

**Hitachi TagmaStore®
Universal Storage Platform and
Network Storage Controller**

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

© 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. See 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.

Trademarks

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.

Hitachi TagmaStore is a registered trademark of Hitachi Data Systems Corporation.

HP-UX is a registered trademark of the Hewlett-Packard Development Company, L.P.

ESCON, FlashCopy, IBM, and z/OS are registered trademarks of International Business Machines Corporation in the United States.

Windows is a registered trademark of Microsoft Corporation.

Java and Solaris are trademarks of Sun Microsystems, Inc.

UNIX is a registered trademark of The Open Group in the United States and other countries.

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.

Notice of Export Controls

Export of technical data contained in this document may require an export license from the United States government and/or the government of Japan. Please contact the Hitachi Data Systems Legal Department for any export compliance questions.

Document Revision Level

Revision	Date	Description
MK-94RD218-P	July 2004	Preliminary Release
MK-94RD218-00	August 2004	Initial Release, supersedes and replaces MK-94RD218-P
MK-94RD218-01	October 2004	Revision 1, supersedes and replaces MK-94RD218-00
MK-94RD218-02	March 2005	Revision 2, supersedes and replaces MK-94RD218-01
MK-94RD218-03	April 2005	Revision 3, supersedes and replaces MK-94RD218-02
MK-94RD218-04	May 2005	Revision 4, supersedes and replaces MK-94RD218-03
MK-94RD218-05	June 2005	Revision 5, supersedes and replaces MK-94RD218-04
MK-94RD218-06	September 2005	Revision 6, supersedes and replaces MK-94RD218-05
MK-94RD218-07	November 2005	Revision 7, supersedes and replaces MK-94RD218-06
MK-94RD218-08	February 2006	Revision 8, supersedes and replaces MK-94RD218-07
MK-94RD218-09	April 2006	Revision 9, supersedes and replaces MK-94RD218-08
MK-94RD218-10	May 2006	Revision 10, supersedes and replaces MK-94RD218-09
MK-94RD218-11	October 2006	Revision 11, supersedes and replaces MK-94RD218-10
MK-94RD218-12	January 2008	Revision 12, supersedes and replaces MK-94RD218-11
MK-94RD218-13	January 2008	Revision 13, supersedes and replaces MK-94RD218-12
MK-94RD218-14	April 2008	Revision 14, supersedes and replaces MK-94RD218-13
MK-94RD218-15	June 2008	Revision 15, supersedes and replaces MK-94RD218-14

Source Document(s) for this Revision

- Pm50d09i (RSD Source Document).

Changes in this Revision

- In Table A.15, Files with Statistics about RC Operations by UR/URz (Journal Groups), revised the field names in the table.

Preface

This document describes and provides instructions for using the Performance Monitor, Volume Migration, and Server Priority Manager features of the Hitachi TagmaStore® Universal Storage Platform (USP) and Network Storage Controller (NSC) storage subsystems.

- Hitachi Performance Monitor collects usage statistics about physical resources in one storage system and feeds them to Hitachi Volume Migration and Hitachi Server Priority Manager software to enable load balancing and performance tuning.
- Hitachi Volume Migration monitors, analyzes, and balances your system, either automatically or by making tuning recommendations to administrators to eliminate performance bottlenecks.
- Hitachi Server Priority Manager optimizes the data-access performance of servers.

This document assumes the following:

- The user is familiar with the operating system and web browser software on the system hosting the Storage Navigator software.
- The user has a background in data processing and understands RAID storage subsystems and their basic functions.
- The user is familiar with the TagmaStore USP and/or NSC subsystem and has read and understands the User and Reference Guide for the subsystem and the *Storage Navigator User's Guide* (MK-94RD206).

Notes:

- In this document the term “TagmaStore” refers to both the Universal Storage Platform and the Network Storage Controller, unless otherwise noted.
- In this document the term “Universal Storage Platform (USP)” refers to all models of the TagmaStore USP subsystem, unless otherwise noted.
- In this document the term “Network Storage Controller (NSC)” refers to all models of the TagmaStore NSC subsystem, unless otherwise noted.
- This product includes software developed by the Apache Group for use in the Apache HTTP server project (<http://www.apache.org/>).
- The Storage Navigator screens shown in this document were captured on a Windows® system with the Internet Explorer web browser. The Storage Navigator screens may display differently on other operating systems and browsers. For more information, see the *Storage Navigator Users Guide* (MK-94RD206).
- The functions described in this document are for storage administrators only. For more information on user access, see the *Storage Navigator User's Guide* (MK-94RD206).

