displays the usage statistics about resources in the storage system.
- **LDEV** tab
Displays the statistics about workload on disks.
- **Port-LUN** tab
Displays the statistics about traffic at ports and LU paths in the storage system.
- **WWN** tab
Displays the statistics about traffic at path between host bus adapters and ports.

This section explains the contents of the tabs in the **Performance Management** window and the contents of the **Monitoring Options** window.

Performance Management Window, Physical Tab

When you click **Go, Performance Manager** and then **Performance Management** on the menu bar of the Storage Navigator main window, Performance Monitor starts and the **Performance Management** window is active. The **Performance Management** window includes the **Physical** tab, which lets you view usage rates for parity groups, volumes, channel processors, disk processors, etc. In addition, when you use Universal Volume Manager to map volumes in an external storage system (storage system other than USP V/VM) to the internal volumes, the **Physical** tab also lets you view usage conditions of volumes in the external storage system (i.e., external volumes) and the groups of external volumes.

For details on how to use this window, see Monitoring Resources in the Storage System.




Figure 4-1 Physical Tab of the Performance Management Window

When the **Physical** tab is active, the **Performance Management** window contains the following items:

- When **Monitoring Switch** is **Enable**, Performance Monitor is monitoring the storage system (a **Disable** setting indicates that the storage system is not being monitored).

- **Gathering Interval** indicates the interval of collecting statistics in short range monitoring. For example, if the number of the CUs to be monitored is 64 or less, and **1 min.** is displayed and **short-range** is selected as the storing period of statistics, the list and graph in the **Physical** tab display the statistics obtained every one minute. In case 65 or more CUs are monitored, the statistics are displayed every 5, 10 or 15 minutes.



Notes: The gathering interval in long range monitoring is fixed to 15 minutes. If you select **long-range** as the storing period of statistics, the list and graph display the statistics obtained every 15 minutes regardless of the value of **Gathering Interval**. For details on the storing period of statistics (short range and long range), see Understanding Statistical Storage Ranges.

- In the drop-down list on the right of **Monitoring Data**, you can select the storing period of statistics from **short-range** and **long-range**.
Select which range of statistics you want to view in the window: statistics in long range or short range.




Figure 4-2 Drop-Down List to Select Storing Period of Statistics (Physical Tab)

The storing period of statistics is the range of monitoring data (statistics collected by monitoring) that can be displayed. You can specify a part of term within the selected range to narrow the statistics to be displayed in the list and graph on the **Performance Management** window.

The differences in selecting **short-range** and **long-range** are described in When you select Long-Range Storage and When you Select Short-Range Storage.

- **Monitoring Term** lets you narrow the range of usage statistics that should be displayed in the window.

Starting and ending times for collecting statistics are displayed on both sides of the slide bar. Performance Monitor stores the monitoring data between these times, and you can specify the desired term within this range as the target of display in lists and graphs.

For example, if you want to view usage statistics within the range of 10:30 July 1 2006 to 22:30 July 31 2006, you set 2006/07/01 10:30 to the **From** box, set 2006/07/31 22:30 to the **To** box, and then click the **Apply** button.

To set a date and time to the **From** and **To** boxes, do either of the following:

- Move the slider to the left or to the right.
- In the text box, select the number that you want to change. Next, click the upward or downward arrow button.

When you specify dates and time in **From** and **To** boxes, Performance Monitor calculates the length of the specified period and displays the calculated length. The length of the period is displayed in days when you select **long-range**, and it is displayed in minutes when you select **short-range**.



Note:

- The From and To boxes are grayed out if the monitoring data (that is, usage statistics) is not stored in the storage system.
 - The Real Time option is grayed out when the Physical tab is active.
-
- In the **Monitoring Data** area, the drop-down list on the upper right of the list specifies the type of statistics to be displayed in the window. When the **Physical** tab is active, the drop-down list contains only one entry (i.e. **Usage**).
 - The tree lists items such as parity groups, channel adapters (CHAs). The tree can display the following items.
 - Icon displayed below the **Parity Group** or **External Group** folder:
 a parity group or an external volume group
 - Icons displayed below the **CHA** folder:
 -  CHA-2T (Serial 8Port) * an ESCON channel adapter
 -  CHA-1G (Mfibre 4Port) * a FICON channel adapter
 -  CHA-1E (Fibre 16Port) * a Fibre Channel adapter in **Standard** mode
 -  CHA-1E (Fibre 16Port) * a Fibre Channel adapter in **High Speed** mode
 -  CHA-1E (Fibre 16Port) * a Fibre Channel adapter in **Initiator/External MIX** mode

* The channel adapter number and number of ports that displayed on the right side of the icon are examples.
 - Icons displayed below the **DKA** folder:
 -  a disk processor (DKP)
 -  a data recovery and reconstruction processor (DRR)
 - Icon displayed below the **Access Path Usage** folder:
 an access path

No icon is displayed below the **Cache** folder.

The numbers on the right of icons (🗄️) displayed below the **Parity Group** or **External Group** folder are IDs of parity groups or external volume groups. The letter "E" at the beginning of an ID indicates that the group is an external volume group.



Notes: A volume existing in an external storage system and mapped to a volume in the USP V/VM storage system by using Universal Volume Manager is called an *external volume*. An *external volume group* is a quantity of external volumes grouped together for managing and that do not contain any parity information, unlike a parity group. However, **Performance Management** window treats external volume groups same as parity groups for convenience.

The parity group icon (🗄️) can represent a single parity group. The parity group icon can also represent two or more parity groups that are concatenated. If two or more parity groups are concatenated, volumes can be striped across two or more drives. Therefore, concatenated parity groups provide faster access (particularly, faster sequential access) to data.

For example, if the parity group icon (🗄️) indicates a single parity group 1-3, the text 1-3 appears on the right of the icon. If the parity group icon indicates two or more parity groups that are connected together, all the connected parity groups appear on the right of the icon. For example, if the parity group 1-3 is connected with the parity group 1-4, the text 1-3[1-4] appears on the right of the parity group icon. (All the parity groups connected with 1-3 are enclosed by square brackets).



Note: Storage Navigator does not allow you to connect two or more parity groups. If you want to use connected parity groups, contact the maintenance personnel.

- The list displays statistics about parity group usage, processor usage, etc.

The list displays up to a maximum of 4,096 resources at a time. If the number of resources exceeds 4,096, the **Previous** and **Next** buttons allows you to display the remaining resources. If you select connected parity groups in the tree, the list displays usage statistics for all the connected parity groups.

If you select the Array Control Processor (**ACP**) folder in the tree, the list displays a list of disk adapters (see the figure below), so that you are able to confirm whether each disk adapter is located in *Cluster-1* or *Cluster-2*. For example, if the **Cluster-1** column displays **0** and the **Cluster-2** column displays a hyphen, the disk adapter is located in Cluster-1.

Adapter	Cluster-1	Cluster-2
DKA-1B	0	-
DKA-1C	0	-
DKA-1D	0	-
DKA-1E	0	-
DKA-2H	-	0
DKA-2J	-	0
DKA-2K	-	0
DKA-2L	-	0

Figure 4-3 List of Disk Adapters



Note: When you select the parity group icon and the list displays the icon of **concatenated parity groups**, the list only displays the ID of the parity group at the top of the concatenated parity groups.

For details on the list contents, see Monitoring Resources in the Storage System.

- The **Page** area displays the number of the current page and the following items are being used to change pages of list.
 - **Previous** button allows you to display the previous 4,096 resources.
 - **N/M** drop-down list: The **N** displays the number of the current page. The **M** displays total number of pages. You can click the drop-down list, and choose the number of the page you want to display.
 - **Next** button allows you to display the next 4,096 resources.
- The **Volume Migration** button starts the Volume Migration if that program is enabled and **long-range** is specified for the display range. Volume Migration lets you optimize hard disk drive performance. For details, see Overview of Volume Migration and Using the Volume Migration Windows.
- The **Draw** button displays a line graph that illustrates changes in usage statistics. The graph can display up to eight lines simultaneously.
- The line graph illustrates changes in usage statistics. The vertical axis indicates the usage rates (in percentage). The horizontal axis indicates dates and/or times.

If you select connected parity groups in the tree, the graph displays changes in usage statistics for all the connected parity groups.

When you illustrate a graph of the following information with specifying **short-range**, you can select an item to be displayed in the graph from the drop-down list on the upper-right of the graph:

- Information of external volume groups or external volumes.
- Information of cache memory (write pending rates or usage statistics about cache memory).

In addition, when you display information of external volume groups or external volumes, you can select the highest value of the Y-axis (the vertical axis) of the graph at the **Chart Y Axis Rate** drop-down list on the upper left of the graph.

Figure 4-4 shows the drop-down lists to select an item displayed in the graph and to select the highest value of the Y-axis.




Figure 4-4 Selecting an Item Displayed in the Graph and the Highest Value of the Y-Axis (Physical Tab)

If you select an item to be displayed from the drop-down list before selecting the **Draw** button, the graph shows the selected item. After drawing, if you select another item from the drop-down list, the graph will be updated without reselecting the **Draw** button. Depending upon the values of the selected item, arrange the graph by changing the highest value of the Y-axis.

When you select Long-Range Storage

The **Performance Management** window displays the statistics collected and stored in long range. The usage statistics about resources for three months (i.e. 93 days) collected every 15 minutes can be displayed in the window.

When you select **long-range**, the **Volume Migration** button is activated. The system administrator can start Volume Migration to migrate volumes for balancing workloads based on the monitoring results displayed in the **Performance Management** window.

**Notes:**

- When you select **long-range**, the value displayed at **Gathering Interval** is ineffective. The gathering interval is fixed to 15 minutes regardless of the displayed value.
 - When **long-range** is selected, you cannot view the statistics of external volume groups and external volumes, and hyphens (-) appear in the list instead of these values. In this case, you cannot draw the graph. To view this data, select **short-range**. For details on viewing the usage statistics of external volume groups and external volumes, see Viewing Usage Statistics on External Volume Groups and Viewing Usage Statistics on External Volumes in External Volume Groups.
 - When **long-range** is selected, you cannot view the ratio of ShadowImage processing and so on to all processing, and the usage statistics about cache memory. For details on "the ratio of ShadowImage processing and so on to all processing", see the description about the **ShadowImage** column in Viewing Usage Statistics on Volumes in Parity Groups.
-

When you Select Short-Range Storage

The **Performance Management** window displays the statistics collected and stored in short range. The usage statistics about resources are collected with the interval indicated by **Gathering Interval**. The storing period of statistics, which is equivalent to the range of monitoring data that can be displayed, changes between 8 hours and 15 days depending on the gathering interval.



Note: When you select **short-range**, the **Volume Migration** button is deactivated. In addition, estimated usage rates of volumes after migration by Volume Migration cannot be displayed.

The usage statistics for a same term might be slightly different between selecting **short-range** and **long-range** because the monitoring precision of these two interval types differs.

For details on the types of storing period of statistics, see Understanding Statistical Storage Ranges. For details on the relationship between collection interval and the storing period of the statistics, see Monitoring Options Window.

LDEV Tab of the Performance Monitor Window

When you click **Go, Performance Manager** and then **Performance Management** on the menu bar of the Storage Navigator main window, Performance Monitor starts and the **Performance Management** window is active. The **Performance Management** window includes the **LDEV** tab, which lets you view statistics about disk access performance. For example, the window displays the I/O rate (the number of I/Os per second), the transfer rate (the size of data transferred per second), the average response time for parity groups and volumes.

For details on how to use this window, see *Monitoring Hard Disk Drives*.




Figure 4-5 LDEV Tab of the Performance Management Window

When the **LDEV** tab is active, the **Performance Management** window contains the following items:

- When **Monitoring Switch** is **Enable**, Performance Monitor is monitoring the storage system (a **Disable** setting indicates that the system is not being monitored).

- **Gathering Interval** displays a number between 1 and 15 to indicate how often data collection is performed. If the number of the CUs to be monitored is 64 or less, the value between 1 and 15 appears as a gathering interval by minutes. For example, if **1 min.** is displayed, the information obtained every one minute is displayed in the list and the graph. In case 65 or more CUs are monitored, the statistics are displayed every 5, 10 or 15 minutes.
- The drop-down list on the right of **Monitoring Data** indicates storing period of statistics (monitoring data). The statistics displayed in the **LDEV** tab are stored only in short range, therefore, **short-range** is displayed in this drop-down list and you cannot change it. The range of monitoring data that can be displayed in the window is between 8 hours and 15 days depending on the gathering interval.

For details on the types of storing period of statistics, see Understanding Statistical Storage Ranges.

For details on the relationship between collection interval and the storing period of the statistics, see Monitoring Options Window.

- **Monitoring Term** lets you narrow the range of statistics that should be displayed in the window.

Starting and ending times for collecting statistics are displayed on both sides of the slide bar. Performance Monitor stores the monitoring data between these times, and you can specify the desired term within this range as statistics for the specified term will be displayed the target of display in lists and graphs formats.

For example, if you want to view statistics within the range of 10:30 July 1 2006 to 10:30 July 2 2006, you set 2006/07/01 10:30 to the **From** box, set 2006/07/02 10:30 to the **To** box, and then click the **Apply** button.

To set a date and time to the **From** and **To** boxes, do either of the following:

- Move the slider to the left or to the right.
- In the text box, select the number that you want to change. Next, click the upward or downward arrow button.

When you specify dates and time in **From** and **To** boxes, Performance Monitor calculates the length (in minutes) of the specified period and displays the calculated length. When calculating the length in minutes, Performance Monitor rounds up to the nearest minute.



Notes:

- The **From** and **To** boxes are grayed out if the monitoring data (that is, obtained statistics) is not stored in the storage system.
 - The Real Time option is grayed out when the **LDEV** tab is active.
-

- In the **Monitoring Data** area, the drop-down list on the upper right of the list specifies the type of statistics to be displayed in the window. If you want to view I/O rate, select **IOPS** (I/Os per second) from the drop-down list. If you want to view transfer rate, select **MB/s** (megabytes per second) from the drop-down list.
- The tree lists parity groups, external volume groups, and V-VOL groups.

Box folders (for example, **Box 1**, **Box E1**, **Box V1**, and **Box X1**) are displayed below the storage system folder. The number at the end of a **Box** folder name indicates the number at the beginning of parity group ID, external volume group ID or V-VOL group ID. For example, if you double-click the **Box 1** folder, the tree displays a list of parity groups whose IDs begin with **1** (for example, **1-1** and **1-2**). At the right of each ID, the RAID level appears. When you select a parity group, the list on the right lists volumes in the parity group.

The IDs of external volume groups are beginning with the letter "**E**". Therefore, a folder whose number is beginning with the letter "**E**", such as **Box E1**, contains external volume groups. For example, if you double-click the **Box E1** folder, the tree displays a list of external volume groups whose IDs begin with **E1** (for example, **E1-1** and **E1-2**). The external volume groups do not have parity information, and no RAID level appears at the right of IDs.



Note: A volume existing in an external storage system and mapped to a volume in the USP V/VM storage system by using Universal Volume Manager is called an *external volume*. An *external volume group* is a quantity of external volumes grouped together for managing and that does not contain any parity information, unlike a parity group. However, **Performance Management** window treats external volume groups same as parity groups for convenience.

When Copy-on-Write Snapshot is used the IDs of V-VOL groups begin with the letter "**V**". Dynamic Provisioning is used the IDs of V-VOL groups begin with the letter "**X**". Therefore, if a folder number ends with the letter "**V**" or "**X**", such as **Box V1** or **Box X1**, it contains V-VOL groups. For example, if you double-click the **Box V1** folder, the tree displays **V1-1**, which is a V-VOL group ID beginning with **V1**. If you double-click the **Box V2** folder, the tree displays **V2-1**, which is a V-VOL group ID beginning with **V2**. The V-VOL groups do not have parity information, and no RAID level appears at the right of IDs.

Unlike a parity group, a V-VOL group is a group of virtual volumes and does not contain any parity information. However, **Performance Management** windows treat V-VOL groups the same as parity groups for convenience.

The parity group icon () can represent a single parity group. The parity group icon can also represent two or more parity groups that are connected together. If two or more parity groups are connected together, volumes can be striped across two or more drives. Therefore, connected parity groups provide faster access (particularly, faster sequential access) to data.

For example, if the parity group icon () indicates a single parity group 1-3, the text 1-3 appears on the right of the icon. If the parity group icon indicates two or more parity groups that are connected together, all the connected parity groups appear on the right of the icon. For example, if the parity group 1-3 is connected with the parity group 1-4, the text 1-3 [1-4] appears on the right of the parity group icon. (All the parity groups connected with 1-3 are enclosed by square brackets).



Notes:

- Storage Navigator does not allow you to connect two or more parity groups. If you want to use connected parity groups, contact the maintenance personnel.
- When the ShadowImage or ShadowImage for IBM z/OS quick restore operation is being performed, a Storage Navigator window may display old information (status before the quick restore operation) on volume (LDEV) configurations.
- In this case, wait until the quick restore operation completes, and then click File, Refresh on the menu bar of the Storage Navigator main window to update the Storage Navigator window.

-
- The list displays statistics about disk access performance (for example, I/O rate, transfer rate, read hit ratio, write hit ratio, and average response time). For details on the list contents, see Monitoring Hard Disk Drives.

If you select connected parity groups in the tree, the list displays statistics about disk access performance for all the connected parity groups.

The list displays up to a maximum of 4,096 resources at a time. If the number of resources exceeds 4,096, the **Previous** and **Next** buttons allows you to display the remaining resources.



Note: When you select a folder icon such as **Box 1** or **Box 3**, the **concatenated parity group** icon is displayed in a list. The **only** parity group ID displayed will be that of the group at the **top** of the concatenated parity groups. The icon used for **concatenated parity groups** is the same as a regular parity group icon. There is no specific icon for concatenated parity groups.

-
- The **Page** area displays the number of the current page and the following items are being used to change pages of list.
 - **Previous** button allows you to display the previous 4,096 resources.

- **N/M** drop-down list: The **N** displays the number of the current page. The **M** displays total number of pages. You can click the drop-down list, and choose the number of the page you want to display.
- **Next** button allows you to display the next 4,096 resources.
- The **Draw** button displays a line graph that illustrates changes in the I/O rate, the transfer rate, etc. The graph can display up to eight lines simultaneously.
- The line graph illustrates changes in the I/O rate, the transfer rate, etc. The vertical axis indicates the usage rates (in percentage). The horizontal axis indicates dates and/or times.

If you select connected parity groups in the tree, the graph illustrates changes in disk access performance (e.g., the I/O rate and the transfer rate) for all the connected parity groups.

Before clicking the **Draw** button, you should use the drop-down list at the right corner of the list (just below the drop-down list for specifying the type of statistics) to specify the type of information that will be displayed in the graph. The drop-down list to specify the item to be displayed is shown below.

	LDEV	Emulation	IO Rate(IOPS)	Read(IOPS)
	00:02:05	OPEN-V	1	2
	00:02:06	OPEN-V	2	3
	00:02:07	OPEN-V	3	4
	00:03:00	OPEN-V	4	5
	00:03:01	OPEN-V	5	6
	00:03:02	OPEN-V	6	7
	00:03:03	OPEN-V	7	8

IO Rate
 IO Rate
 Read
 Write
 Read Hit
 Write Hit
 Back Trans.
 Draw

Figure 4-6 Selecting an Item to be Displayed in the Graph (LDEV Tab)

The items that can be selected in the drop-down list changes depending on the type of statistics you selected. Some items can be selected only when you select I/O rate (**IOPS**) or the transfer rate (**MB/s**). The items you can select in the drop-down list and the corresponding type of statistics are shown in Table 4-1.

Table 4-1 The Items You Can Select in the Drop-Down List and the Types of Statistics (LDEV Tab)

Item selected in the drop-down list	Meaning	Type of statistics	
		I/O rate	transfer rate
IO Rate	The I/O rate. The number of I/O accesses per second).	✓	
Read	The number of read accesses per second.	✓	
Write	The number of write accesses per second.	✓	
Read Hit	The read hit ratio.	✓	✓
Write Hit	The write hit ratio.	✓	✓
Back Trans.	Backend transfer. The number of data transfers between the cache memory and the hard disk drive.	✓	✓
Trans.	The transfer rate. The size of data transferred per second.		✓

✓: Can be selected (blank): Cannot be selected

When you draw a graph, use the **Detail** check box to illustrate the desired information and the **Chart Y Axis Rate** drop-down list to arrange the graph as you like.




Figure 4-7 Chart Y Axis Rate Drop-Down List and Detail Check Box (LDEV Tab)

The **Chart Y Axis Rate** drop-down list lets you select the highest value of the Y-axis (the vertical axis) of the graph.

If you select the **Detail** check box and then the **Draw** button, the graph displays detailed statistics as explained in Table 4-2. The information in the graph depends on the item selected in the drop-down list on the right of the list.

Table 4-2 Detailed Information that can be Displayed in the Graph (LDEV Tab)

Select Detail and this Item in the Drop-Down List	The Graph Contains
IO Rate	- statistics in sequential access mode - statistics in random access mode - statistics in CFW (cache fast write) mode ⚠ Note: If the read hit ratio or the write hit ratio is high, random access mode is used for transferring data instead of sequential access mode. For example, random access mode is likely to be used for transferring data to disk areas to which the Cache Residency Manager function is applied.
Read	
Write	
Read Hit	
Write Hit	
Back Trans.	- the number of data transfers from the cache memory to hard disk drives ("Cache to Drive") - the number of data transfers from hard disk drives to the cache memory in sequential access mode ("Drive to Cache Sequential") - the number of data transfers from hard disk drives to the cache memory in random access mode ("Drive to Cache Random")
Trans.	The graph does not display detailed information.

Port-LUN Tab of the Performance Monitor Window

When you click **Go, Performance Manager** and then **Performance Management** on the menu bar of the Storage Navigator main window, Performance Monitor starts and the **Performance Management** window is active. The **Performance Management** window includes the **Port-LUN** tab, which lets you view statistics about I/O rates, transfer rates, and average response time at storage system ports, host groups, LU paths, etc.



Note: Performance Monitor can obtain statistics about traffics of ports connected to open-system host groups only. The statistics about traffics of ports connected to mainframe host groups cannot be obtained.

For details on how to use this window, see Monitoring Ports and Monitoring LU Paths.




Figure 4-8 Port-LUN tab of the Performance Management Window

When the **Port-LUN** tab is active, the **Performance Management** window contains the following items:

- When **Monitoring Switch** is **Enable**, Performance Monitor is monitoring the storage system (a **Disable** setting indicates that the system is not being monitored).
- **Gathering Interval** indicates that the statistics are collected at the interval displayed here. If the number of the CUs monitored is 64 or less, the value between 1 and 15 appears as a gathering interval by minutes. For example, if **1 min.** is displayed, the information obtained every one minute is displayed in the list and the graph. In case 65 or more CUs are monitored, the statistics are displayed every 5, 10 or 15 minutes.
- The drop-down list on the right of **Monitoring Data** indicates storing period of statistics (monitoring data). The statistics displayed in the **Port-LUN** tab are stored only in short range. The range of monitoring data that can be displayed in the window is between 8 hours and 15 days depending on the gathering interval.

For details on the types of storing period of statistics, see Understanding Statistical Storage Ranges.

For details on the relationship between collection interval and the storing period of the statistics, see Monitoring Options Window.

- **Monitoring Term** let you narrow the range of statistics that should be displayed in the window.

The starting and ending time for collecting statistics is displayed on both sides of the slide bar. Performance Monitor stores the monitoring data between these times, and you can specify desired term within this range as the target to display in lists and graphs.

For example, if you want to view statistics within the range of 10:30 July 1 2007 to 10:30 July 2 2007, you set 2007/07/01 10:30 to the **From** box, set 2007/07/02 10:30 to the **To** box, and then click the **Apply** button.

To set a date and time to the **From** and **To** boxes, do either of the following:

- Move the slider to the left or to the right.
- In the text box, select the number that you want to change. Next, click the upward or downward arrow button.

When you specify dates and time in **From** and **To** boxes, Performance Monitor calculates the length (in minutes) of the specified period and displays the calculated length. When calculating the length in minutes, Performance Monitor rounds up the fraction.



Note: The **From** and **To** boxes are grayed out if the monitoring data (that is, obtained statistics) is not stored in the storage system.

- The **Real Time** option lets you view statistics in real-time mode, where statistics are updated at a gathering interval you specify between 1 and 15 minutes.

When you select the **Real Time** option, use the drop-down list to select the number of recent collections of statistics which should be displayed in the graph. You can select the number of times from 5 or 90. This setting determines the range of recent statistics to be displayed in the graph. For example, suppose the gathering interval is 1 minute. In this case, if you select **90** from the drop-down list, the graph displays statistics obtained in the last 90 minutes (multiplying 1 minute by 90 times).

- In the **Monitoring Data** area, the drop-down list on the upper right of the list specifies the type of statistics to be displayed in the window. If you want to view I/O rates, select **IOPS** (I/Os per second) from the drop-down list. If you want to view transfer rates, select **MB/s** (megabytes per second) from the drop-down list.
- The tree contains the **Subsystem** folder. Below the **Subsystem** folder are ports (such as  and ):



This icon indicates the attribute of the stored port is **Target**, or **Initiator/External**.



This port icon indicates either of the following:

A Fibre Channel port in **Standard** mode. LUN security is applied to this port.

If the port name is followed by its fibre address, the port is a Fibre Channel port. For example, CL1-A(EF) indicates that the CL1-A port is a Fibre Channel port.



This port icon indicates either of the following:

A Fibre Channel port in **Standard** mode. LUN security is *not* applied to this port.



A Fibre Channel port in **High Speed** mode. LUN security is applied to this port.



A Fibre Channel port in **High Speed** mode. LUN security is *not* applied to this port.



A Fibre Channel port in **Initiator/External MIX** mode. LUN security is *not* applied to this port.



A Fibre Channel port in **Initiator/External MIX** mode. LUN security is applied to this port.

When you double-click a port, the host groups () that correspond to that port are displayed. The host group ID appears on the left of the colon (:). The host group name appears on the right of the colon.

When you double-click a host group, an item named **LUN** () appears. When you select **LUN**, the list on the right lists LU paths.

For details about LUN security, host groups, and LU paths, see the *LUN Manager User's Guide*.

- The list displays statistics (that is, I/O rates, transfer rates, or average response time). For details on the list contents, see Monitoring Ports and Monitoring LU Paths.
- The **SPM** button starts the Server Priority Manager program product if that program has been enabled. For details of Server Priority Manager, see Overview of Server Priority Manager and [Server Priority Manager Operation](#).

The **SPM** button is deactivated in real-time mode. To start Server Priority Manager, activate the **From** and **To** boxes and release Performance Monitor from real-time mode.

- If the **Current Control** label displays **Port Control**, the system is controlled by the upper limits and the threshold specified in the **Port** tab of the **Server Priority Manager** window.
 - If the **Current Control** label displays **WWN Control**, the system is controlled by the upper limits and the threshold specified in the **WWN** tab of the **Server Priority Manager** window.
 - If the **Current Control** label displays **No Control**, the system performance is not controlled by Server Priority Manager.
- The **Draw** button displays a line graph that illustrates changes in the I/O rate or the transfer rate. The graph can display up to eight lines simultaneously.
- The line graph illustrates changes in the I/O rate or the transfer rate. The vertical axis indicates the I/O rate or the transfer rate. The horizontal axis indicates dates and/or times.

When the graph displays I/O rates or the transfer rates for a port controlled by an upper limit or a threshold, the graph also displays a line that indicates the upper limit or the threshold.

When you draw a graph, use the **Detail** check box and the drop-down list to illustrate the desired information, and use the **Chart Y Axis Rate** drop-down list to arrange the graph convenient to work.




Figure 4-9 Chart Y Axis Rate Drop-Down List, Detail Check Box, and the Drop-Down List to Select the Item to be Displayed (Port-LUN Tab)

The **Chart Y Axis Rate** drop-down list lets you select the highest value of the Y-axis (the vertical axis) of the graph.

If you select the **Detail** check box after drawing a graph by selecting the **Draw** button, the graph displays detailed statistics as explained. The detailed statistics can be displayed only:

- When you select the **Subsystem** folder (📁) in the tree and select a port in the list.

The graph displays detailed statistics about workloads on the port selected in the list. For details on the graph, see Viewing Port Workload Statistics.

- When you select **LUN** (📁) in the tree and select a LUN (an address of a volume) in the list.

The graph displays detailed statistics about workloads on the LU paths selected in the list. The information in the graph depends on the item selected in the drop-down list on the right of the **Detail** check box. For details on the graph see Viewing Workload Statistics on LU Paths.

Viewing Port Workload Statistics

Figure 4-10 shows an example of a graph displaying about workload of a port. In this example, the user selected the port CL1-A in the list and then select **Draw**. In this example, 1 minute is specified as the gathering interval. The graph contents changes depending on the selection of the **Detail** check box. The figure shows the following:

- The workload on the port CL1-A is 200 IO/s at 8:00, and 300 IO/s at 10:00 (refer to the graph on the left).

- For the period of 7:59 to 8:00, the maximum workload on CL1-A is 300 IO/s. The average workload on CL1-A is 200 IO/s. The minimum workload on CL1-A is 100 IO/s (refer to the graph on the right).
- For the period of 9:59 to 10:00, the maximum workload on CL1-A is 400 IO/s. The average workload on CL1-A is 300 IO/s. The minimum workload on CL1-A is 200 IO/s (refer to the graph on the right).
- When **Detail** is deselected, only one line appears in the graph. This line is equivalent to the line Ave. (1 min), which appears when **Detail** is selected.




Figure 4-10 Graphs Illustrating Changes in Workloads on a Port

Viewing Workload Statistics on LU Paths

When you select **LUN** () in the tree, select a LUN in the list. select the **Draw** button, and then select the **Detail** check box, the graph displays detailed statistics about workload on the LU paths. The detailed statistics that can be displayed differ depending on items you select in the drop-down list, as explained in Table 4-3.

Table 4-3 Detailed Information that can be Displayed in the Graph (Port-LUN Tab)

Select Detail and this Item in the Drop-Down List	The Graph Contains
IO (the number of I/Os per second)*	- statistics in sequential access mode - statistics in random access mode Note: If the read hit ratio or the write hit ratio is high, random access mode is used for transferring data instead of sequential access mode. For example, random access mode is likely to be used for transferring data to disk areas to which the Cache Residency Manager function is applied.
Read (the number of read accesses per second)*	
Write (the number of write accesses per second)*	
Read Hit (the read hit ratio)	
Write Hit (the write hit ratio)	
Back Trans. (backend transfer; the number of I/Os between the cache memory and hard disk drives)	- the number of data transfers from the cache memory to hard disk drives ("Cache to Drive") - the number of data transfers from hard disk drives to the cache memory in sequential access mode ("Drive to Cache Seq.") - the number of data transfers from hard disk drives to the cache memory in random access mode ("Drive to Cache Rnd.")
* You can select this item only when I/O rates are displayed.	

WWN Tab of the Performance Monitor Window

When you click **Go, Performance Manager** and then **Performance Management** on the menu bar of the Storage Navigator main window, Performance Monitor starts and the **Performance Management** window is active. The **Performance Management** window includes the **WWN** tab, which lets you view statistics (I/O rates, transfer rates, and average response time) about traffic between host bus adapters in the hosts and ports on the storage system.



Caution: The **WWN** tab is unavailable if Server Priority Manager is not enabled.



Note: A **WWN** (**Worldwide Name**) is a 16-digit hexadecimal number used as the unique identifier for a host bus adapter. Host bus adapters are contained in hosts and serve as ports for connecting the hosts and the storage system. One WWN represents one host bus adapter.

For details on how to use this window, see *Monitoring Paths between Host Bus Adapters and Ports*.




Figure 4-11 WWN tab of the Performance Management Window

When the **WWN** tab is active, the **Performance Management** window contains the following items:

- When **Monitoring Switch** is **Enable**, Performance Monitor is monitoring the storage system (a **Disable** setting indicates that the system is not being monitored).
- **Gathering Interval** indicates that the statistics are collected at the interval displayed here. If the number of the CUs to be monitored is 64 or less, the value between 1 and 15 appears as a gathering interval by minutes. For example, if **1 min.** is displayed, the information obtained every one minute is displayed in the list and the graph. In case 65 or more CUs are monitored, the statistics are displayed every 5, 10 or 15 minutes.
- The drop-down list on the right of **Monitoring Data** indicates storing period of statistics (monitoring data). The statistics displayed in the **WWN** tab are stored only in short range, therefore, **short-range** is displayed in this drop-down list and you cannot change it. The range of monitoring data that can be displayed in the window is between 8 hours and 15 days depending on the gathering interval.

For details on the types of storing period of statistics, see Understanding Statistical Storage Ranges.

For details on the relationship between collection interval and the storing period of the statistics, see Monitoring Options Window.

- **Monitoring Term** let you narrow the range of statistics that should be displayed in the window.

The starting and ending time for collecting statistics are displayed on both sides of the slide bar. Performance Monitor stores the monitoring data between these times, and you can specify the desired term within this range. Statistics for the specified term will be displayed in list and graph formats.

For example, if you want to view statistics within the range of 10:30 July 1 2007 to 10:30 July 2 2007, you set 2007/07/01 10:30 to the **From** box, set 2007/07/02 10:30 to the **To** box, and then click the **Apply** button.

To set a date and time to the **From** and **To** boxes, do either of the following:

- Move the slider to the left or to the right.
- In the text box, select the number that you want to change. Next, click the upward or downward arrow button.

When you specify dates and time in **From** and **To** boxes, Performance Monitor calculates the length (in minutes) of the specified period and displays the calculated length. When calculating the length in minutes, Performance Monitor rounds up the fraction.



Note: The **From** and **To** boxes are grayed out if the monitoring data (that is, obtained statistics) is not stored in the storage system.

- The **Real Time** option lets you view statistics in real-time mode, where statistics are updated at a gathering interval you specify between 1 and 15 minutes.

When you select the **Real Time** option, use the drop-down list to select the number of recent collections of statistics which should be displayed in the graph. You can select the number of times from 5 or 90. This setting determines the range of recent statistics to be displayed in the graph. For example, suppose the gathering interval is 1 minute. In this case, if you select **90** from the drop-down list, the graph displays statistics obtained in the last 90 minutes (multiplying 1 minute by 90 times).

- In the **Monitoring Data** area, the drop-down list on the upper right of the list specifies the type of statistics to be displayed in the window. If you want to view I/O rates, select **IOPS** (I/Os per second) from the drop-down list. If you want to view transfer rates, select MB/s (megabytes per second) from the drop-down list.
- The tree contains the **Subsystem** folder. Below the **Subsystem** folder are *SPM groups*, which are groups of multiple WWNs. When you double-click an SPM group (🖨️), the host bus adapters (🖨️) belonging to that SPM group are displayed. The WWN and the SPM name of the host bus adapter are displayed to the right of the 🖨️ icon. If you double-click **Not Grouped** in the tree, host bus adapters (WWNs) that do not belong to any SPM group are displayed.



Note: If the WWN of a host bus adapter (HBA) is displayed in red in the tree, the host bus adapter is connected to two or more ports, but the traffic between the HBA and some of the ports are *not* monitored by Performance Monitor. When many-to-many connections are established between HBAs and ports, make sure that all the traffic between HBAs and ports is monitored (see Monitoring All Traffic between HBAs and Ports for instructions).

- The list displays statistics (that is, I/O rates, transfer rates, or average response time). For details on the list contents, see Monitoring Paths between Host Bus Adapters and Ports.
- The **SPM** button starts the Server Priority Manager program product if Server Priority Manager is enabled. For details of Server Priority Manager, see Overview of Server Priority Manager and Server Priority Manager .
The **SPM** button is deactivated in real-time mode. To start Server Priority Manager, activate the **From** and **To** boxes and release Performance Monitor from real-time mode.

- If the **Current Control** label displays **Port Control**, the system is controlled by the upper limits and the threshold specified in the **Port** tab of the **Server Priority Manager** window.
 - If the **Current Control** label displays **WWN Control**, the system is controlled by the upper limits and the threshold specified in the **WWN** tab of the **Server Priority Manager** window.
 - If the **Current Control** label displays **No Control**, the system performance is not controlled by Server Priority Manager.
- The **Draw** button displays a line graph that illustrates changes in the I/O rate or the transfer rate. The graph can display up to eight lines simultaneously.
- The line graph illustrates changes in the I/O rate or the transfer rate. The vertical axis indicates the usage rates (in percentage). The horizontal axis indicates dates and/or times.

When the graph displays I/O rates or the transfer rates for a host bus adapter or an SPM group controlled by an upper limit, the graph also displays a line that indicates the upper limit.

When you draw a graph, the **Chart Y Axis Rate** drop-down list lets you select the highest value of the Y-axis (the vertical axis) of the graph.




Figure 4-12 Chart Y Axis Rate Drop-Down List, Detail Check Box, and the Drop-Down List to Select the Item to be Displayed (Port-LUN tab)

Monitoring Options Window

When you click **Go, Performance Manager** and then **Performance Management** on the menu bar of the Storage Navigator main window, Performance Monitor starts. When you click the **Monitoring Options** tab, the **Monitoring Options** window is displayed. Use it to make settings for obtaining usage rates about hard disk drives, channel processors, disk processors, etc.



Note: This note explains the following statistics to be displayed in tabs of **Performance Management** windows.

- The statistics of LUs that are displayed in the **Port-LUN** tab.
- The statistics of volumes that are displayed in the **LDEV** tab.

In the above tabs, performance statistics of unused volumes are displayed as hyphens (-), if the range of monitored CUs does not match the range of CUs used in the disk storage system or registered as external volumes. In addition, depending on your disk subsystem configuration, the list may display performance statistics for some volumes and not display performance statistics for other volumes.

To correctly display performance statistics, you must specify CUs to be monitored as follows:

- To display performance statistics of a LUSE volume in the **Port-LUN** tab, you must specify all volumes that make up the LUSE volume as the monitoring targets.
 - To display performance statistics of parity group to be displayed in the **LDEV** tab, you must specify all volumes that belong to the parity group as the monitoring target.
-




Figure 4-13 Monitoring Options Window of Performance Monitor

The **Monitoring Switch** area in the **Monitoring Options** window contains the following items:

- **Current Status**

Select **Enable** to start obtaining statistics from the storage system (that is, monitoring). To stop monitoring, select **Disable**. The default setting is **Disable**.

- **Gathering Interval**

Specify the interval to obtain usage statistics about the storage system for short range monitoring. This option is activated when you specify **Enable** for **Current Status**. If CUs to be monitored are 64 or less, you can specify the value between 1 and 15 minutes by minutes, and the default setting is 1 minute. For example, if you specify 1 minute for the gathering interval, Performance Monitor collect statistics (such as I/O rates and transfer rates) every one minute.

If CUs to be monitored are 65 or more, the gathering interval can be specified to the value 5, 10 or 15 minutes (in the 5-minuted interval), and default is 5 minutes. For example, if you specify the gathering interval to 5 minutes, Performance Monitor gathers statistics (such as I/O rate and transfer rate) every 5 minutes.

This option is effective only for:

- Statistics displayed in the **LDEV**, **Port-LUN**, and **WWN** tabs
- Statistics displayed in the **Physical** tab with selecting **short-range** for the storing period

When viewing the **Physical** tab with **long-range** selected for the storing period, the statistics collected every 15 minutes are displayed regardless of the value of the **Gathering Interval**.

The **Monitoring Target** area in the **Monitoring Options** window contains the following items:

- **PG Selection Support**

If you click this button, all parity groups are listed in the **PG** list.



Note: If you click this button when a large-sized configuration is being used, it may take long time to gather the **PG** list.

- **PG** list:

Indicates the list of IDs of the parity groups that can be monitored. Each cell displayed in the list is accompanied by an icon . Click **PG** on the header to sort the parity groups by ID. If you click the cell for an individual parity group, the CUs that belong to that parity group are displayed in the **CU** table on the right.

- **LDKC** list:

Indicates LDKC numbers. To select an LDKC as a monitoring target, click the LDKC number. All the CUs that belong to that LDKC are displayed.

- **CU** table:

The **CU** table consists of cells representing CUs. Each row consists of 16 cells (CUs). A set of 16 rows represents CUs for one LDKC. The table header row displays the last digit of each CU number in the form of +n (n is an integer between 0 and 9, or a letter from A to F).

To select a CU, click a cell to invert its color. To restore the cell to its original color, click the inverted cell. To select all the (16) CUs of the same number in the second last digits, click the CU number (00 to F0). By dragging the cursor over multiple cells, you can select all the cells from the source to the destination.

For your information, one cell corresponds to one CU. The relation of the display of a cell to the CU status is shown below:

Table 4-4 Relationship between the Display of a Cell and the CU Status in the CU Table

CU Exists	CU Monitoring Status	Letter Displayed in The Cell
Yes	CU is being monitored	S: in black
	CU is to be released from monitoring	<i>R</i> : in blue italics
	CU is not being monitored	None
No	CU will be monitored when it exist	N: in black
	CU is to be released from monitoring	-: Hyphen in black bold
	CU will not be monitored when it exist	-: Hyphen in black

- **Monitoring Target CUs:**
Indicates the number of existing and newly added CUs to be monitored.
- The **Select** button adds the CU selected in the CU table to the CUs to be monitored.

Table 4-5 Characters Indicated in a Cell When the Select Button is Clicked

Before Click	After Click
-: Hyphen in black	<i>N</i> : in black italics
<i>R</i> : in blue italics	S: in black

- The **Release** button removes the CU from monitoring targets.

Table 4-6 Characters Indicated in a Cell When Release Button is Clicked

Before Click	After Click
N: in black	-: Hyphen in black bold
S: in black	<i>R</i> : slanted in blue italics

- The **Apply** button applies settings in the **Monitoring Options** window to the storage system.
- The **Reset** button resets the settings in the **Monitoring Options** window.

Performance Monitor has two kinds of periods (ranges) for collecting and storing statistics: short range and long range.

The storing period of statistics in *short range* is determined by the settings of **Gathering Interval** option. Performance Monitor saves the statistics obtained up to 1440 times in SVP. Therefore, you can estimate the storing period of statistics with "gathering interval multiplied by 1440". For example, if you specify one minute for the gathering interval, the statistics for one day can be stored at the maximum from the following formula:

$$1 \text{ minute} \times 1440 = 1440 \text{ minutes} = 24 \text{ hours} = 1 \text{ day}$$

This storing period is the range of display in the **Performance Management** windows. When you specify one minute for the gathering interval like the example above, Performance Monitor can display the statistics for one day (i.e. 24 hours) in the list and graph at the maximum. Also, when you specify 15 minutes for the gathering interval, Performance Monitor can display the statistics for 15 days in the list and graph at the maximum.

However, the value of the **Gathering Interval** option has nothing to do with the storing period of statistics in *long range* monitoring. The gathering interval in long range is fixed to 15 minutes. Therefore, when you select the **long-range** storing period in the **Physical** tab of Performance Monitor, the display range in the window is always three months (i.e. 93 days).



Note: The monitoring switch settings in the **Monitoring Options** window work in conjunction with the monitoring switch settings in the **Usage Monitor** windows of TrueCopy, TrueCopy for IBM z/OS, Universal Replicator, and Universal Replicator for IBM z/OS. Therefore, if you use the **Monitoring Options** window to start or stop monitoring, or to change the gathering interval, the monitoring settings of these remote copy functions will also change. Conversely, if the settings in each **Usage Monitor** window are changed, the settings in the **Monitoring Options** window of Performance Monitor also change automatically.

Other Windows

Apart from the Performance Monitor window and its tabs, the Performance Monitor interface also includes the following windows:

- **TC Monitor** window

This window displays the information about remote copy operation of TrueCopy. The contents in this window are the same as those of the **Usage Monitor** window of TrueCopy. For details on the contents in this window, see the *TrueCopy User's Guide*.

- **TCz Monitor** window

This window displays the information about remote copy operation of TrueCopy for IBM z/OS. The contents in this window are the same as those of the **Usage Monitor** window of TrueCopy for IBM z/OS. For details on the contents in this window, see the *TrueCopy for IBM z/OS User's Guide*.

- **UR Monitor** window

This window displays the information about remote copy operation of Universal Replicator. The contents in this window are the same as those of the **Usage Monitor** window of Universal Replicator. For details on the contents in this window, see the *Universal Replicator User's Guide*.

- **URz Monitor** window

This window displays the information about remote copy operation of Universal Replicator for IBM z/OS. The contents in this window are the same as those of the **Usage Monitor** window of Universal Replicator for IBM z/OS. For details on the contents in this window, see the *Universal Replicator for IBM z/OS User's Guide*.



Note: If each remote copy function (TrueCopy, TrueCopy for IBM z/OS, Universal Replicator, or Universal Replicator for IBM z/OS) is not installed in your environment, the corresponding tabs (**TC Monitor** tab, **TCz Monitor** tab, **UR Monitor** tab, or **URz Monitor** tab) are unavailable.

Using the Volume Migration Windows

Volume Migration can be used to balance workloads among hard disk drives and thus remove bottlenecks from the system. Volume Migration can migrate volumes from the current parity group to another parity group. Volume migrations can be performed manually (by the system administrator) or automatically by Volume Migration.

In addition, Volume Migration can also migrate external volumes in external volume groups. However, external volumes can be migrated manually by the system administrator, but they cannot be migrated automatically by Volume Migration. For details on external volume groups, see Understanding Statistical Storage Ranges. For details on external volumes, see External Volume Group Usage Statistics.



Notes:

- From an open-system host, users can also use Command Control Interface (CCI) to perform manual volume migration by using commands. To perform volume migration by CCI, Volume Migration should be installed on the storage system. For details on manual volume migration by using CCI, see Sample Volume Migration. For details on CCI, see the *Command Control Interface (CCI) User and Reference Guide*.
 - The Volume Migration windows also display the statuses of volume migration performed by programs other than Volume Migration. Volumes being operated by these other programs cannot be operated by Volume Migration.
 - If the user type of your user ID is storage partition administrator, you cannot use Volume Migration. For details on the limitations when using Performance Manager logged in as a storage partition administrator, see Storage Partition Administrators Limitations.
-

This section describes options, buttons and other components in Volume Migration windows.

Manual Plan Window of the Volume Migration

When you click the **Volume Migration** button in the **Performance Management** window, Volume Migration starts and the **Manual Plan** window of the Volume Migration is displayed. The **Manual Plan** window lets you create and execute manual migration plans.

For details on operations in this window, see Migrating Volumes Manually.




Figure 4-14 Manual Plan Window of the Volume Migration

The **Manual Plan** window displays the following:

- **Monitoring Term:** Displays the monitoring period specified in the **Performance Management** window.
Volume Migration analyzes disk usage information collected by Performance Manager during the monitoring period, and then calculates estimated usage rates of the source and target parity groups after a proposed volume migration. The program does not calculate estimated usage rates of external volume groups.
- **Gathering Interval:** Displays the collection interval in long range used in Performance Monitor. The interval in long range is fixed to 15 minutes.


- **Tree:** The tree lists parity groups and external volume groups in the storage system. The **Parity Group** folder contains icons of parity groups, and the **External Group** folder contains icons of external volume groups.

The graphics of parity group icons and external volume group icons indicate the RAID levels, and HDD classes.


The types of icons for parity groups and external volume groups are as follows.

Parity group icons:

-  indicates a fixed parity group when it appears below the **Parity Group** folder.
-  indicates a RAID-1 parity group. The HDD class for this parity group is A.
-  indicates a RAID-1 parity group. The HDD class for this parity group is B.
-  indicates a RAID-1 parity group. The HDD class for this parity group is C.
-  indicates a RAID-1 parity group. The HDD class for this parity group is D.
-  indicates a RAID-5 parity group. The HDD class for this parity group is A.
-  indicates a RAID-5 parity group. The HDD class for this parity group is B.
-  indicates a RAID-5 parity group. The HDD class for this parity group is C.
-  indicates a RAID-5 parity group. The HDD class for this parity group is D.
-  indicates a RAID-6 parity group. The HDD class for this parity group is A.
-  indicates a RAID-6 parity group. The HDD class for this parity group is B.
-  indicates a RAID-6 parity group. The HDD class for this parity group is C.
-  indicates a RAID-6 parity group. The HDD class for this parity group is D.



- ① The parity group ID. The number before the hyphen is the frame number.
The number after the hyphen is the parity group number.
- ② The number before the slash is the average usage rate.
The number after the slash is the maximum usage rate.
If an exclamation mark (!) is displayed before a usage rate,
the reported parity group usage rate is likely to be inaccurate
If a plus mark (+) is displayed before a usage rate 0, such as "+0",
the usage rate is not completely 0 but less than 1.
- ③ The number and name of the CLPR corresponding to the parity group.
The number before the colon is the CLPR number. The name after the
colon is the CLPR name.




- ④ The parity group 1-3 is connected with the parity group 1-4.

Figure 4-15 Information about Parity Groups

External volume group icons:

-  indicates an external volume group when it appears below the **External Group** folder.
No volumes in this external volume group can be used as target
volumes for manual migration.




- ① The letter "E" indicates an external volume group.
- ② The external volume group ID. The number before the hyphen is a
frame number. The number after the hyphen is an external volume
group number.
- ③ The number and name of the CLPR corresponding to the external
volume group.
The number before the colon is the CLPR number. The name after the
colon is the CLPR name.