Notice: The use of Performance Monitor, Volume Migration, Server Priority Manager, and all other Hitachi Data Systems products is governed by the terms of your agreement(s) with Hitachi Data Systems.

Microcode Level

This document revision applies to TagmaStore USP/NSC microcode 50-09-7x and higher.

Convention for Storage Capacity Values

Storage capacity values for logical devices (LDEVs) on the TagmaStore USP and NSC are calculated based on the following values:

1 KB (kilobyte) = 1,024 bytes
1 MB (megabyte) = 1,024² bytes
1 GB (gigabyte) = 1,024³ bytes
1 TB (terabyte) = 1,024⁴ bytes

Referenced Documents

Hitachi TagmaStore Universal Storage Platform documentation:

- *USP User and Reference Guide*, MK-94RD231
- *ShadowImage for z/OS® User's Guide*, MK-94RD212
- *Compatible PAV for z/OS® User's Guide*, MK-94RD211
- *Volume Security User's Guide*, MK-94RD216
- *Volume Retention Manager User's Guide*, MK-94RD219
- *TrueCopy for z/OS® User and Reference Guide*, MK-94RD214
- *Universal Replicator for z/OS® User and Reference Guide*, MK-94RD224
- *Copy-on-Write Snapshot User's Guide*, MK-95RD277

Hitachi TagmaStore Network Storage Controller documentation:

- *NSC User and Reference Guide*, MK-95RD279

Hitachi TagmaStore Universal Storage Platform and Network Storage Controller documentation:

- *Storage Navigator™ User's Guide*, MK-94RD206
- *Storage Navigator™ Error Codes*, MK-94RD202
- *Cross-System Copy User's Guide*, MK-94RD252
- *ShadowImage™ User's Guide*, MK-94RD204
- *LUN Expansion and Virtual LVI/LUN User's Guide*, MK-94RD205
- *LUN Manager User's Guide*, MK-94RD203
- *Cache Residency Manager User's Guide*, MK-94RD208
- *Performance Manager User's Guide*, MK-94RD218

- *ShadowImage™ User's Guide*, MK-94RD204
- *Virtual Partition Manager (VPM) User's Guide*, MK-94RD259
- *Data Retention Utility User's Guide*, MK-94RD210
- *TrueCopy™ User and Reference Guide*, MK-94RD215
- *Universal Volume Manager User's Guide*, MK-94RD220
- *Universal Replicator User and Reference Guide*, MK-94RD223

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.)

Contents

Chapter 1	Overview of Performance Management.....	1
Chapter 2	Performance Monitor Overview.....	5
2.1	Viewing Usage Statistics for Disk Subsystem Resources	6
2.1.1	Viewing Parity Group Usage Statistics	7
2.1.2	Viewing Logical Volume Usage Statistics	7
2.1.3	Viewing External Volume Group Usage Statistics	7
2.1.4	Viewing External Volume Usage Statistics.....	8
2.1.5	Viewing Channel Processor Usage Statistics	8
2.1.6	Viewing Disk Processor Usage Statistics	9
2.1.7	Viewing DRR Processor Usage Statistics	11
2.1.8	Viewing Write Pending Rate and Cache Memory Usage Statistics	11
2.1.9	Viewing Access Path Usage Statistics	12
2.2	Viewing Disk Workload Statistics.....	13
2.3	Viewing Port Traffic Statistics.....	14
2.4	Viewing LU Paths Traffic Statistics	15
2.5	Viewing Traffic between HBAs and Disk Subsystem Ports.....	15
2.6	Understanding Statistical Storage Ranges	16
2.7	Overview of Performance Monitor Operations	17
2.7.1	Start Monitoring	17
2.7.2	View the Monitoring Results	18
2.7.3	Notes on Performance Monitor.....	19
Chapter 3	Volume Migration Overview.....	21
3.1	Using Volume Migration to Balance Work Loads	22
3.2	About Migrating Logical Volumes Manually	23
3.2.1	Migrating Logical Volumes Manually.....	23
3.2.2	About the Manual Volume Migration Mechanism	26
3.3	Automating Volume Migrations	29
3.3.1	Migrating Logical Volumes Automatically.....	29
3.3.2	About the Auto Migration Mechanism	30
3.4	Migration Requirements and Restrictions.....	33
3.4.1	Calculating the Number of Concurrent Migration Plans	40
3.4.2	Determining the Number of Differential Tables for Mainframe Volume Migration.....	41
3.4.3	Determining the Number of Differential Tables for Open System Volume Migration.....	43
Chapter 4	Server Priority Manager Overview	45
4.1	Safeguarding Performance of High-Priority Hosts	46
4.2	Disabling Upper-Limit Control Automatically	47
4.3	Using Server Priority Manager	48
4.3.1	If One-to-One Connections Link HBAs and Ports.....	48
4.3.2	If Many-to-Many Connections Link HBAs and Ports.....	51
4.4	Notes on Using Server Priority Manager.....	56
4.4.1	Selecting CU Groups	56

4.4.2	Calculating I/O and Transfer Rates	56
4.4.3	Using Server Priority Manager with Remote Copy Functions	57
Chapter 5	Preparing for Performance Manager Operations	59
5.1	System Requirements.....	60
5.2	Understanding the Role of Your Web Client Computer	61
Chapter 6	Using Performance Monitor	63
6.1	About the Performance Monitor Window	64
6.2	About Other Windows in the Interface	65
6.3	About the Physical Tab of the Performance Monitor Window.....	66
6.3.1	When You Select Long-Range Storage	70
6.3.2	When You Select Short-Range Storage	71
6.4	About the LDEV Tab of the Performance Monitor Window	72
6.5	About the Port-LUN Tab of the Performance Monitor Window	77
6.5.1	Viewing Port Workload Statistics	80
6.5.2	Viewing Statistics about Workload on LU Paths.....	81
6.6	About the WWN Tab of the Performance Monitor Window	82
6.7	About the Monitoring Options Window	85
6.8	Starting and Stopping Disk Subsystem Monitoring	87
6.9	Monitoring Resources in the Disk Subsystem.....	88
6.9.1	Viewing Usage Statistics about Parity Groups	88
6.9.2	Viewing Usage Statistics about Logical Volumes in Parity Groups	89
6.9.3	Viewing Usage Statistics about External Volume Groups	92
6.9.4	Viewing Usage Statistics about External Volumes in EV Groups	93
6.9.5	Viewing Usage Statistics about Channel Processors	95
6.9.6	Viewing Usage Statistics about Disk Processors	96
6.9.7	Viewing Data Recovery and Reconstruction Processor Usage	97
6.9.8	Viewing Write Pending and Cache Memory Usage Statistics	98
6.9.9	Viewing Usage Statistics about Access Paths	100
6.10	Monitoring Disk Drives	103
6.10.1	Viewing I/O Rates for Disks	103
6.10.2	Viewing Transfer Rates for Disks	106
6.11	Monitoring Ports.....	109
6.11.1	Viewing I/O Rates for Ports	109
6.11.2	Viewing Transfer Rates for Ports	111
6.11.3	Viewing Details about the I/O and Transfer Rates	114
6.12	Monitoring LU Paths	115
6.12.1	Viewing LU Paths I/O Rates	115
6.12.2	Viewing LU Paths Transfer Rates	116
6.13	Monitoring Paths between Host Bus Adapters and Ports.....	119
6.13.1	Viewing I/O Rates between HBAs	119
6.13.2	Viewing Transfer Rates between HBAs	121
Chapter 7	Using Volume Migration	125
7.1	About the Volume Migration Window	126
7.1.1	About the Manual Plan Tab in the Volume Migration Window	126
7.1.2	About the Auto Plan Tab in the Volume Migration Window	134
7.1.3	About the Attribute Tab in the Volume Migration Window.....	137