Figure 4-16 Information about External Volume Groups

Dynamic Provisioning group icons:

-  indicates a Dynamic Provisioning virtual volume group when it
appears below the **Dynamic Provisioning** folder.



- ① The letter "X" indicates a Dynamic Provisioning virtual volume group.
- ② The virtual volume group ID. The number before the hyphen is a frame number. The number after the hyphen is a virtual volume group number.
- ③ The number and name of the CLPR corresponding to the virtual volume group.
The number before the colon is the CLPR number. The name after the colon is the CLPR name.

Figure 4-17 Information about Dynamic Provisioning Virtual Volume Groups

- **LDEVs:** A list of volumes in the parity groups you selected in the tree is indicated. The following volumes are not displayed in this list:
 - Pool volumes
 - Journal volumes
 - System disks.

Select a row to specify the **Source LDEV**.

LDEV	Indicates the status of the volume, in this format: LDKC:CU:LDEV. An icon showing the status of the volume is displayed on the left of the ID.
Emulation	Indicates the emulation type of the volume (LVI or LUN).
Capacity	Indicates the capacity of the volume.
Ave. (%)	Indicates the average usage rate for the parity group. If an exclamation mark (!) is displayed, the reported usage rate is likely to be inaccurate, because the volumes may be moved by Volume Migration or formatted by Virtual LVI or Open Volume Management. If a plus mark (+) is displayed before 0, such as "+0", the usage rate is not completely 0 but less than 1.
Max (%)	Indicates the maximum usage rate for the parity group. If an exclamation mark (!) is displayed, the reported usage rate is likely to be inaccurate, because the volumes may be moved by Volume Migration or formatted by Virtual LVI or Open Volume Management.
Used Capacity	Indicates the capacity of the pool related to the Dynamic Provisioning virtual volume. The indicated pool capacity might be bigger than the LDEV capacity of the Dynamic Provisioning virtual volume. The maximum possible capacity of the used pool is 42 MB larger than the LDEV capacity. This is because the size of each page of a pool used by Dynamic Provisioning is 42 MB. If you want to verify the data size of the volume whether manual migration can be done or not, you must confirm the data size based on the indicated pool capacity related to the virtual volume.
Pool ID	Indicates the pool ID related to the Dynamic Provisioning virtual volume.
Pool Remainder	Indicates the remaining capacity of the pool related to the Dynamic Provisioning virtual volume.
Pool Capacity	Indicates the capacity of the pool related to the Dynamic Provisioning virtual volume.

The following are **volume icons**, contained in parity groups, external volume groups and Dynamic Provisioning virtual volume groups, which can appear in the tree.

The volumes contained in an external volume group are called *external volumes*. In the tree of **Manual Plan** tab, volumes in a parity group and external volumes are expressed by the same icons. Also, the volumes contained in a Dynamic Provisioning virtual volume group are called *Dynamic Provisioning virtual volumes*. In the tree of **Manual Plan** tab, volumes in a parity group and Dynamic Provisioning virtual volumes are expressed by the same icons.

Volume icons:

-  indicates a normal volume that can be migrated to a different parity group but cannot be used as a migration destination.
-  (blue) indicates a reserved volume of Volume Migration. This volume is reserved as a migration destination by Volume Migration.

When you specify a source volume in the **Manual Plan** window, Volume Migration checks these reserved volumes to find volumes that can be used as target volumes. If Volume Migration finds a reserved volume that can be used as target volumes, Volume Migration changes the volume icon to .

-  (green) indicates a reserved volume of other programs. This volume is reserved by a program other than Volume Migration, therefore, Volume Migration cannot operation it.
-  indicates a volume that can be used as a migration destination.
-  indicates a volume that is being migrated, or a normal volume that is registered in a migration plan. A volume that is being migrated by a program other than Volume Migration is also indicated by this icon.

,  (blue),  (green), and . Among these icons, you can specify  as a source volume.

- **Target** Button: After selecting the volume to migrate in the **LDEVs**, click the **Target** button. The list of LDEVs in the **Target (Reserved) LDEV** is displayed.
- **Target (Reserved) LDEV:** The list appears of volumes that volume migration is possible corresponding to the **Source LDEV**. Select a row to specify the **Target LDEV**. Specify the target volume which is the destination volume in the text box. The usage rate for parity groups before volume migration and the estimated usage rate for parity groups after volume migration are displayed on the right of the text box.

Items displayed in the list are as follows.

LDEV	Indicates the volume ID, in this format: LDKC:CU:LDEV. An icon showing the status of the volume is displayed on the left of the ID.
Emulation	Indicates the emulation type of the volume (LVI or LUN).
Capacity	Indicates the capacity of the volume.
Ave. (%)	Indicates the average usage rate for the parity group. If an exclamation mark (!) is displayed, the reported usage rate is likely to be inaccurate, because the volumes may be moved by Volume Migration or formatted by Virtual LVI or Open Volume Management. If a plus mark (+) is displayed before 0, such as "+0", the usage rate is not completely 0 but less than 1.
Max (%)	Indicates the average usage rate for the parity group. If an exclamation mark (!) is displayed, the reported usage rate is likely to be inaccurate, because the volumes may be moved by Volume Migration or formatted by Virtual LVI or Open Volume Management.
CLPR	indicates the number and name of the CLPR corresponding to the parity group in the format "CLPR-number:CLPR-name". For details on CLPRs, see the <i>Virtual Partition Manager User's Guide</i> .
Pool Remainder	Indicates the remaining capacity of the pool related to the Dynamic Provisioning virtual volume.
Pool Capacity	Indicates the capacity of the pool related to the Dynamic Provisioning virtual volume.
Pool ID	Indicates the pool ID related to the Dynamic Provisioning virtual volume.

The displayed usage rates in the **Target (Reserved) LDEV** list are not usage rates of volumes but usage rates of parity groups that volumes are belonged. Comparing the usage rates of volumes in the **Target (Reserved) LDEV** list before volume migrations enables you easy to select the target volumes.

- **Graph** button displays the **Graph** window, which indicates the estimated usage rates of the specified **Source LDEV** and the **Target (Reserved) LDEV**.




Figure 4-18 Graph Window

- **Source LDEV** lets you specify the source volume (that is, the volume that you want to migrate), On the right of the text box appear the following:
 - The current usage rate for the parity group that the source volume belongs to.
 - An estimated usage rate for the parity group after a proposed volume migration.

If you specify an external volume as a source volume, Volume Migration cannot calculate the estimated usage rate after migration. In this case, a hyphen (-) appears instead of a value.

- **Target (Reserved) LDEV** lets you specify the target volume (that is, the location to which you want to migrate the source volume). On the right of the text box appear the following:
 - The current usage rate for the parity group that the target volume belongs to.
 - An estimated usage rate for the parity group after a proposed volume migration.

If you specify an external volume as a target volume, Volume Migration cannot calculate the estimated usage rate after migration. In this case, a hyphen (-) appears instead of a value.

- The line graph illustrates the past usage rates of the source volume and the target volume.
- The **Close** button closes the Graph window.
- The **Set** button adds a new manual migration plan to the list. The new manual migration plan consists of the volumes selected in **Source LDEV** and **Target (Reserved) LDEV**.
- List: Each row in the list indicates a manual migration plan, which is a pair consisting of a source volume and the corresponding target volume, created in the **Manual Plan** window. A migration plan before being applied to the storage system is displayed in blue, and a migration plan after being applied (or during execution) to the storage system is displayed in black. Volume Migration performs multiple migration plans concurrently. When a migration plan was finished, the plan will disappear from the list.

A migration plan being executed by a program other than Volume Migration (such as CCI) is also displayed in the list all together. Migration plans executed by other programs are not erased from the list automatically until the user operates for finish by using the program that performed the volume migration.

The **Stat** column displays PSUE if the execution of the migration plan which is created by the application software other than Storage Navigator failed. In this case, you must check the status of the application software which created the migration plan.

The **DEL** column displays the character "D" if you attempt to delete a migration plan that is waiting for execution, or already in progress. Migration plans indicated by "D" are actually deleted when you click the **Apply** button.

The **%** column indicates the progress of the manual migration. The progress value is updated once per minute.

The **Source LDEV** columns display information about source volumes:

- **LDEV** indicates the volume ID, in this format: LDKC:CU:LDEV. To the left of the device ID, the  icon is displayed indicating that the volume is being migrated or is registered in a migration plan.
- **Emulation** indicates the emulation type of the volume.
- **Capacity** indicates the capacity of the volume.
- **RAID** indicates the RAID type.
- **PG** indicates the parity group ID or the external volume group ID. The number before the hyphen (-) is the frame number. The number after the hyphen is the group number.



Note: If the indicated parity groups are the **concatenated parity groups**, this item only displays the ID of the parity group at the top of the concatenated parity groups.

- **HDD** indicates the type of the hard disk drive.

- **CLPR** indicates the number and name of the CLPR corresponding to the parity group in the format "CLPR-number:CLPR-name". For details on CLPRs, see the *Virtual Partition Manager User's Guide*.

The **Target LDEV** columns display information about target volumes:

- **LDEV** indicates the volume ID, in this format: LDKC:CU:LDEV. The blue  icon indicates a reserved volume of Volume Migration. The green  icon indicates a reserved volume of other programs.
- **RAID** indicates the RAID type.
- **PG** indicates the parity group ID or the external volume group ID. The number before the hyphen (-) is the frame number. The number after the hyphen is the group number.
- **HDD** indicates the type of the hard disk drive.
- **CLPR** indicates the number and name of the CLPR corresponding to the parity group in the format "CLPR-number:CLPR-name". For details on CLPRs, see the *Virtual Partition Manager User's Guide*.
- **Owner** indicates the program which uses this volume as a target volume. If "USP v" is displayed, this is the target volume which is reserved by Storage Navigator. If "Other[XX]" is displayed, this is the target volume which is reserved by CCI or another program. XX is an id which is given by this program. For example, 01 indicates CCI, and 99 indicates HiCommand Tiered Storage Manager.

A migration plan executed by a program other than Volume Migration is indicated by the  (green) icon displayed in the **LDEV** column and by "Other[XX]" displayed in the **Owner** column. You cannot delete such a migration plan by the **Delete** button.

- **Display usage rate.** Displays the usage rate of the parity groups in the tree and displays the usage rate of the volumes in the **LDEVs** and the **Target (Reserved) LDEV**. If you click this button when usage rates are displayed, usage rates are not displayed in this window.
- The **Delete** button deletes the migration plan (that is, a row in the list) selected in the list.

If you select a migration plan in blue (i.e. a migration plan that is not applied to the storage system yet) and click the **Delete** button, the migration plan is deleted immediately.

If you select a migration plan in black (i.e. a migration plan that is waiting for execution or is being executed) and click the **Delete** button, the **DEL** column displays the character "D" and the selected migration plan will not be deleted until you click the **Apply** button.

The **Delete** button does not take effect when both blue and black migration plans are selected. Confirm that only blue or black migration plans are selected before you select **Delete**.



Note: You cannot use the **Delete** button to delete a migration plan executed by a program other than Volume Migration (such as CCI).



Caution: If you delete a migration plan which is being executed, the data in the target volume is not guaranteed.

- The **Apply** button applies manual migration plans in the list to the storage system to start volume migrations.
- The **Close** button closes the **Manual Plan** window.
- The **Refresh** button updates information in the **Manual Plan** window. The button updates the displaying of the tree after the execution of manual migration plan.

Auto Plan Window of the Volume Migration

When you click the **Volume Migration** button in the **Performance Management** window, Volume Migration starts. The **Auto Plan** window lets you make settings for automating volume migrations.

For details on operations in this window, see Automating Volume Migrations.

Volume Migration Monitoring Term : 2006/05/08 16:45 - 2006/05/09 16:45
Gathering Interval : 15 Minute(s)

Auto Migration Plans

Auto Migration : ☐ Disable Migration Status : Not planned yet Auto Refresh ☐ Enable ☒ Disable
Plan Creation : 2006/12/04 08:45 Next Migration : 2006/12/04 08:45

%	Source LDEV						Target LDEV			
	LDEV	Emulation	Capacity	RAID	PG	HDD	LDEV	RAID	PG	HDD
10	00:00:00	OPEN-V	2.28 GB	N/A	2-1	DKR1B-J047FC	00:02:00	N/A	2-1	DKR1B-J047...
-	00:00:01	OPEN-V	2.28 GB	N/A	2-1	DKR1B-J047FC	00:02:01	N/A	2-1	DKR1B-J047...

Buttons: Make New Plan, Delete, Delete All

Auto Plan Parameters

Auto Migration Function: ☒ Enable ☐ Disable

Sampling Term: Date ☐ None ☒ Every day ☐ Once every X days: 2 ☐ Once a week: Sunday ☐ Once a month: 1
Time From: 21 : 45 To: 00 : 45

Number of Sampling Points: ☒ All sampling points ☐ X highest sampling points: 1

Migration Time: 00 : 45

Auto Migration Conditions: Max. Migration duration: 40 min. Max. disk utilization: 55 % Max. number of vols for migration: 24

Buttons: Default, Set, Reset, Close

Figure 4-19 Auto Plan Window of the Volume Migration



Note: The migration of volumes in a fixed parity group, external volumes, and volumes reserved by a program other than Volume Migration cannot be automated.



Note: Auto migration is not directed to volumes in a fixed parity group, external volumes, and volumes reserved by a program other than Volume Migration.



Note: The **Auto Plan** window of the **Volume Migration** displays the following:

- **Monitoring Term:** Displays the monitoring period specified in the **Performance Management** window.

Volume Migration analyzes disk usage information collected by Performance Monitor during the monitoring period, and then creates auto migration plans based on that analysis.

- **Gathering Interval:** Displays the collection interval in long range used in Performance Monitor. The interval in long range is fixed to 15 minutes.
- **Auto Migration Plans:** Above the list, the following items are displayed:
Auto Migration indicates whether auto migrations will take place. If this option is **Enable**, auto migrations will take place at the time displayed at **Next Migration**. If this option is **Disable**, auto migrations will not take place.

Migration Status indicates the progress of auto migration plans.

- If **Not planned yet** is displayed, Volume Migration has not yet created auto migration plans.
- If **Not performed yet** is displayed, Volume Migration has created auto migration plans but has not executed them yet.
- If **Failed to make a plan** is displayed, Volume Migration failed to create auto migration plans.
- If **Under migration** is displayed, Volume Migration is executing an auto migration plan.
- If **Last migration has been canceled** is displayed, Volume Migration has canceled the last migration plan.

Plan Creation: Indicates when the auto migration plans were created.

Next Migration: Indicates the next time that Volume Migration will perform auto migration operations.

Auto Refresh: Enables you to specify whether to automate updating of information about migration plans, which should be displayed in the list.

- If you select **Enable**, information in the list will automatically be updated, so that you will be able to check the progress of migrations in real time. If a new plan is created, the new plan will automatically be added to the list. A message appears whenever information in the list is updated.
- If you select **Disable**, information in the list will not automatically be updated.



Note: The time values displayed on this window are retrieved from the SVP, and may not correspond to the time displayed on the Storage Navigator computer.

The list displays auto migration plans (that is, pairs of the source volume and the target volume). The plans are executed from the top of the list when you perform auto migration.

- The percentage **%** column indicates the progress of the auto migration. A hyphen appears if Volume Migration has not executed the auto migration plan yet.
- **Source LDEV** columns display information about source volumes.
 - **LDEV** indicates the volume ID, in this format: LDKC:CU:LDEV.
 - **Emulation** indicates the emulation type of the volume.
 - **Capacity** indicates the capacity of the volume.
 - **RAID** indicates the RAID type.
 - **PG** indicates the parity group ID. The number before the hyphen (-) is the frame number. The number after the hyphen is the group number.



Note: If the indicated parity groups are the **concatenated parity groups**, the only parity group ID displayed will be that of the group at the top of the concatenated parity groups.

- **HDD** indicates the type of the hard disk drive.
- **Target LDEV** columns display information about target volumes.
 - **LDEV** indicates the volume ID, in this format: LDKC:CU:LDEV.
 - **RAID** indicates the RAID type.
 - **PG** indicates the parity group ID. The number before the hyphen (-) is the frame number. The number after the hyphen is the group number.
 - **HDD** indicates the type of the hard disk drive.

On the right of the list appear the following buttons:

- **Make New Plan** button discards the existing auto migration plans and then creates new auto migration plans.
- **Delete** button deletes the auto migration plan selected in the list.
- **Delete All** button deletes all the auto migration plans in the list.
- **Auto Plan Parameters** lets you specify parameters that affect automatic creation and execution of migration plans.
 - **Auto Migration Function:** Specifies whether to automate volume migrations. If you select **Enable**, volumes are automatically migrated. If you select **Disable**, volumes are not automatically migrated.
 - **Sampling Term:** The Volume Migration auto migration function analyzes resource usage within the specified Sampling Term and then creates auto migration plans based on that analysis. Use Sampling Term to narrow resource usage statistics to be analyzed; Volume Migration does not analyze resource usage statistics outside of the specified sampling term.
 - **Date:**
If this option is **None**, Volume Migration does not analyze resource usage statistics and does not create auto migration plans.

If this option is **Every day**, Volume Migration creates auto migration plans every day.

If this option is **Once every X days**, Volume Migration creates auto migration plans every X days. Use the drop-down list to specify the number of days (that is, the value of X).

If this option is **Once a week**, Volume Migration creates auto migration plans on the selected day of the week. Use the drop-down list to select a day.

If this option is **Once a month**, Volume Migration creates auto migration plans on the Xth day of the month. Use the drop-down list to specify a day.

- **Time** lets you specify a timeframe. Volume Migration analyzes resource usage statistics within the specified period.
- **Number of Sampling Points:**
 - **All sampling points:** If this option is selected, Volume Migration uses all the average usage rates during the **Sampling Term** to determine disk usage.
 - **X highest sampling points:** If this option is selected, Volume Migration uses the Xth highest average usage rate during the **Sampling Term** to determine disk usage. Use the drop-down list to specify the value of X.
- **Migration Time** lets you specify the time when auto migration starts. If the **Auto Migration Function** option is **Enable**, the auto migration operation starts at the specified time.
- **Auto Migration Condition** lets you specify the following auto migration parameters:
 - **Max. migration duration** lets you specify the time limit (10-120 minutes) for an auto migration plan. If the auto migration plan is not complete within this limit, the remaining migration operation(s) are performed at the next scheduled execution time (if a new plan is not made).
 - **Max. disk utilization** lets you specify the disk usage limit (10-100%) during auto migration. If the most recent usage for any source or target parity group is over this limit when auto migration starts, Volume Migration cancels the auto migration plan and will retry the plan at the next execution time (if a new plan is not made before then).
 - **Max. number of vols for migration** lets you specify the maximum number of volumes (1-40) that can be migrated during one auto migration plan.
- The **Default** button sets each parameter in **Auto Plan Parameters** to the default value and then issues the parameter to the storage system.
- The **Set** button applies each parameter in Auto Plan Parameters to the storage system.

- The **Reset** button restores each parameter in Auto Plan Parameters to the value before change.
- The **Close** button closes the **Auto Plan** window.

Attribute Window of the Volume Migration

When you click the **Volume Migration** button in the **Performance Management** window, Volume Migration starts. The **Attribute** window of Volume Migration lets you find out which parity group belongs to which HDD class. The **Attribute** window also lets you view detailed information about parity groups, external volume groups, and Dynamic Provisioning virtual volume groups.

You must use the **Attribute** window when you reserve target volumes.

For details on operations in this window, see Reserving Migration Destinations for Volumes, Setting a Fixed Parity Group, and Changing the Maximum Disk Usage Rate for Each HDD Class.




Figure 4-20 Attribute Window of the Volume Migration

The **Attribute** window displays the following information:

- **Monitoring Term** displays the monitoring period specified in the **Performance Management** window.
- **Gathering Interval** displays the collection interval in long range used in Performance Monitor. The interval in long range is fixed to 15 minutes.

- Tree: The tree lists HDD classes, external volume groups, and Dynamic Provisioning virtual volume groups. On the right of HDD classes, the upper limit (threshold) of the disk usage rate is displayed.

When you double-click an HDD class folder, a list of parity groups in that HDD class appears. When you double-click the **External Group** folder, a list of external volume groups appears. When you double-click the **Dynamic Provisioning** folder, a list of Dynamic Provisioning virtual volume groups appears.

The icons of parity groups, external volume groups and Dynamic Provisioning virtual volume groups are as follows:

- The icon  indicates a fixed parity group when it appears below an HDD class folder. This icon also indicates an external volume group when it appears below the **External Group** folder.
Volume Migration cannot perform auto migration for the volumes in a fixed parity group or an external volume group.
-  indicates a RAID-1 parity group. The HDD class for this parity group is A.
-  indicates a RAID-1 parity group. The HDD class for this parity group is B.
-  indicates a RAID-1 parity group. The HDD class for this parity group is C.
-  indicates a RAID-1 parity group. The HDD class for this parity group is D.
-  indicates a RAID-5 parity group. The HDD class for this parity group is A.
-  indicates a RAID-5 parity group. The HDD class for this parity group is B.
-  indicates a RAID-5 parity group. The HDD class for this parity group is C.
-  indicates a RAID-5 parity group. The HDD class for this parity group is D.
-  indicates a RAID-6 parity group. The HDD class for this parity group is A.
-  indicates a RAID-6 parity group. The HDD class for this parity group is B.
-  indicates a RAID-6 parity group. The HDD class for this parity group is C.
-  indicates a RAID-6 parity group. The HDD class for this parity group is D.

Parity group icons can represent a single parity group. Parity group icons can also represent two or more parity groups that are connected together. For example, if the text 1-3 appears on the right of an icon, the icon represents the parity group 1-3. If the text 1-3[1-4] appears on the right of an icon, the icon represents two parity groups (that is, 1-3 and 1-4) that are connected together.

- List: Information in the list differs depending on the item that you select in the tree.

When you select a class in the tree, the list displays information about parity groups in that class.

PG	HDD	Ave.(%)(Total)	Max.(%)	Total	Reserved	CLPR
 1-1	DKR1B-J047FC	13	14	2	0	00:CLPR0
 1-2	DKR1B-J047FC	!12	!13	2	0	01:CLPR1
 1-3	DKR1B-J047FC	13	14	2	0	02:CLPR2
 1-4	DKR1B-J047FC	!14	!15	2	0	03:CLPR3

Figure 4-21 Information about Parity Groups

- **PG** indicates the parity group ID. The number before the hyphen (-) is the frame number. The number after the hyphen is the group number. The meaning of icons displayed here is the same as the icons in the tree.



Note: If the indicated parity groups are the **concatenated parity groups**, this item only displays the ID of the parity group at the top of the concatenated parity groups.

- **HDD** indicates the type of the hard disk drive.
- **Ave.** indicates the average usage rate for the parity group.
If an exclamation mark (!) is displayed, the reported usage rate is likely to be inaccurate, because the configuration has changed (e.g., volumes have been moved by Volume Migration or changed by Virtual LVI or Open Volume Management).
If a plus mark (+) is displayed before a usage rate 0, such as "+0", the usage rate is not completely 0 but less than 1.
- **Max.** indicates the maximum usage rate for the parity group.
If an exclamation mark (!) is displayed, the reported usage rate is likely to be inaccurate, because the configuration has changed (e.g., volumes have been moved by Volume Migration or changed by Virtual LVI or Open Volume Management).
- **Total** indicates the total number of volumes in the parity group.
- **Reserved** indicates the number of reserved volumes.
- **CLPR** indicates the number and name of the CLPR corresponding to the parity group in the format "CLPR-number:CLPR-name". For details on CLPRs, see the *Virtual Partition Manager User's Guide*.

When you select a parity group in the tree, the list displays information about volumes in that parity group.

	LDEV	Emulation	Capacity	Ave.(%)(T...	Max.(%)	CLPR	Owner
	00:02:05	OPEN-V	2.28 GB	+0	1	00:CLPR0	Other [05]
	00:02:06	OPEN-V	2.28 GB	11	12	00:CLPR0	USP
	00:02:07	OPEN-V	2.28 GB	2	3	00:CLPR0	USP
	00:03:00	OPEN-V	2.28 GB	13	14	00:CLPR0	-
	00:03:01	OPEN-V	2.28 GB	4	5	00:CLPR0	USP
	00:03:02	OPEN-V	2.28 GB	15	16	00:CLPR0	-
	00:03:03	OPEN-V	2.28 GB	6	7	00:CLPR0	USP

Figure 4-22 Information about Volumes in a Parity Group

- **LDEV** indicates the volume ID, in this format: LDKC:CU:LDEV.

You cannot change the attribute of the following volumes.

- LUSE volumes
- Pool volumes
- System disks
- Journal volumes
- Virtual volumes of Dynamic Provisioning that are not associated with pools

 indicates a normal volume that can be migrated but cannot be used as a migration destination. You cannot change this volume to a reserved volume (.

 indicates a normal volume that can be migrated. If you change the attribute of this volume to a reserved volume (, this volume can be used as a migration destination (that is, target volume).

 (blue) indicates a *reserved volume* of Volume Migration. This volume is reserved as a migration destination by Volume Migration. A reserved volume cannot be migrated. You can change this volume to a normal volume (, and then you can use this volume as a source volume.

 (green) indicates a *reserved volume* of other programs. This volume is reserved by a program other than Volume Migration (such as CCI), therefore, Volume Migration cannot operation it. You cannot change the attribute of this volume.

- **Emulation** indicates the emulation type of the volume.
- **Capacity** indicates the capacity of the volume.
- **Ave.** indicates the average usage rate for the volume.

If an exclamation mark (!) is displayed, the reported usage rate is likely to be inaccurate, because the configuration has changed (e.g., volumes have been moved by Volume Migration or changed by Virtual LVI or Open Volume Management).

If a plus mark (+) is displayed before a usage rate 0, such as "+0", the usage rate is not completely 0 but less than 1.

- **Max.** indicates the maximum usage rate for the volume.

If an exclamation mark (!) is displayed, the reported usage rate is likely to be inaccurate, because the configuration has changed (e.g., volumes have been moved by Volume Migration or changed by Virtual LVI or Open Volume Management).

- **CLPR** indicates the number and name of the CLPR corresponding to the parity group which the volume belongs to, in the format "*CLPR-number:CLPR-name*". For details on CLPRs, see the *Virtual Partition Manager User's Guide*.

- **Owner** indicates the program which reserved this volume.

If "USP v" is displayed, this is the reserved volume which is reserved by Storage Navigator.

If "Other[XX]" is displayed, this is the reserved volume which is reserved by CCI or another program. XX is an id which is given by this program. For example, 01 indicates CCI, and 99 indicates HiCommand Tiered Storage Manager.

If a hyphen (–) is displayed, this volume is not reserved and this is a normal volume.

When you select **External Group** in the tree, the list displays information about external volume groups.

ExG	HDD	Total	Reserved	CLPR
 E1-1	External	2	0	02:CLPR2
 E1-2	External	2	0	03:CLPR3
 E1-3	External	2	0	00:CLPR0
 E1-4	External	2	0	01:CLPR1

Figure 4-23 Information about External Volume Groups

- **ExG** indicates the ID of the external volume group. The letter "E" stands for an external volume group. The number before the hyphen (–) is the frame number. The number after the hyphen is the group number.

External volume groups are displayed with the  icon. Volume Migration cannot perform auto migration for the volumes in an external volume group.

- **HDD** indicates the type of the hard disk drive.
- **Total** indicates the total number of volumes in the external volume group.
- **Reserved** indicates the number of reserved volumes.
- **CLPR** indicates the number and name of the CLPR corresponding to the external volume group in the format "*CLPR-number:CLPR-name*". For details on CLPRs, see the *Virtual Partition Manager User's Guide*.

When you select an external volume group in the tree, the list displays information about external volumes (i.e., volumes in the external volume groups).

	ExLDEV	Emulation	Capacity	CLPR	Owner
	00:07:FA #	OPEN-V	2.28 GB	03:CLPR3	-
	00:07:FB #	OPEN-V	2.28 GB	03:CLPR3	-
	00:07:FC #	OPEN-V	2.28 GB	03:CLPR3	-
	00:07:FD #	OPEN-V	2.28 GB	03:CLPR3	-

Figure 4-24 Information about External Volumes

- **ExLDEV** indicates the ID of the external volume, in this format: LDKC:CU:LDEV.

You cannot change the attribute of the following volumes.

- LUSE volumes
- Pool volumes
- System disks
- Journal volumes
- Virtual volumes of Dynamic Provisioning that are not associated with pools

The meaning of icons (, ,  (blue), and  (green)) displayed at the left of IDs have the same meaning as the icons for volumes in parity groups.

- **Emulation** indicates the emulation type of the volume.
- **Capacity** indicates the capacity of the volume.
- **CLPR** indicates the number and name of the CLPR corresponding to the external volume group which the external volume belongs to, in the format "CLPR-number:CLPR-name". For details on CLPRs, see the *Virtual Partition Manager User's Guide*.
- **Owner** indicates the program which reserved this volume, if it is reserved. The displayed item is the same as that of volumes in parity groups.

When you select **Dynamic Provisioning** in the tree, the list displays information about Dynamic Provisioning virtual volume groups.

Dynamic Provisioning	Total	Reserved	CLPR
 X1-1	10	0	00:CLPR0
 X2-1	11	1	01:CLPR1
 X3-1	12	2	02:CLPR2
 X4-1	13	3	03:CLPR3
 X5-1	14	4	00:CLPR0
 X6-1	15	5	01:CLPR1
 X7-1	16	6	02:CLPR2
 X8-1	17	7	03:CLPR3
 X9-1	18	8	00:CLPR0
 X10-1	19	9	01:CLPR1
 X11-1	20	10	02:CLPR2
 X12-1	21	11	03:CLPR3
 X13-1	22	12	00:CLPR0
 X14-1	23	13	01:CLPR1
 X15-1	24	14	02:CLPR2
 X16-1	25	15	03:CLPR3
 X63489-1	26	16	01:CLPR1

Figure 4-25 Information about Dynamic Provisioning Virtual Volume Groups

- **Dynamic Provisioning** indicates the ID of the Dynamic Provisioning virtual volume group. The letter "**X**" stands for a Dynamic Provisioning virtual volume group. The number before the hyphen (-) is the frame number. The number after the hyphen is the group number.

Dynamic Provisioning virtual volume groups are displayed with the  icon. Volume Migration cannot perform auto migration for the virtual volumes in a Dynamic Provisioning virtual volume group.
- **Total** indicates the total number of volumes in the Dynamic Provisioning virtual volume group.
- **Reserved** indicates the number of reserved Dynamic Provisioning virtual volumes.
- **CLPR** indicates the number and name of the CLPR corresponding to the Dynamic Provisioning virtual volume group in the format "*CLPR-number:CLPR-name*". For details on CLPRs, see the *Virtual Partition Manager User's Guide*.

When you select a Dynamic Provisioning virtual volume group in the tree, the list displays information about Dynamic Provisioning virtual volumes (i.e., virtual volumes in the virtual volume group).

	LDEV	Emulation	Capacity	CLPR	Owner	Pool Remainder	Pool Capacity	Pool ID
	00:14:00 X	OPEN-V	2.28 GB	00:CLPR0	USP V	10.06 GB	12.05 GB	0
	00:14:01 X	OPEN-V	2.28 GB	00:CLPR0	USP V	20.12 GB	24.11 GB	1
	00:14:02 X	OPEN-V	2.28 GB	00:CLPR0	-	10.06 GB	12.05 GB	0
	00:14:03 X	OPEN-V	2.28 GB	00:CLPR0	-	20.12 GB	24.11 GB	1
	00:14:04 X	OPEN-V	2.28 GB	00:CLPR0	USP V	10.06 GB	12.05 GB	0
	00:14:05 X	OPEN-V	2.28 GB	00:CLPR0	Other [05]	20.12 GB	24.11 GB	1
	00:14:06 X	OPEN-V	2.28 GB	00:CLPR0	USP V	10.06 GB	12.05 GB	0
	00:14:07 X	OPEN-V	2.28 GB	00:CLPR0	-	20.12 GB	24.11 GB	1
	00:14:08 X	OPEN-V	2.28 GB	00:CLPR0	Other [08]	10.06 GB	12.05 GB	0
	00:14:09 X	OPEN-V	2.28 GB	00:CLPR0	USP V	20.12 GB	24.11 GB	1
	00:14:0A X	OPEN-V	2.28 GB	00:CLPR0	USP V	10.06 GB	12.05 GB	0
	00:14:0B X	OPEN-V	2.28 GB	00:CLPR0	Other [0B]	20.12 GB	24.11 GB	1
	00:14:0C X	OPEN-V	2.28 GB	00:CLPR0	USP V	10.06 GB	12.05 GB	0
	00:14:0D X	OPEN-V	2.28 GB	00:CLPR0	USP V	20.12 GB	24.11 GB	1
	00:14:0E X	OPEN-V	2.28 GB	00:CLPR0	Other [0E]	10.06 GB	12.05 GB	0
	00:14:0F X	OPEN-V	2.28 GB	00:CLPR0	USP V	20.12 GB	24.11 GB	1

Figure 4-26 Information about Dynamic Provisioning Virtual Volumes

- **LDEV** indicates the ID of the Dynamic Provisioning virtual volume, in this format: LDKC:CU:LDEV.
The icons (, ,) (blue), and (green)) displayed at the left of IDs have the same meaning as the icons for volumes in parity groups.
- **Emulation** indicates the emulation type of the Dynamic Provisioning virtual volume.
- **Capacity** indicates the capacity of the Dynamic Provisioning virtual volume.
- **CLPR** indicates the number and name of the CLPR corresponding to the Dynamic Provisioning virtual volume group which the Dynamic Provisioning virtual volume belongs to, in the format "CLPR-number:CLPR-name". For details on CLPRs, see the *Virtual Partition Manager User's Guide*.
- **Owner** indicates the program which reserved this volume, if it is reserved. The displayed item is the same as that of volumes in parity groups.
- **Pool Remainder** indicates the remaining capacity of the pool related to the Dynamic Provisioning virtual volume. If the Dynamic Provisioning virtual volume is not related to the pool, a hyphen (-) appears instead of a value.
- **Pool Capacity** indicates the capacity of the pool related to the Dynamic Provisioning virtual volume. If the Dynamic Provisioning virtual volume is not related to the pool, a hyphen (-) appears instead of a value.
- **Pool ID** indicates the pool ID related to the Dynamic Provisioning virtual volume. If the Dynamic Provisioning virtual volume is not related to the pool, a hyphen (-) appears instead of a value.
- The **Apply** button applies the settings in the **Attribute** window to the storage system.
- The **Close** button closes the **Attribute** window.

History Window of the Volume Migration

When you click the **Volume Migration** button in the **Performance Management** window, Volume Migration starts and the **Volume Migration** window is displayed. The **Volume Migration History** window includes the History window, which displays information about auto migration operations and manual migration operations in the past. For example, you can find out when migration operations took place and also find out whether migration operation finished successfully.

For details on operations in this window, see Viewing the Migration History Log.




Figure 4-27 History Window of the Volume Migration

The **History** window displays the following information:

- **Monitoring Term** displays the monitoring period specified in the **Performance Management** window.
- **Gathering Interval** displays the collection interval in long range used in Performance Monitor. The interval in long range is fixed to 15 minutes.
- **Auto Migration History** displays logs of creation and execution of auto migration plans. The **Erase** button lets you erase the entire information.

- **Migration History** displays logs of all the events about the manual and auto volume migration. This list also displays logs of migration plans executed by programs other than Volume Migration.

The list displays the following information:

- **Date** displays the date when the event occurs.
- **Time** displays the time when the event occurs.
- **Action** displays the event.
- **Source[Parity Pr.]** displays the source volume and the parity group (or the external volume group).
- **Target[Parity Pr.]** displays the target (destination) volume and the parity group (or the external volume group)
- **Owner** displays the program which executed this migration plan.
If "USP v" is displayed, this is the migration plan which is executed by Storage Navigator.
If "Other[XX]" is displayed, this is the migration plan which is executed by a program product other than Storage Navigator (such as CCI). XX is an id which is given by this program product. For example, 01 indicates CCI, and 99 indicates HiCommand Tiered Storage Manager.
- The **Erase** button lets you erase the entire information.



Notes:

- The **Source[Parity Gr.]** and **Target[Parity Gr.]** columns also display the information about external volume group migration, though these groups are not parity groups.
 - If the source parity groups and target parity groups are the **concatenated parity groups**, the list only displays the ID of the parity group at the top of the concatenated parity groups.
-
- The Page drop-down list displays the pages of the history, up to 256 pages. 2,408 or less cases are displayed in one page.
 - The "<<" button displays the page of the latest migration history that are already gathered.
 - The "<" button Displays the previous page of the displayed migration history.
 - The Drop-down list displays the pages of the history, up to 20 pages.
 - The ">" button displays the next page of the displayed migration history.
 - The ">>" button display the page of the oldest migration history.
 - The **Refresh** button updates information in the **History** window.
 - The **Close** button closes the **History** window.

Using the Server Priority Manager Windows

The first section in this chapter explains **Server Priority Manager** windows. The remaining sections explain procedure for monitoring performance and setting the upper limit and threshold for I/O rates and transfer rates.



Note: If the user type of your user ID is *storage partition administrator*, you cannot use Server Priority Manager. For details on the limitations when using Performance Manager logged in as a storage partition administrator, see Storage Partition Administrators Limitations.

The **Server Priority Manager** window has two tabs: **Port** tab and **WWN** tab. If one-to-one connections are established between host bus adapters and ports, use the **Port** tab. If many-to-many connections are established between host bus adapters and ports, use the **WWN** tab.

This section explains these tabs in the **Server Priority Manager** window.

Port Tab of the Server Priority Manager Window

The **Port** tab lets you set the limit on the performance of non-prioritized ports and set the threshold on the performance of prioritized ports.

For operations in this tab, see Port Tab Operations.




Figure 4-28 Port Tab in the Server Priority Manager Window

The **Port** tab displays the following:

- **Current Control Status** can display either **Port Control** or **WWN Control**.
 - If **Port Control** is displayed, the system is controlled by the upper limits and threshold specified in the Port tab.
 - If **WWN Control** is displayed, the system is controlled by the upper limits and threshold specified in the WWN tab.
 - If **No Control** is displayed, the system performance is *not* controlled by Server Priority Manager.

Tip: If **WWN Control** is displayed when the **Port** tab is active, you can click the **Apply** button to switch control so that **Port Control** is displayed.

Tip: To return the control status to **No Control**, specify **Prio.** for attributes of all the ports and then click the **Apply** button.

- The drop-down list near the upper right corner of the **Server Priority Manager** window enables you to narrow ports in the list:
 - If **All** is selected, all the ports appear in the list.
 - If **Prioritize** is selected, only the prioritized ports appear in the list.
 - If **Non-Prioritize** is selected, only the non-prioritized ports appear in the list.

If you change settings of a port, that port remains in the list regardless of the selection in the drop-down list.

- The drop-down list near the upper left corner of the **Port** tab enables you to change the type of performance statistics to be displayed in the list.
 - If **IOPS** (I/Os per second) is selected, the list displays I/O rates for ports. The I/O rate indicates the number of I/Os per second.
 - If **MB/s** (megabytes per second) is selected, the list displays the transfer rates for ports. The transfer rate indicates the size of data transferred via a port in one second.
- The list displays a list of ports and indicates the I/O rate or the transfer rate for each port. This list also enables you to specify the port attributes, and the threshold and upper limit of the port traffic.

The measurement unit for the values in the list can be specified by the drop-down list above the trees. The port traffic (I/O rate and transfer rate) is monitored by Performance Monitor. To specify the monitoring period, use the **Monitoring Term** area of Performance Monitor.

The list shows the following items:

- The **Port** column indicates ports on the storage system.
- The **Ave.** column indicates the average I/O rate or the average transfer rate for the specified period.
- The **Peak** column indicates the peak I/O rate or the peak transfer rate of the ports for the specified period. This value means the top of the Max. line in the detailed port-traffic graph drawn in the Port-LUN tab of Performance Monitor. For details, see Viewing I/O Rates for Disks and Viewing Transfer Rates for Disks.
- The **Attribute** column indicates the priority of each port. **Prio** indicates a prioritized port. **Non-Prio** indicates a non-prioritized port.
- The **Threshold** columns let you specify the threshold for the I/O rate and the transfer rate for each prioritized port. Either of the **IOPS** or **MB/s** column in the list is activated depending on the selection of the drop-down list above the list.

The **IOPS** column lets you specify the threshold for I/O rates. The **MB/s** column lets you specify the threshold for transfer rates. To specify a threshold, double-click a cell to display the cursor in the cell. If you specify a value in either of the **IOPS** or **MB/s** column, the other column is deactivated. You can specify thresholds for I/O rates and transfer rates all together for different prioritized ports.

Even if you use the different type of rate for the threshold as that used for the upper limit values, the threshold control can work for all the ports.

- The **Upper** columns let you specify the upper limit on the I/O rate and the transfer rate for each non-prioritized port. Either of the **IOPS** or **MB/s** column in the list is activated depending on the selection of the drop-down list above the list.

The **IOPS** column lets you specify the upper limit for I/O rates. The **MB/s** column lets you specify the upper limit for transfer rates. To specify an upper limit, double-click a cell to display the cursor in the cell. If you specify a value in either of the **IOPS** or **MB/s** column, the other column is deactivated. You can specify upper limit values for I/O rates and transfer rates all together for different non-prioritized ports.

- If you select the **All Thresholds** check box and enter a threshold value in the text box, the threshold value will be applied to the entire storage system. If you want to specify the threshold for the I/O rate, select **IOPS** from the drop-down list on the right of the text box. If you want to specify the threshold for the transfer rate, select **MB/s** from the drop-down list. For example, if you specify 128 IOPS in **All Thresholds**, the upper limits on non-prioritized ports are disabled when the sum of I/O rates for all the prioritized ports is below 128 IOPS.

Even if you use the different type of rate (**IOPS** or **MB/s**) for the threshold as that used for the upper limit values, the threshold control can work for all the ports.

- If you check the **Delete ports if CHA is removed** check box, Server Priority Manager will delete, from SVP, the setting information of Server Priority Manager on ports in channel adapters that have been removed.

When a channel adapter is removed, the port and its settings are removed from the **Server Priority Manager** window automatically, but they remain in SVP. This may cause that the old setting for Server Priority Manager to be applied to a different channel adapter than the one newly-installed on the same location.

The **Delete ports if CHA is removed** check box is available only when the following Server Priority Manager settings on ports in a removed channel adapter remains on SVP:

- The setting of prioritized ports or non-prioritized ports.
- The setting of prioritized WWNs or non-prioritized WWNs.
- The **Apply** button applies the settings in this window to the storage system.
- The **Reset** button restores the last applied settings in the window. When you click this button, all the changes displayed with the blue text in the window are canceled.
- The **Initialize** button changes the settings in this window as explained below, and then applies the resulting settings to the storage system:
 - All the ports become prioritized ports.
 - The threshold value for all the ports becomes 0 (zero).



Note: The window will display a hyphen (-) instead of 0 (zero).

-
- If the **All Thresholds** checkbox is checked, the check mark disappears.

- The **Close** button closes the **Server Priority Manager** window.

WWN Tab of the Server Priority Manager Window

The **WWN** tab lets you set the limit on the performance of non-prioritized WWNs and set the threshold on the performance of prioritized WWNs.

For operations in this tab, see WWN Tab Operations and Grouping Host Bus Adapters.




Figure 4-29 WWN Tab in the Server Priority Manager Window

The **WWN** tab displays the following:

- **Current Control Status** can display either **Port Control** or **WWN Control**.
 - If **Port Control** is displayed, the system is controlled by the upper limits and threshold specified in the Port tab.
 - If **WWN Control** is displayed, the system is controlled by the upper limits and threshold specified in the WWN tab.
 - If **No Control** is displayed, the system performance is *not* controlled by Server Priority Manager.

Tip: If **Port Control** is displayed when the **WWN** tab is active, you can click the **Apply** button to switch control so that **WWN Control** is displayed.

Tip: To return the control status to **No Control**, specify **Prio.** for attributes of all the host bus adapters and then click the **Apply** button.

- The drop-down list near the upper right corner of the **Server Priority Manager** window enables you to narrow WWNs (host bus adapters) in the list:
 - If **All** is selected, all the WWNs appear in the list.
 - If **Prioritize** is selected, only the prioritized WWNs appear in the list.
 - If **Non-Prioritize** is selected, only the non-prioritized WWNs appear in the list.
- The upper-left tree lists ports and the host bus adapters connected to these ports in the storage system.

Ports on the storage system are shown below the **Subsystem** folder. The ports are indicated by icons such as  and .

When you double-click on a port, the tree is expanded to display two items: **Monitor** and **Non-Monitor**. The host bus adapters that are connected to the specified port are displayed below **Monitor** or **Non-Monitor**.

- If you double-click **Monitor**, the host bus adapters () whose traffics with the specified port are monitored are displayed below **Monitor**.
- If you double-click **Non-Monitor**, the host bus adapters whose traffics with the specified port are not monitored are displayed below **Non-Monitor**.

The WWN and SPM names of the host bus adapters are displayed on the right of the host bus adapter icon () below **Monitor**. *WWNs* (Worldwide Name) are 16-digit hexadecimal numbers used to uniquely identify host bus adapters. *SPM names* are nicknames assigned by the system administrator so that they can easily identify each host bus adapter.

Only the WWN is displayed on the right of the host bus adapter icon () below **Non-Monitor**.



Note: When many-to-many connections are established between host bus adapters (HBAs) and ports, make sure that all the traffics between HBAs and ports monitored. Therefore, make sure that all the connected HBAs are displayed below **Monitor**. For details on how to move an HBA displayed below **Non-Monitor** to below **Monitor**, see *Monitoring All Traffic between HBAs and Ports*.

The list on the right of the tree changes depending on the item you selected in the tree as follows.

- When you selected a port or **Monitor** icon, the list displays the information of host bus adapters that are connected to the ports(s) and monitored by Performance Monitor.
- When you selected the **Monitor** icon or the **Subsystem** folder, the list becomes blank.
- The lower-left tree lists *SPM groups*. The tree also lists host bus adapters (WWNs) in each SPM group:
 - SPM groups (📁), which contain one or more WWNs, are displayed below the **Subsystem** folder. For details on the SPM groups, see Grouping Host Bus Adapters.
 - If you double-click an SPM group, host bus adapters in that group are displayed in the tree. The WWN and SPM name are displayed to the right of the host bus adapter icon (📁).



Note: If the WWN of a host bus adapter (HBA) is displayed in red in the tree, the host bus adapter is connected to two or more ports, but the traffic between the HBA and some of the ports is *not* monitored by Performance Monitor. When many-to-many connections are established between HBAs and ports, you should make sure that all the traffic between HBAs and ports is monitored. For details on the measures when a WWN is displayed in red, see Monitoring All Traffic between HBAs and Ports.

The list on the right of the tree changes depending on the item you selected in the tree as follows:

- When you selected the **Subsystem** folder, the list displays the information of SPM groups.
- When you selected an SPM group icon (📁), the list displays the information of host bus adapters (📁) contained in that SPM group.
- The **Add WWN** button lets you add a host bus adapter to an SPM group. Before using this button, you must select a host bus adapter (📁) from the upper-left tree and also select an SPM group (📁) from the lower-left tree. You can add a host bus adapter that appears below **Monitor** and does not registered on any other SPM group, yet. If you select a host bus adapter below **Non-Monitor** or a host bus adapter already registered on an SPM group, the **Add NNW** button is deactivated.
- The drop-down list at the upper left corner of the list enables you to change the type of performance statistics to be displayed in the list.
 - If **IOPS** (I/Os per second) is selected, the list displays I/O rates for WWNs (host bus adapters). The I/O rate indicates the number of I/Os per second.
 - If **MB/s** (megabytes per second) is selected, the list displays the transfer rates for WWNs (host bus adapters). The transfer rate indicates the size of data transferred in one second.

- The list displays a list of WWNs and indicates the I/O rate or the transfer rate for each host bus adapter corresponding to the selection in the upper-left tree or lower-left tree. This list also enables you to specify the host bus adapter attributes and the upper limit of the host bus adapter traffic.

The measurement unit for the values in the list can be specified by the drop-down list at the upper left corner of the list. The displayed items will change depending on the selected tree and item. The host bus adapter traffic (I/O rate and transfer rate) is monitored by Performance Monitor. To specify the monitoring period, use the **Monitoring Term** area of Performance Monitor.

On the right side of the list appear total number of WWNs, the number of prioritized WWNs, and the number of non-prioritized WWNs.

The list shows the following items:

- The **WWN** column indicates WWNs of host bus adapters. This column does not appear when you select the **Subsystem** folder in the lower-left tree.
- The **SPM Name** column indicates SPM names of host bus adapters. Server Priority Manager allows you assign an *SPM name* to each host bus adapter so that you can easily identify each host bus adapters in the **Server Priority Manager** windows. This column does not appear when you select the **Subsystem** folder in the lower-left tree.
- The **Group** column indicates the SPM group to which the host bus adapter belongs. This column appears when a port is selected in the upper-left tree and does not appear when an SPM group is selected in the lower-left tree.
- The **Per Port** column indicates the traffic (I/O rate or transfer rate) between the host bus adapter and the port selected in the upper-left tree. This item is displayed only when you select an icon in the upper-left tree. The **Per Port** column contains the following columns:
 - Ave.:** Indicates the average I/O rate or the average transfer rate for the specified period.
 - Max.:** Indicates the maximum I/O rate or the maximum transfer rate for the specified period.
- The **WWN Total** column indicates the sum of the traffic (I/O rate or transfer rate) between the host bus adapter and all the ports connected to the host bus adapter. This value means the total traffic of that host bus adapter. This item is displayed only when you select an icon in the upper-left tree. Whichever port you select in the tree, the **WWN Total** column shows the sum of the traffic to all the ports.

The **WWN Port** column contains the following columns:

- Ave.:** Indicates the average I/O rate or the average transfer rate for the specified period.
- Max.:** Indicates the maximum I/O rate or the maximum transfer rate for the specified period.

- The **Ave.** column is also displayed when you select an icon in the lower-left tree. In this case, the **Ave.** column shows the average value same as that of **WWN Total**.

When you select the **Subsystem** folder in the lower-left tree, the **Ave.** column shows the sum of the traffic of the host bus adapters registered on each SPM group.

- The **Max.** column is also displayed when you select an icon in the lower-left tree. In this case, the **Max.** column shows the maximum value same as that of **WWN Total**.

When you select the **Subsystem** folder in the lower-left tree, the **Max.** column shows the sum of the traffic of the host bus adapters registered on each SPM group.

- The **Attribute** column indicates the priority of each WWN. **Prio.** indicates a prioritized WWN. **Non-Prio.** indicates a non-prioritized WWN. For details on how to change the priority, see Setting Priority for Host Bus Adapters.

If one host bus adapter connects to multiple ports, the attribute setting of the host bus adapter is common to all the ports. Therefore, if you specify a host bus adapter as a prioritized WWN or a non-prioritized WWN for one port, the setting is applied to all the other connected ports automatically.

- The **Upper** columns let you specify the upper limit on the I/O rate and the transfer rate for each host bus adapter. Either of the **IOPS** or **MB/s** column in the list is activated depending on the selection of the drop-down list above the list.

The **IOPS** column lets you specify the upper limit for I/O rates. The **MB/s** column lets you specify the upper limit for transfer rates. To specify an upper limit, double-click a cell to display the cursor in the cell.

If you specify a value in either of the **IOPS** or **MB/s** column, the other column is deactivated. You can specify upper limit values for I/O rates and transfer rates all together for different non-prioritized WWNs.



Notes:

- If one host bus adapter connects to multiple ports, the setting of the upper limit value for a non-prioritized WWN is common to all the ports. Therefore, if you specify a upper limit value of non-prioritized WWN for one port, the setting is applied to all the other connected ports automatically.
 - You cannot change the upper limit value of a host bus adapter that has registered on an SPM group. The upper limit value of such a host bus adapter is defined by the setting of the SPM group to which the host bus adapter is registered. For details on setting the upper limit value of an SPM group, see Setting an Upper-Limit Value to HBAs in an SPM Group.
 - The **Upper** columns will not display if an SPM group () or a host bus adapter () is selected in the lower-left tree.
-

- If you select the **All Thresholds** check box and enter a threshold value in the text box, the threshold value will be applied to the entire storage system. If you want to specify the threshold for the I/O rate, select **IOPS** from the drop-down list on the right of the text box. If you want to specify the threshold for the transfer rate, select **MB/s** from the drop-down list. For example, if you specify 128 IOPS in **All Thresholds**, the upper limits on non-prioritized WWNs are disabled when the sum of I/O rates for all the prioritized WWNs is below 128 IOPS.

Even if you use the different type of rate (**IOPS** or **MB/s**) for the threshold as that used for the upper limit values of the non-prioritized WWNs, the threshold control can work for all the WWNs.



Note: In the **WWN** tab, you cannot specify individual thresholds for each host bus adapter.

- If you check the **Delete ports if CHA is removed** checkbox, Server Priority Manager will delete, from SVP, the setting information of Server Priority Manager on ports in channel adapters that have been removed.

When a channel adapter is removed, the port and its settings will disappear on the **Server Priority Manager** window automatically, but that remains in SVP. This may cause that the old setting for Server Priority Manager are applied to a different channel adapter that is installed on the same location newly.

The **Delete ports if CHA is removed** check box is available only when the following Server Priority Manager settings on ports in a removed channel adapter remains on SVP:
 - The setting of prioritized ports or non-prioritized ports.
 - The setting of prioritized WWNs or non-prioritized WWNs.

- The **Apply** button applies the settings in this window to the storage system.
- The **Reset** button restores the last applied settings in the window. When you click this button, all the changes displayed with the blue text in the window are canceled.
- The **Initialize** button changes the settings in this window as explained below, and then applies the resulting settings to the storage system:
 - All the host bus adapters become prioritized WWNs.
 - If the **All Thresholds** checkbox is checked, the check mark disappears.
- The **Close** button closes the **Server Priority Manager** window.

Performance Monitor Operations

This chapter explains the following performance monitor operations:

- ☐ [Overview of Performance Monitor Operations](#)
- ☐ [Monitoring Resources in the Storage System](#)
- ☐ [Monitoring Hard Disk Drives](#)
- ☐ [Monitoring Ports](#)
- ☐ [Monitoring LU Paths](#)
- ☐ [Viewing HBA Information](#)

Overview of Performance Monitor Operations

This section briefly describes how to use Performance Monitor to monitor the storage system.

Start Monitoring

To start monitoring the storage system, you start Performance Monitor and display the **Monitoring Options** window.



Note: This note explains the following statistics to be displayed in tabs of **Performance Management** windows.

- The statistics of LUs that are displayed in the **Port-LUN** tab.
- The statistics of volumes that are displayed in the **LDEV** tab.

In the above tabs, performance statistics of unused volumes are displayed as hyphens (-), if the range of monitored CUs does not match the range of CUs used in the disk storage system or registered as external volumes. In addition, depending on your disk subsystem configuration, the list may display performance statistics for some volumes and not display performance statistics for other volumes.

To correctly display performance statistics, you must specify CUs to be monitored as follows:

- To display performance statistics of a LUSE volume in the **Port-LUN** tab, you must specify all volumes that make up the LUSE volume as the monitoring targets.
 - To display performance statistics of parity group to be displayed in the **LDEV** tab, you must specify all volumes that belong to the parity group as the monitoring target.
-




Figure 5-1 Monitoring Options Window of Performance Monitor

To start monitoring, select **Enable** for the **Current Status** option in **Monitoring Switch**, and specify the **Gathering Interval** option to set the interval for collecting information. Next, select or release the CUs to be monitored in **Monitoring Target** area. You can specify target CUs either by choosing PG numbers, LDKC numbers, CU numbers or the cells in the **CU** table. For the details on how to select CUs, see Monitoring Options Window. After the setting is completed, select the **Apply** button. Performance Monitor starts to obtain statistics about the storage system and saves the statistics at the specified interval.

If the number of CUs to be monitored is 64 or less, you can select a gathering interval between 1 and 15 minutes. The gathering interval you selected here determines the storing period of the statistics in short range, which is up to 15 days. For example, if you specify the gathering interval as 1 minute, the statistics can be stored 1 day, and if you specify it as 15 minutes, the statistics can be stored 15 days. If the number of the CUs to be monitored is 65 or more, you can choose the gathering interval from among 5, 10 or 15 minutes. The gathering interval you selected determines the storing period (between 8 hours and 1 day) of the statistics. In case of 510 CUs, for example, the statistics can be stored for 8 hours if you specify the gathering interval to 5 minutes, and the statistics can be stored for 1 day if the interval is specified to 15 minutes.

Moreover, the resource usage in the storage system can also be stored in a long range up to 3 months. In this case, however, the gathering interval is fixed to 15 minutes regardless of the value set for the **Gathering Interval** option. For details on the relationship of collection interval and storing period of the statistics, see Monitoring Options Window.



Note: The monitoring switch settings in the **Monitoring Options** window work in conjunction with the monitoring switch settings in the **Usage Monitor** windows of TrueCopy, TrueCopy for IBM z/OS, Universal Replicator, and Universal Replicator for IBM z/OS. Therefore, if you use the **Monitoring Options** window to start or stop monitoring, or to change the gathering interval, the monitoring settings of these remote copy functions will also change. Conversely, if the settings in each **Usage Monitor** window are changed, the settings in the **Monitoring Options** window of Performance Monitor will also change automatically.

View the Monitoring Results

To view the results of monitoring, use the **Performance Management** window of Performance Monitor. When you select the **Performance Management** tab, the **Performance Management** window is displayed. In that window, select a desired tab in the tree left side in the window and the icons which information you can view are displayed in the tree. Select the icon with information you want to view and the statistics obtained in the storage system are displayed in the list on the right side of the tree.




Figure 5-2 Performance Management Window

The information you can view by selecting each tab for the tree is explained below.

- **Physical** tab
Enables you to view usage statistics about resources in the storage system.
- **LDEV** tab
Enables you to view statistics about workload on disks.
- **Port-LUN** tab
Enables you to view statistics about traffic at ports and LU paths in the storage system.
- **WWN** tab
Enables you to view statistics about traffic at path between host bus adapters (HBAs) and ports.

Two sliders are displayed on the upper right of the **Performance Management** window. To the left and right of the two sliders, dates and times are displayed. The term between these dates and times is the period in which statistics are stored. If the date and time on the left is 2007/03/27 16:00 and the date and time on the right is 2007/03/28 16:00, you can view usage statistics for the period of 2007/03/27 16:00 to 2007/03/28 16:00.

If you change the dates and times in the **From** and **To** boxes, you can specify the range of statistics that should be displayed in the window. For example, if the window displays statistics for the last one month, you can change the values in the **From** and **To** boxes to display the statistics for the last week only or to display the statistics for the last three days only.

When monitoring ports, LU paths or host bus adapters in the **Port-LUN** tab or the **WWN** tab, you can view the monitoring results in near-real time. To view the monitoring results in real time, select the **Real Time** option which is included in the **Monitoring Term** area of the **Performance Management** window. The information in the window will be updated at the specified gathering interval (every 1 to 15 minutes). You cannot view the monitoring results displayed in the **Physical** tab and the **LDEV** tab in real time.

Starting and Stopping Storage System Monitoring

To monitor the storage system, start Performance Monitor at first and then start obtaining statistics. You can also stop the monitoring from Performance Monitor. Each procedure is explained below.

To start Performance Monitor, take the following steps:

To start Performance Monitor:

1. Log onto Storage Navigator. The Storage Navigator main window is displayed.
2. Click **Go, Performance Manager** and then **Performance Management** on the menu bar of the Storage Navigator main window.

The **Performance Management** window is displayed.

To start monitoring the storage system:

1. Ensure that the Storage Navigator main window is in **Modify** mode.

The Storage Navigator main window is in **Modify** mode if the background color of the  icon is light yellow. If the background color is gray, the window is in **View** mode and you must change it to **Modify** mode by taking the following steps:

- a. Check to see if the background color of the lock icon  is blue.

If the background color is red , you will not be able to switch from **View** mode to **Modify** mode. Wait for a while and click the  button. If the background color turns blue, you can go to the next step.

- b. Click the  icon.

A message appears, asking whether you want to change the mode.

- c. Click **OK** to close the message.

The background color of the icon changes to light yellow (). The mode changes to **Modify** mode. The background color of the lock icon becomes red ().



Note: Even in **View** mode, you can operate the **Performance Management** window, but you cannot change the settings in the **Monitoring Options** window.

2. Start Performance Monitor and select the **Monitoring Options** tab.
The **Monitoring Options** window is displayed.
3. In **Monitoring Switch**, select **Enable** for the **Current Status** option.
4. Use the drop-down lists in **Gathering Interval** to specify the interval for obtaining the statistics.

5. In the **Monitoring Target** area, click the cells representing the CUs to be monitored, and then click **Select**.
6. Select **Apply**.
Performance Monitor starts monitoring the storage system.



Note: When statistics are collected, a heavy workload is likely to be placed on servers. Therefore, the client processing might slow down.

To stop monitoring the storage system:

1. Start Performance Monitor and select the **Monitoring Options** tab.
The **Monitoring Options** window is displayed.
2. In **Monitoring Switch**, select **Disable** for the **Current Status** option.
The **Gathering Interval** drop-down list is grayed out.
3. Select **Apply**.
Performance Monitor stops monitoring the storage system.

Monitoring Resources in the Storage System

This section describes how to view usage statistics about resources in the storage system.

Before taking the following steps, you need to start monitoring in accordance with the procedure described in Monitoring Options Window and obtain the usage statistics.

Viewing Usage Statistics on Parity Groups

Performance Monitor monitors parity groups and lets you view the average and the maximum usage rate in a specified period. Performance Monitor also displays a graph that illustrates changes in parity group usage within that period.

To view usage statistics about parity groups:

1. Ensure that the Performance Management window is displayed.
2. In the tree, click the Physical tab.
3. In the drop-down list above the tree, select the storing period of statistics from **long-range** and **short-range** for display.

For information on storing period of statistics, see Understanding Statistical Storage Ranges.

When you view usage statistics about parity groups, the items displayed in the list by selecting **long-range** and **short-range** are the same.

4. In the tree, select the **Parity Group** folder.

The list on the right displays usage statistics about parity groups. The displayed statistics are the average and the maximum usage rates for the period specified in the **From** and **To** boxes.



Note: The list displays up to a maximum of 4,096 resources at a time. If the number of resources exceeds 4,096, the **Previous** and **Next** buttons allows you to display the remaining resources.

- If you change the date and time in the **From** and **To** boxes and then click the **Apply** button, Performance Monitor recalculates usage rates and updates information in the list. To change the date and time in **From** and **To** boxes, use the arrow buttons and the sliders (for details, see Performance Management Window, Physical Tab).
- If an exclamation mark (!) is displayed before a usage rate, the reported parity group usage rate is likely to be inaccurate, because the configuration has changed (e.g., volumes have been moved by Volume Migration or ShadowImage, or formatted by Virtual LVI or Open Volume Management).

- If a plus mark (+) is displayed before a usage rate 0, such as "+0", the usage rate is not completely 0 but less than 1.
5. To display a graph that illustrates changes in usage rate for parity groups, select the desired parity groups in the list and then click the **Draw** button.



Note: The range of monitoring and the gathering interval affects the time period represented by a graduation on the horizontal axis.

	PG	RAID	Drive Type	Usage(%)		CLPR
				Ave.(Total)	Max.	
	1-1	5(3D+1P)	DKR1B-J047FC	13	14	00:CLPR0
	1-2	5(3D+1P)	DKR1B-J047FC	112	113	01:CLPR1
	1-3	5(3D+1P)	DKR1B-J047FC	13	14	02:CLPR2
	1-4	5(3D+1P)	DKR1B-J047FC	114	115	03:CLPR3

Figure 5-3 Example of Parity Group Usage Rates Displayed in the List



Notes:

- It is possible that the usage rate for a parity group is not equal to the sum of the usage rate for each volume in that parity group (see Viewing Usage Statistics on Volumes in Parity Groups). This is because the Performance Management window rounds off fractions below the decimal point to the nearest whole number when displaying the usage rate for each volume.
- If there is no volume in a parity group, hyphens(-) are displayed in place of performance statistics on a parity group.

The list displays the following items:

- : When the green checkmark icon  is displayed to the left of a parity group, the graph illustrates changes in usage rate for the parity group.
- **PG:** This column indicates IDs of parity groups.
- **RAID:** This column indicates the RAID level (RAID-1, RAID-5, or RAID-6).
- **Drive Type:** This column indicates types of HDDs (hard disk drives).
- **Usage:**
 - The **Ave. (Total)** column displays the average usage rate in the specified period.
 - The **Max.** column displays the maximum usage rate in the specified period.
- **CLPR:** This column indicates numbers and names of cache logical partitions (CLPRs) corresponding to each parity group in the format "CLPR-number:CLPR-name". For details on CLPRs, see the *Virtual Partition Manager User's Guide*.

Viewing Usage Statistics on Volumes in Parity Groups

Performance Monitor monitors volumes in parity groups and lets you view the average and the maximum usage rate in a specified period. Performance Monitor also displays a graph that illustrates changes in volume usage within that period.

To view usage statistics about volumes in a parity group:

1. Ensure that the **Performance Management** window is displayed.
2. In the tree, click the **Physical** tab
3. In the drop-down list above the tree, select the storing period of statistics from **long-range** and **short-range** for display.

For details on the types of storing period of statistics, see Understanding Statistical Storage Ranges.

When you view usage statistics about volumes, some items displayed in the list differ depending on the selection of the storing period: **long-range** or **short-range**. If you want to use Volume Migration to migrate volumes, select **long-range**. This enables you to confirm the estimated usage rate of parity groups after migration.

If you want to examine the ratio of ShadowImage processing and so on to all the processing in the physical drive, select **short-range**.



Note: The usage statistics for a same term might be slightly different between selecting **short-range** and **long-range** because the monitoring precision of these two interval types differs. Especially, differences in read rates and write rates between the interval types are larger than other usage statistics.

4. In the tree, double-click the **Parity Group** folder.
The folder opens and a list of parity groups is displayed below the folder.
5. Select the desired parity group.
The list on the right displays usage statistics about volumes in the specified parity group. The displayed statistics are the average and the maximum usage rates for the period specified in the **From** and **To** boxes.

**Notes:**

- The list displays up to a maximum of 4,096 resources at a time. If the number of resources exceeds 4,096, the **Previous** and **Next** buttons allows you to display the remaining resources.
- If you change the date and time in the **From** and **To** boxes and then click the **Apply** button, Performance Monitor recalculates usage rates and updates information in the list. To change the date and time in **From** and **To** boxes, use the arrow buttons and the sliders (for details, see Performance Management Window, Physical Tab).
- If an exclamation mark (!) is displayed before a usage rate, the reported volume usage rate is likely to be inaccurate, because the configuration has changed (e.g., volumes have been moved by Volume Migration or ShadowImage, or formatted by Virtual LVI or Open Volume Management).
- If a plus mark (+) is displayed before a usage rate 0, such as "+0", the usage rate is not completely 0 but less than 1.

- To display a graph that illustrates changes in usage rate for volumes, select the desired volumes in the list and then click the **Draw** button.



Note: The range of monitoring and the gathering interval affects the time period represented by a graduation on the horizontal axis.

When selecting long-range for storing period of statistics:

	LDEV	Emulation	Usage(%)		Read Rate(%)		Write Rate(%)		Parity Gr.Use[Exp](%)		CLPR
			Ave.(Total)	Max.	Rnd.	Seq.	Rnd.	Seq.	Ave.(Total)	Max.	
	00:00:00	OPEN-V	+0	1	2	3	4	5	11 -> 6	12 -> 7	00:CLPR0
	00:00:01	OPEN-V	!1	!2	3	4	5	6	11 -> 7	12 -> 8	00:CLPR0
	00:00:02	OPEN-V	2	3	4	5	6	7	11 -> 8	12 -> 9	00:CLPR0
	00:00:03	OPEN-V	!3	!4	5	6	7	8	11 -> 9	12 -> 10	00:CLPR0
	01:00:04	OPEN-V	4	5	6	7	8	9	11 -> 10	12 -> 11	00:CLPR0
	01:00:05	OPEN-V	!5	!6	7	8	9	10	11 -> 11	12 -> 12	00:CLPR0

When selecting short-range for storing period of statistics:

	LDEV	Emulation	Usage(%)		Read Rate(%)		Write Rate(%)		ShadowImage(%)		CLPR
			Ave.(Total)	Max.	Rnd.	Seq.	Rnd.	Seq.	Ave.(Total)	Max.	
	00:00:00	OPEN-V	10	11	12	13	14	15	16	17	00:CLPR0
	00:00:01	OPEN-V	11	12	13	14	15	16	17	18	00:CLPR0
	00:00:02	OPEN-V	12	13	14	15	16	17	18	19	00:CLPR0
	00:00:03	OPEN-V	13	14	15	16	17	18	19	20	00:CLPR0
	01:00:04	OPEN-V	14	15	16	17	18	19	20	21	00:CLPR0
	01:00:05	OPEN-V	15	16	17	18	19	20	21	22	00:CLPR0

Figure 5-4 Examples of Volume Usage Rates Displayed



Note: It is possible that the sum of the usage rate for each volume in a parity group is not equal to the usage rate for that parity group (see Viewing Usage Statistics on Volumes in Parity Groups). This is because the Performance Management window rounds off fractions below the decimal point to the nearest whole number when displaying the usage rate for each volume.

The list displays the following items:

- : When the green checkmark icon  is displayed on the left of a volume, the graph illustrates changes in usage rate for the volume.
- **LDEV:** This column indicates volumes (LDEVs), in this format: LDKC:CU:LDEV.
- **Emulation:** This column indicates device emulation types.
- **Usage:**
 - The **Ave. (Total)** column displays the average usage rate in the specified period.
 - The **Max.** column displays the maximum usage rate in the specified period.
- **Read Rate:**
 - The **Rnd.** column indicates random read rates. A random read rate is the ratio of random read requests to read and write requests.
 - The **Seq.** column indicates sequential read rates. A sequential read rate is the ratio of sequential read requests to read and write requests.
- **Write Rate:**
 - The **Rnd.** column indicates random write rates. A random write rate is the ratio of random write requests to read and write requests.
 - The **Seq.** column indicates sequential write rates. A sequential write rate is the ratio of sequential write requests to read and write requests.
- **Parity Gr. Use[Exp]:** This item is displayed only when you select **long-range** for the storing period of statistics. **Parity Gr. Use[Exp]** assumes the volume might be migrated out of the parity group (or uninstalled) and indicates expected (estimated) average and maximum usage rates of the parity group. The **Ave. (Total)** column indicates an estimated change in the average usage rate. The **Max.** column indicates an estimated change in the maximum usage rate.

For example, if the **Ave. (Total)** box for the volume 0:01 displays "20 -> 18", the average usage rate of the parity group that the volume belongs to is 20 percent. If the volume were migrated out of the parity group, the average usage rate of that group is expected to drop to 18 percent.

- **ShadowImage:** This item is displayed only when you select **short-range** for the storing period of statistics. **ShadowImage** indicates the percentage of the processing of the following programs to all the processing of the physical drives, for each volume. This value is found by dividing access time to physical drives by the following programs by all the access time to physical drives.
 - ShadowImage for IBM z/OS
 - ShadowImage
 - Compatible FlashCopy
 - Compatible FlashCopy V2
 - Volume Migration
 - Copy-on-Write Snapshot

The **Ave. (Total)** column displays the average percentage of processing of the above programs in the specified period.

The **Max.** column displays the maximum percentage of processing of the above programs in the specified period.

For details on the above programs, see the respective user's guides.

- **CLPR:** This column indicates numbers and names of CLPRs corresponding to each parity group which the volume belongs to, in the format "CLPR-number:CLPR-name". For details on CLPRs, see the *Virtual Partition Manager User's Guide*.

Viewing Usage Statistics on External Volume Groups

Performance Monitor monitors external volume groups and lets you view the usage statistics of external volume groups in a specified period. Performance Monitor also displays a graph that illustrates changes in the usage statistics of external volume groups within that period.



Note: You can view the usage statistics about external volume groups only when you select short-range for the storing period of statistics.

To view usage statistics about external volume groups:

1. Ensure that the **Performance Management** window is displayed.
2. In the tree, click the **Physical** tab.
3. In the drop-down list above the tree, select **short-range** for the storing period of statistics.

If you select **long-range**, no statistics appears in the list.

4. In the tree, select the **External Group** folder.

The list on the right displays usage statistics about external volume groups. The displayed statistics are the usage statistics for the period specified in the **From** and **To** boxes.



Note:

- The list displays up to a maximum of 4,096 resources at a time. If the number of resources exceeds 4,096, the **Previous** and **Next** buttons allows you to display the remaining resources.
- If you change the date and time in the **From** and **To** boxes and then click the **Apply** button, Performance Monitor recalculates usage rates and updates information in the list. To change the date and time in **From** and **To** boxes, use the arrow buttons and the sliders (for details, see Performance Management Window, Physical Tab).

5. To display a graph that illustrates changes in usage statistics for external volume groups, select the desired external volume groups in the list and then click the **Draw** button.



Note: The range of monitoring and the gathering interval affects the time period represented by a graduation on the horizontal axis.

	ExG	Response Time(ms)	Trans.(MB/sec)	CLPR
	E1-1	2000	4.883	02:CLPR2
	E1-2	2001	4.884	03:CLPR3
	E1-3	2002	4.885	00:CLPR0
	E1-4	2003	4.886	01:CLPR1

Figure 5-5 Example of Usage Statistics of an External Volume Group Usages Displayed in the List

The list displays the following items:

- : When the green checkmark icon  is displayed on the left of an external volume group, the graph illustrates changes in usage statistics for the external volume group.
- **ExG:** This column indicates IDs of external volume groups. The letter "E" at the beginning of IDs indicates the group is an external volume group.
- **Response Time:** This column indicates the time for replying from an external volume group when I/O accesses are made from USP V/VM storage system to the external volume group. The unit is milliseconds. The average response time in the period specified at **Monitoring Term** is displayed.
- **Trans.:** This column indicates the sizes of data transferred between the storage system and the external volume group in one second.
- **CLPR:** This column indicates numbers and names of CLPRs corresponding to each external volume group in the format "CLPR-number:CLPR-name". For details on CLPRs, see the *Virtual Partition Manager User's Guide*.

Viewing Usage Statistics on External Volumes in External Volume Groups

Performance Monitor monitors external volumes in external volume groups and lets you view the usage statistics of external volumes. Performance Monitor also displays a graph that illustrates changes in the usage statistics of external volumes within that period.



Note: You can view the usage statistics about external volumes only when you select **short-range** for the storing period of statistics.

To view usage statistics about external volumes in an external volume group:

1. Ensure that the **Performance Management** window is displayed.
2. In the tree, click the **Physical** tab.
3. In the drop-down list above the tree, select **short-range** for the storing period of statistics.

If you select **long-range**, no statistics appears in the list.

4. In the tree, select the **External Group** folder.

The folder opens and a list of external volume groups is displayed below the folder.

5. Select the desired external volume group.

The list on the right displays usage statistics about external volumes in the specified external volume group. The displayed statistics are the average and the maximum usage rates for the period specified in the **From** and **To** boxes.



Notes:

- The list displays up to a maximum of 4,096 resources at a time. If the number of resources exceeds 4,096, the **Previous** and **Next** buttons allows you to display the remaining resources.
- If you change the date and time in the From and To boxes and then click the Apply button, Performance Monitor recalculates usage rates and updates information in the list. To change the date and time in From and To boxes, use the arrow buttons and the sliders (for details, see Performance Management Window, Physical Tab).

-
6. To display a graph that illustrates changes in usage rate for external volumes, select the desired external volumes in the list and then click the **Draw** button.



Note: The range of monitoring and the gathering interval affects the time period represented by a graduation on the horizontal axis.

	ExLDEV	Emulation	Response Time(ms)	Trans.(MB/sec)	CLPR
	00:08:12 #	OPEN-V	2000	2.930	01:CLPR1
	00:08:13 #	OPEN-V	2001	2.931	01:CLPR1
	00:08:14 #	OPEN-V	2002	2.932	01:CLPR1
	00:08:15 #	OPEN-V	2003	2.933	01:CLPR1

Figure 5-6 Example of External Volume Usage Rates Displayed

The list displays the following items:

- : When the green checkmark icon  is displayed on the left of an external volume, the graph illustrates changes in usage rate for the external volume.
- **ExLDEV:** This column indicates external volumes, in this format: LDKC:CU:LDEV. A number ending in # indicates the volume is an external volume.
- **Emulation:** This column indicates device emulation types.
- **Response Time:** This column indicates the time for replying from an external volume when I/O accesses are made from USP V/VM storage system to the external volume. The unit is milliseconds. The average response time in the period specified at **Monitoring Term** is displayed.
- **Trans.:** This column indicates the sizes of data transferred between the storage system and the external volume in one second.
- **CLPR:** This column indicates numbers and names of CLPRs corresponding to each external volume group which the external volume belongs to, in the format "CLPR-number:CLPR-name". For details on CLPRs, see the *Virtual Partition Manager User's Guide*.

Viewing Usage Statistics on Channel Processors

Performance Monitor monitors channel processors in each channel adapter and lets you view the average and the maximum usage rate in a specified period. Performance Monitor also displays a graph that illustrates changes in channel processor usage within that period.

To view usage statistics about channel processors:

1. Ensure that the **Performance Management** window is displayed.
2. In the tree, click the **Physical** tab.
3. In the drop-down list above the tree, select the storing period of statistics from **long-range** and **short-range** for display.

For details on the types of storing period of statistics, see Understanding Statistical Storage Ranges.

When you view usage statistics about channel processors, the items displayed in the list by selecting **long-range** and **short-range** are the same.

4. In the tree, do one of the following:
 - If you want to view usage statistics about all the channel processors in your storage system, select the **Client-Host Interface Processor (CHIP)** folder.
 - If you want to view usage statistics about channel processors in a channel adapter, double-click the **Client-Host Interface Processor (CHIP)** folder and then select the desired channel adapter.

The list on the right displays usage statistics about the channel processors. The displayed statistics are the average and the maximum usage rates for the period specified in the **From** and **To** boxes.



Note: If you change the date and time in the **From** and **To** boxes and then click the **Apply** button, Performance Monitor recalculates usage rates and updates information in the list. To change the date and time in **From** and **To** boxes, use the arrow buttons and the sliders (for details, see Performance Management Window, Physical Tab).

5. To display a graph that illustrates changes in usage rate for channel processors, select the desired channel processors in the list and then click the **Draw** button.



Note: The range of monitoring and the gathering interval affects the time period represented by a graduation on the horizontal axis.

	ID	Usage(%)	
		Ave.(Total)	Max.
	CHP00-1P	20	82
	CHP01-1P	20	82
	CHP02-1P	20	82
	CHP03-1P	20	82

Figure 5-7 Example of Channel Processors Usage Rates Displayed

The list displays the following items:

- : When the green checkmark icon  is displayed on the left of a channel processor, the graph illustrates changes in usage rate for the channel processor.
- **ID**: This column displays ID numbers for channel processors.
- **Usage**:
 - The **Ave. (Total)** column displays the average usage rate in the specified period.
 - The **Max.** column displays the maximum usage rate in the specified period.

Viewing Usage Statistics on Disk Processors

Performance Monitor monitors disk processors (DKPs) and lets you view the average and the maximum usage rate in a specified period. Performance Monitor also displays a graph that illustrates changes in disk processor usage within that period.

To view usage statistics about disk processors:

1. Ensure that the **Performance Management** window is displayed.
2. In the tree, click the **Physical** tab.
3. In the drop-down list above the tree, select the storing period of statistics from **long-range** and **short-range** for display.
For details on the types of storing period of statistics, see Understanding Statistical Storage Ranges.
When you view usage statistics about disk processors, the items displayed in the list by selecting **long-range** and **short-range** are the same.
4. In the tree, double-click the **ACP** folder.
5. Click **DKP** from below the **ACP** folder.

The list on the right displays usage statistics about disk processors. The displayed statistics are the average and the maximum usage rates for the period specified in the **From** and **To** boxes.



Note: If you change the date and time in the **From** and **To** boxes and then click the **Apply** button, Performance Monitor recalculates usage rates and updates information in the list. To change the date and time in **From** and **To** boxes, use the arrow buttons and the sliders (for details, see Performance Management Window, Physical Tab).

- If you **want** to display a graph that illustrates changes in usage rate for disk processors, select the desired disk processors in the list and then click the **Draw** button.



Note: The range of monitoring and the gathering interval affects the time period represented by a graduation on the horizontal axis.

	ID	Usage(%)	
		Ave.(Total)	Max.
	DKP80-1B	30	83
	DKP81-1B	30	83
	DKP82-1B	30	83
	DKP83-1B	30	83
	DKP90-1C	30	83
	DKP91-1C	30	83
	DKP92-1C	30	83
	DKP93-1C	30	83
	DKPA0-1D	30	83

Figure 5-8 Example of Disk Processor Usage Rates Displayed

The list displays the following items:

- : When the green checkmark icon  is displayed on the left of a disk processor, the graph illustrates changes in usage rate for the disk processor.
- ID:** This column displays ID numbers for disk processors.
- Usage:**
 - The **Ave. (Total)** column displays the average usage rate in the specified period.
 - The **Max.** column displays the maximum usage rate in the specified period.

Viewing Usage Statistics on Data Recovery and Reconstruction Processors

Performance Monitor tracks data recovery and reconstruction processors (DRRs) and lets you view the average and the maximum usage rate in a specified period. Performance Manager also displays a graph illustrating changes in DRR usage within that period.

To view usage statistics about DRRs:

1. Ensure that the **Performance Management** window is displayed.
2. In the tree, click the **Physical** tab.
3. In the drop-down list above the tree, select the storing period of statistics from **long-range** and **short-range** for display.

For details on the types of storing period of statistics, see Understanding Statistical Storage Ranges.

When you view usage statistics about DRRs, the items displayed in the list by selecting **long-range** and **short-range** are the same.

4. In the tree, double-click the **ACP** folder.
5. Click **DRR** from below the **ACP** folder.

The list on the right displays usage statistics about DRRs. The displayed statistics are the average and the maximum usage rates for the period specified in the **From** and **To** boxes.



Note: If you change the date and time in the **From** and **To** boxes and then click the **Apply** button, Performance Monitor recalculates usage rates and updates information in the list. To change the date and time in **From** and **To** boxes, use the arrow buttons and the sliders (for details, see Performance Management Window, Physical Tab).

6. If you want to display a graph that illustrates changes in usage rate for DRRs, select the desired DRRs in the list and then click the **Draw** button.



Note: The range of monitoring and the gathering interval affects the time period represented by a graduation on the horizontal axis.

	ID	Usage(%)	
		Ave.(Total)	Max.
	DRR80-1B	40	84
	DRR81-1B	40	84
	DRR82-1B	40	84
	DRR83-1B	40	84
	DRR90-1C	40	84
	DRR91-1C	40	84
	DRR92-1C	40	84
	DRR93-1C	40	84
	DRRA0-1D	40	84

Figure 5-9 Example of DRR Usage Rates Displayed in the List

The list displays the following items:

- : When the green checkmark icon  is displayed on the left of a DRR, the graph illustrates changes in usage rate for the DRR.
- **ID:** This column displays ID numbers for DRRs.
- **Usage:**
 - The **Ave. (Total)** column displays the average usage rate in the specified period.
 - The **Max.** column displays the maximum usage rate in the specified period.

Viewing Write Pending and Cache Memory Usage Statistics

Performance Monitor lets you view write pending rates. A write pending rate indicates the ratio of write-pending data to the cache memory capacity. Performance Monitor displays the average and the maximum write pending rate in a specified period. When you select **short-range** for storing period of statistics, Performance Monitor also displays the average and the maximum usage statistics about the cache memory in a specified period. In addition, Performance Monitor can display a graph that illustrates changes in write pending rate or usage statistics of the cache memory within that period.

To view the write pending rate or usage statistics of the cache memory:

1. Ensure that the **Performance Management** window is displayed.
2. In the tree, click the **Physical** tab.
3. In the drop-down list above the tree, select the storing period of statistics from **long-range** and **short-range** for display.

For details on the types of storing period of statistics, see Understanding Statistical Storage Ranges.

The items displayed in the list are a little different depending on the selection of the storing period: **long-range** or **short-range**. The write pending rates are displayed by both selection, but usage statistics of the cache memory are displayed only when you select **short-range**.

4. In the tree, select the **Cache** folder.

The list on the right displays the write pending rates. If you select **short-range** for the storing period of statistics, the list also displays usage statistics about the cache memory.

Average and maximum usage rates for the period specified in the **From** and **To** boxes are displayed.



Note: If you change the date and time in the **From** and **To** boxes and then click the **Apply** button, Performance Monitor recalculates the write pending rates and usage statistics about the cache memory, and updates information in the list. To change the date and time in **From** and **To** boxes, use the arrow buttons and the sliders (for details, see Performance Management Window, Physical Tab).

5. If you want to display a graph that illustrates changes in the write pending rate or in usage statistics about the cache memory, select the row of the write pending rate in the list and then click the **Draw** button.



Note: The range of monitoring and the gathering interval affects the time period represented by a graduation on the horizontal axis.

6. When you select **short-range** for the storing period of statistics, you can select the item to be illustrated in the graph from the drop-down list at the upper right on the graph.

Select **Write Pending** to view the write pending rate graph and select **Usage** to view that of the usage statistics for cache memory. The graph will be updated without reselecting the **Draw** button. You can also select the item to be displayed before selecting the **Draw** button at first.




Figure 5-10 Selecting an Item to be Displayed in the Graph (Short-range)

Examples of write pending rate displayed in the lists are shown below:

When selecting long-range for storing period of statistics:

	CLPR	Write Pending(%)	
		Ave.(Total)	Moment Max.
	00:CLPR0	1	2
	01:CLPR1	4	5
	02:CLPR2	7	8
	03:CLPR3	10	11

When selecting short-range for storing period of statistics:

	CLPR	Usage(%)		Write Pending(%)	
		Ave.(Total)	Max.	Ave.(Total)	Moment Max.
	00:CLPR0	1	2	3	4
	01:CLPR1	4	5	6	7
	02:CLPR2	7	8	9	10
	03:CLPR3	10	11	12	13

Figure 5-11 Example of Write Pending Rate and Cache Usage Rate Displayed

The list displays the following items:

- : When the green checkmark icon  is displayed on the left of the write pending rate, the graph illustrates changes in the write pending rate and usage statistics about the cache memory.
- CLPR:** This column indicates numbers and names of cache logical partitions (CLPRs) in the format "CLPR-number:CLPR-name". For details on CLPRs, see the *Virtual Partition Manager User's Guide*.
- Usage:** This item is displayed only when you select **short-range** for the storing period of statistics.
 - The **Ave. (Total)** column displays the average usage rate of the cache in the specified period.
 - The **Max.** column displays the maximum usage rate of the cache in the specified period.
- Write Pending:**
 - The **Ave. (Total)** column displays the average write pending rate for the specified period.
 - The **Moment Max.** column displays the maximum write pending rate for the specified period.

Viewing Usage Statistics on Access Paths

The channel adapters (CHAs) and the disk adapters (DKAs) transfer data to the cache switch (CSW) and the shared memory (SM) when I/O requests are issued from hosts to the storage system. In some configurations, DKAs are called array control processors (ACPs). Also, the cache switch transfers data to the cache memory.

Performance Monitor audits these data transfer paths and lets you view the average and the maximum usage rate for the paths in a specified period. Performance Monitor also displays a graph illustrating changes in path usage within that period.

To view usage statistics about paths:

1. Ensure that the **Performance Management** window is displayed.
2. In the tree, click the **Physical** tab.
3. In the drop-down list above the tree, select the storing period of statistics from **long-range** and **short-range** for display.

For details on the types of storing period of statistics, see Understanding Statistical Storage Ranges.

When you view usage statistics about paths, the items displayed in the list by selecting **long-range** and **short-range** are the same.

4. In the tree, double-click the **Access Path Usage** folder.
5. Do one of the following:
 - To check usage statistics about paths between adapters (CHAs and DKAs) and the cache switch, select **Adapter-CSW**, from below the **Access Path Usage** folder.
 - To check usage statistics about paths between adapters (CHAs and DKAs) and the shared memory, select **Adapter-SM**, from below the **Access Path Usage** folder.
 - To check usage statistics about paths between cache switches and the cache memory, select **CSW-Cache**, from below the **Access Path Usage** folder.

The list on the right in Figure 5-12 displays the average and maximum usage rate for the specified paths. The displayed statistics are the average and the maximum usage rates for the period specified in the **From** and **To** boxes.



Note: If you change the date and time in the **From** and **To** boxes and then click the **Apply** button, Performance Monitor recalculates usage rates and updates information in the list. To change the date and time in **From** and **To** boxes, use the arrow buttons and the sliders (for details, see Performance Management Window, Physical Tab).

6. If you want to display a graph that illustrates changes in usage statistics about paths, select the desired paths in the list and then click the **Draw** button.



Note: The range of monitoring and the gathering interval affects the time period represented by a graduation on the horizontal axis.

Paths between adapters and cache switches:

	Adapter	CSW	Usage(%)	
			Ave.(Total)	Max.
	CHA-1P	CSW-1N	60	86
	CHA-1P	CSW-1A	60	86
	CHA-1Q	CSW-1N	60	86
	CHA-1Q	CSW-1A	60	86
	CHA-1R	CSW-1N	60	86
	CHA-1R	CSW-1A	60	86
	CHA-1S	CSW-1N	60	86
	CHA-1S	CSW-1A	60	86

Paths between adapters and shared memories:

	Adapter	SM	Usage(%)	
			Ave.(Total)	Max.
	CHA-1P	Side-A	50	85
	CHA-1P	Side-B	50	85
	CHA-1Q	Side-A	50	85
	CHA-1Q	Side-B	50	85
	CHA-1R	Side-A	50	85
	CHA-1R	Side-B	50	85
	CHA-1S	Side-A	50	85
	CHA-1S	Side-B	50	85

Paths between cache switches and cache memory:

	CSW	Cache	Usage(%)	
			Ave.(Total)	Max.
	CSW-1N	CACHE-1T	70	87
	CSW-1A	CACHE-1T	70	87
	CSW-2M	CACHE-1T	70	87
	CSW-2Z	CACHE-1T	70	87
	CSW-1N	CACHE-1U	70	87
	CSW-1A	CACHE-1U	70	87
	CSW-2M	CACHE-1U	70	87
	CSW-2Z	CACHE-1U	70	87

Figure 5-12 Examples of Usage Statistics Displayed in the List

- : When the green checkmark icon  is displayed on the left of the access path, the graph illustrates changes in usage rate for the access path.
- **Adapter:** This column indicates adapters.
- **CSW:** This column indicates cache switches.
- **SM:** This column indicates shared memories.

- **Cache:** This column indicates the cache memories.
- **Usage:**
 - The **Ave. (Total)** column displays the average usage rate for the specified period.
 - The **Max.** column displays the maximum usage rate for the specified period.

Monitoring Hard Disk Drives

The **LDEV** tab of the **Performance Management** window lets you check workloads on physical hard disk drives (parity groups) or on volumes. The **LDEV** tab displays the *I/O rate* and the *transfer rate*. The I/O rate indicates the number of disk I/Os per second. The transfer rate indicates the size of data transferred to the disk in one second.

In addition, the **LDEV** tab displays the *read hit ratio* and the *write hit ratio*.

- For a read I/O, when the requested data is already in cache, the operation is classified as a *read hit*.
- For a write I/O, when the requested data is already in cache, the operation is classified as a *write hit*.

This section describes how to view the statistics about disk access performance. Before taking the following steps, you need to start monitoring in accordance with the procedure described in Starting and Stopping Storage System Monitoring and obtain the usage statistics.

Viewing I/O Rates for Disks

Performance Monitor monitors hard disk drives and measures I/O rates (that is, the number of disk I/Os per second).

To view I/O rates:

1. Ensure that the **Performance Management** window is displayed.
2. Click the **LDEV** tab.
The tree displays a list of parity groups and external volume groups.
3. Select **IOPS** from the drop-down list on the right side of the window.
4. In the tree, do one of the following:
 - If you want to view the I/O rate for each parity group and external volume group, select the **Subsystem** folder or a **Box** folder.
The list on the right displays the I/O rate for each group.



Note: If you select the **Subsystem** folder, the list displays all parity groups and external volume groups. To narrow the number of groups to be displayed in the list, select a **Box** folder. For example, if you select the **Box 1** folder, the list displays only the parity groups whose IDs start with "1-".

- If you want to view the I/O rate for each volume, select the parity group or external volume group that contains the volumes. The list on the right displays the I/O rate for each volume in the selected group.

The displayed statistics are the average and the maximum I/O rates for the period specified in the **From** and **To** boxes.



Notes:

- The list displays up to a maximum of 4,096 resources at a time. If the number of resources exceeds 4,096, the **Previous** and **Next** buttons allows you to display the remaining resources.
 - If you change the date and time in the **From** and **To** boxes and then click the **Apply** button, Performance Monitor recalculates usage rates and updates information in the list. To change the date and time in **From** and **To** boxes, use the arrow buttons and the sliders (for details, see Performance Management Window, Physical Tab).
-

5. To display a graph to find out how the I/O rate has been changed, take the following steps:

- a. In the list, select the parity group(s), the external volume groups, or the volume(s) that you want.
- b. Use the drop-down list at the right side of the list to select the type of information that you want to view in the graph. For details on the information that can be displayed, see Table 4-1.
- c. Click the **Draw** button.

A graph appears below the list. The horizontal axis indicates the time.



Caution: If the graph does not display changes in the I/O rate (for example, if the line in the graph runs vertically), it is recommended that you change the value in the **Chart Y Axis Rate** drop-down list. For example, if the largest value in the list is 200 and the value in **Chart Y Axis Rate** is 100, you should select a value larger than 200 from **Chart Y Axis Rate**.

- d. To view more detailed information in the graph, use the drop-down list to the right of the list to select the type of information that you want. Next, select the **Detail** check box at the lower right of the list and then click the **Draw** button. The detailed graph contents change as described in Table 4-2.



Caution: If more than one parity group or volume is selected in the list, you cannot select the **Detail** check box to view detailed information.

I/O rate for parity groups and external volume groups (When the Subsystem folder is selected):

	PG	IO Rate(IOPS)	Read(IOPS)	Write(IOPS)	Read Hit(%)	Write Hit(%)	Back Trans.(cou...	Response Time(ms)	CLPR
	1-1	1	2	3	4	5	6	5.008	00:CLPR0
	1-2	2	3	4	5	6	7	5.009	01:CLPR1
	1-3	3	4	5	6	7	8	5.010	02:CLPR2
	1-4	4	5	6	7	8	9	5.011	03:CLPR3
	2-1	5	6	7	8	9	10	5.012	00:CLPR0
	2-2	6	7	8	9	10	11	5.013	01:CLPR1

I/O rate for parity groups or external volume groups (When the Box 1 folder is selected):

	PG	IO Rate(IOPS)	Read(IOPS)	Write(IOPS)	Read Hit(%)	Write Hit(%)	Back Trans.(cou...	Response Time(ms)	CLPR
	1-1	1	2	3	4	5	6	5.008	00:CLPR0
	1-2	2	3	4	5	6	7	5.009	01:CLPR1
	1-3	3	4	5	6	7	8	5.010	02:CLPR2
	1-4	4	5	6	7	8	9	5.011	03:CLPR3

I/O rate for volumes:

	LDEV	Emulation	IO Rate(IOPS)	Read(IOPS)	Write(IOPS)	Read Hit(%)	Write Hit(%)	Back Trans.(count/sec)	Response Time(ms)	CLPR
	00:00:00	OPEN-V	1	2	3	4	5	6	5.008	00:CLPR0
	00:00:01	OPEN-V	2	3	4	5	6	7	5.009	00:CLPR0
	00:00:02	OPEN-V	3	4	5	6	7	8	5.010	00:CLPR0
	00:00:03	OPEN-V	4	5	6	7	8	9	5.011	00:CLPR0
	01:00:04	OPEN-V	5	6	7	8	9	10	5.012	00:CLPR0
	01:00:05	OPEN-V	6	7	8	9	10	11	5.013	00:CLPR0
	01:00:06	OPEN-V	7	8	9	10	11	12	5.014	00:CLPR0

Figure 5-13 I/O Rates for Disks

Note: It is possible that the I/O rate for a parity group or an external volume group is not equal to the sum of the I/O rate for each volume in that group. This is because the Performance Management window omits fractions below the decimal point when displaying the I/O rate for each volume.

Note: If there is no volume in a parity group, hyphens (-) are displayed in place of performance statistics on a parity group.

The list displays the following:

- : When the green checkmark icon is displayed on the left of an item, the graph displays workload statistics about the parity group, the external volume group, or the volume.
- PG:** Indicates the parity group ID or the external volume group ID. If the ID starts with the letter "E", the group is an external volume group. If the ID starts with the letter "V" or "X", the group is a V-VOL group.
- LDEV:** Indicates the volume ID. If the ID ends with the pound or gate symbol (#), the volume is an external volume. If the ID ends with the symbol "V" or "X", the volume is a virtual volume.

- **Emulation:** Indicates the emulation type.
- **IO Rate:** Indicates the number of I/O requests to the parity group, the external volume group, or the volume per second.
- **Read:** Indicates the number of read requests to the parity group, the external volume group, or the volume per second.
- **Write:** Indicates the number of write requests to the parity group, the external volume group, or the volume per second.
- **Read Hit:** Indicates the read hit ratio.
- **Write Hit:** Indicates the write hit ratio.
- **Back Trans.:** Indicates the number of data transfers per second between the parity group (or the external volume group, or the volume) and the cache memory.
- **Response Time:** This column indicates the time for replying from the parity group, the external volume group, or the volume when I/O accesses are made from the host to them. The unit is milliseconds. The average response time in the period specified at **Monitoring Term** is displayed.
- **CLPR:** This column indicates numbers and names of CLPRs corresponding to each parity group, external volume group, or volume in the format "CLPR-number:CLPR-name". For details on CLPRs, see the *Virtual Partition Manager User's Guide*.