7.1.4	About the History Tab in the Volume Migration Window.....	142
7.2	Starting Volume Migration	144
7.3	Reserving Migration Destinations for Logical Volumes.....	145
7.4	Migrating Logical Volumes Manually	146
7.5	Automating Volume Migrations	149
7.5.1	Setting a Fixed Parity Group.....	149
7.5.2	Changing the Maximum Disk Usage Rate for Each HDD Class.....	149
7.5.3	Setting Auto Migration Parameters	150
7.5.4	Remaking Auto Migration Plans	152
7.6	Viewing the Migration History Log	153
Chapter 8	Using Server Priority Manager	161
8.1	About the Server Priority Manager Window	162
8.1.1	About the Port Tab of the Server Priority Manager Window.....	162
8.1.2	About the WWN Tab of the Server Priority Manager Window	164
8.2	Starting Server Priority Manager	170
8.3	If One-to-One Connections Link HBAs and Ports.....	171
8.3.1	Setting Priority for Ports on the Disk Subsystem	171
8.3.2	Measuring Traffic between HBAs and Ports	172
8.3.3	Analyzing Traffic Statistics	172
8.3.4	Setting Upper-Limit Values for Traffic at Non-Prioritized Ports	173
8.3.5	Setting a Threshold	174
8.4	If Many-to-Many Connections Link HBAs and Ports	176
8.4.1	Monitoring All Traffic between HBAs and Ports	176
8.4.2	Setting Priority for Host Bus Adapters	178
8.4.3	Analyzing Traffic Statistics	179
8.4.4	Setting Upper-Limit Values for Non-Prioritized WWNs	180
8.4.5	Setting a Threshold	182
8.4.6	Changing the SPM Name of a Host Bus Adapter	182
8.4.7	Replacing a Host Bus Adapter	183
8.5	Grouping Host Bus Adapters	185
8.5.1	Containing Multiple HBAs in an SPM Group.....	185
8.5.2	Deleting an HBA from an SPM Group	186
8.5.3	Switching the Priority of an SPM Group.....	186
8.5.4	Setting an Upper-Limit Value to HBAs in an SPM Group.....	187
8.5.5	Renaming an SPM Group.....	188
8.5.6	Deleting an SPM Group	188
Chapter 9	Troubleshooting	189
9.1	Troubleshooting Performance Monitor	189
9.2	Troubleshooting Volume Migration	189
9.3	Troubleshooting Server Priority Manager	189
9.4	Calling the Hitachi Data Systems Technical Support Center.....	190
Appendix A	Exporting Monitoring Data to Files	191
A.1	Files to be Exported.....	192
A.2	Preparing to Use the Export Tool	205
A.2.1	Requirements for Using the Export Tool.....	205
A.2.2	Installing the Export Tool on a Windows Computer.....	206

A.2.3	Installing the Export Tool on a UNIX Computer	206
A.3	Using the Export Tool	208
A.3.1	Preparing a Command File	208
A.3.2	Preparing a Batch File.....	211
A.3.3	Running the Export Tool	212
A.4	Command Reference	216
A.4.1	svpip Subcommand	218
A.4.1.1	Syntax	218
A.4.1.2	Description.....	218
A.4.1.3	Operands	218
A.4.1.4	Example	218
A.4.2	Retry Subcommand	219
A.4.2.1	Syntax	219
A.4.2.2	Description.....	219
A.4.2.3	Operands	219
A.4.2.4	Example	219
A.4.3	Login Subcommand	220
A.4.3.1	Syntax	220
A.4.3.2	Description.....	220
A.4.3.3	Operands	220
A.4.3.4	Example	220
A.4.4	Show Subcommand	221
A.4.4.1	Syntax	221
A.4.4.2	Description.....	221
A.4.5	Group Subcommand	223
A.4.5.1	Syntax	223
A.4.5.2	Description.....	223
A.4.5.3	Operands	226
A.4.5.4	Examples	234
A.4.6	Shorrange Subcommand	236
A.4.6.1	Syntax	236
A.4.6.2	Description.....	236
A.4.6.3	Operands	236
A.4.6.4	Examples	238
A.4.7	Longrange Subcommand	239
A.4.7.1	Syntax	239
A.4.7.2	Description.....	239
A.4.7.3	Operands	239
A.4.7.4	Examples	240
A.4.8	Outpath Subcommand.....	242
A.4.8.1	Syntax	242
A.4.8.2	Description.....	242
A.4.8.3	Operands	242
A.4.8.4	Examples	242
A.4.9	Option Subcommand.....	243
A.4.9.1	Syntax	243
A.4.9.2	Description.....	243
A.4.9.3	Operands	243
A.4.9.4	Example	243
A.4.10	Apply Subcommand	244