Viewing Transfer Rates for Disks

Performance Monitor audits hard disk drives and measures transfer rates (that is, the size of data transferred in one second).

To view transfer rates:

1. Ensure that the **Performance Management** window is displayed.
2. Click the **LDEV** tab.
The tree displays a list of parity groups and external volume groups.
3. Select **MB/s** from the drop-down list on the right side of the window.
4. In the tree, do one of the following:
 - To view the transfer rate for each parity group and external volume group, select the **Subsystem** folder or a **Box** folder.
The list on the right displays the transfer rate for each group.

**Notes:**

- If you select the **Subsystem** folder, the list displays all parity groups and external volume groups. To narrow the number of groups to be displayed in the list, select a **Box** folder. For example, if you select the **Box 1** folder, the list displays only the parity groups whose IDs start with "1-".
- To view the transfer rate for each volume, select the parity group or external volume group that contains the volumes.
The list on the right displays the transfer rate for each volume in the selected group.

The displayed statistics are the average and the maximum transfer rates for the period specified in the **From** and **To** boxes.

**Notes:**

- The list displays up to a maximum of 4,096 resources at a time. If the number of resources exceeds 4,096, the **Previous** and **Next** buttons allows you to display the remaining resources.
- If you change the date and time in the **From** and **To** boxes and then click the **Apply** button, Performance Monitor recalculates usage rates and updates information in the list. To change the date and time in **From** and **To** boxes, use the arrow buttons and the sliders (for details, see LDEV Tab of the Performance Monitor Window).

-
5. If you want to display a graph to find out how the transfer rate has been changed, take the following steps:
 - a. In the list, select the parity group(s), the external volume groups, or the volume(s) that you want.
 - b. Use the drop-down list at the right side of the list to select the type of information that you want to view in the graph. For details on the information that can be displayed, see Table 4-1.
 - c. Click the **Draw** button.A graph appears below the list. The horizontal axis indicates the time.



Caution: If the graph does not display changes in the transfer rate (for example, if the line in the graph runs vertically), it is recommended that you change the value in the **Chart Y Axis Rate** drop-down list. For example, if the largest value in the list is 200 and the value in **Chart Y Axis Rate** is 100, you should select a value larger than 200 from **Chart Y Axis Rate**.

- d. If you want to view more detailed information in the graph, use the drop-down list to the right of the list to select the type of information that you want. Next, select the **Detail** check box at the lower right of the list and then click the **Draw** button. The detailed graph contents change as described in Table 4-2.



Caution: If more than one parity group or volume is selected in the list, you cannot select the **Detail** check box to view detailed information.

Transfer rate for parity groups and external volume groups (When the Subsystem folder is selected):

	PG	Trans.(MB/s)	Read Hit(%)	Write Hit(%)	Back Trans.(cou...	Response Time(ms)	CLPR
	1-1	1.960	4	5	6	5.008	00:CLPR0
	1-2	1.961	5	6	7	5.009	01:CLPR1
	1-3	1.962	6	7	8	5.010	02:CLPR2
	1-4	1.963	7	8	9	5.011	03:CLPR3
	2-1	1.964	8	9	10	5.012	00:CLPR0
	2-2	1.965	9	10	11	5.013	01:CLPR1

Transfer rate for parity groups or external volume groups (When the Box 1 folder is selected):

	PG	Trans.(MB/s)	Read Hit(%)	Write Hit(%)	Back Trans.(count/sec)	Response Time(ms)	CLPR
	1-1	1.960	4	5	6	5.008	00:CLPR0
	1-2	1.961	5	6	7	5.009	01:CLPR1
	1-3	1.962	6	7	8	5.010	02:CLPR2
	1-4	1.963	7	8	9	5.011	03:CLPR3

Transfer rate for volumes:

	LDEV	Emulation	Trans.(MB/s)	Read Hit(%)	Write Hit(%)	Back Trans.(count/sec)	Response Time(ms)	CLPR
	00:00:00	OPEN-V	1.960	4	5	6	5.008	00:CLPR0
	00:00:01	OPEN-V	1.961	5	6	7	5.009	00:CLPR0
	00:00:02	OPEN-V	1.962	6	7	8	5.010	00:CLPR0
	00:00:03	OPEN-V	1.963	7	8	9	5.011	00:CLPR0
	01:00:04	OPEN-V	1.964	8	9	10	5.012	00:CLPR0
	01:00:05	OPEN-V	1.965	9	10	11	5.013	00:CLPR0
	01:00:06	OPEN-V	1.966	10	11	12	5.014	00:CLPR0

Figure 5-14 Transfer Rates for Disks



Notes:

- It is possible that the transfer rate for a parity group or an external volume group is not equal to the sum of the transfer rate for each volume in that group. This is because the Performance Management window omits fractions below the decimal point when displaying the transfer rate for each volume.
- If there is no volume in a parity group, hyphens (-) are displayed in place of performance statistics on a parity group.

The list displays the following:

- : When the green checkmark icon  is displayed on the left of an item, the graph displays workload statistics about the parity group, the external volume group, or the volume.
- **PG**: Indicates the parity group ID or the external volume group ID. If the ID starts with the letter "E", the group is an external volume group.
- **LDEV**: Indicates the volume ID. If the ID ends with the symbol "#", the volume is a an external volume.
- **Emulation**: Indicates the emulation type.
- **Trans.**: Indicates the size (in megabytes) of data transferred to the parity group, the external volume group, or the volume per second.
- **Read Hit**: Indicates the read hit ratio.
- **Write Hit**: Indicates the write hit ratio.
- **Back Trans.**: Indicates the number of data transfers per second between the parity group (or the external volume group, or the volume) and the cache memory.
- **Response Time**: This column indicates the time for replying from the parity group, the external volume group, or the volume when I/O accesses are made from the host to them. The unit is milliseconds. The average response time in the period specified at **Monitoring Term** is displayed.



Note: This column displays a hyphen (-) if the I/O rate is 0 (zero).

- **CLPR**: This column indicates numbers and names of CLPRs corresponding to each parity group, external volume group, or volume in the format "CLPR-number:CLPR-name". For details on CLPRs, see the *Virtual Partition Manager User's Guide*.

Monitoring Ports

Performance Monitor monitors ports on the storage system and measures I/O rates (that is, the number of I/Os per second) and transfer rates (that is, the size of data transferred per second). This section describes how to view I/O rates and transfer rates of ports on the storage system. Before taking the following steps, you need to start monitoring in accordance with the procedure described in Starting and Stopping Storage System Monitoring and obtain the usage statistics.



Note: Performance Monitor can obtain statistics about traffics of ports connected to open-system host groups only. The statistics about traffics of ports connected to mainframe host groups cannot be obtained.

Viewing I/O Rates for Ports

Performance Monitor monitors ports on the storage system and measures I/O rates (that is, the number of disk I/Os per second).

To view I/O rates:

1. Ensure that the **Performance Management** window is displayed.
2. Select the **Port-LUN** tab.
The tree displays a **Subsystem** folder.
3. Select the **Subsystem** folder.
The **Target** folder and the **Initiator/External** folder are displayed.
4. Select **IOPS** from the drop-down list on the right side of the window.
5. In **Monitoring Term**, do one of the following:
 - To view the I/O rate in real time, you must select the **Real Time** option, specify the number of recent collections of statistics which should be displayed in the graph, and then click **Apply**.
 - To view I/O rates for a certain period of time in the last 24 hours, you must select the **From** option, change the date and time in the **From** and **To** boxes, and then click **Apply**. Use the arrow button and the sliders when you change the date and time in the **From** and **To** boxes.

For details on the **Real Time** option and the **From** option, see Viewing Port Workload Statistics.

6. In the tree, select the **Target** folder, or the **Initiator/External** folder.
The list on the right displays I/O rates for the ports on the storage system. If you click the **Target** folder the tree displays a list of ports on the storage system. If you click a port icon (for example,  and ) , the host groups () corresponding to that port appear.

Tips:

- If you select a port (for example,  and ) in the tree, the list displays I/O rates for all the host bus adapters connected to the selected port=.
 - If you select a host group () in the tree, the list displays I/O rates for host bus adapters in the host group.
7. To display a graph to find out how the I/O rate has been changed, take the following steps:
- a. In the list, select one or more ports or host bus adapters (WWNs).
 - b. Select the **Draw** button.



Caution: If the graph does not display changes in the I/O rate (for example, if the line in the graph runs vertically), it is recommended that you change the value in the **Chart Y Axis Rate** drop-down list. For example, if the largest value in the list is 200 and the value in **Chart Y Axis Rate** is 100, you should select a value larger than 200 from **Chart Y Axis Rate**.

I/O rate for ports (When the Subsystem folder is selected):

	Port	Current(IOPS)	Ave.(IOPS)(Total)	Peak(IOPS)	Response Time(ms)	Attribute
	All Prio.	-	14	19	1.999	Prio.
	All Non-Prio.	-	13	18	1.998	Non-Prio.

All Prio. indicates all the prioritized ports.

All Non-Prio. indicates all the non-prioritized ports

I/O rate for ports (When the Target folder is selected):

	Port	Current(IOPS)	Ave.(IOPS)(Total)	Peak(IOPS)	Response Time(ms)	Attribute
	CL1-A(E8)	-	15	20	2.000	Prio.
	CL3-A(E4)	-	65	70	2.050	Non-Prio.
	CL5-A(E2)	-	115	120	2.100	Non-Prio.
	CL7-A(E1)	-	165	170	2.150	Prio.
	CL1-B(E0)	-	215	220	2.200	Prio.
	CL3-B(DC)	-	265	270	2.250	Non-Prio.
	CL5-B(DA)	-	315	320	2.300	Non-Prio.

I/O rate for ports (When the Initiator/External folder is selected):

	Port	Current(IOPS)	Ave.(IOPS)(Total)	Peak(IOPS)	Response Time(ms)	Initiator/External
	CL1-E(D6)	-	415	420	2.400	Initiator
	CL3-E[1-E 2nd](D5)	-	465	470	2.450	Initiator
	CL5-E(D4)	-	515	520	2.500	External
	CL7-E[5-E 2nd](D3)	-	565	570	2.550	External

I/O rate for host bus adapters connected to a specified port (When a port ((for example, and) is selected):

	WWN	SPM Name	Current(IOPS)	Ave.(IOPS)(Total)	Max.(IOPS)	Response Time(ms)	Attribute
	0000000000000001	SPM-1	-	15	20	2.000	Non-Prio.
	0000000000000002	SPM-2	-	65	70	2.050	Non-Prio.
	0000000000000003	SPM-3	-	115	120	2.100	Non-Prio.
	0000000000000004	SPM-4	-	165	170	2.150	Non-Prio.
	0000000000000005	SPM-5	-	215	220	2.200	Non-Prio.
	0000000000000006	SPM-6	-	265	270	2.250	Non-Prio.
	Others	-	-	15	20	2.000	-

**I/O rates for host bus adapters in a host group
(When a host group () is selected):**

	WWN	Nickname	Current(IOPS)	Ave.(IOPS)(Total)	Max.(IOPS)	Response Time(ms)	Attribute
	0000000000000001	WWN-1	-	15	20	2.000	Non-Prio.
	0000000000000002	WWN-2	-	65	70	2.050	Non-Prio.

Figure 5-15 I/O Rates for Ports

The list displays the following:

- : When the green checkmark icon  is displayed on the left of an item, the graph displays changes in workload statistics about the item.
- **Port**: Indicates ports on the storage system.
- **WWN**: Indicates WWNs of the host bus adapters.
- **SPM Name**: Indicates SPM names of the host bus adapters. Server Priority Manager allows you assign an SPM name to each host bus adapter so that you can easily identify each host bus adapters in the Server Priority Manager windows.
- **Nickname**: Indicates the nickname for the host bus adapters. LUN Manager allows you to assign a nickname to each host bus adapter so that you can easily identify each host bus adapter in the LUN Manager windows.
- **Current**: Indicates the current I/O rate.
- **Ave.**: Indicates the average I/O rate for the specified period.
- **Peak**: Indicates the peak I/O rate of the ports for the specified period. This value is shown in the list when you select the **Subsystem** folder in the tree.

If you select a port in the list, click the **Draw** button, and select the **Detail** check box, the detailed graph of the port I/O rate is drawn. The **Peak** value means the top of the **Max.** line in this graph.

- **Max.**: Indicates the maximum I/O rate for the specified period. This value is shown in the list when you select a port icon or host group icon in the tree.
- **Response Time**: This column indicates the time for replying from the port or host bus adapter when I/O accesses are made from the host to them. The unit is milliseconds. The average response time in the period specified at **Monitoring Term** is displayed.
- **Attribute**: Indicates the priority of each port. **Prio.** indicates a prioritized port. **Non-Prio.** indicates a non-prioritized port.



Note: In the list of **Port-LUN** tab, two types of aliases appear for host bus adapters: SPM name and nickname. If you select a port icon in the tree, SPM names defined by Server Priority Manager appear. If you select a host group icon in the tree, nicknames defined by LUN Manager appear. We recommend that you specify the same name for an SPM name and a nickname for convenience of host bus adapter management.

- **Initiator/External:** Indicates the port attribute. **Initiator** indicates an initiator port. **External** indicates an external port. Both of the ports are not controlled by Server Priority Manager.

Viewing Transfer Rates for Ports

Performance Monitor monitors ports on the storage system and measures transfer rates (that is, the size of data transferred in one second).

To view transfer rates:

1. Ensure that the **Performance Management** window is displayed.
2. Select the **Port-LUN** tab.
The tree displays a **Subsystem** folder.
3. Select the **Subsystem** folder.
The **Target** folder and the **Initiator/External** folder are displayed.
4. Select **MB/s** from the drop-down list on the right side of the window.
5. In **Monitoring Term**, do one of the following:
 - To view the transfer rate in real time, you must select the **Real Time** option, specify the number of recent collections of statistics which should be displayed in the graph, and then click **Apply**.
 - To view transfer rates for a certain period of time in the last 24 hours, you must select the **From** option, change the date and time in the **From** and **To** boxes, and then click **Apply**. Use the arrow button and the sliders when you change the date and time in the **From** and **To** boxes.

For details on the **Real Time** option and the **From** option, see Viewing Port Workload Statistics.

6. In the tree, select the **Target** folder, or the **Initiator/External** folder.
The list on the right displays transfer rates for the ports on the storage system. If you click the **Target** folder, the tree displays a list of ports on the storage system. If you click a port icon (for example,  and ) , the host groups () corresponding to that port appear.

Tips:

- If you select a port (for example,  and ) in the tree, the list displays transfer rates for all the host bus adapters connected to the selected port.
- If you select a host group () in the tree, the list displays transfer rates for host bus adapters in the host group.



Note: You cannot view information about host bus adapters if the host group is not registered in LUN Manager.

7. If you want to display a graph to find out how the transfer rate has been changed, take the following steps:
 - a. In the list, select one or more ports or host bus adapters (WWNs).
 - b. Select the **Draw** button.



Caution: If the graph does not display changes in the transfer rate (for example, if the line in the graph runs vertically), it is recommended that you change the value in the **Chart Y Axis Rate** drop-down list. For example, if the largest value in the list is 200 and the value in **Chart Y Axis Rate** is 100, you should select a value larger than 200 from **Chart Y Axis Rate**.

Transfer rate for ports (When the Subsystem folder is selected):

	Port	Current(MB/s)	Ave.(MB/s)(Total)	Peak(MB/s)	Response Time(ms)	Attribute
	All Prio.	-	1.499	1.999	1.999	Prio.
	All Non-Prio.	-	1.498	1.998	1.998	Non-Prio.

All Prio. indicates all the prioritized ports.

All Non-Prio. indicates all the non-prioritized ports.

Transfer rate for ports (When the Target folder is selected):

	Port	Current(MB/s)	Ave.(MB/s)(Total)	Peak(MB/s)	Response Time(ms)	Attribute
	CL1-A(E8)	-	1.500	2.000	2.000	Prio.
	CL3-A(E4)	-	1.549	2.049	2.050	Non-Prio.
	CL5-A(E2)	-	1.598	2.098	2.100	Non-Prio.
	CL7-A(E1)	-	1.646	2.146	2.150	Prio.
	CL1-B(E0)	-	1.695	2.195	2.200	Prio.
	CL3-B(DC)	-	1.744	2.244	2.250	Non-Prio.
	CL5-B(DA)	-	1.793	2.293	2.300	Non-Prio.

I/O rate for ports (When the Initiator/External folder is selected):

	Port	Current(MB/s)	Ave.(MB/s)(Total)	Peak(MB/s)	Response Time(ms)	Initiator/External
	CL1-E(D6)	-	1.891	2.391	2.400	Initiator
	CL3-E[1-E 2nd](D5)	-	1.939	2.439	2.450	Initiator
	CL5-E(D4)	-	1.988	2.488	2.500	External
	CL7-E[5-E 2nd](D3)	-	2.037	2.537	2.550	External

**Transfer rate for host bus adapters connected to a specified port
(When a port ((for example, and) is selected):**

	WWN	SPM Name	Current(MB/s)	Ave.(MB/s)(Total)	Max.(MB/s)	Response Time(ms)	Attribute
	0000000000000001	SPM-1	-	1.500	2.000	2.000	Non-Prio.
	0000000000000002	SPM-2	-	1.549	2.049	2.050	Non-Prio.
	0000000000000003	SPM-3	-	1.598	2.098	2.100	Non-Prio.
	0000000000000004	SPM-4	-	1.646	2.146	2.150	Non-Prio.
	0000000000000005	SPM-5	-	1.695	2.195	2.200	Non-Prio.
	0000000000000006	SPM-6	-	1.744	2.244	2.250	Non-Prio.
	Others	-	-	1.500	2.000	2.000	-

Transfer rates for host bus adapters in a host group(When a host group () is selected):

	WWN	Nickname	Current(MB/s)	Ave.(MB/s)(Total)	Max.(MB/s)	Response Time(ms)	Attribute
	0000000000000001	WWN-1	-	1.500	2.000	2.000	Non-Prio.
	0000000000000002	WWN-2	-	1.549	2.049	2.050	Non-Prio.

Figure 5-16 Transfer Rates for Ports

The list displays the following:

- : When the green checkmark icon  is displayed on the left of an item, the graph displays changes in workload statistics about the item.
- **Port**: Indicates ports on the storage system.
- **WWN**: Indicates WWNs of the host bus adapters.
- **SPM Name**: Indicates SPM names of the host bus adapters. Server Priority Manager allows you assign an SPM name to each host bus adapter so that you can easily identify each host bus adapters in the Server Priority Manager windows.
- **Nickname**: Indicates the nickname for the host bus adapters. LUN Manager allows you to assign a nickname to each host bus adapter so that you can easily identify each host bus adapter in the LUN Manager windows.
- **Current**: Indicates the current transfer rate.
- **Ave.**: Indicates the average transfer rate for the specified period.
- **Peak**: Indicates the peak transfer rate of the ports for the specified period. This value is shown in the list when you select the **Subsystem** folder in the tree.

If you select a port in the list, click the **Draw** button, and select the **Detail** check box, the detailed graph of the port transfer rate is drawn. The **Peak** value means the top of the **Max.** line in this graph.

- **Max.**: Indicates the maximum transfer rate for the specified period. This value is shown in the list when you select a port icon or host group icon in the tree.

- **Response Time:** This column indicates the time for replying from the port or host bus adapter when I/O accesses are made from the host to them. The unit is milliseconds. The average response time in the period specified at **Monitoring Term** is displayed.



Note: This column displays a hyphen (-) if the I/O rate is 0 (zero).

- **Attribute:** Indicates the priority of each port. **Prio.** indicates a prioritized port. **Non-Prio.** indicates a non-prioritized port.



Note: In the list of **Port-LUN** tab, two types of aliases appear for host bus adapters: SPM name and nickname. If you select a port icon in the tree, SPM names defined by Server Priority Manager appear. If you select a host group icon in the tree, nicknames defined by LUN Manager appear. We recommend that you specify the same name for an SPM name and a nickname for convenience of host bus adapter management.

- **Initiator/External:** Indicates the port attribute. **Initiator** indicates an initiator port. **External** indicates an external port. Both of the ports are not controlled by Server Priority Manager.

Viewing Details about the I/O and Transfer Rates

To view detailed information about the I/O rate or the transfer rate for a port:

1. From the drop-down list at the right side of the window, select the type of statistics to be displayed.
 - To view I/O rates, select **IOPS** (I/Os per second) from the drop-down list.
 - To view transfer rates, select **MB/s** (megabytes per second) from the drop-down list.
2. Select the **Subsystem** folder in the tree.
3. Select a port from the list.
4. Select the **Draw** button.

The graph that is not in detail is displayed.

5. Select the **Detail** check box

The graph contents change as described in Figure 4-10.

Monitoring LU Paths

Performance Monitor monitors LU paths and measures I/O rates (that is, the number of I/Os per second) and transfer rates (that is, the size of data transferred per second). This section describes how to view I/O rates and transfer rates of LU paths on the storage system. Before taking the following steps, you need to start monitoring in accordance with the procedure described in Starting and Stopping Storage System Monitoring and obtain the usage statistics.



Note: The traffic statistics reported for an LU is aggregated across all LU paths defined for an LU.

- I/O rate is the sum of I/Os across all LU paths defined for an LU.
 - Transfer rate is the total transfer rate across all LU paths defined for an LU.
 - **Response time** is the average response time across all LU paths defined for an LU.
-

Viewing LU Paths I/O Rates

Performance Monitor monitors LU paths and measures I/O rates (that is, the number of disk I/Os per second).

To view I/O rates:

1. Ensure that the **Performance Management** window is displayed.
2. Select the **Port-LUN** tab.
The tree displays a list of ports on the storage system.
3. Select **IOPS** from the drop-down list on the right side of the window.
4. In **Monitoring Term**, do one of the following:
 - To view the I/O rate in real time, you must select the **Real Time** option, specify the number of recent collections of statistics which should be displayed in the graph, and then click **Apply**.
 - To view I/O rates for a certain period of time in the last 24 hours, you must select the **From** option, change the date and time in the **From** and **To** boxes, and then click **Apply**. Use the arrow button and the sliders when you change the date and time in the **From** and **To** boxes.

For details on the **Real Time** option and the **From** option, see Viewing Port Workload Statistics.

5. In the tree, double-click a port (for example,  or ) and then select a host group (.

The **LUN** icon () appears.

6. Select **LUN**.

The list on the right displays a list of LU paths and I/O rates.

7. If you want to display a graph to find out how the I/O rate has been changed, take the following steps:
- In the list, select one or more LUNs.
 - Select the **Draw** button.



Note: If the graph does not display changes in the I/O rate (for example, if the line in the graph runs vertically), it is recommended that you change the value in the **Chart Y Axis Rate** drop-down list. For example, if the largest value in the list is 200 and the value in **Chart Y Axis Rate** is 100, you should select a value larger than 200 from **Chart Y Axis Rate**.

8. If you want to view more detailed information in the graph, select the **Detail** check box at the lower right of the list. The graph contents change as described in Table 4-3.



Note: If more than one row is selected in the list, you cannot select the **Detail** check box.

	LUN	LDEV	Emulation	Paths	Current(IOPS)	Ave.(IOPS)(Total)	Max.(IOPS)	Response Time(ms)
	0000	00:00:00	OPEN-V	7	-	15	20	2.000
	0001	00:00:01	OPEN-V	7	-	65	70	2.050
	0002	00:00:02	OPEN-V	7	-	115	120	2.100

Figure 5-17 I/O Rates for LU Paths

The list displays the following:

- : When the green checkbox icon is displayed on the left of an item, the graph displays changes in workload statistics about the item.
- LUN**: Indicates LUNs (logical unit numbers).
- LDEV**: Indicates IDs of volumes, in the following format: LDKC:CU:LDEV. An ID ending in **#** indicates the volume is an external volume. An ID ending in **V** or **X** indicates the volume is a virtual volume.
- Emulation**: Indicates emulation types.
- Paths**: Indicates the number of LU paths (i.e., paths to volumes).
- Current**: Indicates the current I/O rate.
- Ave.**: Indicates the average I/O rate for the specified period.
- Max.**: Indicates the maximum I/O rate for the specified period.

- **Response Time:** This column indicates the time for replying from the LU paths when I/O accesses are made from the host to the LU paths. The unit is milliseconds. The average response time in the period specified at **Monitoring Term** is displayed.

Viewing LU Paths Transfer Rates

Performance Monitor monitors LU paths and measures transfer rates (that is, the size of data transferred in one second).

To view transfer rates:

1. Ensure that the **Performance Management** window is displayed.
2. Select the **Port-LUN** tab.
The tree displays a list of ports on the storage system.
3. Select **MB/s** from the drop-down list on the right side of the window.
4. In **Monitoring Term**, do one of the following:
 - To view the transfer rate in real time, you must select the **Real Time** option, specify the number of recent collections of statistics which should be displayed in the graph, and then click **Apply**.
 - To view transfer rates for a certain period of time in the last 24 hours, you must select the **From** option, change the date and time in the **From** and **To** boxes, and then click **Apply**. Use the arrow button and the sliders when you change the date and time in the **From** and **To** boxes.

For details on the **Real Time** option and the **From** option, see Port-LUN Tab of the Performance Monitor Window.

5. In the tree, double-click a port (for example,  or ) and then select a host group ().

The **LUN** icon () appears.

6. Select **LUN**.
The list on the right displays a list of LU paths and transfer rates.
7. If you want to display a graph to find out how the transfer rate has been changed, take the following steps:
 - a. In the list, select one or more LUNs.
 - b. Click the **Draw** button.



Note: If the graph does not display changes in the transfer rate (for example, if the line in the graph runs vertically), it is recommended that you change the value in the **Chart Y Axis Rate** drop-down list. For example, if the largest value in the list is 200 and the value in **Chart Y Axis Rate** is 100, you should select a value larger than 200 from **Chart Y Axis Rate**.

8. If you want to view more detailed information in the graph, select the **Detail** check box at the lower right of the list. The graph contents change as described in Table 4-3.



Note: If more than one row is selected in the list, you cannot select the **Detail** check box.

	LUN	LDEV	Emulation	Paths	Current(MB/s)	Ave.(MB/s)(Total)	Max.(MB/s)	Response Time(ms)
☒	0000	00:00:00	OPEN-V	7	-	1.500	2.000	2.000
☐	0001	00:00:01	OPEN-V	7	-	1.549	2.049	2.050
☐	0002	00:00:02	OPEN-V	7	-	1.598	2.098	2.100

Figure 5-18 Transfer Rates for LU Paths

The list displays the following:

- : When the green checkmark icon is displayed on the left of an item, the graph displays changes in workload statistics about the item.
- **LUN**: Indicates LUNs (logical unit numbers).
- **LDEV**: Indicates IDs of volumes, in the following format: LDKC:CU:LDEV. An ID ending in **#** indicates the volume is an external volume. An ID ending in **V** or **X** indicates the volume is a virtual volume.
- **Emulation**: Indicates emulation types.
- **Paths**: Indicates the number of LU paths (i.e. paths to volumes).
- **Current**: Indicates the current transfer rate.
- **Ave.**: Indicates the average transfer rate for the specified period.
- **Max.**: Indicates the maximum transfer rate for the specified period.
- **Response Time**: This column indicates the time for replying from the LU paths when I/O accesses are made from the host to the LU paths. The unit is milliseconds. The average response time in the period specified at **Monitoring Term** is displayed.



Note: This column displays a hyphen (-) if the I/O rate is 0 (zero).

Monitoring Paths between Host Bus Adapters and Ports

If Server Priority Manager is enabled, Performance Monitor can be used to monitor paths between host bus adapters (HBAs) in host servers and ports on the storage system. HBAs are contained in host servers and serve as ports for connecting the host servers to the storage system. This section describes how to view I/O rates and transfer rates between host bus HBAs and ports.

Before taking the steps described in the following sections, you need to do the following:

1. Perform a Server Priority Manager operation (see Monitoring All Traffic between HBAs and Ports and Setting Priority for Host Bus Adapters).
2. Start the monitoring by using Performance Monitor in accordance with the procedure described in Starting and Stopping Storage System Monitoring and obtain the traffic data between host bus HBAs and ports.

Viewing HBA Information

Viewing I/O Rates between HBAs

Performance Monitor monitors traffic between HBAs in the hosts and ports on the storage system, and measures I/O rates (that is, the number of disk I/Os per second).

To view I/O rates:

1. Ensure that the Performance Management window is displayed.
2. Click the **WWN** tab.

The tree displays a list of SPM groups (). Below that list, an item named **Not Grouped**.

- To display host bus adapters () in the SPM group, double-click an SPM group.
 - To display host bus adapters () that do not belong to any SPM group, double-click **Not Grouped**.
3. Select **IOPS** from the drop-down list on the right side of the window.
 4. In **Monitoring Term**, do one of the following:
 - To view the I/O rate in real time, you must select the **Real Time** option, specify the number of recent collections of statistics which should be displayed in the graph, and then click **Apply**.
 - To view I/O rates for a certain period of time in the last 24 hours, you must select the **From** option, change the date and time in the **From** and **To** boxes, and then click **Apply**. Use the arrow button and the sliders when you change the date and time in the **From** and **To** boxes.

For details on the **Real Time** option and the **From** option, see WWN Tab of the Performance Monitor Window.

5. In the tree, do one of the following:

- To view the I/O rate for host bus adapters in an SPM group, select the SPM group.
The list on the right displays the I/O rate.
- To view the I/O rate for host bus adapters that do not belong to any SPM group, select **Not Grouped**.
The list on the right displays the I/O rate.

Tips:

- If you select the **Subsystem** folder in the tree, the list displays the I/O rate at each SPM group.
 - If you select a host bus adapter () in the tree, the list displays the I/O rate at each port connected to the selected host bus adapter.
6. If you want to display a graph to find out how the I/O rate has been changed, take the following steps:
- a. In the list, select one or more SPM groups, WWNs, or ports.
 - b. Select the **Draw** button.



Caution: If the graph does not display changes in the I/O rate (for example, if the line in the graph runs vertically), it is recommended that you change the value in the **Chart Y Axis Rate** drop-down list. For example, if the largest value in the list is 200 and the value in **Chart Y Axis Rate** is 100, you should select a value larger than 200 from **Chart Y Axis Rate**.



Note: If the WWN of a host bus adapter (HBA) is displayed in red in the tree, the host bus adapter is connected to two or more ports, but the traffic between the HBA and some of the ports are *not* monitored by Performance Monitor. When many-to-many connections are established between HBAs and ports, you should make sure that all the traffic between HBAs and ports is monitored. For details on the measures when a WWN is displayed in red, see Monitoring All Traffic between HBAs and Ports.

**I/O rate at host bus adapters
(Displayed when an SPM group or Not Grouped is selected):**

	WWN	SPM Name	Current(IOPS)	Ave.(IOPS)(Total)	Max.(IOPS)	Response Time(ms)	Attribute
	0000000000000001	SPM-1	-	65	70	2.050	Non-Prio.
	0000000000000009	SPM-9	-	465	470	2.450	Non-Prio.
	0000000000000011	SPM-17	-	865	870	2.850	Non-Prio.
	0000000000000019	SPM-25	-	1265	1270	3.250	Non-Prio.

I/O rate at each SPM group

(Displayed when the Subsystem folder is selected):

	Group	Current(IOPS)	Ave.(IOPS)(Total)	Max.(IOPS)	Response Time(ms)	Attribute
	group-1	-	65	70	2.050	Non-Prio.
	group-2	-	115	120	2.100	Non-Prio.
	group-3	-	165	170	2.150	Non-Prio.
	group-4	-	215	220	2.200	Non-Prio.
	<Not Grouped>	-	15	20	2.000	-

I/O rate at each port connected to a specified host bus adapter

(Displayed when a host bus adapter is selected):

	Port	Current(IOPS)	Ave.(IOPS)(Total)	Max.(IOPS)	Response Time(ms)	Attribute
	CL1-A	-	15	20	2.000	Prio.
	CL1-B	-	1615	1620	3.600	Non-Prio.

Figure 5-19 I/O Rates for Host Bus Adapters (WWN Tab)

The list displays the following:

- : When the green checkmark icon  is displayed on the left of an item, the graph displays changes in workload statistics about the item.
- **Group**: Indicates SPM groups.
- **WWN**: Indicates WWNs of the host bus adapters.
- **SPM Name**: Indicates SPM names of host bus adapters. Server Priority Manager allows you assign an SPM name to each host bus adapter so that you can easily identify each host bus adapters in the Server Priority Manager windows.
- **Port**: Indicates ports on the storage system.
- **Current**: Indicates the current I/O rate.
- **Ave.**: Indicates the average I/O rate for the specified period.
- **Max.**: Indicates the maximum I/O rate for the specified period.
- **Response Time**: This column indicates the time for replying from the host bus adapter, SPM group, or port when I/O accesses are made from the host to them. The unit is milliseconds. The average response time in the period specified at **Monitoring Term** is displayed.
- **Attribute**: Indicates the priority of each host bus adapter (HBA). **Prio.** indicates a high-priority HBA (a prioritized WWN). **Non-Prio.** indicates a low-priority HBA (a non-prioritized WWN).

Viewing Transfer Rates between HBAs

Performance Monitor monitors traffic between HBAs in the hosts and ports on the storage system, and measures transfer rates (that is, the size of data transferred in one second).

To view transfer rates:

1. Ensure that the **Performance Management** window is displayed.
2. Select the **WWN** tab.

The tree displays a list of SPM groups (). Below SPM groups, an item named **Not Grouped** is displayed.

- To display host bus adapters () in the SPM group, double-click an SPM group.
 - To display host bus adapters () that do not belong to any SPM group, double-click **Not Grouped**.
3. Select **MB/s** from the drop-down list on the right side of the window.
 4. In **Monitoring Term**, do one of the following:
 - To view the transfer rate in real time, you must select the **Real Time** option, specify the number of recent collections of statistics which should be displayed in the graph, and then click **Apply**.
 - To view transfer rates for a certain period of time in the last 24 hours, you must select the **From** option, change the date and time in the **From** and **To** boxes, and then click **Apply**. Use the arrow button and the sliders when you change the date and time in the **From** and **To** boxes.

For details on the **Real Time** option and the **From** option, see WWN Tab of the Performance Monitor Window.

5. In the tree, do one of the following:
 - To view the transfer rate for host bus adapters in an SPM group, select the SPM group.
The list on the right displays the transfer rate.
 - To view the transfer rate for host bus adapters that do not belong to any SPM group, select **Not Grouped**.
The list on the right displays the transfer rate.

Tips:

- If you select the **Subsystem** folder in the tree, the list displays the transfer rate at each SPM group.
- If you select a host bus adapter () in the tree, the list displays the transfer rate at each port connected to the selected host bus adapter.

6. If you want to display a graph to find out how the transfer rate has been changed, take the following steps:
 - a. In the list, select one or more SPM groups, WWNs, or ports.
 - b. Select the **Draw** button.



Notes:

- If the graph does not display changes in the transfer rate (for example, if the line in the graph runs vertically), it is recommended that you change the value in the **Chart Y Axis Rate** drop-down list. For example, if the largest value in the list is 200 and the value in **Chart Y Axis Rate** is 100, you should select a value larger than 200 from **Chart Y Axis Rate**.
- If the WWN of a host bus adapter (HBA) is displayed in red in the tree, the host bus adapter is connected to two or more ports, but the traffic between the HBA and some of the ports are *not* monitored by Performance Monitor. When many-to-many connections are established between HBAs and ports, you should make sure that all the traffic between HBAs and ports is monitored. For details on the measures when a WWN is displayed in red, see Monitoring All Traffic between HBAs and Ports.

Transfer rate at host bus adapters

(Displayed when an SPM group or Not Grouped is selected):

	WWN	SPM Name	Current(MB/s)	Ave.(MB/s)(Total)	Max.(MB/s)	Response Time(ms)	Attribute
	0000000000000001	SPM-1	-	1.549	2.049	2.050	Non-Prio.
	0000000000000009	SPM-9	-	1.939	2.439	2.450	Non-Prio.
	0000000000000011	SPM-17	-	2.330	2.830	2.850	Non-Prio.
	0000000000000019	SPM-25	-	2.721	3.221	3.250	Non-Prio.

Transfer rate at each SPM group

(Displayed when the Subsystem folder is selected):

	Group	Current(MB/s)	Ave.(MB/s)(Total)	Max.(MB/s)	Response Time(ms)	Attribute
	group-1	-	1.549	2.049	2.050	Non-Prio.
	group-2	-	1.598	2.098	2.100	Non-Prio.
	group-3	-	1.646	2.146	2.150	Non-Prio.
	group-4	-	1.695	2.195	2.200	Non-Prio.
	<Not Grouped>	-	1.500	2.000	2.000	-

Transfer rate at each port connected to a specified host bus adapter

(Displayed when a host bus adapter is selected):

	Port	Current(MB/s)	Ave.(MB/s)(Total)	Max.(MB/s)	Response Time(ms)	Attribute
	CL1-A	-	1.500	2.000	2.000	Prio.
	CL1-B	-	3.063	3.563	3.600	Non-Prio.

Figure 5-20 Transfer Rates for Host Bus Adapters (WWN Tab)

The list displays the following:

- : When the green checkmark icon  is displayed on the left of an item, the graph displays changes in workload statistics about the item.
- **Group:** Indicates SPM groups.
- **WWN:** Indicates WWNs of host bus adapters.
- **SPM Name:** Indicates SPM names of host bus adapters. Server Priority Manager allows you assign an *SPM name* to each host bus adapter so that you can easily identify each host bus adapters in the **Server Priority Manager** windows.
- **Port:** Indicates ports on the storage system.
- **Current:** Indicates the current transfer rate.
- **Ave.:** Indicates the average transfer rate for the specified period.
- **Max.:** Indicates the maximum transfer rate for the specified period.
- **Response Time:** This column indicates the time for replying from the host bus adapter, SPM group, or port when I/O accesses are made from the host to them. The unit is milliseconds. The average response time in the period specified at **Monitoring Term** is displayed.



Note: This column displays a hyphen (-) if the I/O rate is 0 (zero).

- **Attribute:** Indicates the priority of each host bus adapter (HBA). **Prio.** indicates a high-priority HBA (a prioritized WWN). **Non-Prio.** indicates a low-priority HBA (a non-prioritized WWN).

Volume Migration Operations

This chapter explains the following volume migration operations:

- [Starting Volume Migration](#)
- [Reserving Migration Destinations for Volumes](#)
- [Migrating Volumes Manually](#)
- [Automating Volume Migrations](#)
- [Viewing the Migration History Log](#)

Starting Volume Migration

To start Volume Migration, take the following steps:

1. Ensure that the Storage Navigator main window is in **Modify** mode.

The Storage Navigator main window is in **Modify** mode if the background color of the  (pen tip) icon is light yellow. If the background color is gray, the window is in **View** mode and you must change it to **Modify** mode by taking the following steps:

- a. Verify whether the background color of the lock icon  is blue.

If the background color is red , you will not be able to switch from **View** mode to **Modify** mode. Wait for a while and click the  button. If the background color turns blue, you can go to the next step.

- b. Select the  icon.

A message asks whether you want to change the mode.

- c. Select **OK** to close the message.

The background color of the icon changes to light yellow (). The mode changes to **Modify** mode. The background color of the lock icon becomes red ().

2. Ensure that the Physical tab is active in the **Performance Management** window.
3. In the drop-down list at the right of **Monitoring Data**, select **long-range**.
The **Volume Migration** button is activated.
4. Select the **Volume Migration** button.
The Volume Migration window is displayed.

Reserving Migration Destinations for Volumes

To be able to migrate volumes, you must reserve disk spaces (actually, volumes) that can be used for migration destinations. You must reserve volumes regardless of migration methods (manual migration or automatic migration). Migration will not take place when no volume is reserved.

Throughout the USP V/VM documentation set, the term "reserved volume" is used to refer to a volume that is reserved for use as a migration destination. Volumes reserved by Volume Migration are indicated by blue  icons. Also, the term "normal volume" is sometimes used to refer to a volume that is not reserved as a migration destination.

To reserve a volume as a migration destinations:

1. Start Volume Migration and select the **Attribute** tab in the **Volume Migration** window.
2. In the tree, locate a volume that you want to use as a migration destination.

To locate a volume, first double-click an HDD class (or the **External Group** folder) in the tree to display a list of parity groups (or external volume groups). Next, click a parity group (or an external volume group).

The list on the right displays a list of volumes in that group.

3. In the list on the right of the tree, locate volumes preceded by the  icon. From those volumes, select the desired volume and right-click it.



Note: If you want to reserve two or more volumes, you can select the desired volumes and right-click the selection.

4. Select **Reserved LDEV** from the pop-up menu.

The icon changes to a blue  icon.

5. Select the **Apply** button.

The settings in the window are applied to the storage system. The specified volume is reserved as a migration destination.

Tip: If you right-click a volume preceded by the blue  icon (reserved volume of Volume Migration) and then choose **Normal LDEV** from the pop-up menu, the selected icon changes to  and you will not be able to use the volume as a migration destination.

**Notes:**

- When you click the **Apply** button one time, up to 64 volumes can be reserved as migration destinations. If you click the **Apply** button to reserve more than 64 volumes, an error occurs. If you want to reserve more than 64 volumes, you must follow the above procedure at least twice. For example, if you want to reserve 100 volumes, you need to reserve 64 or fewer volumes and then reserve the remaining volumes.
 - A volume preceded by the green  icon is reserved by a program other than Volume Migration and you cannot change the attribute of the volume by using the pop-up menu.
-

Migrating Volumes Manually

To migrate volumes manually, you must create migration plans. A migration plan is a pair of a source volume and the corresponding target (destination) volume. A migration plan is prepared for each volume that you want to migrate.

In manual migration, you can specify an external volume in an external volume group as a source volume or a target volume of migration.



Caution: Do not perform manual migration operations while the auto migration function is active. If you perform a manual migration operation, all the auto migration plans created until then will be canceled.

Before migrating volumes manually, stop the auto migration function first by selecting **Auto Plan** in the **Volume Migration** window, and then following the procedure below:

1. In **Auto Migration Plans**, select the **Delete All** button to remove all the auto migration plans.
2. In **Auto Migration Function**, select Disable.
3. Select the **Set** button to apply the settings to the storage system.



Caution: When performing a series of manual migration operations, remember that the configuration changes after each migration is complete.

To migrate volumes manually:

1. Start Volume Migration and select the **Manual Plan** tab in the Volume Migration window.
2. In the tree, double-click a folder icon of parity group (or an external volume group).

A list of parity groups or external volumes is displayed.

A list of volumes in that parity group (or external volume group) is displayed. Volumes that can be migrated is indicated by the  icon. Usage rate of the parity groups are displayed if you select the **Display usage rate.** button.



Note: the **Display usage rate.** button is toggle-able. If the button is held down, usage rates are displayed. If the button is not held down, usage rates are not displayed. Usage rates are not displayed in initial state of this tab.

3. Select the parity group icon that includes the volume you want to migrate.
The list of volumes that are belonged to the parity group is displayed in the **LDEVs** list.
4. Select the volume you want to migrate from the **LDEVs** and click the **Target** button.
Volume Migration finds the volumes that can be used as target volumes and displays the list of the volumes in the **Target(Reserved) LDEV** list.
5. Select the target volume from the volumes displayed in the **Target LDEV** and click the **Graph** button.
The **Graph** window appears and displays the estimated value of the performance after migration of the selected volume.
6. Check the estimated results in **Source LDEV**, **Target LDEV** on the **Graph** window and click the **Close** button.
The **Graph** window is closed.
 - If you anticipate getting the desired disk performance, go to the next step:
 - If not, return to step 2 or 3 to change the source and target volumes.
 - If you want to migrate the volumes that are already decided, you can skip the operations step 5 and step 6.
7. Select the **Set** button.
A new migration plan, consisting of the specified source and target volumes, is added to the list. The new migration plan is displayed in blue.



Note: Throughout the USP V/VM documentation set, the term "migration plan" is used to refer to a pair of a source volume and the corresponding target volume.

8. If you want to create and execute more than one migration plan, repeat steps 2 through 7 to add migration plans. The newly added migration plan will appear below the existing migration plans.
9. Select the **Apply** button.

The migration plans are applied to the storage system and the migrations start. Volume Migration performs multiple migration plans concurrently. The font color of the migration plans during execution change from blue to black.

When a volume starts migration, the percent % column displays the progress of the migration. When the migration finished, the migration plan of that volume will disappear from the list.

If execution of a migration plan created by an application software other than Storage Navigator failed, PSUE is displayed in the **Stat** column. In this case, you must check the status of the application software which created the migration plan.

**Notes:**

- Up to eight migration plans can be applied with each click of the **Apply** button. To create and apply more than eight migration plans, use the **Apply** button several times. While executing the migration plans, you can apply additional migration plans.
 - The number of the migration plans that can be executed concurrently might be restricted depending on the usage of other USP V/VM programs, and also depending on the emulation types and sizes of the migrated volumes.
-

Related topics:

- To delete a migration plan, select a migration plan in the list and then click the **Delete** button.
- If you select a migration plan in black (a migration plan that is waiting for execution or being executed) and then click the **Delete** button, the character "D" appears in the **DEL** column, and the migration plan will not be deleted until you click the **Apply** button.
- You cannot use the **Delete** button to delete migration plans executed by programs other than Volume Migration. These migration plans are indicated by the  (green) icon displayed in the **LDEV** column and by "Other [XX]" displayed in the **Owner** column in the **Target LDEV** list. These migration plans also disappear from the list when the migration finished.

**Cautions:**

- If you delete a migration plan that is being executed, the data in the target volume will not be guaranteed.
 - When an error condition exists in the storage system, resource usage can increase or become unbalanced. Do not use usage statistics collected during an error condition as the basis for planning Volume Migration operations.
-

Automating Volume Migrations

You can set volume migrations automatically take place every day at the specified time. This section explains the procedure for automating volume migrations.



Note: Volume Migration cannot perform auto migration using external volumes and volumes reserved by a program other than Volume Migration.

Setting a Fixed Parity Group

A fixed parity group is a parity group that cannot be used by the auto migration function. If you want to exclude a parity group (all volumes in the group) from auto migration operations, you need to fix the parity group. Volumes in fixed parity groups will not be used for auto migration operations, but can be used as needed for manual migration operations.

To fix a parity group so that it cannot be migrated manually:

1. Start Volume Migration and select the **Attribute** window of the Volume Migration.
2. In the tree, double-click an HDD class folder.

The list on the right displays a list of parity groups in the HDD class. The  icon indicates "fixed" parity groups, which cannot include source nor target volumes. Other icons indicate parity groups that can include source or target volumes.

3. From the parity groups that are **not** displayed by the  icon, select a parity group that you want to exclude from auto migration operations. Next, right click the  parity group and select **Fixed PG** from the pop-up menu. The selected parity group is indicated by the  icon and becomes "fixed".
4. Select the **Apply** button in the **Attribute** window.

Tip: If you want to enable volumes in a fixed parity group to be moved to another parity group, you must change the fixed parity group to a "normal" parity group. To do this, right-click the fixed parity group (indicated by the  icon) and then select **Normal PG** from the pop-up menu.



Note: When you perform manual migration, you can specify volumes in the fixed parity groups as the source or target volumes.

Changing the Maximum Disk Usage Rate for Each HDD Class

By default, each HDD class is assigned a certain limit on disk usage rate. As explained earlier, this limit affects auto migration behavior. If necessary, you can change the limit on disk usage rate.

To change the limit on disk usage rate for an HDD class:

1. Start Volume Migration and select the **Attribute** tab in the **Volume Migration** windows.
2. In the tree, right-click an HDD class folder and then select **Change Class Threshold** from the pop-up menu.
The Change Threshold window is displayed.
3. From the **Threshold** drop-down list, select the value that you want to use as the limit and then click **OK**.
In the tree, the selected value appears at the right of the specified class.
4. Select the **Apply** button in the **Attribute** window.
The specified limit is applied to the storage system.




Figure 6-1 Change Threshold Dialog Box

Setting Auto Migration Parameters

To be able to automate volume migration, you must specify migration parameters such as the migration time.

To configure auto migration parameters:

1. Start Volume Migration and select the **Auto Plan** tab in the **Volume Migration** windows.
2. In **Auto Refresh**, specify **Enable** or **Disable**.
 - If you specify **Enable**, a new plan will be automatically added to the **Auto Migration Plans** list when the plan is automatically created. If a migration plan executes, the progress of execution is displayed and will automatically be updated in the **Auto Migration Plans** list. Volume Migration displays a message notifying you whenever the **Auto Migration Plans** list is updated.

- If you specify **Disable**, the **Auto Migration Plans** list will not be updated when a migration plan is automatically created or executed.
3. Use the **Date** and **Time** boxes in **Sampling Term** to specify a period so that Volume Migration analyzes resource usage statistics for that period.

For example, if you want to improve disk access speed every day from 9:00 a.m. to 5:00 p.m., you should specify the period as "every day from 9:00 a.m. to 5:00 p.m.". Auto migration plans will not be created if you select **None** in the **Date** box.

Important: Make sure that the data term and the auto migration execution time are separated by at least an hour. For example, if the sampling term is from 10:00 to 14:00, auto migration should not be performed between 09:00 and 15:00.