A.4.10.1 Syntax	244
A.4.10.2 Description	244
A.4.11 Set Subcommand	245
A.4.11.1 Syntax	245
A.4.11.2 Description	245
A.4.11.3 Operands.....	245
A.4.11.4 Examples.....	246
A.4.12 Help Subcommand.....	247
A.4.12.1 Syntax	247
A.4.12.2 Description	247
A.4.12.3 Example.....	247
A.4.13 Java Command for Exporting Data in Files	248
A.4.13.1 Syntax	248
A.4.13.2 Description	248
A.4.13.3 Operands.....	248
A.4.13.4 Examples.....	249
A.5 Troubleshooting the Export Tool.....	250

Appendix B Limitations for Storage Partition Administrators.....253

B.1 Limitations of Performance Monitor	254
B.2 Limitations of the Export Tool	256

Appendix C Manual Volume Migration Using CCI.....257

C.1 Migrating Volumes Using CCI.....	258
--------------------------------------	-----

Acronyms and Abbreviations

Index

List of Figures

Figure 1.1	Performance Management Solution Components	2
Figure 1.2	Performance Management Process Flow Diagram	3
Figure 2.1	Access Paths	12
Figure 2.2	Monitoring Options Window of Performance Monitor	17
Figure 2.3	Performance Monitor Window	18
Figure 3.1	Sample Migration Plan in Volume Migration Window	23
Figure 3.5	Specifying Migration Destination	25
Figure 3.6	Estimated Usage Rate for Parity Groups	26
Figure 3.7	Data Flow during a Volume Migration Operation	27
Figure 3.8	Data Flow after a Volume Migration Operation	28
Figure 3.9	Migrating to Improve Disk Usage	31
Figure 3.10	Keeping Reserve Volumes in High HDD Class Available	31
Figure 3.11	Migrating a Volume Automatically to a Lower HDD Class.....	32
Figure 3.12	Migrating a Volume Automatically within Same HDD Class.....	32
Figure 4.1	Network A (One-to-One connections between HBAs and Ports).....	49
Figure 4.4	Network B (Many-to-Many connections between HBAs and Ports)	51
Figure 4.9	Line Graphs indicating Port Traffic Changes	57
Figure 6.1	The Physical Tab of the Performance Monitor Window	66
Figure 6.4	Selecting an Item Displayed on the Graph	70
Figure 6.5	LDEV tab of Performance Monitor Window	72
Figure 6.7	Chart Y Axis Rate Drop-Down List and Detail Check Box (LDEV Tab)	76
Figure 6.8	Port-LUN Tab of the Performance Monitor Window	77
Figure 6.10	Graphs Illustrating Changes Port Workloads.....	81
Figure 6.13	The Monitoring Options Window of Performance Monitor	85
Figure 6.22	Examples of Write Pending and Cache Usage Rates Displayed	100
Figure 6.23	Examples of Usage Statistics Displayed	102
Figure 6.24	I/O Rates for Disks.....	105
Figure 6.25	Transfer Rates for Disks.....	107
Figure 6.26	I/O Rates for Ports.....	110
Figure 6.27	Transfer Rates for Ports	113
Figure 6.30	I/O Rates for Host Bus Adapters (WWN Tab)	121
Figure 6.31	Transfer Rates for Host Bus Adapters (WWN Tab)	123
Figure 7.1	Manual Plan Tab in the Volume Migration Window	126
Figure 7.5	Auto Plan Tab in the Volume Migration Window	134
Figure 7.6	Attribute Tab in the Volume Migration Window	138
Figure 7.11	History Tab in the Volume Migration Window	142
Figure 8.1	Port Tab of the Server Priority Manager Window	162
Figure 8.2	WWN Tab of the Server Priority Manager Window.....	165
Figure 8.3	Add WWN Dialog Box	177
Figure A.1	Example of a Text File	191
Figure A.2	Example of a Command File	208

Figure A.3	Scripts in Batch Files (runWin.bat and runUNIX.bat)	211
Figure A.4	Scripts in Batch Files (when specifying host name of a proxy host).....	212
Figure A.5	Scripts in Batch Files (when specifying IP address of a proxy host).....	212
Figure A.6	Example of Executing a Batch File (on a Windows computer)	213
Figure A.7	Example of Command Prompt Outputs when Export Tool Runs.....	213
Figure A.8	Example of a Batch File, including a Reference to Termination Code.....	215
Figure A.9	Examples of Comments.....	217
Figure A.10	Information output by Show Subcommand	222
Figure A.11	Starting and Ending Times for Collecting Monitoring Data	237
Figure A.12	Starting and Ending Times for Collecting Monitoring Data	240
Figure B.1	Example of a Text File	255
Figure C.1	Status Transition Chart of a Pair during Volume Migration by CCI	259

List of Tables

Table 2.1	Migrating Volumes when Disk Processor Usage Rate is High (not applicable to NSC)	9
Table 3.1	Combinations of Volumes That Can Be Migrated.....	33
Table 3.2	Non-Cascade Volumes and Volume Migration Source Volumes	35
Table 3.3	Cascade Volumes and Volume Migration Source Volumes	36
Table 3.4	Migrated Volume Capacity and Estimated Delay in Copying Speed	39
Table 3.5	Number of Control Cylinders According to Emulation Types.....	42
Table 3.6	Total Number of Differential Tables per Migration Plan.....	43
Table 3.7	Number of Control Cylinders by Emulation Type	43
Table 6.1	Cascade Volumes and Volume Migration Source Volumes	75
Table 6.2	Information Displayed in the LDEV Tab Graph	76
Table 6.3	Information Displayed in the Port-LUN Tab Graph	81
Table 7.1	Messages in the Auto Migration Logs (1 of 5)	154
Table 7.2	Messages in the Auto Migration Logs (2 of 5)	154
Table 7.3	Messages in the Auto Migration Logs (3 of 5)	156
Table 7.4	Messages in the Auto Migration Logs (4 of 5)	158
Table 7.5	Messages in the Auto Migration Logs (5 of 5)	159
Table 7.6	Messages in the Migration Logs (Displayed in the <i>Action</i> Column)	160
Table A.1	Performance Monitor Windows and Monitoring Data Saved by Export Tool ...	192
Table A.2	Files with Statistics about Resource Usage and Write Pending Rates	193
Table A.3	Files with Statistics about Parity Groups, External Volume Groups, or V-VOL Groups	194
Table A.4	Files with Statistics about Logical Volumes in Parity/External Volume Groups, or V-VOL Groups.....	195
Table A.5	Files with Statistics about Ports	196
Table A.6	Files with Statistics about Host Bus Adapters Connected to Ports.....	196
Table A.7	Files with Statistics about LUs	197
Table A.8	Files with Statistics about SPM Groups	197
Table A.9	Files with Statistics about Host Bus Adapters Belonging to SPM Groups	198
Table A.10	Files with Statistics about RC Operations by TC/TCz (Disk Subsystem)	198
Table A.11	Files with Statistics about RC Operations by TC/TCz (LUs)	199
Table A.12	Files with Statistics about RC Operations by TC/TCz (LUs controlled by CU) ..	200
Table A.13	Files Containing Statistics about Remote Copy Operations by TrueCopy and TrueCopy for z/OS (At CLPR)	201
Table A.14	Files with Statistics about RC Operations by UR/URz (Disk Subsystem).....	201
Table A.15	Files with Statistics about RC Operations by UR/URz (Journal Groups).....	202
Table A.16	Files with Statistics about RC Operations by UR/URz (LUs)	203
Table A.17	Files with Statistics about RC Operations by UR/URz (LUs controlled by CU) ..	203
Table A.18	Estimates of Time Required for Exporting Files.....	214
Table A.19	Termination Codes Returned by Export Tool	215
Table A.20	Subcommand List	216
Table A.21	Syntax Descriptions	217
Table A.22	Operands of the Group Subcommand and Saved Monitoring Data	224
Table A.23	Monitoring Items to which Longrange Subcommand can be Applied	239

Table A.24	Troubleshooting the Export Tool.....	250
Table A.25	Messages to be Issued by Export Tool	251
Table B.1	Limitations for Storage Partition Administrators (Performance Monitor)	254
Table C.1	Errors during Volume Migration or Migration Cancellation using CCI	259

Chapter 1 Overview of Performance Management

The Hitachi TagmaStore® Universal Storage Platform (USP) includes a suite of performance management software products. These products allow you to monitor and tune disk subsystem performance. The Performance Manager suite of software consists of the following:

- **Performance Monitor**
- **Volume Migration**
- **Server Priority Manager (SPM)**

Performance Monitor lets you obtain usage statistics about the physical hard disk drive, logical volumes, processors, or other resources in your disk subsystem. Performance Monitor also lets you obtain statistics about workloads on disk drives and traffic between hosts and the disk subsystem. The **Performance Monitor** 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. For an overview, see Chapter 2.