Auto migration plans will not be created if you select **None** in the **Date** box.
 4. In **Number of Sampling Points**, do one of the following:
 - To enable Volume Migration to analyze all the average usage rates in the sampling term, select **All sampling points**.
 - To enable Volume Migration to analyze only the X highest average usage rates in the sampling term, select **X highest sampling points** and then select the value of X from the drop-down list.
 5. In **Migration Time**, specify when to start auto migration operations.

**Cautions:**

- Volume migrations might impose heavy workloads on disks. For this reason, it is recommended that you specify a time when disk usage rates are low.
- Auto migration is likely to place a heavy workload on servers. Therefore, auto migration might slow down processing on the clients.
- The migration time must be at least 15 minutes plus the specified **Max. migration duration** earlier than the specified **Sampling Term** start time.
- The migration time must be at least 60 minutes later than the specified **Sampling Term** end time.

-
6. If necessary, specify auto migration parameters as follows:
 - a. Use the **Max. migration duration** drop-down list to set a time limit for auto migrations.



Note: If auto migration is not complete within this limit, Volume Migration cancels the incomplete operations.

- b. Use the Max. disk utilization drop-down list to set a disk usage limit.

- c. Use the Max. number of vols for migration drop-down list to specify the maximum number of volumes that can be auto migrated at the same time.
7. In Auto Migration Function select **Enable**.
8. Select the **Set** button.

The settings in this window will be applied to the storage system.

Migration plans will be automatically created 30 minutes later than the specified sampling term. For example, if you specified 13:00 - 15:00 for the sampling term, plan creation will be performed at 15:30.
9. Select the **Close** button to close the **Auto Plan** window.



Note: Plans are executed from the top of the list downward when Volume Migration performs auto migration.

Related topics

- To delete all the auto migration plans, click the **Delete All** button.
- To delete an auto migration plan, select the plan in the list and then select the **Delete** button.

Remaking Auto Migration Plans

If auto migration operations did not produce enough improvements in disk access performance, you can discard the current migration plans and make new auto migration plans.

To make new auto migration plans:

1. If necessary, do the following:
 - Change the limit on disk usage rate for HDD classes. For details, see Changing the Maximum Disk Usage Rate for Each HDD Class.
 - Change parameters in **Auto Plan Parameters**. For details see Setting Auto Migration Parameters.
2. In the **Auto Plan** window of the **Volume Migration**, select the **Make New Plan** button.

All the existing auto migration plans are deleted and new auto migration plans are created.
3. Select the **Close** button to close the **Auto Plan** window.

Viewing the Migration History Log

The **History** window of the **Volume Migration** displays logs of auto migration operations and manual migration operations.

To view migration logs:

1. Start Volume Migration to open the Volume Migration windows.
2. Select the **History** tab of the Volume Migration windows.
 - To view logs of auto migration operations, look at **Auto Migration History**. For messages that might appear in **Auto Migration History**, refer to Table 6-1.
 - To view logs of manual migration operations, look at **Migration History**. For messages that might appear in **Migration History**, refer to
 - Table 6-2. In **Migration History**, logs of auto migration operations are also displayed together with that of manual migration operations.



Note: Messages in the **History** window may not be arranged in the order that corresponding events take place.

Table 6-1 Messages in the Auto Migration Logs

A. Messages about starting and ending auto migration		
yyyy/mm/dd hh:min The auto migration started. (a) Volume Migration started to execute auto migration plans. (a) Start time		
yyyy/mm/dd hh:min All the auto migration plans were executed normally. (a) The volume migration ended. All the auto migration plans were executed normally. (a) Execution end time		
yyyy/mm/dd hh:min The auto migration stopped because XXXXXXXXXXXX (a) (b) The execution of auto migration plans stopped due to the cause shown in XXXXXXXXXXXX. The auto migration plans executed before yyyy/mm/dd hh:min were ended successfully. The remaining auto migration plans were not executed. (a) Time when auto migration was stopped (b) Cause of stop:		
	the statistics could not be obtained.	Volume Migration failed to obtain the monitoring data.
	the SVP was rebooted.	The SVP was rebooted.

A. Messages about starting and ending auto migration			
yyyy/mm/dd hh:min The auto migration plans remaining after the migration error were deleted. (a) Volume Migration failed to execute an auto migration plan. So, Volume Migration deleted the remaining plans. (a) Time when the remaining auto migration plans were deleted			
B. Messages about executing individual auto migration plans			
yyyy/mm/dd hh:min The auto migration plan (LDKC:CU:LDEV->LDKC:CU:LDEV) was executed normally. (a) (b) (c) The execution of the auto migration plan (LDKC:CU:LDEV->LDKC:CU:LDEV) was finished normally. (a) Execution end time (b) Source volume (c) Target volume			
yyyy/mm/dd hh:min The auto migration plan (LDKC:CU:LDEV->LDKC:CU:LDEV) was canceled because (a) (b) (c) of timeout. The auto migration plan (LDKC:CU:LDEV->LDKC:CU:LDEV) was canceled because the duration of the migration exceeded the time specified in Max. Migration duration of the Auto Plan window. (a) Canceled time (b) Source volume (c) Target volume			
yyyy/mm/dd hh:min The auto migration plan (LDKC:CU:LDEV (X-X)->LDKC:CU:LDEV (X-X)) was (a) (b) (c) (d) (e) canceled because of invalid parity group information. The auto migration plan (LDKC:CU:LDEV (X-X)->LDKC:CU:LDEV (X-X)) was canceled because the specified parity group information is incorrect. The correspondence between the source or target volume and the parity group is invalid. That volume might have been migrated manually after the creation of the auto migration plan. (a) Canceled time (b) Source volume (c) Parity group of the source volume (d) Target volume (e) Parity group of the target volume			
yyyy/mm/dd hh:min The auto migration plan (LDKC:CU:LDEV->LDKC:CU:LDEV) was stopped. (a) (b) (c) Cause: XXXXXXXXXXXX (d) The execution of the auto migration plan (LDKC:CU:LDEV->LDKC:CU:LDEV) was stopped due to the cause shown in XXXXXXXXXXXX. The status of the volumes returned to that before the migration. This auto migration plan is not valid. If necessary, create auto migration plans again. (a) Stopped time (b) Source volume (c) Target volume (d) Cause of stop:			

B. Messages about executing individual auto migration plans		
	The specified target volume is not a reserved volume.	The specified target volume is not a reserved volume. This volume was changed to a normal volume after the creation of the auto migration plan.
	Emulation type mismatch between source and target volumes.	The emulation type of the source volume is different from that of the target volume. The volume configuration might have been changed after the creation of the auto migration plan.
	The usage rate of the parity group is over the specified maximum disk utilization, or the usage rate is unknown.	The disk usage rate of the parity group including the source or target volume is over the specified limit. Or, the disk usage rate of the parity group including the source or target volume is unknown because no monitoring data exist about them. If necessary, change the limit of the disk usage rate and create auto migration plans again. You can change the limit at Max. disk utilization in Auto Plan Parameters of the Auto Plan window.
	The migration failed. Error code: XXXXXXXX.	Failed to migrate the volume due to an error indicated by Error code XXXXXXXX. For details and measures of the error, see the <i>Storage Navigator Messages</i> .
	Volume size mismatch between source and target volumes.	The size of the source volume is different from that of the target volume. The volume configuration might have been changed after the creation of the auto migration plan.
	The emulation type of the specified target volume is not supported.	The emulation type of the specified target volume is not supported. The volume configuration might have been changed after the creation of the auto migration plan.
	Failed to check the usage rate of the parity groups.	The check of the disk usage rate of the parity group including the source or target volume ended abnormally. For the measures, contact the service engineer.
	Failed to check the target volume.	The check of the target volume ended abnormally. For the measures, contact the service engineer.
	An unknown error occurred in checking the target volume.	An unknown error occurred in checking the target volume. For the measures, contact the service engineer.

C. Messages about creating auto migration plans (output when the creation started or ended)		
	yyyy/mm/dd hh:min The creation of auto migration plans started. (a) The creation of auto migration plans started. (a) Start time	
	The auto migration plan XXX was created. Source = LDKC:CU:LDEV in Group X-X, (a) (b) (c) Target = LDKC:CU:LDEV in Group X-X (d) (e) The auto migration plan XXX was created. (a) The number of the created auto migration plan (b) Source volume (c) Parity group of the source volume (d) Target volume (e) Parity group of the target volume	

C. Messages about creating auto migration plans (output when the creation started or ended)	
X auto migration plan(s) were created. (a) X auto migration plans were created. (a) Number of auto migration plans	
yyyy/mm/dd hh:min The creation of auto migration plans finished. (a) Creating auto migration plans finished. (a) Creation end time	
yyyy/mm/dd hh:min Failed to create auto migration plans. The statistics could not be obtained. (a) Volume Migration failed to create auto migration plans because the monitoring data could not be obtained. (a) Time when the creation of auto migration plans was failed.	
yyyy/mm/dd hh:min Failed to create auto migration plans. Cause: XXXXXXXXXXXX (a) (b) Error location: Class X Group X-X (c) (d) The creation of auto migration plans failed due to the cause shown in XXXXXXXXXXXX. Perform the appropriate measure described in "(b) Cause of failure" below, and then create auto migration plans again. (a) Time when creating auto migration plans failed (b) Cause of failure:	
Failed to specify a proper target volume.	Volume Migration could not specify a proper target volume for creating an auto migration plan. To find the cause of the failure, see the message output during the creation of auto migration plans. For details on the measure, see <i>"(d) Cause why a proper target volume could not be found"</i> in the second message of the <i>"D. Messages about creating each auto migration plan (output during the creation)"</i> section.
Insufficient statistics data in the sampling term.	There are insufficient statistics data for creating auto migration plans. Check the condition of monitoring data acquisition, and then collect the statistics data again. Otherwise, check the sampling term.
Insufficient reserved volumes.	There are insufficient reserved volumes for creating auto migration plans. If Volume Migration failed to obtain the specified information, check the reserve volumes in the Attribute window.
Failed to obtain XXXXXXXXXXXX.	Volume Migration failed to obtain the specified information or collected statistics data which are required for creating auto migration plans. XXXXXXXXXXXX shows the type of the information specified for migration-plan creation or collected statistics data. If Volume Migration failed to obtain the specified information, check the values specified in Auto Plan Parameters of the Auto Plan window. If Volume Migration failed to obtain the statistics data, check the status of the collected statistics data.

C. Messages about creating auto migration plans (output when the creation started or ended)		
	Failed to get XXXXXXXXXXXX.	Volume Migration failed to get the information from XXXXXXXXXXXX. XXXXXXXXXXXX shows the type of the information specified for migration-plan creation. For the measures, contact the service engineer.
	Failed to write in XXXXXXXXXXXX.	Volume Migration failed to write in the information of XXXXXXXXXXXX. XXXXXXXXXXXX shows the type of the information specified for migration-plan creation. For the measures, contact the service engineer.
	Invalid XXXXXXXXXXXX.	Volume Migration could not use XXXXXXXXXXXX because it was invalid. XXXXXXXXXXXX shows the type of the information specified for migration-plan creation or collected statistics data. If the specified information is invalid, check the values specified in Auto Plan Parameters of the Auto Plan window. If the statistics data are invalid, check the status of the collected statistics data.
	Memory allocation error.	Volume migration failed to allocate the memory for creating auto migration plans. For the measures, contact the service engineer.
(c) Class in which the failure occurred (d) Parity group in which the failure occurred		
Failed to create auto migration plans. The C:\dkc200\others\hihsmatm.dat file does not exist. The creation of auto migration plans failed because there was not a hihsmatm.dat file under the C:\dkc200\others folder. For the measures, contact the service engineer.		
In Group X-X, there are X invalid samples. These samples were not used for creating the auto migration plans. (a) (b) There are some invalid samples in the parity group X-X. Therefore, Volume Migration created auto migration plans without using these samples. (If there is one invalid sample, this message is shown as "...there is invalid sample. This sample was ...") (a) Parity group (b) Number of invalid samples		
There is too much invalid data in Group X-X LDKC:CU:LDEV. (a) (b) All the data about the volume were not used for creating the auto migration plans. There is too much invalid data in the volume LDKC:CU:LDEV of the parity group X-X. Therefore, Volume Migration created auto migration plans without using this volume. (a) Parity group (b) Volume		
yyyy/mm/dd hh:min The creation of auto migration plans has already been started. (a) Volume Migration has already started the creation of auto migration plans. (a) Creation start time		

D. Messages about creating each auto migration plan (output during the creation)	
yyyy/mm/dd hh:min Source candidate = LDKC:CU:LDEV in Group X-X (a) (b) (c) Target candidate = LDKCCU:LDEV in Group X-X XXXXXXXXXXXX (d) (e) (f) Volume Migration found a proper candidate for the target volume corresponding to the specified candidate for the source volume. (a) Start time of creating this auto migration plan. (b) Candidate for the source volume (c) Parity group of the candidate for the source volume (d) Candidate for the target volume (e) Parity group of the candidate for the target volume (f) Migration type	
Auto migration of a volume to a higher HDD class	Migrates the volume to a parity group in a higher HDD class
Auto migration of a volume to a lower HDD class	Migrates the volume to a parity group in a lower HDD class
Auto migration of a volume in the same HDD class	Migrates the volume to a parity group in the same HDD class
yyyy/mm/dd hh:min Source candidate = LDKC:CU:LDEV in Group X-X (a) (b) (c) Failed to specify a proper target volume for this source candidate. Cause: XXXXXXXXXXXX (d) Volume Migration could not find a proper candidate for the target volume corresponding to the specified candidate for the source volume. (a) Start time of creating this auto migration plan. (b) Candidate for the source volume (c) Parity group of the candidate for the source volume (d) Cause why a proper target volume could not be found:	
The parity group of the target candidate includes a volume that was migrated in the specified sampling term.	The parity group which the candidate for the target volume belongs to includes a volume that was migrated in the specified sampling term.
The disk usage rate of the parity group might exceed the class threshold.	If Volume Migration migrates the candidate for the source volume, the disk usage rate of the parity group to which the source volume will be migrated might exceed the HDD class threshold. Check the values of the HDD class thresholds in the Attribute window, and change the value, if required.
Insufficient reserved volumes.	There are insufficient reserved volumes. Check the number of reserved volumes and their locations.
E. Other messages	
yyyy/mm/dd hh:min The parameters for creating auto migration plans were initialized. (a) The parameters for creating auto migration plans were initialized. (a) Time when the parameters were initialized	
yyyy/mm/dd hh:min Failed to initialize the parameters for creating auto migration plans. (a) Volume Migration failed to initialize the parameters for creating auto migration plans. (a) Time when the initialization failed	

E. Other messages	
yyyy/mm/dd hh:min The auto migration plan(s) were deleted. (a) The auto migration plan(s) were deleted. (a) Deleted time	
yyyy/mm/dd hh:min Auto migration is not available. All the parity groups cannot be used for auto migration. (a) migration. All the parity groups in the storage system cannot be used as the auto-migration source or target. Auto migration of volumes is not available. This message is output when you start Volume Migration in the following conditions: 1. - All the volumes in the storage system are external volumes 2. - All the parity groups in the storage system have On-Demand settings (a) Output time of this message	
An internal logic error occurred. An internal logic error occurred in the SVP. For the measures, contact the service engineer.	

Table 6-2 Messages in the Migration Logs (displayed in the Action column)

Message	Meaning
Migration Started	The migration operation started.
Migration Completed	The migration operation completed successfully.
Migration Canceled by User	The migration operation was canceled by the user.
Migration Failed	The migration operation failed.
Migration Canceled by Controller	The migration operation was canceled by Volume Migration (e.g., volume usage rate exceeded specified maximum).
Migration Canceled by Controller (ShadowImage) Migration Canceled by Controller (TC (*1)) Migration Canceled by Controller (UR (*2)) Migration Canceled by Controller (XRC(*3)) Migration Canceled by Controller (CC(*4))	Because each program product started processing, the migration operation was canceled by Volume Migration.
*1: It indicates TrueCopy. *2: It indicates Universal Replicator. *3: It indicates Compatible XRC. *4: It indicates IBM 3990 Concurrent Copy.	

Server Priority Manager Operation

This chapter explains the following server priority manager operations:

- [Overview of Server Priority Manager Operations](#)
- [Port Tab Operations](#)
- [WWN Tab Operations](#)

Overview of Server Priority Manager Operations

Procedures for using Server Priority Manager depend on connection between host bus adapters and storage system ports.

If one-to-one connections are established between host bus adapters and ports, you specify the priority of I/O operations, upper limit value, and threshold value on each port. Because one port connects to one HBA, you can define the server priority by the port.

However, if many-to-many connections are established between host bus adapters and ports, you cannot define the server priority by the port, because one port can connect to multiple host bus adapters, and also multiple ports can connect to one host bus adapter. Therefore, in the many-to-many connection environment, you specify the priority of I/O operations and upper limit value on each host bus adapter. In this case, you specify one threshold value for the entire storage system.

If one-to-one connections are established between host bus adapters and ports, you use the **Port** tab of the **Server Priority Manager** window. If many-to-many connections are established between host bus adapters and ports, you use the **WWN** tab of the **Server Priority Manager** window. This section explains the operation procedures in each tab.



Note: *Host bus adapters* (HBAs) are adapters contained in hosts and serve as host ports for connecting the hosts and the storage system.

If One-to-One Connections Link HBAs and Ports

Figure 7-1 shows an example of a network in which each host bus adapter is connected to only one port on the storage system (Henceforth, this network is referred to as *network A*). Host bus adapters and the storage system ports are directly connected and are not connected via hubs and switches.




Figure 7-1 Network A (One-to-one Connections between HBAs and Ports)

If one-to-one connections are established between HBAs and ports, take the following major steps:

1. Set priority to ports on the storage system using the **Port** tab of the **Server Priority Manager** window.

In *network A*, the ports 1A and 1C are connected to high-priority production servers. The port 2A is connected to a low-priority development server. Therefore, the ports 1A and 1C should be given high priority, and the port 2A should be given low priority.

Figure 7-2 shows a portion of the **Server Priority Manager** window where the abbreviation **Prio.** indicates that the associated port is given high priority, and the abbreviation **Non-Prio.** indicates that the port is given low priority.



Note: Throughout this manual, the term *prioritized port* is used to refer to a high-priority port. The term *non-prioritized port* is used to refer to a low-priority port.

Port	Ave.[IOPS]	Peak[IOPS]	Attribute
CL1-A(E2)	-	-	Prio.
CL1-C(DC)	-	-	Prio.
CL2-A(EF)	-	-	Non-Prio.

Prio. indicates a prioritized port.
Non-Prio. indicates a non-prioritized port.

Figure 7-2 Priority Specified in the Server Priority Manager Window

2. Monitor traffic at ports. You must obtain statistics about traffic at each port on the storage system.

There are two types of traffic statistics: the *I/O rate* and the *transfer rate*. The I/O rate is the number of I/Os per second. The transfer rate is the size of data transferred between a host and the storage system. When you view traffic statistics in the window, you select either the I/O rate or the transfer rate. The **Port-LUN** tab of the **Performance Management** window lets you view a line graph that illustrates changes in traffic.

Figure 7-3 is a graph illustrating changes in the I/O rate for the three ports (1A, 1C, and 2A). According to the graph, the I/O rate for 1A and 1C was approximately 400 IO/s at first. The I/O rate for 2A was approximately 100 IO/s at first. However, as the I/O rate for 2A gradually increased from 100 IO/s to 200 IO/s, the I/O rate for 1A and 1C decreased from 400 IO/s to 200 IO/s. This fact indicates that the high-priority production servers have suffered lowered performance. If you were the network administrator, you probably would like to maintain the I/O rate for prioritized ports (1A and 1C) at 400 IO/s. To maintain the I/O rate at 400 IO/s, you must set an upper limit to the I/O rate for the port 2A.

For detailed information about monitoring traffic, see *Setting Priority for Ports on the Storage System and Analyzing Traffic Statistics*.




Figure 7-3 Traffic at Ports

3. Set an upper limit to traffic at the non-prioritized port. To prevent decline in I/O rates at prioritized ports, you set upper limit values to the I/O rate for non-prioritized ports.

When you set an upper limit for the first time, it is recommended that the upper limit be approximately 90 percent of the peak traffic. In *network A*, the peak I/O rate for the non-prioritized port (2A) is 200 IO/s. So, the recommended upper limit for 2A is 180 IO/s.

For details on how to set an upper limit, see Setting Priority for Ports on the Storage System.

4. Check the result of applying upper limit values. After applying upper limit values, you must measure traffic at ports. You must view traffic statistics for prioritized ports 1A and 1C to check whether the host performance is improved to a desirable level.

In *network A*, the desirable I/O rate for ports 1A and 1C is 400 IO/s. If the I/O rate reaches 400 IO/s, production server performance has reached to a desirable level. If production server performance is not improved to a desirable level, you can change the upper limit to a smaller value and then apply the new upper limit to the storage system. In *network A*, if the upper limit is set to 180 IO/s but the I/O rate for 1A and 1C is still below 400 IO/s, the administrator needs to change the upper limit until the I/O rate reaches 400 IO/s.

5. If necessary, apply a threshold. If you want to use threshold control, set threshold values in the **Port** tab in the **Server Priority Manager** window. You can set threshold values in either of the following ways:

- Set one threshold to each prioritized port

In *network A*, if you set a threshold of 200 IO/s to the port 1A and set a threshold of 100 IO/s to the port 1C, the upper limit on the non-prioritized port (2A) is disabled when either of the following conditions is satisfied:

- the I/O rate for the port 1A is 200 IO/s or lower
- the I/O rate for the port 1C is 100 IO/s or lower

- Set only one threshold to the entire storage system

In *network A*, if you set a threshold of 500 IO/s to the storage system, the upper limit on the non-prioritized port (2A) is disabled when the sum of the I/O rates for all prioritized ports (1A and 1C) goes below 500 IO/s.

For details on how to set a threshold, see Upper-Limit Control.

If Many-to-Many Connections Link HBAs and Ports

Figure 7-4 gives an example of a network in which a production server and a development server are connected to the storage system (Henceforth, this network is referred to as *network B*). The host bus adapter (wwn01) in the production server is connected to four ports (1A, 1C, 2A and 2C). The host bus adapters (wwn02 and wwn03) in the development server are also connected to the four ports.




Figure 7-4 Network B (Many-to-Many Connections are Established between HBAs and Ports)

If many-to-many connections are established between HBAs and ports, take the following major steps:

1. Find WWNs of host bus adapters. Before using Server Priority Manager, you must find the WWN (Worldwide Name) of each host bus adapter in host servers. WWNs are 16-digit hexadecimal numbers used to identify host bus adapters. For details on how to find WWNs, see the *LUN Manager User's Guide*.
2. Ensure that all host bus adapters connected to ports in the storage system are monitored. Use the **WWN** tab of the **Server Priority Manager** window to define which port is connected to which host bus adapter. Place host bus adapters connected to each port below the **Monitor** icons.

In *network B*, each of the four ports is connected to three host bus adapters (wwn01, wwn02, and wwn03). Place the host bus adapter icons of wwn01, wwn02, and wwn03 below the **Monitor** icons for all the four port icons.

The resulting definitions on the window are as follows:




Figure 7-5 Specifying Host Bus Adapters to be Monitored

For more detailed instruction, see Setting Priority for Ports on the Storage System.



Note: Server Priority Manager is unable to monitor and control the performance of hosts whose host bus adapters are placed below the **Non-Monitor** icon.

3. Set priority to host bus adapters using the **WWN** tab of the **Server Priority Manager** window.

In *network B*, the production server is given high priority and the development server is given low priority. If your network is configured as in Figure 7-4, you must give high priority to wwn01 and also give low priority to wwn02 and wwn03.

To give priority to the three host bus adapters, take the following steps:

- a. In the **WWN** tab, select one of the four ports that the HBAs are connected to (i.e. ports 1A, 1C, 2A, and 2C).
- b. Set **Prio.** to wwn01. Also, set **Non-Prio.** to wwn02 and wwn03.




Figure 7-6 Priority Specified in the Server Priority Manager Window



Note: Throughout this manual, the term *prioritized WWN* refers to a high-priority host bus adapter (for example, wwn01). The term *non-prioritized port* refers to a low-priority host bus adapter (for example, wwn02 and wwn03).

4. Monitor traffic between host bus adapter and ports. You must obtain statistics about traffic between host bus adapter and ports.

There are two types of traffic statistics: the *I/O rate* and the *transfer rate*. The I/O rate is the number of I/Os per second. The transfer rate is the size of data transferred between a host and the storage system. When you view traffic statistics in the window, you select either the I/O rate or the transfer rate.

If your network is configured as *network B*, you must do the following:

- Measure traffic between the port 1A and the three host bus adapters (wwn01, wwn02 and wwn03).
- Measure traffic between the port 1C and the three host bus adapters (wwn01, wwn02 and wwn03).
- Measure traffic between the port 2A and the three host bus adapters (wwn01, wwn02 and wwn03).
- Measure traffic between the port 2C and the three host bus adapters (wwn01, wwn02 and wwn03).

Figure 7-7 illustrates a graph that describes the I/O rate at the paths between each port and the host bus adapters. According to the graph, the I/O rate at the path between 1A and the prioritized WWN (wwn01) was approximately 400 IO/s at first. The I/O rate at the path between 1A and the non-prioritized WWNs (wwn02 and wwn03) was approximately 100 IO/s at first. However, as the I/O rate for non-prioritized WWNs (wwn02 and wwn03) gradually increased from 100 IO/s to 200 IO/s, the I/O rate for the prioritized WWN (wwn01) decreased from 400 IO/s to 200 IO/s. This indicates that the high-priority production server has degraded. If you were the network administrator, you probably would like to maintain the I/O rate for the prioritized WWN (wwn01) at 400 IO/s.

For more information about monitoring traffic, see *Setting Priority for Host Bus Adapters and Analyzing Traffic Statistics*.




Figure 7-7 Traffic at Ports

5. Set an upper limit to traffic between ports and the non-prioritized WWN to prevent decline in I/O rates at prioritized WWNs.

When you set an upper limit for the first time, the upper limit should be approximately 90 percent of the peak traffic level.

In *network B*, the peak I/O rate at the paths between port 1A and the non-prioritized WWNs (wwn02 and wwn03) is 200 IO/s. The peak I/O rate at the paths between port 1C and the non-prioritized WWNs is also 200 IO/s. The peak I/O rate at the paths between port 2A and the non-prioritized WWNs is also 200 IO/s. The peak I/O rate at the paths between port 2C and the non-prioritized WWNs is also 200 IO/s. So, the recommended upper limit for the non-prioritized WWNs is 720 IO/s (= 200 × 4 × 0.90).

If your network is configured as in Figure 7-4, you must do the following in this order:

- a. In the **WWN** tab, select one of the four ports that the HBAs are connected to (i.e. ports 1A, 1C, 2A, and 2C).
- b. Set an upper limit to the non-prioritized WWNs (wwn02 and wwn03).

Figure 7-8 is the result of setting the upper limit of 720 IO/s to the paths between 1A and the non-prioritized WWNs. For details on how to set an upper limit, see Setting Priority for Ports on the Storage System.

Port WWN		IOPS									
		WWN	SPM Name	...	Pe...	WV...	Attribute	Upper			
								IOPS	MB/s		
		0000000000000001	wwn01	...	15	20	Prio.	0	-		
		0000000000000002	wwn02	...	16	21	Non-Prio.	720	-		
		0000000000000003	wwn03	...	17	22	Non-Prio.	720	-		

Figure 7-8 Setting Upper Limits

6. Check the result of applying upper limit values. After applying upper limit values, you must measure traffic at ports. View traffic statistics for the prioritized WWN to check whether the host performance is improved to a desirable level.

In *network B*, the desirable I/O rate for the prioritized WWN is 400 IO/s. If the I/O rate reaches 400 IO/s, production server performance has reached to a desirable level. If production server performance is not improved to a desirable level, you can change the upper limit to a smaller value and then apply the new upper limit to the storage system. In *network B*, if the upper limit is set to 720 IO/s but the I/O rate for wwn01 is still below 400 IO/s, the administrator needs to change the upper limit until the I/O rate reaches 400 IO/s.



Note: If an upper limit of the non-prioritized WWN is set to zero or nearly zero, I/O performance might be lowered. If I/O performance is lowered, the host cannot be connected to the storage system in some cases.

7. If necessary, apply a threshold. If you want to use threshold control, set a threshold in the **WWN** tab in the **Server Priority Manager** window.

In the **WWN** tab, you can specify only one threshold for the entire storage system, regardless of the number of prioritized WWNs. For example, if there are three prioritized WWNs in the network and the threshold is 100 IO/s, the upper limit on the non-prioritized WWNs is disabled when the sum of the I/O rates for all prioritized WWNs goes below 100 IO/s.

For details on how to set a threshold, see [Setting a Threshold](#).



Caution: If you enter zero (0) in a cell to disable the upper limit, the cell displays a hyphen (-) and the threshold for the prioritized port becomes ineffective. If the thresholds of all the prioritized ports are ineffective, threshold control will not be performed but upper limit control will be performed. If you set thresholds for multiple prioritized ports and the I/O rate or transfer rate becomes below the threshold at one of the prioritized ports, threshold control takes effect and the upper limits of the non-prioritized ports are disabled. The following table shows the relationship between the thresholds and the upper limits.

Table 7-1 Relationship between the Thresholds of the Prioritized WWN and the Upper Limits of the Non-prioritized WWN

Threshold Settings	A Number Other Than Zero Is Set to The Upper Limit of The Non-prioritized WWN	Zero Is Set to The Upper Limit of The Non-prioritized WWN
Threshold Is Set to The Prioritized WWN	When thresholds are set to multiple prioritized WWNs, depending on the transfer rate, the following controls are executed. ▪ If I/O rate or transfer rate exceeds the threshold in any prioritized WWN, upper limits of all the non-prioritized WWNs take effect. ▪ If I/O rate or transfer rate goes below the threshold in any prioritized WWN, upper limits of all the non-prioritized WWNs do not take effect.	The threshold control of the prioritized WWN is not executed.
Threshold Is Not Set to The Prioritized WWN	The specified upper limit always takes effect.	

Starting Server Priority Manager

To start Server Priority Manager, take the following steps.

To start Server Priority Manager:

1. Ensure that the Storage Navigator main window is in **Modify** mode (i.e., with the background color of the pen tip icon  showing light yellow). If the background color is gray, the window is in **View** mode and you must change it to **Modify** mode by taking the following steps:
 - a. Verify whether the background color of the lock icon  is blue.
 If the background color is red , you will not be able to switch from **View** mode to **Modify** mode. Wait for a while and click the  button. When the background color turns blue, you can go to the next step.
 - b. Select the  (pen tip) icon.
 A prompt asks whether you want to change the mode.
 - c. Select **OK** to close the message.
 The background color of the icon changes to light yellow (). The mode changes to **Modify** mode. The background color of the lock icon becomes red ().
2. Ensure that the **WWN** tab or the **Port-LUN** tab is active in the **Performance Management** window.
3. Ensure that the **Real Time** option is cleared.
 You cannot start Server Priority Manager in real-time mode.

4. Select the **SPM** button.

The **Server Priority Manager** window is displayed.

- Use the **Port** tab if one-to-one connection is established between host bus adapters and storage system ports. For details on operations in the **Port** tab, see Port Tab Operations.
- Use the **WWN** tab if many-to-many connection is established between host bus adapters and storage system ports. For details on operations in the **WWN** tab, see WWN Tab Operations.

Port Tab Operations

If one-to-one connections are established between host bus adapters (HBAs) and storage system ports, use the **Port** tab in the **Server Priority Manager** window to do the following:

- Analyze traffic statistics
- Measure traffic between host bus adapters and storage system ports
- Set priority to ports on the storage system
- Set an upper limit to traffic at each non-prioritized port
- Set a threshold to the storage system or to each prioritized port, if necessary

If one-to-one connections are established between host bus adapters and ports, you should specify the priority of I/O operations on each port. You can specify the upper limit values on the non-prioritized ports, and if necessary, you can also specify the threshold values on the prioritized ports. Moreover, you can use one threshold value applied for the entire storage system.

For details on the system configuration of one-to-one connections between host bus adapters and ports, see *If One-to-One Connections Link HBAs and Ports*. This section explains operation procedures you can perform for ports and the entire storage system.

Analyzing Traffic Statistics

The traffic statistics reveal the number of I/Os that have been made via ports. The traffic statistics also reveal the amount of data that have been transferred via ports. You must analyze the traffic statistics to determine upper limit values that should be applied to I/O rates or transfer rates for non-prioritized ports.

The following is the procedure for using the **Server Priority Manager** window to analyze traffic statistics. You can also use the **Performance Management** window to analyze traffic statistics. Performance Monitor can display a line graph that indicates changes in traffic (For details, see *Viewing I/O Rates for Ports* and *Viewing Transfer Rates for Ports*).

To analyze traffic statistics:

1. Start Server Priority Manager.
The **Server Priority Manager** window is displayed.
2. Select the Port tab.
3. Select **All** from the drop-down list at the top right corner of the window.
4. Do one of the following:

- To analyze I/O rates, select **IOPS** from the drop-down list at the upper left corner of the list.
- To analyze transfer rates, select **MB/s** from the drop-down list at the upper left corner of the list.

The list displays traffic statistics (i.e. the average and peak I/O rates or transfer rates) of the ports.

5. Analyze the information in the list and then determine upper limit values that should be applied to non-prioritized ports. If necessary, determine threshold values that should be applied to prioritized ports. For details on the upper limit values and threshold values, see If One-to-One Connections Link HBAs and Ports.

Setting Priority for Ports on the Storage System

If one-to-one connection is established between HBAs and ports, you need to measure traffic between high-priority HBAs and prioritized ports. You also need to measure traffic between low-priority HBAs and non-prioritized ports.

Prioritized ports are ports on which the processing has high priority and *non-prioritized ports* are ports on which the processing has low priority. Specify a port that connects to a high-priority host bus adapter as a prioritized port. Specify a port that connects to a low-priority host bus adapter as a non-prioritized port.

To set priority to ports on the storage system:

1. Start Server Priority Manager.

The **Server Priority Manager** window is displayed.

2. Ensure that the Port tab is displayed.
3. Select **All** from the drop-down list at the top right corner of the window.
4. Right-click a high-priority port and then select **Non-Prio ->> Prio** from the pop-up menu. If there is more than one high-priority port, repeat this operation.

The **Attribute** column displays **Prio**.

5. Right-click a low-priority port and then select **Prio ->> Non-Prio** from the pop-up menu. If there is more than one low-priority port, repeat this operation.

The **Attribute** column displays **Non-Prio**.

You must set upper limit values for the **Non-prio.** specified ports. For detail about the setting of upper limit values, see Setting Upper-Limit Values to Traffic at Non-prioritized Ports.

6. Click the **Apply** button.

The settings on the window are applied to the storage system.

After priority has been set, you can implement the procedure for measuring traffic (I/O rates and transfer rates) (see Starting and Stopping Storage System Monitoring)

Setting Upper-Limit Values to Traffic at Non-prioritized Ports

After you analyze traffic statistics, you must set upper limit values to I/O rates or transfer rates for non-prioritized ports. Upper limit values for I/O rates are used to suppress the number of I/Os from the low priority host servers and thus provide better performance for high-priority host servers. Upper limit values for transfer rates are used to suppress the amount of data that should be transferred between the storage system and the low priority ports, and thus provide better performance for high-priority host servers.

To limit the I/O rate or transfer rate of a non-prioritized port:

1. Start Server Priority Manager.

The **Server Priority Manager** window is displayed.

2. Select the **Port** tab.
3. Do one of the following:
 - To limit the I/O rate for the non-prioritized port, select **IOPS** from the drop-down list at the upper left corner of the list.
 - To limit the transfer rate for the non-prioritized port, select **MB/s** from the drop-down list at the upper left corner of the list.
4. Locate the non-prioritized port in the list.



Notes:

- The **Attribute** column of the list indicates whether ports are prioritized or non-prioritized.
- If you cannot find any non-prioritized port in the list, check the drop-down list at the top right corner of the window. If the drop-down list displays **Prioritize**, select **All** or **Non-Prioritize** from the drop-down list.

-
5. Do one of the following:
 - To limit the I/O rate for the non-prioritized port, double-click the desired cell in the **IOPS** column in **Upper**. Next, enter the upper limit value in the cell.
 - To limit the transfer rate for the non-prioritized port, double-click the desired cell in the **MB/s** column in **Upper**. Next, enter the upper limit value in the cell.

In the list, either of **IOPS** or **MB/s** column is activated depending on the rate selected at step 3 above. You can use either of them to specify the upper limit value for one port. You can specify different types of rates (**IOPS** or **MB/s**) for the upper limit values of different non-prioritized ports.

The upper limit value that you entered is displayed in blue.

6. Select the **Apply** button.

The settings in the window are applied to the storage system. The upper limit value that you entered turns black.



Note: If an upper limit of the non-prioritized WWN is set to zero or nearly zero, I/O performance might be lowered. If I/O performance is lowered, the host cannot be connected to the storage system in some cases.

Setting a Threshold

If threshold control is used, upper limit control is automatically disabled when traffic between production servers and the storage system is reduced to a specified level. For details, see Upper-Limit Control and If One-to-One Connections Link HBAs and Ports.

If one-to-one connections are established between HBAs and ports, you can set the threshold value by the following two ways:

- Set a threshold value for each prioritized port
- Set one threshold value for the entire storage system

The procedures for these operations are explained below.

To set threshold values to traffic at prioritized ports:

1. Start Server Priority Manager.

The **Server Priority Manager** window is displayed.

2. Select the Port tab.
3. To set a threshold value for each prioritized port, select the type of rates for the threshold value from the drop-down list at the upper left corner of the list.
 - To use the I/O rates for the threshold value, select **IOPS**.
 - To use the transfer rates for the threshold value, select **MB/s**.



Note: If you want to set one threshold value for the entire storage system, this step is unnecessary.

4. Do one of the following:
 - To set a threshold to each prioritized port, locate the desired prioritized port, which is indicated by **Prio.** in the **Attribute** column. Next, double-click the cell in the **IOPS** or **MB/s** column in **Threshold**, and then enter the threshold value. In the list, either of **IOPS** or **MB/s** column is activated depending on the rate selected at step 3 above. Repeat this operation to set the thresholds for all the prioritized ports. You can use different types of rates (**IOPS** or **MB/s**) for thresholds of different prioritized ports.



Caution: If you enter zero (0) in a cell to disable the upper limit, the cell displays a hyphen (-) and the threshold for the prioritized port becomes ineffective. If the thresholds of all the prioritized ports are ineffective, threshold control will not be performed but upper limit control will be performed. If you set thresholds for multiple prioritized ports and the I/O rate or transfer rate becomes below the threshold at one of the prioritized port, threshold control works in the entire storage system and the upper limits of the non-prioritized ports are disabled. The following table shows the relationship between the thresholds and the upper limits.

Table 7-2 Relationship between the Thresholds of the Prioritized Port and the Upper Limits of the Non-prioritized Port

Thresholds Settings	A Number Other Than Zero Is Set to The Upper Limit of The Non-prioritized port	Zero Is Set to The Upper Limit of The Non-prioritized port
Threshold Is Set to The Prioritized port	When thresholds are set to multiple prioritized ports, depending on the transfer rate, following controls are executed. ▪ If I/O rate or transfer rate exceeds the threshold in any prioritized port, upper limits of all the non-prioritized ports take effect. ▪ If I/O rate or transfer rate goes below the threshold in any prioritized port, upper limits of all the non-prioritized ports do not take effect.	The threshold control of the prioritized port is not executed.
Threshold Is Not Set to The Prioritized port	The specified upper limit always takes effect.	

- To set one threshold to the entire storage system, select the **All Thresholds** check box. Next, select **IOPS** or **MB/s** from the drop-down list of right side in **All Thresholds** and enter the threshold value in the text box. Even if the types of rates for upper limit values and the threshold are different, the threshold control can work for all the non-prioritized ports.
5. Select the **Apply** button.
- The settings in the window are applied to the storage system.

WWN Tab Operations

If many-to-many connections are established between host bus adapters (HBAs) and storage system ports, you use the **WWN** tab in the **Server Priority Manager** window to do the following:

- Make all the traffics between host bus adapters and ports monitored
- Analyze traffic statistics
- Measure traffic between host bus adapters and storage system ports
- Set priority to host bus adapters
- Set an upper limit on traffic at non-prioritized WWNs
- Set a threshold, if necessary

If many-to-many connections are established between host bus adapters and ports, you should specify the priority of I/O operations on each host bus adapter. You can specify the upper limit values on the non-prioritized WWNs. If necessary, you can set one threshold value applied for the entire storage system. When many-to-many connections are established between host bus adapters and ports, you cannot set individual thresholds for prioritized WWNs.

For the system configuration of many-to-many connections between host bus adapters and ports, see [If Many-to-Many Connections Link HBAs and Ports](#).

For details on the system configuration of many-to-many connections between host bus adapters and ports, see [If Many-to-Many Connections Link HBAs and Ports](#). This section explains operation procedures you can perform for host bus adapters and the entire storage system.

Monitoring All Traffic between HBAs and Ports

When many-to-many connections are established between host bus adapters (HBAs) and ports, you should make sure that all the traffic between HBAs and ports is monitored.

To make all the traffics between host bus adapters and ports monitored:

1. Start Server Priority Manager.
The **Server Priority Manager** window is displayed.
2. Ensure that the **WWN** tab is displayed.
The two trees are displayed in the left side of the **WWN** tab. The upper-left tree lists ports in the storage system.
3. Select **All** from the drop-down list at the top right corner of the window.
4. In the upper-left tree, double-click a port.

5. Double-click **Non-Monitor** below the specified port.
If there are any host bus adapters whose traffics with the specified port are not monitored, those host bus adapters are displayed below **Non-Monitor**.
6. Right-click **Monitor** and then select **Add WWN** from the pop-up menu.
The **Add WWN** window is displayed. This window lets you add a WWN of a host bus adapter to **Monitor**.
7. In the **Add WWN** window, specify the WWN and the SPM name.
Expand the **WWN** drop-down list to show the WWNs of the host bus adapters that are connected to the port but are not monitored. These host bus adapters are the same as that displayed in step 5. From that drop-down list, select a WWN and specify the SPM name. You can specify up to 16 characters for an SPM name.



Note: We recommend that you specify the same names for the SPM names and the nicknames of the host bus adapters for convenience of host bus adapter management. *Nicknames* are aliases of host bus adapters defined by LUN Manager. In the **Port-LUN** tab of Performance Monitor, not only SPM names but also nicknames are displayed as the aliases of host bus adapters (WWNs) in the list. Therefore, if you specify both the same aliases, the management of the host bus adapters is easier.

8. Select the **OK** button.
The selected WWN (of the host bus adapter) is moved from **Non-Monitor** to **Monitor**.



Note: If the specified host bus adapter is also connected to other ports, after selecting the **OK** button, a message will appear that asks you whether to change the settings of that host bus adapter for other ports, too. Make the same setting for all the ports.

9. Repeat step 6 to 8 to move all the host bus adapters displayed below **Non-Monitor** to below **Monitor**.



Note: If you disconnect a host that has been connected via a cable to your storage system or change the port to the another port of the host, the WWN for the host will remain in the WWN list of the WWN tab. If you want to delete the WWN from the WWN list, you can delete the WWN by using LUN Manager. For detail information of the deleting old WWNs from the WWN list, see the *LUN Manager User's Guide*.

10. Click the **Apply** button in the **Server Priority Manager** window.
The settings on the window are applied to the storage system.



Note: If you are using Windows, you can drag and drop the desired WWNs from **Non-Monitor** to **Monitor**. When you drop a WWN to **Monitor**, the **Add WWN** window is displayed in which you can specify the SPM name only.




Figure 7-9 Add WWN Window

If you add a port or host bus adapter to the storage system after the settings above, the traffics about connections to the newly added port or host bus adapter will not be monitored. In this case, follow the procedure above again to make all the traffics between host bus adapters and ports monitored.

Up to 32 host bus adapters (WWNs) can be monitored for one port. If more than 32 host bus adapters are connected to one port, the traffics about some host bus adapters will be obliged to be excluded from the monitoring target. Consider the intended use of each host and move the host bus adapters which you think not necessary to be monitored to **Non-Monitor** by the following steps.

To exclude traffic between a host bus adapter and a port from the monitoring target:

1. Start Server Priority Manager and ensure that the **WWN** tab is displayed.
2. Select **All** from the drop-down list at the top right corner of the window.
3. In the upper-left tree, double-click a port to which more than 32 host bus adapters are connected.
4. Double-click **Monitor** below the specified port.
5. Right-click the WWN of a host bus adapter you want to exclude from the monitoring target and then select **Delete WWN** from the pop-up menu.

If you are using Windows, you can drag and drop the desired WWNs from **Monitor** to **Non-Monitor**.



Note:

- If the selected host bus adapter is connected to multiple ports, when you select the host bus adapter and select the **Delete WWN** pop-up menu, a message will appear that asks you whether to move the host bus adapter from **Monitor** to **Non-Monitor** below all other ports, too.
 - If the selected host bus adapter is contained in an SPM group, a message will appear that tell you to delete the host bus adapter from the SPM group on ahead. You cannot move a host bus adapter which is contained in an SPM group from **Monitor** to **Non-Monitor**. For details on how to delete a host bus adapter from an SPM group, see Deleting an HBA from an SPM Group.
-

6. Click **OK** for the confirmation message that asks you whether to delete the host bus adapter.

The deleted host bus adapter (WWN) is moved from **Monitor** to **Non-Monitor**.

7. Click the **Apply** button in the **Server Priority Manager** window.
The settings on the window are applied to the storage system.

Analyzing Traffic Statistics

The traffic statistics reveal the number of I/Os that have been made via ports from HBAs. They also reveal the amount of data that have been transferred between ports and HBAs. You must analyze the traffic statistics to determine upper limit values that should be applied to I/O rates or transfer rates for low-priority HBAs.

The following is the procedure for using the **Server Priority Manager** window to analyze traffic statistics. You can also use the **Performance Management** window to analyze traffic statistics. Performance Monitor can display a line graph that indicates changes in traffic (For details, see Viewing I/O Rates for Disks and Viewing Transfer Rates for Disks).

To analyze traffic statistics:

1. Start Server Priority Manager.

The **Server Priority Manager** window is displayed.

2. Select the WWN tab.
3. Select **All** from the drop-down list at the top right corner of the window.
4. Do one of the following:
 - To analyze I/O rates, select **IOPS** from the drop-down list at the upper left corner of the list.
 - To analyze transfer rates, select **MB/s** from the drop-down list at the upper left corner of the list.
5. Below the **Subsystem** folder in the upper-left tree, click the icon of the port whose traffic statistics you want to collect.

The list displays traffic statistics (I/O rates or transfer rates) about the host bus adapters that connect to the selected port.

The following two types of traffic are shown. The traffic has attributes including the average and maximum values.

- Traffic between the host bus adapter and the selected port (shown in **Per Port**)
- Sum of the traffic between the host bus adapter and all the ports connected to the host bus adapter (shown in **WWN Total**)



Notes:

- The traffic statistics only about the host bus adapters below **Monitor** appear in the list.
 - The **WWN Total** traffic statistics will also be displayed in the list when you click an icon in the lower-left tree. If you click the **Subsystem** folder in the lower-left tree, the sum of the traffic of the host bus adapters registered on each SPM group is displayed. For details on SPM groups, see Grouping Host Bus Adapters.
-

6. Analyze the information in the list and then determine upper limit values that should be applied to non-prioritized WWNs. If necessary, determine threshold values that should be applied to prioritized WWNs. For details, see If Many-to-Many Connections Link HBAs and Ports.

Setting Priority for Host Bus Adapters

If many-to-many connection is established between host bus adapters (HBAs) and ports, you need to define the priority of WWNs, measure traffic between each HBA and the port that the HBA is connected to, and analyze the traffics.

The host bus adapters (HBAs) are divided into two types: *Prioritized WWNs* and *non-prioritized WWNs*. Prioritized WWNs are the host bus adapters that are used for the high-priority processing, and *non-prioritized WWNs* are the host bus adapters that are used for the low-priority processing. Specify a host bus adapter existed in a server, on which the high-priority processing is performed, as a prioritized WWNs. Specify a host bus adapter existed in a server, on which the low-priority processing is performed, as a non-prioritized WWNs.

To set priority to host bus adapters:

1. Start Server Priority Manager.
The **Server Priority Manager** window is displayed.
2. Select the **WWN** tab.
3. Select **All** from the drop-down list at the top right corner of the window.
4. In the upper-left tree, double-click a port.
5. Double-click **Monitor**, which is displayed below the specified port.
6. Check to see if all the WWNs of the host bus adapters to be controlled by using Server Priority Manager appear below **Monitor**.

If some of the WWNs are missing, use the procedure in Monitoring All Traffic between HBAs and Ports to move all WWNs to below **Monitor**.

7. Click **Monitor** to display the information of the host bus adapters that are monitored in the list on the right of the tree.
8. Right-click a host bus adapter (WWN) in the list and then select **Non-Prio ->> Prio** from the pop-up menu.