Volume Migration lets you balance workload among hard disk drives, logical volumes, and processors to remove bottlenecks from your system. If Performance Monitor indicates that a lot of I/O reads and writes are made to certain hard disk drives, you can use Volume Migration to distribute workloads to other disk drives. For details, see Chapter 3.

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 Chapter 4.

Operating Notes:

- 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.

Figure 1.1 illustrates the performance management solution from Hitachi Data Systems.

Performance Monitoring

- Base component of performance management
- Total performance monitoring




Figure 1.1 Performance Management Solution Components

Using Storage Navigator along with Performance Monitor, Volume Migration, and Server Priority Manager provides a simple and straightforward performance management solution.

Figure 1.2 illustrates and simplifies the performance management process.




Figure 1.2 Performance Management Process Flow Diagram

Chapter 2 Performance Monitor Overview

Performance Monitor tracks your disk subsystem and lets you obtain statistics about the following:

- Resources in your disk subsystem
- Workloads on disks and ports

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

Performance Monitor can also display the status of remote copies by TrueCopy, TrueCopy for z/OS®, Universal Replicator, and Universal Replicator for z/OS . The displayed contents are the same as those displayed in the **Usage Monitor** windows of each remote copy function.

An Export Tool enables you to save information on the **Performance Monitor** window into files, so you can use spreadsheet or database software to analyze the monitoring results. For detailed information about the Export Tool, see Appendix A.

This chapter is an overview of performance monitoring organized as follows:

- Section 2.1 Viewing Usage Statistics for Disk Subsystem Resources
- Section 2.2 Viewing Disk Workload Statistics
- Section 2.3 Viewing Port Traffic Statistics
- Section 2.4 Viewing LU Path Traffic Statistics
- Section 2.5 Viewing Traffic Between HBAs and Disk Subsystem Ports
- Section 2.6 Understanding Statistical Storage Ranges

2.1 Viewing Usage Statistics for Disk Subsystem Resources

If disk subsystem resources (e.g., hard disk drives) become overloaded, your system might require a lot of time for processing data and, in worst cases, might go down. To avoid such a situation, you must monitor resource utilization and, if necessary, take preventive measures.

Performance Monitor monitors your disk subsystem and obtains usage statistics about resources such as front- and back-end processors, hard disk drives, and logical volumes at an interval specified beforehand. You can specify any interval between 1 and 15 minutes. The **Performance Monitor** window allows you to view statistics about the resources and detect the resources on which workloads converge. The **Performance Monitor** window displays a line graph that helps you detect the peak time of the resource usage.

Note: The collected statistics that expired after a specific storing period are lost from the disk subsystem and are not displayed in the **Performance Monitor** window. Statistics about resources in the disk subsystem are stored by two methods with different collecting intervals and storage periods. Short-range information is stored between 1 and 15 days, depending on the collecting interval specified by the user. Long-range information is stored 93 days. For more information about data storage periods, see section 2.6.

If workloads converge on some particular resources, you need to balance the workloads by installing new hard disk drives or by using Volume Migration. For information about Volume Migration, see Chapter 3.

The Performance Monitor windows let you view statistical information about the following:

- Parity group Usage
- Logical Volume Usage
- External Volume Group Usage
- External Volume Usage
- Channel Processor Usage
- Disk Processor Usage
- Access Path Usage
- Data Recovery and Reconstruction Processor Usage
- Write Pending Rate and Cache Memory Usage
- Disk Workload
- Port Traffic
- LU Path Traffic
- Traffic Between HBAs and Disk Subsystem Ports

Each information type is briefly described in the following sections.

2.1.1 Viewing Parity Group Usage Statistics

A *parity group* is a group of hard disk drives (HDDs) that form the basic unit of storage for the TagmaStore USP and NSC subsystem. All HDDs in a parity group must have the same physical capacity. TagmaStore USP and NSC support the following three types of parity groups:

- **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 section 6.9.1.