The **Attribute** column of the selected WWN in the list displays **Prio**. If you want to specify more than one prioritized WWN, repeat this operation.



Note: You cannot change the priority of a WWN which is contained in an SPM group. For details on how to change the attribute of a WWN contained in an SPM group, see Switching Priority of an SPM Group.

9. Right-click a host bus adapter (WWN) in the list and then select **Prio ->> Non-Prio** from the pop-up menu.

The **Attribute** column of the selected WWN in the list displays **Non-Prio**. If you want to specify more than one non-prioritized WWN, repeat this operation.



Note: You cannot change the priority of a WWN which is contained in an SPM group. For details on how to change the attribute of a WWN contained in an SPM group, see Switching Priority of an SPM Group.

You must set upper limit values for the **Non-prio**. specified ports. For details, see Setting Upper-Limit Values for Non-Prioritized WWNs.

10. Repeat steps 4 to 9 for ports (except for the port selected in step 9).

If one host bus adapter is connected to multiple ports and you specify the priority of the host bus adapter for one port, the specified priority will be also applied to the host bus adapter settings for other connected ports automatically.

11. Select the **Apply** button in the **Server Priority Manager** window.
The settings on the window are applied to the storage system.

Follow the instructions in Starting and Stopping Storage System Monitoring to measure traffic (i.e.. I/O rates and transfer rates).

Setting Upper-Limit Values for Non-Prioritized WWNs

After you analyze traffic statistics about prioritized WWNs and non-prioritized WWNs, you must set upper limit values to I/O rates or transfer rates for non-prioritized WWNs. Upper limit values for I/O rates are used to suppress the number of I/Os from the low priority host servers and thus provide better performance for high-priority host servers. Upper limit values for transfer rates are used to suppress the amount of data that should be transferred between the storage system and the low priority ports, thus providing better performance for high-priority host servers.

Tip: To set the same upper limit value to more than one non-prioritized WWN, use an SPM group. For details on SPM groups, see Grouping Host Bus Adapters.

To limit the I/O rate or transfer rate of a non-prioritized WWN:

1. Start Server Priority Manager.

The **Server Priority Manager** window is displayed.

2. Ensure that the WWN tab is displayed.

3. Do one of the following:

- To limit the I/O rate of the non-prioritized WWN, select **IOPS** from the drop-down list at the upper left corner of the list.
- To limit the transfer rate of the non-prioritized WWN, select **MB/s** from the drop-down list at the upper left corner of the list.

4. In the upper-left tree, click the icon of the port whose traffic you want to limit below the **Subsystem** folder.

The information about the host bus adapters which connect to the selected port is displayed in the list.

5. Locate the non-prioritized WWN in the list.



Notes:

- The **Attribute** column of the list indicates whether WWNs are prioritized or non-prioritized. The **Attribute** column of a non-prioritized WWN displays **Non-Prio**.
- If you cannot find any non-prioritized WWN in the list, check the drop-down list at the top right corner of the window. If the drop-down list displays **Prioritize**, select **All** or **Non-Prioritize** from the drop-down list.

-
6. Do one of the following:

- To limit the I/O rate of the non-prioritized WWN, double-click the desired cell in the **IOPS** column in **Upper**. Next, enter the upper limit value in the cell.
- To limit the transfer rate of the non-prioritized WWN, double-click the desired cell in the **MB/s** column in **Upper**. Next, enter the upper limit value in the cell.

In the list, either of the **IOPS** cells or **MB/s** cells are activated depending on the rate you specified in step 3. You can specify the limit value by using either of the I/O rate or transfer rate for each host bus adapter. The upper limit value that you entered is displayed in blue. It is allowed that you specify upper limit values by using the I/O rate for some host bus adapters and specify them by using the transfer rate for the other host bus adapters.



Notes:

- You cannot specify or change the upper limit value of a host bus adapter which is contained in an SPM group. The upper limit value of such a host bus adapter is defined by the SPM group settings. For details on how to specify an upper limit value for an SPM group, see Setting an Upper-Limit Value to HBAs in an SPM Group.
- If one host bus adapter is connected to multiple ports and you specify an upper limit value of the host bus adapter for one port, the specified upper limit value will be applied to the host bus adapter settings for other connected ports automatically.

-
7. Select the **Apply** button.

The settings in the window are applied to the storage system. The upper limit value that you entered turns black.

Setting a Threshold

If threshold control is used, upper limit control is automatically disabled when traffic between production servers and the storage system is reduced to a specified level. For details, see Upper-Limit Control and If Many-to-Many Connections Link HBAs and Ports.

If many-to-many connections are established between host bus adapters and storage system ports, you can set one threshold value for the entire storage system. In this environment, you cannot set individual threshold values for each prioritized WWN.

To set a threshold value:

1. Start Server Priority Manager.

The **Server Priority Manager** window is displayed.

2. Select the **WWN** tab.
3. Select the **All Thresholds** check box.
4. Select **IOPS** or **MB/s** from the drop-down list of **All Thresholds**, and do one of the following:
 - To specify the threshold value by using the I/O rate, select **IOPS** from the drop-down list below the check box.
 - To specify the threshold value by using the transfer rate, select **MB/s** from the drop-down list below the check box.

Even if the types of rates differ between the upper limit values and the threshold value, the threshold control is effective for all the non-prioritized WWNs.

5. Enter the threshold in the text box of **All Thresholds**.
6. Select the **Apply** button.

The settings in the window are applied to the storage system.

Changing the SPM Name of a Host Bus Adapter

The **Server Priority Manager** window allows you to assign an *SPM name* to a host bus adapter (HBA). Although you can identify HBAs by WWNs (Worldwide Names), you will be able to identify HBAs more easily if you assign SPM names. WWNs are 16-digit hexadecimal numbers and cannot be changed. However, SPM names should not necessarily be 16-digit hexadecimal numbers and can be changed.

The following is the procedure for changing an already assigned SPM name. For details on how to assign an SPM name, see *Monitoring All Traffic between HBAs and Ports*.

To change an SPM name:

1. Start Server Priority Manager.

The **Server Priority Manager** window is displayed.

2. Ensure that the WWN tab is displayed.
3. In the upper-left tree, select a host bus adapter (HBA) from below **Monitor** and then right-click the selection.
4. From the pop-up menu, select **Change WWN and SPM Name**.

The **Change WWN and SPM Name** window is displayed.

5. Enter a new SPM name in the **SPM Name** box and then select **OK**.



Note: You can use up to 16 characters for an SPM name.

6. In the **Server Priority Manager** window, select the **Apply** button.

The settings in the window are applied to the storage system.




Figure 7-10 Change WWN and SPM Name Window

Replacing a Host Bus Adapter

If a host bus adapter fails, you will need to replace the adapter with a new one. After you finish replacement, you will need to delete the old host bus adapter from the **Server Priority Manager** window and then register the new host bus adapter.



Note: When you add a new host bus adapter rather than replacing an old one, the WWN of the added host bus adapter is automatically displayed below **Non-Monitor** for the connected port in the list.

Follow the procedure below to remove the old adapter and register a new adapter quickly and easily.

To register a new host bus adapter after replacement:

1. Start Server Priority Manager.
The **Server Priority Manager** window is displayed.
2. Select the WWN tab.
3. In the upper-left tree, select the old host bus adapter (🖨️) from below **Monitor** and then right-click the selection.
4. From the pop-up menu, select **Change WWN and SPM Name**.
The **Change WWN and SPM Name** window is displayed.
5. Enter the WWN of the new host bus adapter in the **WWN** combo box.
You can select the WWN of the newly connected host bus adapter in the **WWN** combo box. If you are using Windows, you can drag the WWN of the new host bus adapter displayed below **Non-Monitor** and drop it to **Monitor**.
6. If necessary, enter a new SPM name in the **SPM Name** box.



Note: You can use up to 16 characters for an SPM name.

7. Select **OK** to close the **Change WWN and SPM Name** window.
8. In the **Server Priority Manager** window, select the **Apply** button.
The settings in the window are applied to the storage system.

Grouping Host Bus Adapters

Server Priority Manager allows you to create an *SPM group* to contain multiple host bus adapters. All the host bus adapters (HBAs) in one SPM group must be of the same priority. Prioritized WWNs (i.e. high-priority HBAs) and non-prioritized WWNs (i.e. low-priority HBAs) cannot be mixed in the same group.

You can use an SPM group to switch priority of multiple HBAs from *prioritized* to *non-prioritized*, or vice versa. You can also use an SPM group to set the same upper limit value to all the HBAs in the group.

Containing Multiple HBAs in an SPM Group

A host bus adapter can be contained in only one SPM group. To create an SPM group and contain multiple host bus adapters in the group, take the following steps:

To contain multiple HBAs in an SPM group:

1. Start Server Priority Manager.
The **Server Priority Manager** window is displayed.
2. Select the **WWN** tab.
3. In the lower-left tree, select and right-click the Subsystem folder.
4. From the pop-up menu, select **Add New SPM Group**.
5. In the **Add New SPM Group** window, enter the name of the SPM group and then select **OK**.

An SPM group is created. An SPM group icon () is added to the lower-left tree.

6. Select an HBA  from the upper-left tree and select an SPM group from the lower-left tree. Next, select the **Add WWN** button.

Repeat this operation until all the desired HBAs are added to the SPM group.



Notes:

- Select a host bus adapter from below **Monitor**. You cannot add HBAs from below **Non-Monitor** to SPM groups.
- When you select a host bus adapter which is already contained in some SPM group from the upper-left tree, the Add WWN button is not activated. Select a host bus adapter which is not contained in any SPM groups.
- Windows users can add a host bus adapter to an SPM group by dragging the host bus adapter to the SPM group.

-
7. Select the **Apply** button.

The settings in the window are applied to the storage system.




Figure 7-11 Add New SPM Group Window

Deleting an HBA from an SPM Group

To delete a host bus adapter from the SPM group, take the following steps.

1. Start Server Priority Manager.

The **Server Priority Manager** window is displayed.

2. Select the **WWN** tab.
3. In the lower-left tree, double-click the SPM group (🗄️) that contains the host bus adapter to be deleted.
4. Below the SPM icon, right-click the icon the host bus adapter (🗄️) you want to delete.
5. Select **Delete WWN** from the pop-up menu.

The selected host bus adapter icon is deleted from the tree.

6. Select the **Apply** button.

The settings on the window are applied to the storage system.

Switching Priority of an SPM Group

All the host bus adapters (HBAs) in one SPM group must be of the same priority. Prioritized WWNs (i.e. high-priority HBAs) and non-prioritized WWNs (i.e. low-priority HBAs) cannot be mixed in one SPM group.

You can use an SPM group to switch priority of multiple HBAs from *prioritized* to *non-prioritized*, or vice versa.

To switch priority:

1. Start Server Priority Manager.

The **Server Priority Manager** window is displayed.

2. Select the WWN tab.
3. In the lower-left tree, select and right-click an SPM group (🗄️).
4. Do one of the following:
 - To switch priority from *prioritized* to *non-prioritized*, select **Prio ->> Non-Prio** from the pop-up menu.

- To switch priority from *non-prioritized* to *prioritized*, select **Non-Prio - >> Prio** from the pop-up menu.
- 5. Select the **Apply** button.

The settings in the window are applied to the storage system.

Setting an Upper-Limit Value to HBAs in an SPM Group

If all the host bus adapters in an SPM group are *non-prioritized WWNs* (i.e. low-priority HBAs), you can set an upper limit value to HBA performance (i.e. I/O rate or transfer rate). You can assign one upper limit value for one SPM group.

For example, suppose that the upper limit value **100 IOPS** is assigned to an SPM group consisting of four host bus adapters. If the sum of the I/O rate of the four HBAs reaches 100 IOPS, Server Priority Manager controls the system so that the sum of the I/O rates will not exceed 100 IOPS.

To set an upper limit value to HBAs in an SPM group:

1. Start Server Priority Manager.
The **Server Priority Manager** window is displayed.
2. Select the WWN tab.
3. In the lower-left tree, select and right-click the Subsystem folder or an SPM group (.
4. If you selected the **Subsystem** folder, take the following steps:
 - a. Select **IOPS** or **MB/s** from the drop-down list at the upper-left corner of the list. Select **IOPS** if you want to assign an upper limit to the I/O rate. Select **MB/s** if you want to assign an upper limit to the transfer rate.
 - b. To assign an upper limit to the I/O rate, enter the upper limit value in the **IOPS** column of the list.
To assign an upper limit to the transfer rate, enter the upper limit value in the **MB/s** column of the list.
Tips: If you cannot see the **IOPS** or **MB/s** column, scroll the list to the left. The column is located at the right side of the list.

If you selected an SPM group () , take the following steps:

- a. Right-click the selected SPM group and then select **Change Upper Limit** from the pop-up menu.
The **Change Upper Limit** window is displayed.
- b. To assign an upper limit to the I/O rate, enter the upper limit value and then select **IOPS** from the drop-down list. Next, select **OK**.
To assign an upper limit to the transfer rate, enter the upper limit value and then select **MB/s** from the drop-down list. Next, select **OK**.

5. In the **Server Priority Manager** window, select the **Apply** button.
The settings in the window are applied to the storage system.




Figure 7-12 Change Upper Limit Window



Note: To confirm an upper limit value specified for each SPM group, select the **Subsystem** folder in the lower-left tree of the **WWN** tab. The SPM groups are displayed in the list and you can confirm each upper limit value.

Renaming an SPM Group

To rename an SPM group, take the following steps.

1. Start Server Priority Manager.
The **Server Priority Manager** window is displayed.
2. Select the WWN tab.
3. In the lower-left tree, select and right-click an SPM group (📁).
4. Select **Rename SPM Group** from the pop-up menu.
The **Rename SPM Group** window is displayed.
5. Enter the new name and select **OK**.
6. In the **Server Priority Manager** window, select the **Apply** button.
The settings in the window are applied to the storage system.




Figure 7-13 Rename SPM Group Window

Deleting an SPM Group

If you want to delete an SPM group, take the following steps:

To delete an SPM group:

1. Start Server Priority Manager.
The **Server Priority Manager** window is displayed.
2. Select the WWN tab.
3. In the lower-left tree, select and right-click an SPM group ().
4. Select **Delete SPM Group** from the pop-up menu.
5. In the **Server Priority Manager** window, select the **Apply** button.
The settings in the window are applied to the storage system.

Using the Export Tool

This chapter explains using the export tool.

- ☐ [Files to be Exported](#)
- ☐ [Preparing for Using the Export Tool](#)
- ☐ [Using the Export Tool](#)
- ☐ [Command Reference](#)

Files to be Exported

The Export Tool enables you to save monitoring data in the **Performance Management** window into files. The Export Tool also enables you to save monitoring data about remote copy operations into files. The Export Tool usually compresses monitoring data in ZIP files. To use a text editor or spreadsheet software to view or edit the monitoring data, you usually need to decompress the ZIP files to extract CSV files. However, if you want the Export Tool to save monitoring data in CSV files instead of ZIP files, you can force the Export Tool to do so.

Table 8-1 shows the correspondence between the windows of Performance Monitor and the monitoring data that can be saved by the Export Tool. For details on the ZIP files and CSV files that are saved, refer to the tables indicated in the **See** column.

Table 8-1 Performance Management Windows and Monitoring Data Saved by the Export Tool

Window	Monitoring Data	See
Physical tab in the Performance Management Window	Statistics about resource usage and write pending rates	Table 8-2
LDEV tab in the Performance Management window	Statistics about parity groups, external volume groups, or V-VOL groups	Table 8-3
	Statistics about volumes in parity groups, in external volume groups, or in V-VOL groups	Table 8-4
Port-LUN tab in the Performance Management window	Statistics about ports	Table 8-5
	Statistics about host bus adapters connected to ports	Table 8-6
	Statistics about volumes(LUs)	Table 8-7
WWN tab in the Performance Management window	Statistics about SPM groups	Table 8-8
	Statistics about host bus adapters belonging to SPM groups	Table 8-9
TC Monitor window and TCz Monitor window	Statistics about remote copy operations by TrueCopy and TrueCopy for IBM z/OS (in the whole volume)	Table 8-10
	Statistics about remote copy operations by TrueCopy and TrueCopy for IBM z/OS (for each volume (LU))	Table 8-11
	Statistics about remote copy operations by TrueCopy and TrueCopy for IBM z/OS (at volumes controlled by a particular CU)	Table 8-12
	Statistics about remote copy operations by TrueCopy and TrueCopy for IBM z/OS (At CLPR)	Table 8-13
UR Monitor window and URz Monitor window	Statistics about remote copy operations by Universal Replicator and Universal Replicator for IBM z/OS (in the whole volume)	Table 8-14
	Statistics about remote copy operations by Universal Replicator and Universal Replicator for IBM z/OS (at journal groups)	Table 8-15
	Statistics about remote copy operations by Universal Replicator and Universal Replicator for IBM z/OS (for each volume (LU))	Table 8-16

Window	Monitoring Data	See
	Statistics about remote copy operations by Universal Replicator and Universal Replicator for IBM z/OS (at volumes controlled by a particular CU)	Table 8-17

Table 8-2 Files with Statistics about Resource Usage and Write Pending Rates

ZIP File	CSV File	Data Saved in the File
PhyPG_dat.ZIP	PHY_Long_PG.csv	Usage rates for parity groups in long range
	PHY_Short_PG.csv	Usage rates for parity groups in short range
PhyLDEV_dat.ZIP	PHY_Long_LDEV_x-y.csv	Usage rates for volumes in a parity group in long range
	PHY_Short_LDEV_x-y.csv	Usage rates for volumes in a parity group in short range
	PHY_Short_LDEV_SI_x-y.csv	Usage rates for ShadowImage volumes in a parity group in short range
PhyExG_dat.ZIP	PHY_ExG_Response.csv	Average response time for external volume groups (The unit is millisecond*1)
	PHY_ExG_Trans.csv	Amount of transferred data for external volume groups (The unit is KB/sec)
PhyExLDEV_dat.ZIP	PHY_ExLDEV_Response_x-y.csv	Average response time for volumes in an external volume group (The unit is millisecond*1)
	PHY_ExLDEV_Trans_x-y.csv	Amount of data transferred for volumes in an external volume group (The unit is KB/sec)
PhyProc_dat.ZIP	PHY_Long_CHP.csv	Usage rates for channel processors in long range
	PHY_Short_CHP.csv	Usage rates for channel processors in short range
	PHY_Long_DKP.csv	Usage rates for disk processors in long range
	PHY_Short_DKP.csv	Usage rates for disk processors in short range
	PHY_Long_DRR.csv	Usage rates for DRRs (data recovery and reconstruction processors) in long range
	PHY_Short_DRR.csv	Usage rates for DRRs (data recovery and reconstruction processors) in short range
PhyCSW_dat.ZIP	PHY_Long_MPA_CSW.csv	Usage rates for access paths between channel adapters and cache memories in long range Usage rates for access paths between disk adapters and cache memories in long range

ZIP File	CSV File	Data Saved in the File
	PHY_Short_MPA_CSW.csv	Usage rates for access paths between channel adapters and cache memories in short range Usage rates for access paths between disk adapters and cache memories in short range
	PHY_Long_MPA_SMA.csv	Usage rates for access paths between channel adapters and the shared memory in long range Usage rates for access paths between disk adapters and the shared memory in long range
	PHY_Short_MPA_SMA.csv	Usage rates for access paths between channel adapters and the shared memory in short range Usage rates for access paths between disk adapters and the shared memory in short range
	PHY_Long_CSW_CMA.csv	Usage rates for access paths between cache switches and cache memories in long range
	PHY_Short_CSW_CMA.csv	Usage rates for access paths between cache switches and cache memories in short range
	PHY_Long_Write_Pending_Rate.csv	Write pending rates in long range
	PHY_Short_Write_Pending_Rate.csv	Write pending rates in short range
	PHY_Short_Cache_Usage_Rate.csv	Usage rates for cache memory
* 1 millisecond is one-thousandth of 1 second.  Note: The letters "x-y" in CSV filenames indicate a parity group or external volume group. Both the statistics in long range and in short range are stored about resource usage and write pending rates		

Table 8-3 Files with Statistics about Parity Groups, External Volume Groups or V-VOL Groups

ZIP File	CSV File	Data Saved in the File
PG_dat.ZIP	PG_IOPS.csv	The number of read and write operations per second
	PG_TransRate.csv	The size of data transferred per second (The unit is KB/sec)
	PG_Read_IOPS.csv	The number of read operations per second
	PG_Seq_Read_IOPS.csv	The number of sequential read operations per second
	PG_Rnd_Read_IOPS.csv	The number of random read operations per second
	PG_CFW_Read_IOPS.csv	The number of read operations in "cache-fast-write" mode per second
	PG_Write_IOPS.csv	The number of write operations per second
	PG_Seq_Write_IOPS.csv	The number of sequential write operations per second
	PG_Rnd_Write_IOPS.csv	The number of random write operations per second

ZIP File	CSV File	Data Saved in the File
	PG_CFW_Write_IOPS.csv	The number of write operations in "cache-fast-write" mode per second
	PG_Read_Hit.csv	The read hit ratio
	PG_Seq_Read_Hit.csv	The read hit ratio in sequential access mode
	PG_Rnd_Read_Hit.csv	The read hit ratio in random access mode
	PG_CFW_Read_Hit.csv	The read hit ratio in "cache-fast-write" mode
	PG_Write_Hit.csv	The write hit ratio
	PG_Seq_Write_Hit.csv	The write hit ratio in sequential access mode
	PG_Rnd_Write_Hit.csv	The write hit ratio in random access mode
	PG_CFW_Write_Hit.csv	The write hit ratio in "cache-fast-write" mode
	PG_BackTrans.csv	The number of data transfer operations between cache memories and hard disk drives (i.e., parity groups, external volume groups, or V-VOL groups) per second
	PG_C2D_Trans.csv	The number of data transfer operations from cache memories and hard disk drives (i.e., parity groups, external volume groups, or V-VOL groups)
	PG_D2CS_Trans.csv	The number of data transfer operations from hard disk drives (i.e., parity groups, external volume groups, or V-VOL groups) to cache memories in sequential access mode
	PG_D2CR_Trans.csv	The number of data transfer operations from hard disk drives (i.e., parity groups, external volume groups, or V-VOL groups) to cache memories in random access mode
	PG_Response.csv	The average response time at parity groups, external volume groups, or V-VOL groups (The unit is microsecond*1)
* 1 microsecond is one-millionth of 1 second.		

Table 8-4 Files with Statistics about Volumes in Parity / External Volume Groups, or in V-VOL Groups

ZIP File	CSV File	Data Saved in the File
LDEV_IOPS.ZIP	LDEV_IOPS_x-y.csv	The number of read and write operations per second
LDEV_TransRate.ZIP	LDEV_TransRate_x-y.csv	The size of data transferred per second (The unit is KB/sec)
LDEV_Read_IOPS.ZIP	LDEV_Read_IOPS_x-y.csv	The number of read operations per second
LDEV_Seq_Read_IOPS.ZIP	LDEV_Seq_Read_IOPS_x-y.csv	The number of sequential read operations per second
LDEV_Rnd_Read_IOPS.ZIP	LDEV_Rnd_Read_IOPS_x-y.csv	The number of random read operations per second
LDEV_CFW_Read_IOPS.ZIP	LDEV_CFW_Read_IOPS_x-y.csv	The number of read operations in "cache-fast-write" mode per second

ZIP File	CSV File	Data Saved in the File
LDEV_Write_IOPS.ZIP	LDEV_Write_IOPS_x-y.csv	The number of write operations per second
LDEV_Seq_Write_IOPS.ZIP	LDEV_Seq_Write_IOPS_x-y.csv	The number of sequential write operations per second
LDEV_Rnd_Write_IOPS.ZIP	LDEV_Rnd_Write_IOPS_x-y.csv	The number of random write operations per second
LDEV_CFW_Write_IOPS.ZIP	LDEV_CFW_Write_IOPS_x-y.csv	The number of write operations in "cache-fast-write" mode per second
LDEV_Read_Hit.ZIP	LDEV_Read_Hit_x-y.csv	The read hit ratio
LDEV_Seq_Read_Hit.ZIP	LDEV_Seq_Read_Hit_x-y.csv	The read hit ratio in sequential access mode
LDEV_Rnd_Read_Hit.ZIP	LDEV_Rnd_Read_Hit_x-y.csv	The read hit ratio in random access mode
LDEV_CFW_Read_Hit.ZIP	LDEV_CFW_Read_Hit_x-y.csv	The read hit ratio in "cache-fast-write" mode
LDEV_Write_Hit.ZIP	LDEV_Write_Hit_x-y.csv	The write hit ratio
LDEV_Seq_Write_Hit.ZIP	LDEV_Seq_Write_Hit_x-y.csv	The write hit ratio in sequential access mode
LDEV_Rnd_Write_Hit.ZIP	LDEV_Rnd_Write_Hit_x-y.csv	The write hit ratio in random access mode
LDEV_CFW_Write_Hit.ZIP	LDEV_CFW_Write_Hit_x-y.csv	The write hit ratio in "cache-fast-write" mode
LDEV_BackTrans.ZIP	LDEV_BackTrans_x-y.csv	The number of data transfer operations between cache memories and hard disk drives (i.e., volumes) per second
LDEV_C2D_Trans.ZIP	LDEV_C2D_Trans_x-y.csv	The number of data transfer operations from cache memories and hard disk drives (i.e., volumes)
LDEV_D2CS_Trans.ZIP	LDEV_D2CS_Trans_x-y.csv	The number of data transfer operations from hard disk drives (i.e., volumes) to cache memories in sequential access mode
LDEV_D2CR_Trans.ZIP	LDEV_D2CR_Trans_x-y.csv	The number of data transfer operations from hard disk drives (i.e., volumes) to cache memories in random access mode
LDEV_Response.ZIP	LDEV_Response_x-y.csv	The average response time at volumes (The unit is microsecond*1)
* 1 microsecond is one-millionth of 1 second.  Note: The letters "x-y" in CSV filenames indicate a parity group. For example, if the filename is LDEV_IOPS_1-2.csv, the file contains the I/O rate for each volume in the parity group 1-2.		

Table 8-5 Files with Statistics about Ports

ZIP File	CSV File	Data Saved in the File
Port_dat.ZIP	Port_IOPS.csv	The number of read and write operations per second at ports
	Port_KBPS.csv	The size of data transferred per second at ports (The unit is KB/sec)
	Port_Response.csv	The average response time at ports(The unit is microsecond*1)
	Port_Initiator_IOPS.csv	The number of read and write operations per second at Initiator/External ports
	Port_Initiator_KBPS.csv	The size of data transferred per second at Initiator/External ports (The unit is KB/sec)
	Port_Initiator_Response.csv	The average response time at Initiator/External ports(The unit is microsecond*1)
* 1 microsecond is one-millionth of 1 second.		

Table 8-6 Files with Statistics about Host Bus Adapters Connected to Ports

ZIP File	CSV File	Data Saved in the File
PortWWN_dat.ZIP	PortWWN_xx_IOPS.csv	The I/O rate (that is, the number of read and write operations per second) for HBAs that are connected to a port
	PortWWN_xx_KBPS.csv	The size of data transferred per second between a port and the HBAs connected to that port (The unit is KB/sec)
	PortWWN_xx_Response.csv	The average response time between a port and the HBAs connected to that port (The unit is microsecond*1)
* 1 microsecond is one-millionth of 1 second. Notes: The letters "xx" in CSV filenames indicate a port name. For example, if the filename is PortWWN_1A_IOPS.csv, the file contains the I/O rate for each host bus adapter connected to the CL1-A port. If files are exported to a Windows computer, CSV filenames may end with numbers (for example, PortWWN_1A_IOPS-1.csv and PortWWN_1a_IOPS-2.csv).		

Table 8-7 Files with Statistics about Volumes(LUs)

ZIP File	CSV File	Data Saved in the File
LU_dat.ZIP	LU_IOPS.csv	The number of read and write operations per second
	LU_TransRate.csv	The size of data transferred per second (The unit is KB/sec)
	LU_Seq_Read_IOPS.csv	The number of sequential read operations per second
	LU_Rnd_Read_IOPS.csv	The number of random read operations per second
	LU_Seq_Write_IOPS.csv	The number of sequential write operations per second
	LU_Rnd_Write_IOPS.csv	The number of random write operations per second
	LU_Seq_Read_Hit.csv	The read hit ratio in sequential access mode
	LU_Rnd_Read_Hit.csv	The read hit ratio in random access mode
	LU_Seq_Write_Hit.csv	The write hit ratio in sequential access mode
	LU_Rnd_Write_Hit.csv	The write hit ratio in random access mode
	LU_C2D_Trans.csv	The number of data transfer operations from cache memories and hard disk drives (i.e., LUs)
	LU_D2CS_Trans.csv	The number of data transfer operations from hard disk drives (i.e., LUs) to cache memories in sequential access mode
	LU_D2CR_Trans.csv	The number of data transfer operations from hard disk drives (i.e., LUs) to cache memories in random access mode
	LU_Response.csv	The average response time at volumes(LUs) The unit is microsecond*1
* 1 microsecond is one-millionth of 1 second.		

Table 8-8 Files with Statistics about SPM Groups

ZIP File	CSV File	Data Saved in the File
PPCG_dat.ZIP	PPCG_IOPS.csv	The number of read and write operations per second
	PPCG_KBPS.csv	The size of data transferred per second (The unit is KB/sec)
	PPCG_Response.csv	The average response time at SPM groups (The unit is microsecond*1)
* 1 microsecond is one-millionth of 1 second.		

Table 8-9 Files with Statistics about Host Bus Adapters Belonging to SPM Groups

ZIP File	CSV File	Data Saved in the File
PPCGWWN_dat.ZIP	PPCGWWN_xx_IOPS.csv	The I/O rate (that is, the number of read and write operations per second) for HBAs belonging to an SPM group
	PPCGWWN_xx_KBPS.csv	The transfer rate for HBAs belonging to an SPM group (The unit is KB/sec)
	PPCGWWN_xx_Response.csv	The average response time for HBAs belonging to an SPM group (The unit is microsecond*1)
	PPCGWWN_NotGrouped_IOPS.csv	The I/O rate (that is, the number of read and write operations per second) for HBAs that do not belong to any SPM group
	PPCGWWN_NotGrouped_KBPS.csv	The transfer rate for HBAs that do not belong to any SPM group (The unit is KB/sec)
	PPCGWWN_NotGrouped_Response.csv	The average response time for HBAs that do not belong to any SPM group (The unit is microsecond*1)
* 1 microsecond is one-millionth of 1 second.  Notes: The letters "xx" in CSV filenames indicate the name of an SPM group. If files are exported to a Windows computer, CSV filenames may end with numbers (for example, PPCGWWN_mygroup_IOPS-1.csv and PPCGWWN_MyGroup_IOPS-2.csv).		

Table 8-10 Files with Statistics about Remote Copy Operations by TrueCopy and TrueCopy for IBM z/OS (In the Whole Volumes)

ZIP File	CSV File	Data Saved in the File
RemoteCopy_dat.ZIP	RemoteCopy.csv	The following data in the whole volumes are saved: The usage rate for sidefile cache The total number of remote I/Os (read and write operations). The total number of remote write I/Os. The number of errors that occur during remote I/O The number of initial copy remote I/Os. The average transfer rate (kB/sec) for initial copy remote I/Os. The average response time (millisecond*1) for initial copy. The number of update copy remote I/Os. The average transfer rate (kB/sec) for update copy remote I/Os. The average response time (millisecond*1) for update copy The number of restore copy remote I/Os The number of hits of restore copy remote I/Os The number of asynchronous update copy remote I/Os The number of asynchronous recordsets The average transfer rate (kB/sec) for asynchronous update copy remote I/Os The average response time (millisecond*1) for asynchronous update copy The number of scheduled recordsets The number of recordsets that do not arrive during the schedule The number of remaining recordsets when the schedule is completed The number of job activations of consistency manager The percentage of completion of copy operations (i.e., number of synchronized pairs / total number of pairs) The number of tracks that have not yet been copied by the initial copy or resync copy operation
* 1 millisecond is one-thousandth of 1 second.		

Table 8-11 Files with Statistics about Remote Copy Operations by TrueCopy and TrueCopy for IBM z/OS (for each Volume (LU))

ZIP File	CSV File	Data Saved in the File
RCLU_dat.ZIP	RCLU_All_RIO.csv	The total number of remote I/Os (read and write operations)
	RCLU_All_Read.csv	The total number of remote read I/Os
	RCLU_All_Write.csv	The total number of remote write I/Os
	RCLU_RIO_Error.csv	The number of errors that occur during remote I/O
	RCLU_Initial_Copy_RIO.csv	The number of initial copy remote I/Os
	RCLU_Initial_Copy_Hit.csv	The number of hits of initial copy remote I/Os
	RCLU_Initial_Copy_Transfer.csv	The average transfer rate (kB/sec) for initial copy remote I/Os
	RCLU_Initial_Copy_Response.csv	The average response time (millisecond*1) for the initial copy of each volume (LU)
	RCLU_Migration_Copy_RIO.csv	The number of migration copy remote I/Os
	RCLU_Migration_Copy_Hit.csv	The number of hits of migration copy remote I/Os
	RCLU_Update_Copy_RIO.csv	The number of update copy remote I/Os
	RCLU_Update_Copy_Hit.csv	The number of hits of update copy remote I/Os
	RCLU_Update_Copy_Transfer.csv	The average transfer rate (kB/sec) for update copy remote I/Os
	RCLU_Update_Copy_Response.csv	The average response time (millisecond*1) for the update copy of each volume (LU)
	RCLU_Restore_Copy_RIO.csv	The number of restore copy remote I/Os
	RCLU_Restore_Copy_Hit.csv	The number of hits of restore copy remote I/Os
	RCLU_Asynchronous_RIO.csv	The number of asynchronous update copy remote I/Os
	RCLU_Recordset.csv	The number of asynchronous recordsets
	RCLU_Asynchronous_Copy_Transfer.csv	The average transfer rate (kB/sec) for asynchronous update copy remote I/Os
	RCLU_Asynchronous_Copy_Response.csv	The average response time (millisecond*1) for the asynchronous update copy of each volume (LU)
	RCLU_Scheduling_Recordset.csv	The number of scheduled recordsets
	RCLU_Scheduling_Miss_Recordset.csv	The number of recordsets that do not arrive during the schedule
	RCLU_Remained_Recordset.csv	The number of remaining recordsets when the schedule is completed
	RCLU_Scheduling_Attempt.csv	The number of job activations of consistency manager
	RCLU_Pair_Synchronized.csv	The percentage of completion of copy operations (i.e., number of synchronized pairs / total number of pairs)

ZIP File	CSV File	Data Saved in the File
	RCLU_Out_of_Tracks.csv	The number of tracks that have not yet been copied by the initial copy or resync copy operation
* 1 millisecond is one-thousandth of 1 second.		

Table 8-12 Files with Statistics about Remote Copy Operations by TrueCopy and TrueCopy for IBM z/OS (At Volumes Controlled by a Particular CU)

ZIP File	CSV File	Data Saved in the File
RCLDEV_All_RIO.ZIP	RCLDEV_All_RIO_xx.csv	The total number of remote I/Os (read and write operations)
RCLDEV_All_Read.ZIP	RCLDEV_All_Read_xx.csv	The total number of remote read I/Os
RCLDEV_All_Write.ZIP	RCLDEV_All_Write_xx.csv	The total number of remote write I/Os
RCLDEV_RIO_Error.ZIP	RCLDEV_RIO_Error_xx.csv	The number of errors that occur during remote I/O
RCLDEV_Initial_Copy_RIO.ZIP	RCLDEV_Initial_Copy_RIO_xx.csv	The number of initial copy remote I/Os
RCLDEV_Initial_Copy_Hit.ZIP	RCLDEV_Initial_Copy_Hit_xx.csv	The number of hits of initial copy remote I/Os
RCLDEV_Initial_Copy_Transfer.ZIP	RCLDEV_Initial_Copy_Transfer_xx.csv	The average transfer rate (kB/sec) for initial copy remote I/Os
RCLDEV_Initial_Copy_Response.ZIP	RCLDEV_Initial_Copy_Response_xx.csv	The average response time(millisecond*1) for initial copy at volumes
RCLDEV_Migration_Copy_RIO.ZIP	RCLDEV_Migration_Copy_RIO_xx.csv	The number of migration copy remote I/Os
RCLDEV_Migration_Copy_Hit.ZIP	RCLDEV_Migration_Copy_Hit_xx.csv	The number of hits of migration copy remote I/Os
RCLDEV_Update_Copy_RIO.ZIP	RCLDEV_Update_Copy_RIO_xx.csv	The number of update copy remote I/Os
RCLDEV_Update_Copy_Hit.ZIP	RCLDEV_Update_Copy_Hit_xx.csv	The number of hits of update copy remote I/Os
RCLDEV_Update_Copy_Transfer.ZIP	RCLDEV_Update_Copy_Transfer_xx.csv	The average transfer rate (kB/sec) for update copy remote I/Os
RCLDEV_Update_Copy_Response.ZIP	RCLDEV_Update_Copy_Response_xx.csv	The average response time (millisecond*1) for the update copy at volumes
RCLDEV_Restore_Copy_RIO.ZIP	RCLDEV_Restore_Copy_RIO_xx.csv	The number of restore copy remote I/Os
RCLDEV_Restore_Copy_Hit.ZIP	RCLDEV_Restore_Copy_Hit_xx.csv	The number of hits of restore copy remote I/Os
RCLDEV_Asynchronous_RIO.ZIP	RCLDEV_Asynchronous_RIO_xx.csv	The number of asynchronous update copy remote I/Os
RCLDEV_Recordset.ZIP	RCLDEV_Recordset_xx.csv	The number of asynchronous recordsets

ZIP File	CSV File	Data Saved in the File
RCLDEV_Asynchronous_Copy_Transfer.ZIP	RCLDEV_Asynchronous_Copy_Transfer_xx.csv	The average transfer rate (kB/sec) for asynchronous update copy remote I/Os
RCLDEV_Asynchronous_Copy_Response.ZIP	RCLDEV_Asynchronous_Copy_Response_xx.csv	The average response time (millisecond*1) for the asynchronous update copy at volumes
RCLDEV_Scheduling_Recordset.ZIP	RCLDEV_Scheduling_Recordset_xx.csv	The number of scheduled recordsets
RCLDEV_Scheduling_Miss_Recordset.ZIP	RCLDEV_Scheduling_Miss_Recordset_xx.csv	The number of recordsets that do not arrive during the schedule
RCLDEV_Remained_Recordset.ZIP	RCLDEV_Remained_Recordset_xx.csv	The number of remaining recordsets when the schedule is completed
RCLDEV_Scheduling_Attempt.ZIP	RCLDEV_Scheduling_Attempt_xx.csv	The number of job activations of consistency manager
RCLDEV_Pair_Synchronized.ZIP	RCLDEV_Pair_Synchronized_xx.csv	The percentage of completion of copy operations (i.e., number of synchronized pairs / total number of pairs)
RCLDEV_Out_of_Tracks.ZIP	RCLDEV_Out_of_Tracks_xx.csv	The number of tracks that have not yet been copied by the initial copy or Resync copy operation
* 1 millisecond is one-thousandth of 1 second.  Note: The letters "xx" in CSV filenames indicate a CU image number. For example, if the filename is RCLDEV_All_RIO_10.csv, the file contains the total number of remote I/Os of the volumes controlled by the CU whose image number is 10.		

Table 8-13 Files with Statistics about Remote Copy Operations by TrueCopy and TrueCopy for IBM z/OS (At CLPR)

ZIP File	CSV File	Data Saved in the File
RCCLPR_dat.ZIP	RCCLPR_SideFile.csv	Sidefile usage rate per CLPR (The unit is %)

Table 8-14 Files with Statistics about Remote Copy Operations by Universal Replicator and Universal Replicator for IBM z/OS (In the Whole Volumes)

ZIP File	CSV File	Data Saved in the File
UniversalReplicator.ZIP	UniversalReplicator.csv	- The following data in the whole volumes are saved: - The number of write I/Os per second. - The amount of data that are written per second (The unit is KB/sec) - The initial copy hit rate (The unit is percent) - The average transfer rate for initial copy operations (The unit is KB/sec) - The number of asynchronous remote I/Os per second at the primary storage system - The number of journals at the primary storage system - The average transfer rate for journals in the primary storage system (The unit is KB/sec) - The remote I/O average response time on the primary storage system (The unit is millisecond*1) - The number of asynchronous remote I/Os per second at the secondary storage system - The number of journals at the secondary storage system - The average transfer rate for journals in the secondary storage system (The unit is KB/sec) - The remote I/O average response time on the secondary storage system (The unit is millisecond*1)
* 1 millisecond is one-thousandth of 1 second.		

Table 8-15 Files with Statistics about Remote Copy Operations by Universal Replicator and Universal Replicator for IBM z/OS (At Journal Groups)

ZIP File	CSV File	Data Saved in the File
URJNL_dat.ZIP	URJNL_Write_Record.csv	The number of write I/Os per second
	URJNL_Write_Transfer.csv	The amount of data that are written per second (The unit is KB/sec)
	URJNL_Initial_Copy_Hit.csv	The initial copy hit rate (The unit is percent)
	URJNL_Initial_Copy_Transfer.csv	The average transfer rate for initial copy operations (The unit is KB/sec)
	URJNL_MCU_Asynchronous_RIO.csv	The number of asynchronous remote I/Os per second at the primary storage system
	URJNL_MCU_Asynchronous_Journal.csv	The number of journals at the primary storage system

ZIP File	CSV File	Data Saved in the File
	URJNL_MCU_Asynchronous_Copy_Transfer.csv	The average transfer rate for journals in the primary storage system (The unit is KB/sec)
	URJNL_MCU_Asynchronous_Copy_Response.csv	The remote I/O average response time on the primary storage system (The unit is millisecond*1)
	URJNL_RCU_Asynchronous_RIO.csv	The number of asynchronous remote I/Os per second at the secondary storage system
	URJNL_RCU_Asynchronous_Journal.csv	The number of journals at the secondary storage system
	URJNL_RCU_Asynchronous_Copy_Transfer.csv	The average transfer rate for journals in the secondary storage system (The unit is KB/sec)
	URJNL_RCU_Asynchronous_Copy_Response.csv	The remote I/O average response time on the secondary storage system (The unit is millisecond*1)
	URJNL_MCU_Journal_Filling.csv	Data usage rate for master journals (The unit is percent)
	URJNL_MCU_JNCB_Filling.csv	Metadata usage rate for master journals (The unit is percent)
	URJNL_RCU_Journal_Filling.csv	Data usage rate for restore journals (The unit is percent)
	URJNL_RCU_JNCB_Filling.csv	Metadata usage rate for restore journals (The unit is percent)
* 1 millisecond is one-thousandth of 1 second.		

Table 8-16 Files with Statistics about Remote Copy Operations by Universal Replicator and Universal Replicator for IBM z/OS (For Each Volume (LU))

ZIP File	CSV File	Data Saved in the File
URLU_dat.ZIP	URLU_Read_Record.csv	The number of read I/Os per second
	URLU_Read_Hit.csv	The number of read hit records per second
	URLU_Write_Record.csv	The number of write I/Os per second
	URLU_Write_Hit.csv	The number of write hit records per second
	URLU_Read_Transfer.csv	The amount of data that are read per second (The unit is KB/sec)
	URLU_Write_Transfer.csv	The amount of data that are written per second (The unit is KB/sec)
	URLU_Initial_Copy_Hit.csv	The initial copy hit rate (The unit is percent)
	URLU_Initial_Copy_Transfer.csv	The average transfer rate for initial copy operations (The unit is KB/sec)

Table 8-17 Files with Statistics about Remote Copy Operations by Universal Replicator and Universal Replicator for IBM z/OS (At Volumes Controlled by a Particular CU)

ZIP File	CSV File	Data Saved in the File
URLDEV_Read_Record.ZIP	URLDEV_Read_Record_xx.csv	The number of read I/Os per second
URLDEV_Read_Hit.ZIP	URLDEV_Read_Hit_xx.csv	The number of read hit records per second
URLDEV_Write_Record.ZIP	URLDEV_Write_Record_xx.csv	The number of write I/Os per second
URLDEV_Write_Hit.ZIP	URLDEV_Write_Hit_xx.csv	The number of write hit records per second
URLDEV_Read_Transfer.ZIP	URLDEV_Read_Transfer_xx.csv	The amount of data that are read per second (The unit is KB/sec)
URLDEV_Write_Transfer.ZIP	URLDEV_Write_Transfer_xx.csv	The amount of data that are written per second (The unit is KB/sec)
URLDEV_Initial_Copy_Hit.ZIP	URLDEV_Initial_Copy_Hit_xx.csv	The initial copy hit rate (The unit is percent)
URLDEV_Initial_Copy_Transfer.ZIP	URLDEV_Initial_Copy_Transfer_xx.csv	The average transfer rate for initial copy operations (The unit is KB/sec)
 Note: The letters "xx" in CSV filenames indicate a CU image number. For example, if the filename is URLDEV_Read_Record_10.csv, the file contains the number of read I/Os (per second) of the volumes controlled by the CU whose image number is 10.		

Preparing for Using the Export Tool

This section explains how to prepare for using the Export Tool.

Requirements for Using the Export Tool

The following components are required to use the Export Tool:

- **a Windows computer or a UNIX computer**

The Export Tool runs on Windows computers and UNIX computers that can run Storage Navigator. If your Windows or UNIX computer is unable to run Storage Navigator, your computer is unable to run the Export Tool. For detailed information about computers that can run Storage Navigator, see the *Storage Navigator User's Guide*.

- **Java Runtime Environment (JRE)**

To be able to use the Export Tool, you must install Java Runtime Environment on your Windows or UNIX computer. If your computer runs Storage Navigator, JRE is already installed on your computer and you can install the Export Tool. If your computer does not run Storage Navigator but contains an appropriate version of JRE, you can install the Export Tool on your computer.



Note: The JRE version required for running the Export Tool is the same as the JRE version required for running Storage Navigator. For detailed information about the JRE version required for running Storage Navigator, see the *Storage Navigator User's Guide*.

A user ID exclusively for use with the Export Tool

If you want to use the Export Tool, you must create a user ID that will be used exclusively with the Export Tool. When you create the user ID, note the following:

- **Permissions of USP V/VM programs**

If you use the Export Tool only to save the monitoring data into files, *do not* assign any permission to the user ID for use with the Export Tool. If a user ID having permission is used with the Export Tool, the storage system configuration might be changed in an unfavorable way by an unidentified user.

If you use the Export Tool not only to save monitoring data but also to start or stop monitoring and to change the gathering interval by the **set** subcommand, the user ID needs to have at least one of permissions for Performance Monitor, TrueCopy, TrueCopy for IBM z/OS, Universal Replicator, and Universal Replicator for IBM z/OS.

– **User types**

You can specify any user type for the user ID for use with the Export Tool.

If you specify "*storage administrator*" for the user ID, all the monitoring data described from Table 8-2 to Table 8-17 can be saved into files.

If you specify "*storage partition administrator*" for the user ID, the monitoring data that can be saved and the functions of the Export Tool are limited. For details, see Using the Export Tool.

For detailed information about how to create a user ID, see the *Storage Navigator User's Guide*.

- **The Export Tool program**

The Export Tool is contained in CD-ROM Disc 2, which is named *Host PP*. For detailed information about how to install the Export Tool on a Windows computer, see Installing the Export Tool on a Windows Computer. For detailed information about how to install the Export Tool on a UNIX computer, see Installing the Export Tool on a UNIX Computer.

Installing the Export Tool on a Windows Computer

To install the Export Tool on a Windows computer:

1. Create a directory on your Windows computer.
In later steps, you will install the Export Tool on the new directory.
2. Insert the Export Tool installation media into the CD-ROM drive.
3. Locate the self-extracting file **export.EXE** in the directory **\program\monitor\win_nt** in your CD-ROM disc, and then copy export.EXE to the new directory that you created earlier.
4. Double-click export.EXE on your computer.

The Export Tool is installed. Also, a new directory named **export** is created.



Notes:

- The **export** directory contains a couple of files, which include **runUnix.bat**. It is recommended that you delete runUnix.bat because this file is no longer needed.
 - The Export Tool program is a Java class file and is located in the **export\lib** directory.
-

Installing the Export Tool on a UNIX Computer

To install the Export Tool on a UNIX computer:

1. Create a directory on your UNIX computer.
In later steps, you will install the Export Tool on the new directory.
2. Mount the Export Tool installation media.
3. Do one of the following:
 - If you are using Solaris, locate the archive file **export.tar** in the directory **/program/monitor/solaris** in your CD-ROM disc, and then copy **export.tar** to the new directory that you created earlier.
 - If you are using HP-UX, locate the archive file **export.tar** in the directory **/program/monitor/HP-UX** in your CD-ROM disc, and then copy **export.tar** to the new directory that you created earlier.
4. Decompress **export.tar** on your computer.
The Export Tool is installed. Also, a new directory named **export** is created.



Notes:

- The **export** directory contains a couple of files, which include **runWin.bat**. It is recommended that you delete **runWin.bat** because this file is no longer needed.
 - The Export Tool program is a Java class file and is located in the **export/lib** directory.
-

Using the Export Tool

To be able to export monitoring data, you must prepare a command file and a batch file. This section explains how to prepare a command file and a batch file, and then explains how to run the Export Tool.

- Preparing a command file
- Preparing a batch file
- Running the Export Tool

Preparing a Command File

To be able to run the Export Tool, you must write scripts for exporting monitoring data. When writing scripts, you need to write several subcommands in a *command file*. When you run the Export Tool, the subcommands in the command file execute sequentially and then the monitoring data are saved in files.

Figure 8-1 gives an example of a command file:

```
svpip 158.214.135.57      ; Specifies IP address of SVP
login expusr passwd      ; Logs user into SVP
show                     ; Outputs storing period to standard output
group PhyPG Long         ; Specifies type of data to be exported and type of
                        ; storing period
group RemoteCopy         ; Specifies type of data to be exported
short-range 200610010850:200610010910
                        ; Specifies term of data to be exported for data stored
                        ; in short range
long-range 200609301430:200610011430
                        ; Specifies term of data to be exported for data stored
                        ; in long range
outpath out              ; Specifies directory in which files will be saved
option compress          ; Specifies whether to compress files
apply                   ; Executes processing for saving monitoring data in files
```

Figure 8-1 Example of a Command File



Note: In the above scripts, the semicolon (;) indicates the beginning of a comment. Characters from a semicolon to the end of the line are regarded as a comment.

The scripts in this command file are explained as follows:

- svpip 158.214.135.57

This script specifies that you are logging into the SVP whose IP address is 158.214.135.57. You must log into the SVP when using the Export Tool.

The `svpip` subcommand specifies the IP address of the SVP. You must include the `svpip` subcommand in your command file. For detailed information about the `svpip` subcommand, see `svpip Subcommand`.

- `login expusr passwd`

This script specifies that you use the user ID **expusr** and the password **passwd** to log into the SVP.

The `login` subcommand logs the specified user into the SVP. You must include the `login` subcommand in your command file. For detailed information about the `login` subcommand, see `login Subcommand`.



Caution: When you write the `login` subcommand in your command file, you must specify a user ID that should be used exclusively for running the Export Tool. See `Requirements for Using the Export Tool` for reference.

- `show`

The `show` subcommand checks the SVP to find the period of monitoring data stored in the SVP and the data collection interval (that is called "gathering interval" in Performance Monitor), and then outputs them to the standard output (for example, the command prompt) and the log file.

Performance Monitor collects statistics by the two types of storing periods: *in short range* and *in long range*. The `show` subcommand displays the storing periods and the gathering intervals for these two types of monitoring data.

The following is an example of information that the `show` subcommand outputs:

Short Range	From: 2006/10/01 01:00 - To: 2006/10/01 15:00 Interval: 1min.
Long Range	From: 2006/09/01 00:00 - To: 2006/10/01 15:00 Interval: 15min.

`Short Range` indicates the storing period and gathering interval of the monitoring data stored in short range. `Long Range` indicates those of the monitoring data stored in long range. In the above example, the monitoring data in short range is stored every 1 minute in the term of 1:00-15:00 on Oct. 1, 2006. Also, the monitoring data in long range is stored every 15 minutes in the term of Sep. 1, 2006, 0:00 through Oct. 1, 2006, 15:00. When you run the Export Tool, you can export monitoring data within these periods into files.

All the monitoring items are stored in short range, but a part of monitoring items is stored in both the short range and long range. For details on monitoring items that can be stored in long range, see `long-range Subcommand`.

The use of the `show` subcommand is not mandatory, but it is recommended that you include the `show` subcommand in your command file. If an error occurs when you run the Export Tool, you might be able to find the error cause by checking the log file for information issued by the `show` subcommand. For detailed information about the `show` subcommand, see `show Subcommand`.

- `group PhyPG Long` and `group RemoteCopy`

The `group` subcommand specifies the type of data that you want to export. Specify an operand following `group` to define the type of data to be exported. Basically, monitoring data stored in short range is exported. But you can direct to export monitoring data stored in long range when you specify some of the operands.

The script `group PhyPG Long` in Figure 8-1 specifies to export usage statistics about parity groups in long range. Also, the script `group RemoteCopy` specifies to export statistics about remote copy operations by TrueCopy and TrueCopy for IBM z/OS in short range. You can describe multiple lines of the `group` subcommand to export multiple monitoring items at the same time.

For detailed information about the `group` subcommand, see `group Subcommand`.

- `short-range 200610010850:200610010910` and
`long-range 200609301430:200610011430`

The `short-range` and `long-range` subcommands specify the term of monitoring data to be exported. Use these subcommands when you want to narrow the `export-target` term within the stored data. You can specify both the `short-range` and `long-range` subcommands at the same time. The difference between these subcommands is as follows:

- The `short-range` subcommand is valid for monitoring data in short range. You can use this subcommand to narrow the `export-target` term for all the monitoring items you can specify by the `group` subcommand.
Specify a term within "Short Range From XXX To XXX" which is output by the `show` subcommand.
- The `long-range` subcommand is valid for monitoring data in long range. You can use this subcommand only when you specify the `PhyPG`, `PhyLDEV`, `PhyProc`, or `PhyCSW` operand with the `Long` option in the `group` subcommand. (The items that can be saved by these operands are the monitoring data displayed in the **Physical** tab of the **Performance Management** window with selecting **long-range**.)
Specify a term within "Long Range From XXX To XXX" which is output by the `show` subcommand.

In Figure 8-1, the script `short-range 200610010850:200610010910` specifies the term 8:50-9:10 on Oct. 1, 2006. This script is applied to the `group RemoteCopy` subcommand in this example. When you run the Export Tool, it will export the statistics about remote copy operations by TrueCopy and TrueCopy for IBM z/OS in the term specified by the `short-range` subcommand.

Also, in Figure 8-1, the script `long-range 200609301430:200610011430` specifies the term from Sep. 30, 2006, 14:30 to Oct. 1, 2006, 14:30. This script is applied to the `group PhyPG Long` subcommand in this example. When you run the Export Tool, it will export the usage statistics about parity groups in the term specified by the `long-range` subcommand.

If you run the Export Tool without specifying the `short-range` or `long-range` subcommand, the monitoring data in the whole storing period (data in the period displayed by the `show` subcommand) will be exported.

- For detailed information about the `short-range` subcommand, see `short-range Subcommand`.
- For detailed information about the `long-range` subcommand, see `long-range Subcommand`.

- `outpath out`

This script specifies that files should be saved in the directory named `out` in the current directory.

The `outpath` subcommand specifies the directory in which files should be saved. For detailed information about the `outpath` subcommand, see `outpath Subcommand`.

- `option compress`

This script specifies that the Export Tool should compress monitoring data in ZIP files.

The `option` subcommand specifies whether to save files in ZIP format or in CSV format. For detailed information about the `option` subcommand, see `option Subcommand`.

- `apply`

The `apply` subcommand saves monitoring data in files. For detailed information about the `apply` command, see `apply Subcommand`.

When you install the Export Tool, the file **command.txt** will be stored in the **export** directory. The **command.txt** file contains sample scripts for your command file. It is recommended that you customize scripts in `command.txt` according to your needs. For detailed information about subcommand syntax, see `Command Reference`.

Preparing a Batch File

To run the Export Tool, you need a batch file. The Export Tool starts and saves monitoring data in files when you execute the batch file.

The installation directory for the Export Tool (that is, the **export** directory) contains two batch files: **runWin.bat** and **runUnix.bat**. If your computer runs Windows, use **runWin.bat**. If your computer runs UNIX, use **runUnix.bat**.

Figure 8-2 illustrates scripts in runWin.bat and runUnix.bat. These batch files include a command line that executes a Java command. When you execute your batch file, the Java command executes subcommands specified in your command file and then saves monitoring data in files.

Batch file for Windows computers (runWin.bat)

```
Java -classpath "./lib/JSanExport.jar;./lib/JSanRmiServerSx.jar" -Xmx536870912 -  
Dmd.command=command.txt -Dmd.logpath=log sanproject.getmondat.RJMdMain ↵  
pause ↵
```

Batch file for UNIX computers (runUnix.bat)

```
#!/bin/sh ↵  
Java -classpath "./lib/JSanExport.jar:./lib/JSanRmiServerSx.jar" -Xmx536870912 -  
Dmd.command=command.txt -Dmd.logpath=log sanproject.getmondat.RJMdMain ↵
```

Figure 8-2 Scripts in Batch Files



Note: The "↵" symbol indicates the end of a command line.

If the computer running the Export Tool communicates directly with the SVP, you usually do not need to change scripts in **runWin.bat** and **runUnix.bat**. However, you might need to edit the Java command script in your text editor in some occasions, for example:

- if the name of your command file is not **command.txt**
- if you moved your command file to a different directory
- if you do not want to save log files in the "**log**" directory
- if you want to name log files as you like

If the computer that runs the Export Tool communicates with the SVP via a proxy host, you need to edit the Java command script in your text editor. When editing the Java command script, you need to specify the host name (or the IP address) and the port number of the proxy host. For example, if the host name is *Jupiter* and the port number is *8080*, the resulting command script as shown in Figure 8-3:

Batch file for Windows computers (runWin.bat)

```
Java -classpath "./lib/JSanExport.jar;./lib/JSanRmiServerSx.jar" -Dhttp.proxyHost=Jupiter
-Dhttp.proxyPort=8080 -Xmx536870912 -Dmd.command=command.txt -Dmd.logpath=log
sanproject.getmondar.RJMdMain ↵
pause ↵
```

Batch file for UNIX computers (runUnix.bat)

```
#!/bin/sh ↵
Java -classpath "./lib/JSanExport.jar;./lib/JSanRmiServerSx.jar" -Dhttp.proxyHost=Jupiter
-Dhttp.proxyPort=8080 -Xmx536870912 -Dmd.command=command.txt -Dmd.logpath=log
sanproject.getmondar.RJMdMain ↵
```

Figure 8-3 Scripts in Batch Files (When Specifying the Host Name of a Proxy Host)



Note: The "↵" symbol indicates the end of a command line.

If the IP address of the proxy host is *158.211.122.124* and the port number is *8080*, the resulting command script is as follows:

Batch file for Windows computers (runWin.bat)

```
Java -classpath "./lib/JSanExport.jar;./lib/JSanRmiServerSx.jar" -Dhttp.proxyHost=158.211.122.124
-Dhttp.proxyPort=8080 -Xmx536870912 -Dmd.command=command.txt -Dmd.logpath=log
sanproject.getmondar.RJMdMain ↵
pause ↵
```

Batch file for UNIX computers (runUnix.bat)

```
#!/bin/sh ↵
Java -classpath "./lib/JSanExport.jar;./lib/JSanRmiServerSx.jar" -Dhttp.proxyHost=158.211.122.124
-Dhttp.proxyPort=8080 -Xmx536870912 -Dmd.command=command.txt -Dmd.logpath=log
sanproject.getmondar.RJMdMain ↵
```

Figure 8-4 Scripts in Batch Files (When Specifying the IP Address of a Proxy Host)



Note: The "↵" symbol indicates the end of a command line.

For detailed information about syntax of the Java command, see Java Command for Exporting Data In Files.

Running the Export Tool

To run the Export Tool and save monitoring data in files, you need to execute your batch file. To execute your batch file, you need to enter the name of the batch file at the command prompt and then press the <Enter> key. If you are using a Windows computer, you can double-click the batch file to execute the batch file.

c:\WINDOWS> cd c:\export	← Go to the directory containing runWin.bat
c:\export> runWin.bat	← Execute runWin.bat

Figure 8-5 Example of Executing a Batch File (on a Windows Computer)

When the Export Tool starts exporting monitoring data, dots (...) are issued to the standard output (for example, the command prompt). The dots increment as export processing continues. If an internal error occurs, the exclamation mark (!) is issued to the standard output and the Export Tool attempts to restart exporting data. If the export processing restarts, dots reappear and increment until export processing finishes.

[2] svpip 158.214.135.57	← Displays the currently-running subcommand
[3] login User = expusr, Passwd = [*****]	← Displays the currently-running subcommand
:	
:	
[6] group Port	← Displays the currently-running subcommand
:	
:	
[20] apply	← Displays the currently-running subcommand
Start gathering port data	← Indicates that the export processing starts
Target = 16, Total = 16	
+-----+-----+-----+-----+-----+-----+	
.....!	← Indicates that an error occurs during the export processing
.....	← Dots appear and increment as the export processing continues
End gathering port data	← Indicates that the export processing ends successfully

Figure 8-6 Example of Command Prompt Outputs from Export Tool

When the Export Tool finishes successfully, monitoring data are usually compressed in ZIP-format archive files. If you want to obtain CSV files, you need to decompress ZIP files and extract CSV files out of the ZIP files. If the operating system on your computer does not include a feature for decompressing ZIP files, you need to obtain software for decompressing files.

For a complete list of files to be saved by the Export Tool, see Using the Export Tool.



Note: When an internal error occurs during export processing, the exclamation mark (!) appears to signal the error. If this happens, the Export Tool will make up to three more attempts at processing.

If export processing does not finish through three retries or if an internal error occurs other than those in Table 8-18, the Export Tool does not retry the processing. In this case, you need to quit the command prompt and then run the Export Tool again.

You can change the maximum number of retries by using the **retry** subcommand. For detailed information about the **retry** subcommand, see retry Subcommand.

Table 8-18 Errors for Which Export Tool Retries Processing

Error Message ID	Cause of Error
0001 4001	An error occurred during SVP processing.
0001 5400	Since the SVP is busy, the monitoring data cannot be obtained.
0001 5508	An administrator is changing a system environment file.
0002 2016	Array is refreshing, or the settings by the user are registered.
0002 5510	The storage system is in internal process, or some other user is changing configuration.
0002 6502	Now processing.
0002 9000	Another user has lock.
0003 2016	A service engineer is accessing the storage system in Modify mode.
0003 2033	The SVP is not ready yet, or an internal processing is being executed.
0003 3006	An error occurred during SVP processing.
0405 8003	The storage system status is invalid.
5205 2003	An internal process is being executed, or maintenance is in progress.
5205 2033	The SVP is now updating the statistics data.
5305 2033	The SVP is now updating the statistics data.
5305 8002	The storage system status is invalid.

- If you specify the `nocompress` operand for the `option` subcommand, the Export Tool saves files in CSV format instead of ZIP format (For detailed information, see option Subcommand).



Note: When files are saved in CSV format instead of ZIP format, the file saving process could take longer and the resulting files could be larger.

- Files saved by the Export Tool are often very large. The total file size for all the files can be as large as approximately 2 GB. For this reason, the exporting process might take a lot of time. If you want to export statistics spanning a long period of time, it is recommended that you run the Export Tool multiple times for different periods, rather than run only once to export the entire time span as a single large file. For example, if you want to export statistics spanning 24 hours, run the tool eight times to export statistics in three hour increments.

Table 8-19 Estimate of Time Required for Exporting Files

Operand for the Group Subcommand	Estimated Time	Remarks
Port	5 minutes	This estimate assumes that the Export Tool should save statistics about 128 ports within a 24-hour period.
PortWWN	5 minutes	This estimate assumes that the Export Tool should save statistics about 128 ports within a 24-hour period.
PPCG	5 minutes	This estimate assumes: There are eight SPM groups, and eight WWNs are registered on each SPM group. There is a WWN that is not registered on any SPM group. The Export Tool should save statistics about SPM groups and WWNs described above within a 24-hour period.
LDEV	60 minutes	This estimate assumes: The Export Tool should save statistics about 8,192 volumes within a 24-hour period. The Export Tool is used eight times. Each time the Export Tool is used, the tool obtains statistics within a 3-hour period.
LU	60 minutes	This estimate assumes: The Export Tool should save statistics about 12,288 LUs within a 24-hour period. The Export Tool is used eight times. Each time the Export Tool is used, the tool obtains statistics within a 3-hour period.

- Whenever the Export Tool runs, the Export Tool creates a new log file on your computer. Therefore, if you run the Export Tool repeatedly, the size of free space on your computer will be reduced. To secure free space on your computer, you are strongly recommended to delete log files regularly. For information about the directory containing log files, see Java Command for Exporting Data In Files.
- For information about how to solve errors with the Export Tool, see Overview of Export Tool.
- The Export Tool returns a termination code when the Export Tool finishes.

Table 8-20 Termination Codes Returned by the Export Tool

Termination Code	Meaning
0	The Export Tool finished successfully.
1	An error occurred when the set subcommand (see set subcommand) executed, because an attempt to switch to Modify mode failed. Some other user might have been logged on in Modify mode.
2	An error occurred due to some reason unrelated to system option modes (i.e., View mode and Modify mode)
3	An error occurred due to more than one reason. One of the reasons is that an attempt to switch to Modify mode failed when the set subcommand (see set subcommand) executed. Some other user might have been logged on in Modify mode.
4	The user ID has none of permissions for Performance Monitor, TrueCopy, TrueCopy for IBM z/OS, Universal Replicator, and Universal Replicator for IBM z/OS.

If you want to use a reference to a termination code in your batch file, do the following:

- To use such a reference in a Windows batch file, write **%errorlevel%** in the batch file.
- To use such a reference in a UNIX Bourne shell script, write **%?** in the shell script.
- To use such a reference in a UNIX C shell script, write **%status** in the shell script.

A reference to a termination code is used in the following example of a Windows batch file. If this batch file executes and the Export Tool returns the termination code 1 or 3, the command prompt displays a message that indicates the **set** subcommand fails.

```
Java -classpath "./lib/JSanExport.jar;./lib/JSanRmiServerSx.jar" -Xmx536870912 -
Dmd.command=command.txt -Dmd.logpath=log sanproject.getmondar.RJMcMain ↵
if %errorlevel%==1 echo THE SET SUBCOMMAND FAILED ↵
if %errorlevel%==3 echo THE SET SUBCOMMAND FAILED ↵
pause ↵
```

Figure 8-7 Example of a Batch File Including References to Termination Codes

Note: The "↵" symbol indicates the end of a command line.

Command Reference

This section provides the syntax of the subcommands that you can write in your command file and the Java command that should be used in your batch file. Table 8-21 lists the subcommands explained in this section. The Java command is explained in Java Command for Exporting Data In Files.

Table 8-21 Subcommand List

Subcommand	Function	See
svpip	Specifies the IP address of the SVP to be logged in.	svpip Subcommand
retry	Makes settings on retries of export processing.	retry Subcommand
login	Logs the specified user into the SVP.	login Subcommand
show	Checks the SVP to find the period of monitoring data stored in the SVP and the data collection interval (that is called "gathering interval"), and then outputs them to the standard output and the log file.	show Subcommand
group	Specifies the type of data that you want to export.	group Subcommand
short-range	Specifies the term of monitoring data to be exported for short-range monitoring data.	short-range Subcommand
long-range	Specifies the term of monitoring data to be exported for long-range monitoring data.	long-range Subcommand
outpath	Specifies the directory in which files should be saved.	outpath Subcommand
option	Specifies whether to save files in ZIP format or in CSV format.	option Subcommand
apply	Saves monitoring data in files.	apply Subcommand
set	Starts or ends monitoring the storage system, and also specifies the gathering interval in short range monitoring.	set subcommand
help	Displays the online help for subcommands.	set subcommand

Command Syntax

This section explains the syntax of subcommands that you can write in your command file. This section also explains the syntax of the Java command that should be used in your batch file.

- **Conventions used in this section**

This section uses the following symbols and typefaces to explain syntax:

△

Indicates a space.

bold

Indicates characters that you must type as they are.

italics

Indicates a type of an operand. You do not need to type characters in italics as they are.

[]

Indicates one or more operands that can be omitted.

If two or more operands are enclosed by these square brackets and are delimited by vertical bars (|), you can select one of the operands.

{ }

Indicates that you must select one operand from the operands enclosed by the braces. Two or more operands are enclosed by the braces and are delimited by vertical bars (|).

...

Indicates that a previously used operand can be repeated.

Table 8-22 Syntax Descriptions

Syntax	The Syntax Indicates that You Can Write the Following Script
connect △ <i>ip-address</i>	connect 123.01.22.33
destination △[<i>directory</i>]	▪ destination ▪ destination c:\temp
compress △[yes no]	▪ compress ▪ compress yes ▪ compress no
answer △{ yes no }	▪ answer yes ▪ answer no
ports △[<i>name</i>][△...]	▪ ports ▪ ports port-1 ▪ ports port-1 port-2

- **Notes on writing script in the command file**

When you write a script in your command file, be aware of the following:

- Ensure that only one subcommand is used in one line.
- Empty lines in any command file will be ignored.
- Use a semicolon (;) if you want to insert a comment in your command file. If you enter a semicolon in one line, the remaining characters in that line will be regarded as a comment.

```
//////////  
;;;      COMMAND FILE: command.txt      ;;;;  
//////////  
svpip 158.214.135.57      ; IP address of SVP  
login expusr "passwd"      ; Log onto SVP
```

Figure 8-8 **Examples of Comments**

- **Viewing the Online Help for subcommands**

You can display the online Help to view the syntax of subcommands when you are working at the command prompt. To be able to view the online Help, you must use the **help** subcommand of the Export Tool. For detailed information about how to use the **help** subcommand, see help Subcommand.

svpip Subcommand

Syntax

svpip△{*ip-address*|*host-name*}

Description

The **svpip** subcommand specifies the IP address or the host name of the SVP.

Operands

ip-address

Specifies the IP address of the SVP. If the SVP is managed with IPv6 (Internet Protocol Version 6), you must specify the *ip-address* operand to match the format of IPv6. If the export tool runs on Windows XP, the interface identifier (for example, "%5") must be added to the end of the specified IP address.

host-name

Specifies the host name of the SVP.

If the host name includes any character that is neither an alphanumeric nor a period, the host name must be enclosed by double quotation marks ("").

Example

The following example specifies the IP address of the SVP as 158.214.127.170.

```
svpip 158.214.127.170
```

retry Subcommand

Syntax

retry△[**time**=*m*]△[**count**=*n*]

Description

The **retry** subcommand makes settings on retries of export processing.

When an internal error occurs during export processing, the Export Tool stops processing and then retries export processing. By default, the Export Tool can retry processing up to three times, but you can change the maximum number of retries by using the **retry** subcommand.

By default, the interval between one retry and the next retry is two minutes. You can change the interval by using the **retry** subcommand.

The **retry** subcommand must execute before the **login** subcommand executes.

Operands

time=*m*

Specifies the interval between retries in minutes.

m is a value within the range of 1 to 59.

If this operand is omitted, the interval between retries is two minutes.

count=*n*

Specifies the maximum number of retries.

If *n* is 0, the number of retries is unlimited.

If this operand is omitted, the maximum number of retries is 3.

Example

If the following command file is used, the interval between retries is five minutes and the maximum number of retries is 10.

```
svpip 158.214.135.57
retry time=5 count=10
login expusr passwd
show
group Port
short-range 200604010850:200604010910
outpath out
option compress
apply
```


login Subcommand

Syntax

login△*userid*△*password*

Description

The **login** subcommand uses a user ID and a password to log the specified user in the SVP.

The **svpip** subcommand must execute before the **login** subcommand executes.

The **login** subcommand fails if monitoring data does not exist in the SVP.

Operands

userid

Specifies the user ID for the SVP.

If the user ID includes any non-alphanumeric character, the user ID must be enclosed by double quotation marks (").



Note: Be sure to specify a user ID that should be used exclusively with the Export Tool. For detailed information, see Requirements for Using the Export Tool.

password

Specifies the password of the user.

If the password includes any non-alphanumeric character, the password ID must be enclosed by double quotation marks (").

Example

This example logs the user **expusr** into the SVP whose IP address is 158.214.127.170. The password is **pswd**.

```
svpip 158.214.127.170  
login expuser pswd
```

show Subcommand

Syntax

show

Description

The **show** subcommand outputs the following information to the standard output (for example, to the command prompt):

- the period during which monitoring data was collected onto the SVP (storing period)
- the interval at which the monitoring data was collected (gathering interval).

Performance Monitor collects statistics by the two types of storing periods: *in short range* and *in long range*. In short-range monitoring, the monitoring data between 8 hours and 15 days is stored in the SVP, and in long-range monitoring, the monitoring data up to three months is stored in the SVP. For details on the two storing periods, see Understanding Statistical Storage Ranges.

The **show** subcommand displays the storing period and the gathering interval for these two types of monitoring data: in short range and in long range. For example, the **show** subcommand outputs the following information:

Short Range	From: 2006/10/01 01:00 - To: 2006/10/01 15:00	Interval: 1min.
Long Range	From: 2006/09/01 00:00 - To: 2006/10/01 15:00	Interval: 15min.

`Short Range` indicates the storing period and gathering interval of the monitoring data stored in short range. `Long Range` indicates those of the monitoring data stored in long range. When you run the Export Tool, you can export the monitoring data within these periods into files. If you use the **short-range** or **long-range** subcommand additionally, you can narrow the term of data to be exported (see short-range Subcommand or long-range Subcommand).

`From` indicates the starting time for collecting monitoring data. `To` indicates the ending time for collecting monitoring data.

`Interval` indicates the interval at which the monitoring data was collected (gathering interval). For example, `Interval 15min.` indicates that monitoring data was collected at 15-minute intervals.

Storing periods output by the **show** subcommand is the same as the information displayed in the **Monitoring Term** area of the **Performance Management** window.

The **show** subcommand outputs the period from May 2, 2006, 03:12 to May 3, 2006, 03:12.

The screenshot shows a web interface for monitoring. On the left is a blue sidebar with the text "Monitoring Term". The main area has a light blue background. It contains two sets of date and time pickers. The first set, on the left, has a "From" field with the value "2006 /05 /03 01 :42" and a "To" field with the value "2006 /05 /03 03 :12". The second set, on the right, has a "From" field with the value "2006/05/02 03:12" and a "To" field with the value "2006/05/03 03:12". Red boxes highlight the "To" field of the first set and the "From" and "To" fields of the second set. A red line connects the text box above to the "From" field of the second set. Below the date pickers, there is a "Real Time" radio button, a dropdown menu showing the number "5", and the text "90 min.". An "Apply" button is located at the bottom right.

Figure 8-9 Information output by the show subcommand

The **login** command must execute before the **show** subcommand executes.

group Subcommand

Syntax

```
group△{PhyPG△[Short|Long]△[[parity-group-id]:[parity-group-id]][△...]|
PhyLDEV△[Short|Long]△[[parity-group-id]:[parity-group-id]][△...]|
PhyExG△[[exg-id]:[exg-id]][△...]|
PhyExLDEV△[[exg-id]:[exg-id]][△...]|
PhyProc△[Short|Long]|
PhyCSW△[Short|Long]|
PG△[[parity-group-id|V-VOL-group-id|exg-id]:
[parity-group-id|V-VOL-group-id|exg-id]][△...]|
LDEV△[[parity-group-id|V-VOL-group-id|exg-id]:
[parity-group-id|V-VOL-group-id|exg-id]][△...]|
Port△[[port-name]:[port-name]][△...]|
PortWWN△[[port-name]:[port-name]][△...]|
LU△[[port-name.host-group-id]:[port-name.host-group-id]][△...]|
PPCG△[[SPM-group-name]:[SPM-group-name]][△...]|
PPCGWWN△[[SPM-group-name]:[SPM-group-name]][△...]|
RemoteCopy△|
RCLU△[[port-name.host-group-id]:[port-name.host-group-id]][△...]|
RCLDEV△[[CU-id]:[CU-id]][△...]|
RCCLPR□|
UniversalReplicator△|
URJNL△[[JNL-group-id]:[JNL-group-id]][△...]|
URLU△[[port-name.host-group-id]:[port-name.host-group-id]][△...]|
URLDEV△[[CU-id]:[CU-id]][△...]|
}
```

Description

The **group** subcommand specifies the type of monitoring data that you want to export. This command uses an operand (such as **PhyPG** and **PhyLDEV** above) to specify a type of monitoring data.

Table 8-23 shows the monitoring data that can be saved into files by each operand, and the saved ZIP files. For details on the monitoring data saved in these files, refer to the tables indicated in the **See** column.

Table 8-23 Operands of the group Subcommand and Saved Monitoring Data

Operand	Window of Performance Monitor	Monitoring Data Saved in the File	Saved ZIP File	See
PhyPG	Physical tab in the Performance Management	Usage statistics about parity groups	PhyPG_dat.ZIP (*1)	Table 8-2
PhyLDEV		Usage statistics about volumes	PhyLDEV_dat.ZIP (*1)	

Operand	Window of Performance Monitor	Monitoring Data Saved in the File	Saved ZIP File	See
PhyExG	window	Usage conditions about external volume groups	PhyExG_dat.ZIP	
PhyExLDEV		Usage conditions about external volumes	PhyExLDEV_dat.ZIP	
PhyProc		Usage statistics about channel processors, disk processors, and data recovery and reconstruction processors	PhyProc_dat.ZIP (*1)	
PhyCSW		Usage statistics about access paths, write pending rate, and cache	PhyCSW_dat.ZIP (*1)	
PG	LDEV tab in the Performance Management window	Statistics about parity groups, external volume groups, or V-VOL groups	PG_dat.ZIP	Table 8-3
LDEV		Statistics about volumes in parity groups, in external volume groups, or in V-VOL groups	LDEV_XXXXX.ZIP (*2)	Table 8-4
Port	Port-LUN tab in the Performance Management window	Statistics about ports	Port_dat.ZIP	Table 8-5
PortWWN		Statistics about host bus adapters connected to ports	PortWWN_dat.ZIP	Table 8-6
LU		Statistics about LUs	LU_dat.ZIP	Table 8-7
PPCG	WWN tab in the Performance Management window	Statistics about SPM groups	PPCG_dat.ZIP	Table 8-8
PPCGWWN		Statistics about host bus adapters belonging to SPM groups	PPCGWWN_dat.ZIP	Table 8-9
RemoteCopy	TC Monitor window and TCz Monitor window	Statistics about remote copy operations by TrueCopy and TrueCopy for IBM z/OS (in complete volumes)	RemoteCopy_dat.ZIP	Table 8-10
RCLU		Statistics about remote copy operations by TrueCopy and TrueCopy for IBM z/OS (for each volume (LU))	RCLU_dat.ZIP	Table 8-11
RCLDEV		Statistics about remote copy operations by TrueCopy and TrueCopy for IBM z/OS (for volumes controlled by a particular CU)	RCLDEV_XXXXX.ZIP (*3)	Table 8-12
RCCLPR		Statistics about remote copy operations by TrueCopy and TrueCopy for IBM z/OS (at CLPR)	RCCLPR_dat.ZIP	Table 8-13

Operand	Window of Performance Monitor	Monitoring Data Saved in the File	Saved ZIP File	See
UniversalReplicator	UR Monitor window and URz Monitor window	Statistics about remote copy operations by Universal Replicator and Universal Replicator for IBM z/OS (for entire volumes)	UniversalReplicator.ZIP	Table 8-14
URJNL		Statistics about remote copy operations by Universal Replicator and Universal Replicator for IBM z/OS (for journal groups)	URJNL_dat.ZIP	Table 8-15
URLU		Statistics about remote copy operations by Universal Replicator and Universal Replicator for IBM z/OS (for each volume (LU))	URLU_dat.ZIP	Table 8-16
URLDEV		Statistics about remote copy operations by Universal Replicator and Universal Replicator for IBM z/OS (for volumes controlled by a particular CU)	URLDEV_XXXXX.ZIP (*4)	Table 8-17

*1: When you specify the **PhyPG**, **PhyLDEV**, **PhyProc**, or **PhyCSW** operand, you can select the storing period of the monitoring data to be exported from short range or long range. When you specify other operands, the monitoring data in short range is exported.

*2: A ZIP file whose name begins with `LDEV_`.

*3: A ZIP file whose name begins with `RCLDEV_`.

*4: A ZIP file whose name begins with `URLDEV_`.

You can use the **group** subcommand more than one time in a command file. For example, you can write the following script:

```
group PortWWN CL1-A:CL1-B
group PPCG spmg01:spmg02
group RemoteCopy
```

If an operand is used more than one time in a command file, the last operand takes effect. In the example below, the first **group** subcommand does not take effect, but the second **group** subcommand takes effect:

```
group PortWWN CL1-A:CL1-B
group PortWWN CL2-A:CL2-B
```

Operands

PhyPG△[**Short**|**Long**]△[[*parity-group-id*]:[*parity-group-id*]](△...]

Use this operand when you want to export statistics about parity group usage rates, which are displayed in the **Physical** tab of the **Performance Management** window. When statistics are exported to a ZIP file, the file name will be `PhyPG_dat.ZIP`. For details on the statistics exported by this operand, see Table 8-2.

You can use the **Short** or **Long** option to select the storing period of the monitoring data to be exported. If you specify **Short**, the exported file will contain statistics in a short range for up to 15 days. If you specify **Long**, the exported file will contain statistics in a long range for up to three months (i.e., up to 93 days). If neither **Short** nor **Long** is specified, statistics in both the short and long range are exported.

When you specify variables *parity-group-id*, you can narrow the range of parity groups whose monitoring data are to be exported. *parity-group-id* is a parity group ID. The colon (:) indicates a range. For example, `1-1:1-5` indicates parity groups from 1-1 to 1-5.

Ensure that the *parity-group-id* value on the left of the colon is smaller than the *parity-group-id* value on the right of the colon. For example, you can specify `PhyPG 1-1:1-5`, but you cannot specify `PhyPG 1-5:1-1`. Also, you can specify `PhyPG 1-5:2-1`, but you cannot specify `PhyPG 2-1:1-5`.

If *parity-group-id* is not specified, the monitoring data of all the parity groups will be exported.

PhyLDEV△[**Short**|**Long**]△[[*parity-group-id*]:[*parity-group-id*]](△...]

Use this operand when you want to export statistics about volume usage rates, which are displayed in the **Physical** tab of the **Performance Management** window. When statistics are exported to a ZIP file, the file name will be `PhyLDEV_dat.ZIP`. For details on the statistics exported by this operand, see Table 8-2.

You can use the **Short** or **Long** option to select the storing period of the monitoring data to be exported. If you specify **Short**, the exported file will contain statistics in short range for up to 15 days. If you specify **Long**, the exported file will contain statistics in long range for up to three months (i.e., up to 93 days). If neither **Short** nor **Long** is specified, statistics in both the short and long range are exported.

When you specify variables *parity-group-id*, you can narrow the range of parity groups whose monitoring data are to be exported. *parity-group-id* is a parity group ID. The colon (:) indicates a range. For example, `1-1:1-5` indicates parity groups from 1-1 to 1-5.

Ensure that the *parity-group-id* value on the left of the colon is smaller than the *parity-group-id* value on the right of the colon. For example, you can specify `PhyLDEV 1-1:1-5`, but you cannot specify `PhyLDEV 1-5:1-1`. Also, you can specify `PhyLDEV 1-5:2-1`, but you cannot specify `PhyLDEV 2-1:1-5`.

If *parity-group-id* is not specified, the monitoring data of all the volumes will be exported.

PhyExG△[[*exg-id*]:[*exg-id*]][△...]

Use this operand when you want to export statistics about external volume groups, which are displayed in the **Physical** tab of the **Performance Management** window. When statistics are exported to a ZIP file, the file name will be `PhyExG_dat.ZIP`. For details on the statistics exported by this operand, see Table 8-2.

When you specify variables *exg-id*, you can narrow the range of external volume groups whose monitoring data are to be exported. *exg-id* is an ID of an external volume group. The colon (:) indicates a range. For example, `E1-1:E1-5` indicates external volume groups from E1-1 to E1-5.

Ensure that the *exg-id* value on the left of the colon is smaller than the *exg-id* value on the right of the colon. For example, you can specify `PhyExG E1-1:E1-5`, but you cannot specify `PhyExG E1-5:E1-1`. Also, you can specify `PhyExG E1-5:E2-1`, but you cannot specify `PhyExG E2-1:E1-5`.

If *exg-id* is not specified, the monitoring data of all the external volume groups will be exported.

PhyExLDEV△[[*exg-id*]:[*exg-id*]][△...]

Use this operand when you want to export statistics about volumes in external volume groups, which are displayed in the **Physical** tab of the **Performance Management** window. When statistics are exported to a ZIP file, the file name will be `PhyExLDEV_dat.ZIP`. For details on the statistics exported by this operand, see Table 8-2.

When you specify variables *exg-id*, you can narrow the range of external volume groups whose monitoring data are to be exported. *exg-id* is an ID of an external volume group. The colon (:) indicates a range. For example, `E1-1:E1-5` indicates external volume groups from E1-1 to E1-5.

Ensure that the *exg-id* value on the left of the colon is smaller than the *exg-id* value on the right of the colon. For example, you can specify `PhyExLDEV E1-1:E1-5`, but you cannot specify `PhyExLDEV E1-5:E1-1`. Also, you can specify `PhyExLDEV E1-5:E2-1`, but you cannot specify `PhyExLDEV E2-1:E1-5`.

If *exg-id* is not specified, the monitoring data of all the external volumes will be exported.

PhyProc△[**Short**|**Long**]

Use this operand when you want to export the following statistics, which are displayed in the **Physical** tab of the **Performance Management** window:

- Usage rates of channel processors
- Usage rates of disk processors
- Usage rates of DRRs (data recovery and reconstruction processors)

When statistics are exported to a ZIP file, the file name will be `PhyProc_dat.ZIP`. For details on the statistics exported by this operand, see Table 8-2.

You can use the **Short** or **Long** option to select the storing period of the monitoring data to be exported. If you specify **Short**, the exported file will contain statistics in short range for up to 15 days. If you specify **Long**, the exported file will contain statistics in long range for up to three months (i.e., up to 93 days). If neither **Short** nor **Long** is specified, statistics in both the short and long range are exported.

PhyCSW△[**Short**|**Long**]

Use this operand when you want to export the following statistics, which are displayed in the **Physical** tab of the **Performance Management** window:

- Usage rates of access paths between channel adapters and cache memories
- Usage rates of access paths between disk adapters and cache memories
- Usage rates of access paths between channel adapters and the shared memory
- Usage rates of access paths between disk adapters and the shared memory
- Usage rates of access paths between cache switches and cache memories
- Write pending rates

When statistics are exported to a ZIP file, the file name will be `PhyCSW_dat.ZIP`. For details on the statistics exported by this operand, see Table 8-2.

You can use the **Short** or **Long** option to select the storing period of the monitoring data to be exported. If you specify **Short**, the exported file will contain statistics in short range for up to 15 days. If you specify **Long**, the exported file will contain statistics in long range for up to three months (i.e., up to 93 days). If neither **Short** nor **Long** is specified, statistics in both the short and long range are exported.

PG△[[*parity-group-id*|*V-VOL-group-id*|*exg-id*]:
[*parity-group-id*|*V-VOL-group-id*|*exg-id*]][△...]

Use this operand when you want to export statistics about parity groups, external volume groups, or V-VOL groups which are displayed in the **LDEV** tab of the **Performance Management** window. When statistics are exported to a ZIP file, the file name will be PG_dat.ZIP. For details on the statistics exported by this operand, see Table 8-3.

When you specify variables *parity-group-id*, *exg-id*, or *V-VOL-group-id*, you can narrow the range of parity groups, external volume groups, or V-VOL groups, whose monitoring data are to be exported. *parity-group-id* is a parity group ID. *exg-id* is an ID of an external volume group. *V-VOL-group-id* is V-VOL group ID. The colon (:) indicates a range. For example, 1-1:1-5 indicates parity groups from 1-1 to 1-5. E1-1:E1-5 indicates external volume groups from E1-1 to E1-5. V1-1:V5-1 indicates V-VOL groups from V1-1 to V5-1. X1-1:X5-1 indicates V-VOL groups from X1-1 to X5-1.

Ensure that the *parity-group-id*, *exg-id*, or *V-VOL-group-id* value on the left of the colon is smaller than the *parity-group-id*, *exg-id*, or *V-VOL-group-id* value on the right of the colon. For example, you can specify PG 1-1:1-5, but you cannot specify PG 1-5:1-1. Also, you can specify PG 1-5:2-1, but you cannot specify PG 2-1:1-5.

If neither of *parity-group-id*, *exg-id*, nor *V-VOL-group-id* is specified, the monitoring data of all the parity groups, external volume groups, and V-VOL group will be exported.

LDEV△[[*parity-group-id*|*V-VOL-group-id*|*exg-id*]:
[*parity-group-id*|*V-VOL-group-id*|*exg-id*]][△...]

Use this operand when you want to export statistics about volumes, which are displayed in the **LDEV** tab of the **Performance Management** window. When statistics are exported to a ZIP file, multiple ZIP files whose names are beginning with LDEV_ will be output. For details on the statistics exported by this operand, see Table 8-5.

When you specify variables *parity-group-id*, *exg-id*, or *V-VOL-group-id*, you can narrow the range of parity groups, external volume groups, or V-VOL groups, whose monitoring data are to be exported. *parity-group-id* is a parity group ID. *exg-id* is an ID of an external volume group. *V-VOL-group-id* is V-VOL group ID. The colon (:) indicates a range. For example, 1-1:1-5 indicates parity groups from 1-1 to 1-5. E1-1:E1-5 indicates external volume groups from E1-1 to E1-5. V1-1:V5-1 indicates V-VOL groups from V1-1 to V5-1. X1-1:X5-1 indicates V-VOL groups from X1-1 to X5-1.

Ensure that the *parity-group-id*, *exg-id*, or *V-VOL-group-id* value on the left of the colon is smaller than the *parity-group-id*, *exg-id*, or *V-VOL-group-id* value on the right of the colon. For example, you can specify LDEV 1-1:1-5, but you cannot specify LDEV 1-5:1-1. Also, you can specify LDEV 1-5:2-1, but you cannot specify LDEV 2-1:1-5.

If neither of *parity-group-id*, *exg-id*, nor *V-VOL-group-id* is specified, the monitoring data of all the volumes (including external volumes and V-VOL groups) will be exported.

Port△[[*port-name*]:[*port-name*]][△...]

Use this operand when you want to export port statistics, which are displayed in the **Port-LUN** tab of the **Performance Management** window. When statistics are exported in a ZIP file, the file name will be Port_dat.ZIP. For details on the statistics exported by this operand, see Table 8-5.

When you specify variables *port-name*, you can narrow the range of ports whose monitoring data are to be exported. *port-name* is a port name. The colon (:) indicates a range. For example, CL3-a:CL3-c indicates ports from CL3-a to CL3-c.

Ensure that the *port-name* value on the left of the colon is smaller than the *port-name* value on the right of the colon. The smallest *port-name* value is CL1-A and the largest *port-name* value is CL4-r. The following formula illustrates which value is smaller than which value:

$$\text{CL1-A} < \text{CL1-B} < \dots < \text{CL2-A} < \text{CL2-B} < \dots < \text{CL3-a} < \text{CL3-b} < \dots < \text{CL4-a} < \dots < \text{CL4-r}$$

For example, you can specify Port CL1-C:CL2-A, but you cannot specify Port CL2-A:CL1-C. Also, you can specify Port CL3-a:CL3-c, but you cannot specify Port CL3-c:CL3-a.

If *port-name* is not specified, the monitoring data of all the ports will be exported.

PortWWN△[[*port-name*]:[*port-name*]][△...]

Use this operand when you want to export statistics about host bus adapters (WWNs) connected to ports, which are displayed in the **Port-LUN** tab of the **Performance Management** window. When statistics are exported in a ZIP file, the file name will be PortWWN_dat.ZIP. For details on the statistics exported by this operand, see Table 8-6.

When you specify variables *port-name*, you can narrow the range of ports whose monitoring data are to be exported. *port-name* is a port name. The colon (:) indicates a range. For example, CL3-a:CL3-c indicates ports from CL3-a to CL3-c.

Ensure that the *port-name* value on the left of the colon is smaller than the *port-name* value on the right of the colon. The smallest *port-name* value is CL1-A and the largest *port-name* value is CL4-r. The following formula illustrates which value is smaller than which value:

$$\text{CL1-A} < \text{CL1-B} < \dots < \text{CL2-A} < \text{CL2-B} < \dots < \text{CL3-a} < \text{CL3-b} < \dots < \text{CL4-a} < \dots < \text{CL4-r}$$

For example, you can specify PortWWN CL1-C:CL2-A, but you cannot specify PortWWN CL2-A:CL1-C. Also, you can specify PortWWN CL3-a:CL3-c, but you cannot specify PortWWN CL3-c:CL3-a.

If *port-name* is not specified, the monitoring data of all the host bus adapters will be exported.

LU△[[*port-name.host-group-id*]:[*port-name.host-group-id*]][△...]

Use this operand when you want to export statistics about LU paths, which are displayed in the **Port-LUN** tab of the **Performance Management** window. When statistics are exported in a ZIP file, the file name will be LU_dat.ZIP. For details on the statistics exported by this operand, see Table 8-7.

When you specify variables *port-name.host-group-id*, you can narrow the range of LU paths whose monitoring data are to be exported. *port-name* is a port name. *host-group-id* is the ID of a host group (that is, a host storage domain). The host group (host storage domain) ID must be a hexadecimal numeral. The colon (:) indicates a range. For example, CL1-C.01:CL1-C.03 indicates the range from the host group #01 of the CL1-C port to the host group #03 of the CL1-C port.

Ensure that the value on the left of the colon is smaller than the value on the right of the colon. The smallest *port-name* value is CL1-A and the largest *port-name* value is CL4-r. The following formula illustrates which *port-name* value is smaller than which *port-name* value:

CL1-A < CL1-B < ... < CL2-A < CL2-B < ... < CL3-a < CL3-b < ... < CL4-a < ... < CL4-r

For example, you can specify LU CL1-C.01:CL2-A.01, but you cannot specify LU CL2-A.01:CL1-C.01. Also, you can specify LU CL1-C.01:CL1-C.03, but you cannot specify LU CL1-C.03:CL1-C.01.

If *port-name.host-group-id* is not specified, the monitoring data of all the LU paths will be exported.

PPCG△[[*SPM-group-name*]:[*SPM-group-name*]][△...]

Use this operand when you want to export statistics about SPM groups, which are displayed in the **WWN** tab of the **Performance Management** window. When statistics are exported in a ZIP file, the file name will be PPCG_dat.ZIP. For details on the statistics exported by this operand, see Table 8-8.

When you specify variables *SPM-group-name*, you can narrow the range of SPM groups whose monitoring data are to be exported. *SPM-group-name* is the name of an SPM group. If the name includes any non-alphanumeric character, the name must be enclosed by double quotation marks (""). The colon (:) indicates a range. For example, Grp01:Grp03 indicates a range of SPM groups from Grp01 to Grp03.

Ensure that the *SPM-group-name* value on the left of the colon is smaller than the *SPM-group-name* value on the right of the colon. Numerals are smaller than letters and lowercase letters are smaller than uppercase letters. In the following formulae, values are arranged so that smaller values are on the left and larger values are on the right:

– 0 < 1 < 2 < ... < 9 < a < b < ... < z < A < B < ... < Z

- cygnus < raid < Cancer < Pisces < RAID < RAID5

Regardless of whether you specify or omit SPM group names, the exported CSV files contain statistics about host bus adapters that do not belong to any SPM groups. The exported CSV files use the heading named `Not Grouped` to indicate statistics about these host bus adapters.

If *SPM-group-name* is not specified, the monitoring data of all the SPM groups will be exported.

PPCGWWN△[[*SPM-group-name*]:[*SPM-group-name*]](△...]

Use this operand when you want to export statistics about host bus adapters (WWNs) belonging to SPM groups, which are displayed in the **WWN** tab of the **Performance Management** window. When statistics are exported in a ZIP file, the file name will be PPCGWWN_dat.ZIP. For details on the statistics exported by this operand, see Table 8-9.

When you specify variables *SPM-group-name*, you can narrow the range of SPM groups whose monitoring data are to be exported. *SPM-group-name* is the name of an SPM group. If the name includes any non-alphanumeric character, the name must be enclosed by double quotation marks ("). The colon (:) indicates a range. For example, Grp01:Grp03 indicates a range of SPM groups from Grp01 to Grp03.

Ensure that the *SPM-group-name* value on the left of the colon is smaller than the *SPM-group-name* value on the right of the colon. Numerals are smaller than letters and lowercase letters are smaller than uppercase letters. In the following formulae, values are arranged so that smaller values are on the left and larger values are on the right:

- 0 < 1 < 2 < ... < 9 < a < b < ... < z < A < B < ... < Z
- cygnus < raid < Cancer < Pisces < RAID < RAID5

If *SPM-group-name* is not specified, the monitoring data of all the host bus adapters will be exported.

RemoteCopy

Use this operand when you want to export statistics about remote copy operations which are displayed in the **TC Monitor** window and the **TCz Monitor** window. By using this operand, you can export monitoring data about remote copy operations performed by TrueCopy and TrueCopy for IBM z/OS in the whole volumes. When statistics are exported to a ZIP file, the file name will be RemoteCopy_dat.ZIP. For details on the statistics exported by this operand, see Table 8-10.

RCLU△[[*port-name.host-group-id*]:[*port-name.host-group-id*]][△...]

Use this operand when you want to export statistics about remote copy operations which are displayed in the **TC Monitor** window and the **TCz Monitor** window. By using this operand, you can export monitoring data about remote copy operations performed by TrueCopy and TrueCopy for IBM z/OS at each volume (LU). When statistics are exported to a ZIP file, the file name will be RCLU_dat.ZIP. For details on the statistics exported by this operand, see Table 8-11.

When you specify variables *port-name.host-group-id*, you can narrow the range of LU paths whose monitoring data are to be exported. *port-name* is a port name. *host-group-id* is the ID of a host group. The host group ID must be a hexadecimal numeral. The colon (:) indicates a range. For example, CL1-C.01:CL1-C.03 indicates the range from the host group #01 of the CL1-C port to the host group #03 of the CL1-C port.

Ensure that the value on the left of the colon is smaller than the value on the right of the colon. The smallest *port-name* value is CL1-A and the largest *port-name* value is CL4-r. The following formula illustrates which *port-name* value is smaller than which *port-name* value:

CL1-A < CL1-B < ... < CL2-A < CL2-B < ... < CL3-a < CL3-b < ... < CL4-a < ... < CL4-r

For example, you can specify RCLU CL1-C.01:CL2-A.01, but you cannot specify RCLU CL2-A.01:CL1-C.01. Also, you can specify RCLU CL1-C.01:CL1-C.03, but you cannot specify RCLU CL1-C.03:CL1-C.01.

If *port-name.host-group-id* is not specified, the monitoring data of all the volumes (LUs) will be exported.

RCLDEV△[[*LDKC-CU-id*]:[*LDKC-CU-id*]][△...]

Use this operand when you want to export statistics about remote copy operations which are displayed in the **TC Monitor** window and the **TCz Monitor** window. By using this operand, you can export monitoring data about remote copy operations performed by TrueCopy and TrueCopy for IBM z/OS at volumes controlled by each CU. When statistics are exported to a ZIP file, multiple ZIP files whose names are beginning with RCLDEV_ will be output. For details on the statistics exported by this operand, see Table 8-12.

When you specify variables *LDKC-CU-id*, you can narrow the range of LDKC:CUs that control the volumes whose monitoring data are to be exported. *LDKC-CU-id* is an ID of a LDKC:CU. The colon (:) indicates a range. For example, 000:105 indicates LDKC:CUs from 00:00 to 01:05.

Ensure that the *LDKC-CU-id* value on the left of the colon is smaller than the *LDKC-CU-id* value on the right of the colon. For example, you can specify RCLDEV 000:105, but you cannot specify RCLDEV 105:000.

If *LDKC-CU-id* is not specified, the monitoring data of all the volumes will be exported.

RCCLPR

Use this operand when you want to export statistics about remote copy operations which are displayed in the **TC Monitor** window and the **TCz Monitor** window. By using this operand, you can export monitoring data about remote copy operations performed by TrueCopy and TrueCopy for IBM z/OS at each CLPR. When statistics are exported to a ZIP file, the file name will be RCCLPR_dat.ZIP. For details on the statistics exported by this operand, see Table 8-13.



Note: Monitoring data are grouped by SLPR, and are exported per CLPR. If there are two SLPRs that are *SLPR0* (corresponding to CLPR0, and CLPR2) and *SLPR1* (corresponding to CLPR1), those CLPRs are arranged as follows: CLPR0, CLPR2, CLPR1.

UniversalReplicator

Use this operand when you want to export statistics about remote copy operations which are displayed in the **UR Monitor** window and the **URz Monitor** window. By using this operand, you can export monitoring data about remote copy operations performed by Universal Replicator and Universal Replicator for IBM z/OS in the whole volume. When statistics are exported to a ZIP file, the file name will be UniversalReplicator.ZIP. For details on the statistics exported by this operand, see Table 8-14.

URJNL△[*JNL-group-id*]:[*JNL-group-id*][△...]

Use this operand when you want to export statistics about remote copy operations which are displayed in the **UR Monitor** window and the **URz Monitor** window. By using this operand, you can export monitoring data about remote copy operations performed by Universal Replicator and Universal Replicator for IBM z/OS at each journal group. When statistics are exported to a ZIP file, the file name will be URJNL_dat.ZIP. For details on the statistics exported by this operand, see Table 8-15.

When you specify variables *JNL-group-id*, you can narrow the range of journal groups whose monitoring data are to be exported. *JNL-group-id* is a journal group number. The colon (:) indicates a range. For example, 00:05 indicates journal groups from 00 to 05.

Ensure that the *JNL-group-id* value on the left of the colon is smaller than the *JNL-group-id* value on the right of the colon. For example, you can specify URJNL 00:05, but you cannot specify URJNL 05:00.

If *JNL-group-id* is not specified, the monitoring data of all the journal volumes will be exported.

URLU△[[*port-name.host-group-id*]:[*port-name.host-group-id*]][△...]

Use this operand when you want to export statistics about remote copy operations which are displayed in the **UR Monitor** window and the **URz Monitor** window. By using this operand, you can export monitoring data about remote copy operations performed by Universal Replicator and Universal Replicator for IBM z/OS at each volume(LU). When statistics are exported to a ZIP file, the file name will be URLU_dat.ZIP. For details on the statistics exported by this operand, see Table 8-16.

When you specify variables *port-name.host-group-id*, you can narrow the range of LU paths whose monitoring data are to be exported. *port-name* is a port name. *host-group-id* is the ID of a host group. The host group ID must be a hexadecimal numeral. The colon (:) indicates a range. For example, CL1-C.01:CL1-C.03 indicates the range from the host group #01 of the CL1-C port to the host group #03 of the CL1-C port.

Ensure that the value on the left of the colon is smaller than the value on the right of the colon. The smallest *port-name* value is CL1-A and the largest *port-name* value is CL4-r. The following formula illustrates which *port-name* value is smaller than which *port-name* value:

CL1-A < CL1-B < ... < CL2-A < CL2-B < ... < CL3-a < CL3-b < ... < CL4-a < ... < CL4-r

For example, you can specify URLU CL1-C.01:CL2-A.01, but you cannot specify URLU CL2-A.01:CL1-C.01. Also, you can specify URLU CL1-C.01:CL1-C.03, but you cannot specify URLU CL1-C.03:CL1-C.01.

If *port-name.host-group-id* is not specified, the monitoring data of all the volumes (LUs) will be exported.

URLDEV△[[*LDKC-CU-id*]:[*LDKC-CU-id*]][△...]

Use this operand when you want to export statistics about remote copy operations which are displayed in the **UR Monitor** window and the **URz Monitor** window. By using this operand, you can export monitoring data about remote copy operations performed by Universal Replicator and Universal Replicator for IBM z/OS at volumes controlled by each CU. When statistics are exported to a ZIP file, multiple ZIP files whose names are beginning with URLDEV_ will be output. For details on the statistics exported by this operand, see Table 8-17.

When you specify variables *LDKC-CU-id*, you can narrow the range of LDKC:CUs that control the volumes whose monitoring data are to be exported. *LDKC-CU-id* is an ID of a LDKC:CU. The colon (:) indicates a range. For example, 000:105 indicates LDKC:CUs from 00:00 to 01:05.

Ensure that the *LDKC-CU-id* value on the left of the colon is smaller than the *LDKC-CU-id* value on the right of the colon. For example, you can specify URLDEV 000:105, but you cannot specify URLDEV 105:000.

If *LDKC-CU-id* is not specified, the monitoring data of all the volumes will be exported.

Examples

The following example exports statistics about host bus adapters and SPM groups:

```
group PortWWN
group PPCG
```

The following example exports statistics about three ports (CL1-A, CL1-B, and CL1-C):

```
group Port CL1-A:CL1-C
```

The following example exports statistics about six ports (CL1-A to CL1-C, and CL2-A to CL2-C)

```
group Port CL1-A:CL1-C CL2-A:CL2-C
```

The following example exports statistics about the parity group 1-3:

```
group PG 1-3:1-3
```

The following example exports statistics about the parity group 1-3 and other parity groups whose ID is larger than 1-3 (for example, 1-4 and 1-5):

```
group PG 1-3:
```

The following example exports statistics about the external volume groups E1-1 to E1-5:

```
group PG E1-1:E1-5
```

The following example exports statistics about the parity group 1-3 and other parity groups whose ID is smaller than 1-3 (for example, 1-1 and 1-2):

```
group LDEV:1-3
```

The following example exports statistics about LU paths for the host group (host storage domain) ID *01* for the port *CL1-A*:

```
group LU CL1-A.01:CL1-A.01
```

short-range Subcommand

Syntax

short-range△[[yyyyMMddhhmm][{+|-}hhmm]:[yyyyMMddhhmm][{+|-}hhmm]]

Description

The **short-range** subcommand enables you to specify a term of monitoring data to be exported into files. Use this subcommand when you want to narrow the export-target term within the stored data.

The **short-range** subcommand is valid for monitoring data in short range. The monitoring data in short range is the contents displayed in the following windows:

- The **Performance Management** window with selecting **short-range** as the storing period
- The **TC Monitor** and **TCz Monitor** windows
- The **UR Monitor** and **URz Monitor** windows

All the monitoring items are stored in short range. Therefore, you can use the **short-range** subcommand whichever operand you specify to the **group** subcommand. If you run the Export Tool without specifying the **short-range** subcommand, the data stored in the whole monitoring term will be exported.

The **login** subcommand must execute before the **short-range** subcommand executes.

Operands

The value on the left of the colon (:) specifies the starting time of the period. The value on the right of the colon specifies the ending time of the period. Specify the term within "Short Range From XXX To XXX" which is output by the **show** subcommand.

If no value is specified on the left of the colon, the starting time for collecting monitoring data is assumed. If no value is specified on the right of the colon, the ending time for collecting monitoring data is assumed. The starting and ending times for collecting monitoring data are displayed in the **Monitoring Term** area in the **Performance Management** window.

Starting time for collecting statistics

Ending time for collecting statistics

Monitoring Term

From 2006/05/03 01:42 To 2006/05/03 03:12

Real Time 5 90 min. Apply

Figure 8-10 Starting and Ending Time for Collecting Monitoring Data

yyyyMMddhhmm

yyyyMMdd indicates the year, the month, and the day. *hhmm* indicates the hour and the minute.

If *yyyyMMddhhmm* is omitted on the left of the colon, the starting time for collecting monitoring data is assumed. If *yyyyMMddhhmm* is omitted on the right of the colon, the ending time for collecting monitoring data is assumed.

+hhmm

Adds time (*hhmm*) to *yyyyMMddhhmm* if *yyyyMMddhhmm* is specified. For example, **200601230000+0130** indicates Jan. 23, 2006. 01:30.

Adds time to the starting time for collecting monitoring data, if *yyyyMMddhhmm* is omitted.

-hhmm

Subtracts time (*hhmm*) from *yyyyMMddhhmm* if *yyyyMMddhhmm* is specified. For example, **200601230000-0130** indicates Jan. 22, 2006. 22:30.

Subtracts time from the ending time for collecting monitoring data, if *yyyyMMddhhmm* is omitted.



Note: If the last two digit of the time on the left or right of the colon (:) is not a multiple of the sampling interval, the time will automatically be changed so that the last two digits is a multiple of the sampling interval. If this change occurs to the time on the left of the colon, the time will be smaller than the original time. If this change occurs to the time on the right of the colon, the time will be larger than the original time. The following are the examples:

- **If the time on the left is 10:15, the time on the right is 20:30, and the sampling interval is 10 minutes:**

The time on the left will be changed to 10:10 because the last two digits of the time is not a multiple of 10 minutes. The time on the right will remain unchanged because the last two digits of the time is a multiple of 10 minutes.

- **If the time on the left is 10:15, the time on the right is 20:30, and the sampling interval is 7 minutes:**

The time on the left will be changed to 10:14 because the last two digits of the time is not a multiple of 7 minutes. The time on the right will be changed to 20:35 because of the same reason.

Examples

The examples below assume that:

- the starting time for collecting monitoring data is Jan. 1, 2006, 00:00,
- the ending time for collecting monitoring data is Jan. 2, 2006, 00:00.

`short-range 200601010930:200601011730`

The Export Tool saves monitoring data within the range of Jan. 1, 9:30-17:30.

`short-range 200601010930:`

The Export Tool saves monitoring data within the range of Jan. 1, 9:30 to Jan. 2, 00:00.

`shortrange:200601011730`

The Export Tool saves monitoring data within the range of Jan. 1, 0:00-17:30.

`short-range +0001:`

The Export Tool saves monitoring data within the range of Jan. 1, 0:01 to Jan. 2, 00:00.

`short-range -0001:`

The Export Tool saves monitoring data within the range of Jan. 1, 23:59 to Jan. 2, 00:00.

`shortrange:+0001`

The Export Tool saves monitoring data within the range of Jan. 1, 0:00-00:01.

`shortrange:-0001`

The Export Tool saves monitoring data within the range of Jan. 1, 0:00-23:59.

`short-range +0101:-0101`

The Export Tool saves monitoring data within the range of Jan. 1, 1:01-22:59.

`short-range 200601010900+0130:200601011700-0130`

The Export Tool saves monitoring data within the range of Jan. 1, 10:30-15:30.

short-range 200601010900-0130:200601011700+0130

The Export Tool saves monitoring data within the range of Jan. 1, 7:30-18:30.

short-range 200601010900-0130:

The Export Tool saves monitoring data within the range of Jan. 1, 7:30 to Jan. 2, 00:00.

long-range Subcommand

Syntax

long-range△[[yyyyMMddhhmm][{+|-}ddhhmm]:[yyyyMMddhhmm][{+|-}ddhhmm]]

Description

The **long-range** subcommand enables you to specify a term of monitoring data to be exported into files. Use this subcommand when you want to narrow the export-target term within the stored data.

The **long-range** subcommand is valid for monitoring data in long range. The monitoring data in long range is the contents displayed in the **Physical** tab of the **Performance Management** window with selecting **long-range** as the storing period. The monitoring items whose data can be stored in long range are limited. Table 8-24 shows the monitoring items to which the **long-range** subcommand can be applied, and also shows the operands to export those monitoring items.

Table 8-24 Monitoring Items To Which the long-range Subcommand Can be Applied

Monitoring Data	Operands of the group subcommand
Usage statistics about parity groups	PhyPG Long
Usage statistics about volumes	PhyLDEV Long
Usage statistics about channel processors, disk processors, and data recovery and reconstruction processors	PhyProc Long
Usage statistics about access paths and write pending rate	PhyCSW Long

If you run the Export Tool without specifying the **long-range** subcommand, the data stored in the whole monitoring term will be exported.

The **login** subcommand must execute before the **long-range** subcommand executes.

Operands

The value on the left of the colon (:) specifies the starting time of the period. The value on the right of the colon specifies the ending time of the period. Specify the term within "Long Range From XXX To XXX" which is output by the **show** subcommand.

If no value is specified on the left of the colon, the starting time for collecting monitoring data is assumed. If no value is specified on the right of the colon, the ending time for collecting monitoring data is assumed. The starting and ending times for collecting monitoring data are displayed in the **Monitoring Term** area in the **Performance Management** window.




Figure 8-11 Starting and Ending Time for Collecting Monitoring Data

yyyyMMddhhmm

yyyyMMdd indicates the year, the month, and the day. *hhmm* indicates the hour and the minute.

If *yyyyMMddhhmm* is omitted on the left of the colon, the starting time for collecting monitoring data is assumed. If *yyyyMMddhhmm* is omitted on the right of the colon, the ending time for collecting monitoring data is assumed.

+ddhhmm

Adds time (*ddhhmm*) to *yyyyMMddhhmm* if *yyyyMMddhhmm* is specified. For example, **200601120000+010130** indicates Jan. 13, 2006. 01:30.

Adds time to the starting time for collecting monitoring data, if *yyyyMMddhhmm* is omitted.

-ddhhmm

Subtracts time (*ddhhmm*) from *yyyyMMddhhmm* if *yyyyMMddhhmm* is specified. For example, **200601120000-010130** indicates Jan. 10, 2006. 22:30.

Subtracts time from the ending time for collecting monitoring data, if *yyyyMMddhhmm* is omitted.



Note: Ensure that *mm* is 00, 15, 30, or 45. If you do not specify *mm* in this way, the value on the left of the colon (:) will be rounded down to one of the four values. Also, the value on the right of the colon will be rounded up to one of the four values. For example, if you specify 200601010013:200601010048, the specified value is regarded as 200601010000:200601010100.

Examples

The examples below assume that:

- the starting time for collecting monitoring data is Jan. 1, 2006, 00:00,
- the ending time for collecting monitoring data is Jan. 2, 2006, 00:00.

`long-range 200601010930:200601011730`

The Export Tool saves monitoring data within the range of Jan. 1, 9:30-17:30.

`long-range 200601010930:`

The Export Tool saves monitoring data within the range of Jan. 1, 9:30 to Jan. 2, 00:00.

`longrange:200601011730`

The Export Tool saves monitoring data within the range of Jan. 1, 0:00-17:30.

`long-range +000015:`

The Export Tool saves monitoring data within the range of Jan. 1, 0:15 to Jan. 2, 00:00.

`long-range -000015:`

The Export Tool saves monitoring data within the range of Jan. 1, 23:45 to Jan. 2, 00:00.

`longrange:+000015`

The Export Tool saves monitoring data within the range of Jan. 1, 0:00-00:15.

`longrange:-000015`

The Export Tool saves monitoring data within the range of Jan. 1, 0:00-23:45.

`long-range +000115:-000115`

The Export Tool saves monitoring data within the range of Jan. 1, 1:15-22:45.

`long-range 200601010900+000130:200601011700-000130`

The Export Tool saves monitoring data within the range of Jan. 1, 10:30-15:30.

long-range 200601010900-000130:200601011700+000130

The Export Tool saves monitoring data within the range of Jan. 1, 7:30-18:30.

long-range 200601010900-000130:

The Export Tool saves monitoring data within the range of Jan. 1, 7:30 to Jan. 2, 00:00.

outpath Subcommand

Syntax

outpath△[*path*]

Description

The **outpath** subcommand specifies the directory to which monitoring data will be exported.

Operands

path

Specifies the directory in which files will be saved.

If the directory includes any non-alphanumeric character, the directory must be enclosed by double quotation marks ("). If you want to specify a back slash (\) in the character string enclosed by double quotation marks, repeat the back slash twice such as \\.

If the specified directory does not exist, this subcommand creates a directory that has the specified name.

If this operand is omitted, the current directory is assumed.

Examples

The following example saves files in the directory **C:\Project\out** in a Windows computer:

```
outpath "C:\\Project\\out"
```

The following example saves files in the **out** directory in the current directory:

```
outpath out
```

option Subcommand

Syntax

option△[**compress**|**nocompress**]△[**ask**|**clear**|**noclear**]

Description

The **option** subcommand specifies the following:

- whether to compress monitoring data in ZIP files
- whether to overwrite or delete existing files and directories when saving monitoring data in files

Operands

The two operands below specify whether to compress CSV files into ZIP files. If none of these operands is specified, **compress** is assumed:

compress

Compresses data in ZIP files. To extract CSV files out of a ZIP file, you will need to decompress the ZIP file.

nocompress

Does not compress data in ZIP files and saves data in CSV files.

The three operands below specify whether to overwrite or delete an existing file or directory when the Export Tool saves files. If none of these operands is specified, **ask** is assumed:

ask

Displays a message that asks whether to delete existing files or directories.

clear

Deletes existing files and directories and then saves monitoring data in files.

noclear

Overwrites existing files and directories.

Example

The following example saves monitoring data in CSV files, not in ZIP files:

```
option nocompress
```

apply Subcommand

Syntax

apply

Description

The **apply** subcommand saves monitoring data specified by the **group** subcommand into files.

The **login** subcommand must execute before the **apply** subcommand executes.

The **apply** subcommand does nothing if the **group** subcommand executes.

The settings made by the **group** subcommand will be reset when the **apply** subcommand finishes.

set subcommand

Syntax

set△[**switch**={*m*|**off**}]

Description

The **set** subcommand starts or ends monitoring the storage system (i.e., starts or ends collecting performance statistics). The **set** subcommand also specifies the gathering interval (interval of collecting statistics) in short range monitoring.

If you want to use the **set** subcommand, you must use the **login** subcommand (see login Subcommand) to log on to the SVP. Ensure that the **set** subcommand executes immediately before the Export Tool finishes.

Executing the **set** subcommand generates an error in the following conditions:

- Some other user is being logged onto the SVP in Modify mode.
- Maintenance operations are being performed at the SVP.

If an error occurs, do the following:

- Ensure that all the users who are logged onto the SVP are not in Modify mode. If any user is logged on in Modify mode, ask the user to switch to View mode.
- Wait until maintenance operations finish at the SVP, so that the **set** subcommand can execute.

**Notes:**

- Your batch files can include script that should execute when an error occurs. For information about writing such a script in your batch file, refer to **Notes** in Running the Export Tool.
 - When the **set** subcommand starts or ends the monitoring or changes the gathering interval after the **Performance Management** window is started, the contents displayed in the **Performance Management** window does not change automatically in conjunction with the **set** subcommand operation. To display the current monitoring status in the **Performance Management** window, click **File**, and then **Refresh** on the menu bar of the Storage Navigator main window.
 - If you change the specified gathering interval during a monitoring, the previously gathered monitoring data will be deleted.
-

Operands

switch={*m*|**off**}

To start monitoring, specify the gathering interval (interval of collecting statistics) of monitoring data at *m*. Specify a value between 1 and 15 in *minutes*. *m* is the gathering interval in short range monitoring by Performance Monitor. The gathering interval in long range is fixed to 15 minutes.

To end monitoring, specify **off**.

If this operand is omitted, the **set** subcommand does not make settings for starting or ending monitoring.

Examples

The following command file saves port statistics and then ends monitoring ports:

```
svpip 158.214.135.57
login expusr passwd
show
group Port
short-range 200604010850:200604010910
apply
set switch=off
```

The following command file starts monitoring remote copy operations. The sampling time interval is 10 minutes:

```
svpip 158.214.135.57
login expusr passwd
set switch=10
```

help Subcommand

Syntax

help

Description

The **help** subcommand displays the online help for subcommands.

If you want to view the online help, It is recommended that you create a batch file and a command file that are exclusively used for displaying the online help. For detailed information, refer to *Example* below.

Example

In this example, a command file (cmdHelp.txt) and a batch file (runHelp.bat) are created in the **C:\export** directory in a Windows computer:

- Command file (c:\export\cmdHelp.txt):

```
help
```

- Batch file (c:\export\runHelp.bat):

```
Java -classpath ".\lib\JSanExport.jar;.\lib\JSanRmiServerSx.jar" -Xmx536870912 -  
Dmd.command=cmdHelp.txt  
-Dmd.logpath=log sanproject.getmondat.RJMdMain ↵  
pause ↵
```



Note: The "↵" symbol in this batch file example indicates the end of a command line.

In this example, you must do one of the following to view the online Help:

- Double-click runHelp.bat with the mouse.
- Go to the **c:\export** directory at the command prompt, enter **runHelp** or **runHelp.bat** and then press the <Enter> key.

Java Command for Exporting Data In Files

Syntax

Java△**-classpath**△*class-path*△*property-parameters*△**sanproject.getmondat.RJMdMain**

Description

This Java command starts the Export Tool.

To start the Export Tool, you must write this Java command in your batch file and then run the batch file.

Operands

class-path

Specifies the path to the class file of the Export Tool.

The path must be enclosed in double quotation marks ("").

property-parameters

You can specify the following parameters. At least you must specify - **Dmd.command**.

- **-Dhttp.proxyHost**=host-name-of-proxy-host, or
-Dhttp.proxyHost=IP-address-of-proxy-host

Specifies the host name or the IP address of a proxy host. You must specify this parameter if the computer that runs the Export Tool communicates with the SVP via a proxy host.

- **-Dhttp.proxyPort**=port-number-of-proxy-host

Specifies the port number of a proxy host. You must specify this parameter if the computer that runs the Export Tool communicates with the SVP via a proxy host.

- **-Xmxmemory-size(bytes)**

Specifies the size of memory to be used by JRE when the Export Tool is being executed. You must specify this parameter. The memory size must be 536870912, as shown in the Example later in this section. If an installed memory size is smaller than the recommended memory size of Storage Navigator PC, you must install more memory, before executing the Export Tool.



Note: If an installed memory size is larger than the recommended memory size of Storage Navigator PC, you can specify a memory size larger than as shown in the Example. However, to prevent lowering of execution speed, you do not set oversized memory size.

- **-Dmd.command**=path-to-command-file

Specifies the path to the command file

- **-Dmd.logpath**=path-to-log-file

Specifies the path to log files. A log file will be created whenever the Export Tool executes.

If this parameter is omitted, log files will be saved in the current directory.

- **-Dmd.logfile**=name-of-log-file

Specifies the name of the log file.

If this parameter is omitted, log files are named **exportMMddHHmmss.log**. MMddHHmmss indicates when the Export Tool executed. For example, the log file **export0101091010.log** contains log information about Export Tool execution at Jan. 1, 09:10:10.

Examples

The following example assumes that the computer running the Export Tool communicates with the SVP via a proxy host. In this example, the host name of the proxy host is **Jupiter**, and the port name of the proxy host is **8080**:

```
Java -classpath "./lib/JSanExport.jar;./lib/JSanRmiServerSx.jar" -Dhttp.proxyHost=Jupiter  
-Dhttp.proxyPort=8080 -Xmx536870912 -Dmd.command=command.txt -Dmd.logpath=log  
sanproject.getmondar.RJMdMain ↵
```

In the following example, a log file named **export.log** will be created in the **log** directory below the current directory when the Export Tool executes:

```
Java -classpath "./lib/JSanExport.jar;./lib/JSanRmiServerSx.jar" -Xmx536870912 -  
Dmd.command=command.txt  
-Dmd.logfile=export.log -Dmd.logpath=log sanproject.getmondar.RJMdMain ↵
```



Note: The "↵" symbol indicates the end of a command line.

Using CCI for Manual Volume Migration

In the environment where Volume Migration is installed on the storage system, users can use CCI to perform manual volume migration from an open-system host. This chapter explains the procedure and notes for using CCI to perform manual volume migration.

- [Sample Volume Migration](#)
- [Volume Migration Errors](#)

Sample Volume Migration

Using an example, this section illustrates the procedure for performing manual volume migration using CCI commands. In this example, group1 indicates the group name in the CCI configuration definition file and pair1 indicates the name of the volume whose pair is the target of the operation. For details on CCI operation, see the *Command Control Interface (CCI) User and Reference Guide*.

To perform manual volume migration by using CCI:

1. Start CCI.
2. Type the following command for a SMPL pair to start volume migration:

```
paircreate -g group1 -d pair1 -m cc -vl
```

When volume migration starts, the status of the pair changes to COPY.
3. Type the following command to check the status of the pair:

```
pairdisplay -g group1 -d pair1 -fcex
```

When the volume migration completed, the status of the pair becomes PSUS. When the volume migration failed, the status of the pair becomes PSUE.
4. When the status of the pair becomes PSUS or PSUE, return the pair status to SMPL by the following command:

```
pairsplit -S -g group1 -d pair1
```



Note: When the migration failed (the status of the pair becomes PSUE), repeat step 2 and 3. When it still becomes PSUE, contact the service engineer.

The status transition chart of a pair when the volume migration is performed by using CCI is shown below. The statuses in this chart can be displayed by typing pairdisplay command of CCI.




Figure 9-1 Status Transition Chart of a Pair during Volume Migration by CCI

Volume Migration Errors

When you perform volume migration by using CCI, note the following:

- A reserved volume of Volume Migration cannot be used when you perform volume migration by using CCI.
- A migration plan being executed by Volume Migration cannot be canceled from CCI.
- A migration plan created by Volume Migration cannot be displayed by using CCI.
- Before checking CCI settings in a Storage Navigator window, you must click **File**, and then **Refresh All** on the menu bar of the Storage Navigator main window.
- Migration cannot be performed when the specified volumes is a LUSE volume.
- While the microprogram is being downgraded or upgraded, users must not perform any operations that the resulting microprogram would not support.
- If you delete a manual migration plan of Volume Migration, the volume status changes from SMPL(PD) to SMPL. You can check the volume status by viewing the list at the bottom of the Manual Plan window and confirming whether the manual migration plan exists.

Although you can check the volume status by using the PairDisplay command of CCI, this command does not allow you to distinguish between SMPL and SMPL(PD).

- After you delete a manual migration plan of Volume Migration, if you want to execute a command to perform such tasks as volume migration operation or event waiting, wait for a certain period of time (recommended time is 10 seconds) until the volume status changes from SMPL(PD) to SMPL. If you do not wait for a certain period of time, the command might end abnormally.
- When you try volume migration or migration cancellation by using CCI, "EX_CMDRJE" might be displayed and the command might be refused depending on the condition in the DKC. The following table shows the error codes, causes of errors, and measures for such a case.

Table 9-1 Errors during Volume Migration or Migration Cancellation by using CCI

Error Code	Cause	Measure
2051	The migration target volume is being used with Data Retention Utility.	Release the volume from Data Retention Utility, and then perform the volume migration.
2055	The migration target volume is being used with Universal Replicator.	Release the volume from Universal Replicator, and then perform the volume migration.

Error Code	Cause	Measure
2056	The migration source volume is being used with Universal Replicator.	Release the volume from Universal Replicator, and then perform the volume migration.
2058	The migration target volume and the migration source volume belong to the same parity group.	This volume migration cannot be executed. Check the configuration definition file.
2301	Volume Migration is not installed on the storage system.	Install Volume Migration, and then perform the volume migration.
2306	The LBA size does not match between the migration target volume and the migration source volume.	This volume migration cannot be executed. Check the configuration definition file.
2309	The maximum number of pairs which can be set in the storage system is exceeded.	Reduce the number of pairs set in the storage system, and then perform the volume migration.
2311	The migration target volume is already reserved by Volume Migration.	Cancel the reservation by Volume Migration, and then perform the volume migration.
2312	The migration source volume is online with the host.	Make the migration source volume off-line from the host, and then perform the volume migration.
232F	The migration source volume is already specified as the target volume of Volume Migration.	Release the volume from Volume Migration, and then perform the volume migration.
2331	Either of the following causes can be considered: 1) The migration source volume is already reserved by Volume Migration. 2) The number of slots does not match between the migration target volume and the migration source volume.	Do the following measures respectively: 1) Cancel the reservation by Volume Migration, and then perform the volume migration. 2) This volume migration cannot be executed. Check the configuration definition file.
2332	You cannot add pair settings any more to the volume specified as the migration source.	Reduce the number of pair settings of the specified volume, and then perform the volume migration.
2333	The volume specified by the migration cancellation is not a migration source volume	It is trying to perform the migration cancellation of the pair which is not migrating. Check the configuration definition file.
2336	The emulation type does not match between the migration target volume and the migration source volume.	This volume migration cannot be executed. Check the configuration definition file.
2337	The migration source volume is a leaf volume of ShadowImage.	Release the ShadowImage pair, and then perform the volume migration.
233B	The migration target volume is a primary volume of ShadowImage.	Release the ShadowImage pair, and then perform the volume migration.
233C	The migration target volume is a secondary volume of ShadowImage.	Release the ShadowImage pair, and then perform the volume migration.
2342	The migration target volume is already specified as the target volume of Volume Migration.	Release the volume from Volume Migration, and then perform the volume migration.

Error Code	Cause	Measure
2344	The volume specified by the migration cancellation is not a migration target volume	It is trying to perform the migration cancellation of the pair which is not migrating. Check the configuration definition file.
2346	The migration target volume is a primary volume of TrueCopy.	Release the TrueCopy pair, and then perform the volume migration.
2347	The migration target volume is a secondary volume of TrueCopy.	Release the TrueCopy pair, and then perform the volume migration.
234B	The migration target volume is already specified as the source volume of Volume Migration.	Release the volume from Volume Migration, and then perform the volume migration.
2355	The setting of VLL is different between the migration target volume and the migration source volume.	This volume migration cannot be executed. Check the configuration definition file.
2362	The migration source volume is a primary volume of XRC.	Release the XRC pair, and then perform the volume migration.
2363	The migration target volume is a primary volume of XRC.	Release the XRC pair, and then perform the volume migration.
2365	The migration source volume is a primary volume of IBM 3990 Concurrent Copy (CC).	Release the IBM 3990 Concurrent Copy pair, and then perform the volume migration.
2366	The migration target volume is a primary volume of IBM 3990 Concurrent Copy.	Release the IBM 3990 Concurrent Copy pair, and then perform the volume migration.
2368	The migration source volume is a primary volume of a TrueCopy pair which is not in the SUSPEND status.	Put the TrueCopy pair in the SUSPEND status or release the TrueCopy pair. Then, perform the volume migration.
236B	The migration source volume is a volume of a ShadowImage pair which is in the SP-PEND status.	Perform the volume migration after the ShadowImage pair is changed to the SPLIT status.
2372	The migration source volume is being formatted.	Perform the volume migration after formatting is completed.
2373	The migration source volume is a command device.	This volume migration cannot be executed. Check the configuration definition file.
237C	The migration source volume is an external volume specified as a primary volume of TrueCopy.	Release the TrueCopy pair, and then perform the volume migration.
2382	The migration target volume is being formatted.	Perform the volume migration after formatting is completed.
2383	The migration target volume is a command device.	This volume migration cannot be executed. Check the configuration definition file.
2389	Cache Residency is set on the migration source volume.	Release the setting of Cache Residency, and then perform the volume migration.
238A	Cache Residency is set on the migration target volume.	Release the setting of Cache Residency, and then perform the volume migration.
238B	The migration source volume and the migration target volume is specified as LUSE volumes.	This volume migration cannot be executed. Check the configuration definition file.
238C	The migration target volume is a reserved volume of ShadowImage.	Cancel the reservation by ShadowImage, and then perform the volume migration.

Troubleshooting

This chapter gives troubleshooting information on Performance Monitor, Volume Migration Server Priority Manager and Export Tool. For troubleshooting information on Storage Navigator, see the *Storage Navigator User's Guide* and *Storage Navigator Messages*.

- ❑ [Troubleshooting Performance Monitor](#)
- ❑ [Troubleshooting Volume Migration](#)
- ❑ [Troubleshooting Server Priority Manager](#)
- ❑ [Troubleshooting the Export Tool](#)
- ❑ [Calling the Hitachi Data Systems Support Center](#)

Troubleshooting Performance Monitor

- When the WWN of a host bus adapter is displayed in red in the tree of the **WWN** tab:
The host bus adapter (HBA) whose WWN is displayed in red is connected to two or more ports, but the traffic between the HBA and some of the ports are not monitored by Performance Monitor. When many-to-many connections are established between HBAs and ports, you should make sure that all the traffic between HBAs and ports is monitored. For details on the measures, see *Monitoring All Traffic between HBAs and Ports*.
- When a part of monitoring data is missing:
While displaying Performance Monitor in short range, if I/O workloads between hosts and the storage system become heavy, the storage system gives higher priority to I/O processing than monitoring processing. If you notice that frequently monitoring data is missing, use the Gathering Interval option located in the Monitoring Options window to change to a longer collection interval. For details, see *Start Monitoring and Monitoring Options Window*.



Note: The display of the **LDEV** tab, **Port-LUN** tab, and **WWN** tab is fixed to in short range.

- Although **Monitoring Switch** is set to **Enable**, the monitoring data is not updated:
Because the time setting of SVP is changed, the monitoring data might not be updated. Set **Monitoring Switch** to **Disable**, and set **Monitoring Switch** to **Enable** again.

Troubleshooting Volume Migration

When an error message is displayed on the Storage Navigator computer:

The Volume Migration software displays error messages on the Storage Navigator computer when error conditions occur during Volume Migration operations.

If you need to call the Support Center, make sure to provide as much information about the problem as possible, including the error codes. For information on other error codes displayed on the Storage Navigator computer, see the *Storage Navigator Messages*.

Troubleshooting Server Priority Manager

When the WWN of a host bus adapter is displayed in red in the lower-left tree of the **WWN** tab:

The host bus adapter (HBA) whose WWN is displayed in red is connected to two or more ports, but the traffic between the HBA and some of the ports are not monitored by Performance Monitor. When many-to-many connections are established between HBAs and ports, make sure that all the traffics between HBAs and ports are monitored. For details on the measures, see Monitoring All Traffic between HBAs and Ports.

Troubleshooting the Export Tool

Table 10-1 explains possible problems with the Export Tool and probable solutions to the problems.

Table 10-1 Troubleshooting the Export Tool

Possible Problems	Probable Causes and Recommended Action
You cannot run the batch file.	The path to the Java Virtual Machine (Java.exe) might not be defined in the PATH environment variable. If this is true, you must add that path to the PATH environment variable. For information about how to add a path to the environment variable, refer to the documentation for your operating system. An incorrect version of Java Runtime Environment (JRE) might be installed on your computer. To check the JRE version, enter the following command at the Windows command prompt or the UNIX console window: <pre>Java -version</pre> If the version is incorrect, install the correct version of JRE.
The Export Tool stops and the processing does not continue.	- The command prompt window might be in pause mode. The command prompt window will be in pause mode if you click the command prompt window when the Export Tool is running. To cancel pause mode, you need to activate the command prompt window and then press the <ESC> key. If a timeout of RMI occurs during pause mode, the login will be canceled and an error will occur when you cancel pause mode after the timeout. The error message ID will be (0001 4011). - If a memory size is not specified in a batch file, the Out Of Memory Error occurs in JRE, the Export Tool might stop and the processing might not continue. Confirm whether the specified memory size is correct or not.
The command prompt window was displaying progress of the export processing, but the window stopped displaying progress before the processing stopped. The progress information does not seem to be updated anymore.	

Possible Problems	Probable Causes and Recommended Action
An error occurs and the processing stops.	If the error message ID is (0001 4011), the user is forcibly logged off and the processing stops because the Export Tool did not issue any request to the SVP. The computer running the Export Tool could be slow. Confirm whether you are using a computer that is not supported, or whether the computer is slow. Run the Export Tool again. If the error persists, contact the maintenance personnel. If the error message ID is (0002 5510), probable error causes and solutions are: - An internal processing is being performed in the disk array. Alternatively, another user is changing configurations. Wait for a while and then run the Export Tool again. - Maintenance operations are being performed on the disk array. Wait until the maintenance operations finish and then run the Export Tool again. If the error message ID is none of the above, see Table 10-2.
The monitoring data in the CSV file includes "-1".	The value "-1" indicates that Performance Monitor failed to obtain monitoring data for some reasons. Probable reasons are: - Performance Monitor attempted to obtain statistics when an operation for rebooting the disk array is in progress. - A heavy workload is imposed on the disk array. - If IOPS is 0 (zero), the Response Time that is included in the monitoring data of LUs, LDEVs, ports, WWNs, or external volumes is -1. Because IOPS is 0 (zero), the average response time becomes invalid. - There is no volume in a parity group. - Just after the time when the CUs to be monitored were added, Performance Monitor failed to obtain related files of remote copy, i.e. TrueCopy, TrueCopy for IBM z/OS, Universal Replicator, or Universal Replicator for IBM z/OS for all volumes or journal groups. For details about the files, see Table 8-10, Table 8-14, and Table 8-15.
When the Export Tool terminated abnormally due to error, the row of Check License is shown as UnmarshalException in the log file.	It might be unsuitable combination of DKCMAIN/SVP program version and export tool version. Confirm whether versions of these programs are correct.
When a CSV file is opened, the parity group ID and volume ID are displayed as following: - The parity group IDs are displayed as dates - The volume IDs are displayed with a decimal point	To display a CSV file correctly, you need to perform following operations: - Start Microsoft® Excel. - On the menu bar, select Data, Import External Data, and Import Text File, and specify a CSV file to import. The Text Import Wizard - Step 1 of 3 dialog box is displayed - In the Text Import Wizard - Step 1 of 3 dialog box, click Next. The Text Import Wizard - Step 2 of 3 dialog box is displayed - In the Text Import Wizard - Step 2 of 3 dialog box, check only Comma in the Delimiter area, and click Next. The Text Import Wizard - Step 3 of 3 dialog box is displayed - In the Text Import Wizard - Step 3 of 3 dialog box, select all columns of Date preview, and check Text in the Column data format area on the upper right of this dialog box. - Click Finish. The imported CSV file is displayed.

Possible Problems	Probable Causes and Recommended Action
When you executed the Export Tool with many volumes specified, the Export Tool terminated abnormally while gathering monitoring data.	Because too many volumes are specified, a timeout error might have occurred due to a heavy workload imposed on the computer where the Export Tool was running. The error message ID is (0001 4011) . Specify fewer volumes. It is recommended that the number of volumes to be specified is 16,384 or less.

If an error occurs when you run the Export Tool, error messages are issued to the standard output (for example, the command prompt) and the log file. Table 10-2 lists the Export Tool messages and recommended actions against errors.

Table 10-2 Messages Issued by the Export Tool

Possible Problems	Probable Causes and Recommended Action
Connection to the server has not been established.	Connection to the server has not been established. Use the login subcommand.
Execution stops.	Execution stops. Remove errors.
Illegal character: " <i>character</i> "	An illegal character is used. Use legal characters.
Invalid length: <i>token</i>	The length is invalid. Specify a value that has a correct length.
Invalid range: <i>range</i>	The specified range is invalid. Specify the correct range.
Invalid value: " <i>value</i> "	The specified value is invalid. Specify a correct value.
Login failed	1. An attempt to log into the SVP failed. Probable causes are: 2. An incorrect operand is used for the <i>svpip</i> subcommand. 3. An incorrect operand is used for the <i>login</i> subcommand. 4. The specified user ID is used by another person, and the person is being logged in. 5. Two users are currently displaying the Performance Management window. Two users are currently executing the Export Tool. Note: If the error cause is step 4 or 5 above, Do one of the following: - ask one of the users to display another window. - ask one of the users to log off. - wait for one of the users to quit the Export Tool.
Missing command file	The command file is not specified. Specify the name of the command file correctly.
Missing group name	No operand is specified in the group subcommand. Specify operands for the subcommand.
Missing host name	No host name is specified. Specify a host name.
Missing output directory	No directory is specified for saving files. Specify the directory for saving files.
Missing password	The Export Tool cannot find the user ID, which is used to log into the SVP. Specify the password.

Possible Problems	Probable Causes and Recommended Action
Missing svpip	The svpip subcommand is not used. Use the svpip command.
Missing time range	Specify the time range.
Missing user ID	The Export Tool cannot find the user ID, which is used to log into the SVP. Specify the user ID.
Out of range: <i>range</i>	The value is outside the range. If the short-range subcommand or the long-range subcommand is used, specify a value within the range from the monitoring start time to the monitoring end time. If the set subcommand is used with the switch operand, specify a value within the range of 1 to 15.
Permission Denied.	The user ID does not have the required permission. The user ID needs to have at least one of permissions for Performance Monitor, TrueCopy, TrueCopy for IBM z/OS, Universal Replicator, and Universal Replicator for IBM z/OS.
RMI server error (<i>part-code, error-number</i>)	An error occurs at the RMI server. For detailed information, see the <i>Storage Navigator Messages</i> .
Unable to display help message	The Export Tool cannot display the online help due to a system error.
Unable to get serial number	The Export Tool cannot obtain the serial number due to a system error.
Unable to get time range for monitoring	The SVP does not contain monitoring data.
Unable to read command file: <i>file</i>	The Export Tool cannot read the command file. Specify the name of the command file correctly.
Unable to use the command: <i>command</i>	The specified subcommand is unavailable because you logged in as a storage partition administrator.
Unable to use the group name: <i>operand</i>	The specified operand of the group subcommand is unavailable because you logged in as a storage partition administrator.
Unknown host: <i>host</i>	The Export Tool cannot resolve the host name. Specify the correct host name.
Unsupported command: <i>command</i>	The Export Tool does not support the specified command. Specify a correct command.
Unsupported operand: <i>operand</i>	The specified operand is not supported. Correct the specified operand.
Unsupported option: <i>option</i>	The specified option is not supported. Correct the specified option.
Some file exists in <i>path</i> . What do you do? clear(c)/update(u)/stop(p) You selected "action". Is it OK? (y/n)	Files exist in <i>path</i> . If you want to clear the files, press the <c> key. If you want to overwrite the files, press the <u> key. If you want to stop the operation, press the <p> key. When you press a key, a message appears and asks whether to perform the specified action. To perform the specified action, press the <y> key. To cancel the specified action, press the <n> key.

Possible Problems	Probable Causes and Recommended Action
Specify the following subcommand before login subcommand: retry	The retry subcommand is written in an incorrect position in the command file. Write the retry subcommand before the login subcommand.
Start gathering group data Target = xxx, Total = yyy End gathering group data	The Export Tool starts collecting data specified by the <i>group</i> subcommand. The number of <i>targets</i> is xxx and the total number is yyy (refer to Note below). The Export Tool ends collecting data. Note: For example, suppose that the storage system contains 100 parity groups and the command file contains the following command line: <pre>group PG 1-1:1-2</pre> The Export Tool displays the message "Target=2, Total=100", which means that the group subcommand specifies two parity groups and that the total number of parity groups in the storage system is 100.
Syntax error: "line"	A syntax error is detected in a command line in your command file. Check the command line for the syntax error and then correct the script. Note: Some operands must be enclosed by double quotation marks ("). Check the command line to find whether double quotation marks are missing.


Calling the Hitachi Data Systems Support Center

If you need to call 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 error messages displayed on the host system(s).
- The Volume Retention Manager (or other) error code(s) displayed by the Storage Navigator computer.
- The USP V/VM Storage Navigator configuration information saved on diskette using the FD Dump Tool (see the *Storage Navigator User's Guide*).
- The remote service information messages (R-SIMs) logged on the Storage Navigator computer and the reference codes and severity levels of the recent R-SIMs.

The Hitachi Data Systems customer support staff is available 24 hours/day, seven days a week. If you need technical support, please call:


- United States: (800) 446-0744
- Outside the United States: (858) 547-4526



Acronyms and Abbreviations

ACP	array control processor
CCI	command control interface
CHA	channel adapter
CHP	channel processor
CLPR	cache logical partition
CSW	cache switch
CU	control unit (logical control unit)
DASD	direct-access storage device
DFW	DASD fast write
DKA	disk adapter
DKCMAIN	Disk Controller Main
DKP	disk processor
DRR	data recovery and reconstruction
ESCON [®]	Enterprise System Connection (IBM trademark for optical channels)
GB	gigabyte (see Convention for Storage Capacity Values)
HBA	host bus adapter
HDD	hard disk drive
JRE	Java Runtime Environment
KB	kilobyte (see Convention for Storage Capacity Values)
LAN	local-area network
LDEV	logical device
LDKC	logical disk controller
LU	logical unit
LUN	logical unit number
LUSE	LU size expansion
MB	megabyte (see Convention for Storage Capacity Values)
NSC	network storage controller
NVS	nonvolatile storage
PB	petabyte (see Convention for Storage Capacity Values)
PC	personal computer

PDEV	physical device
PSUE	pair suspended-error
PSUS	pair suspended-split
P-VOL	primary volume
RAID	redundant array of independent disks
SIM	service information message
SLPR	storage management logical partition
SM	shared memory
SMPL	simplex
SPM	Server Priority Manager
S-VOL	secondary volume
SVP	service processor
TB	terabyte (see Convention for Storage Capacity Values)
TC	TrueCopy
TCz	TrueCopy for IBM z/OS
UR	Universal Replicator
USP	Universal Storage Platform
VLL	Virtual LVI/LUN
V-VOL	virtual volume (Snapshot Image)
WWN	world-wide name



Index

A

- access path, 2-10
- auto migration, 2-13
 - operations, 6-8
 - overview, 2-21

C

- CHA. --- Refer to channel adapter
- channel adapter, 2-6
- channel processor, 2-6
- CHP. --- Refer to channel processor
- class, 2-16

D

- data recovery and reconstruction processor, 2-9
- development server, 2-26
- disk adapter, 2-6
- disk processor, 2-6
- DKA. --- Refer to disk adapter
- DKP. --- Refer to disk processor
- DRR. --- Refer to data recovery and reconstruction processor

E

- Expanded LUs, 2-34

F

- fixed parity group, 2-22

H

- HBA. --- Refer to host bus adapter
- HDD class, 2-16
- host bus adapter, 2-13, 4-23, 5-46, 7-2

I

- I/O rate, 2-11, 2-12, 5-27, 7-4

M

- manual migration, 2-13

- operations, 6-5
- overview, 2-15
- migration operation, 2-19
 - conditions for canceling, 2-20, 2-36
- migration plan, 2-15, 6-6
 - deleting a manual migration plan, 6-7
 - deleting an auto migration plan, 6-11

N

- nickname, 5-36, 5-39
- non- prioritized port, 7-13
- non-prioritized port, 7-3
- non-prioritized WWN, 7-8, 7-21
- normal volume, 6-3

P

- parity group, 2-4
- Performance Manager
 - system requirements, 3-2
- prioritized port, 7-3, 7-13
- prioritized WWN, 7-8, 7-21
- production server, 2-26

R

- read hit ratio, 2-11, 5-27
- real-time mode, 4-18
- requirements, 2-31
- reserved volume, 4-37, 6-3

S

- source volume, 2-19
- SPM group, 7-29
- SPM name, 5-36, 5-39, 5-48, 5-51

T

- target volume, 2-19, 4-37
- threshold, 2-27
- threshold control, 2-27
- transfer rate, 2-11, 2-12, 5-27, 7-4

U

upper limit control, 2-27

W

write hit rate, 2-11

write hit ratio, 5-27

write pending rate, 2-9

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