2.1.2 Viewing Logical Volume Usage Statistics

Performance Monitor displays the average and maximum usage, including sequential and random access, of each logical volume (LDEV) in a parity group. The logical 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 logical volume usage, you should consider installing additional hardware (e.g., HDDs, disk adapters, 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 logical 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 logical volumes, see section 6.9.2.

2.1.3 Viewing External Volume Group Usage Statistics

If the TagmaStore USP and NSC disk subsystem is connected to an external disk subsystem 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 disk subsystem as logical volumes in the TagmaStore USP and NSC disk subsystem, the mapped volumes in the external disk subsystem 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.

Note: 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 windows of Performance Monitor treat external volumes the same as parity groups for convenience.

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 volumes, see section 6.9.3.

Note: Performance Monitor cannot monitor external volumes that are set for use by Cross-System Copy, therefore, the information about these volumes is not displayed in Performance Monitor windows.

2.1.4 Viewing External Volume Usage Statistics

An *external volume* is a volume existing in an external disk subsystem that is mapped to a volume in the TagmaStore USP and NSC disk subsystem 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 logical volume. For details on how to view usage conditions about external volumes, see section 6.9.4.

Note: Performance Monitor cannot monitor external volumes that are set for use by Cross-System Copy, therefore, the information about these volumes is not displayed in Performance Monitor windows.

2.1.5 Viewing 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 section 6.9.5.

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

2.1.6 Viewing 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 logical volumes from high-usage parity groups to low-usage parity groups.

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

1. See 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. Hitachi Data Systems recommends migrating volumes from higher-usage parity groups to lower-usage parity groups.

Table 2.1 Migrating Volumes when Disk Processor Usage Rate is High (not applicable to NSC)

Cluster	Disk Adapter	Disk Processor	How to Migrate Volumes
1	DKA-1A	DKP80-1A DKP81-1A DKP82-1A DKP83-1A DKP84-1A DKP85-1A DKP86-1A DKP87-1A	Migrate volumes from parity groups with the following IDs to another parity group: 1-x 2-x 3-x 4-x 5-x* 6-x* 11-x 12-x
1	DKA-1B	DKP90-1B DKP91-1B DKP92-1B DKP93-1B DKP94-1B DKP95-1B DKP96-1B DKP97-1B	Migrate volumes from parity groups with the following IDs to another parity group: 5-x 6-x 13-x 14-x
1	DKA-1L	DKPA0-1L DKPA1-1L DKPA2-1L DKPA3-1L DKPA4-1L DKPA5-1L	Migrate volumes from parity groups with the following IDs to another parity group: 7-x 8-x 9-x* 10-x*

Cluster	Disk Adapter	Disk Processor	How to Migrate Volumes
		DKPA6-1L DKPA7-1L	15-x 16-x 17-x* 18-x*
1	DKA-1K	DKPB0-1K DKPB1-1K DKPB2-1K DKPB3-1K DKPB4-1K DKPB5-1K DKPB6-1K DKPB7-1K	Migrate volumes from parity groups with the following IDs to another parity group: 9-x 10-x 17-x 18-x
2	DKA-2M	DKPC0-2M DKPC1-2M DKPC2-2M DKPC3-2M DKPC4-2M DKPC5-2M DKPC6-2M DKPC7-2M	Migrate volumes from parity groups with the following IDs to another parity group: 1-x 2-x 3-x 4-x 5-x* 6-x* 11-x 12-x
2	DKA-2N	DKPD0-2N DKPD1-2N DKPD2-2N DKPD3-2N DKPD4-2N DKPD5-2N DKPD6-2N DKPD7-2N	Migrate volumes from parity groups with the following IDs to another parity group: 5-x 6-x 13-x 14-x
2	DKA-2X	DKPE0-2X DKPE1-2X DKPE2-2X DKPE3-2X DKPE4-2X DKPE5-2X DKPE6-2X DKPE7-2X	Migrate volumes from parity groups with the following IDs to another parity group: 7-x 8-x 9-x* 10-x* 15-x 16-x 17-x* 18-x*
2	DKA-2W	DKPF0-2W DKPF1-2W DKPF2-2W DKPF3-2W	Migrate volumes from parity groups with the following IDs to another parity group: 9-x 10-x

Cluster	Disk Adapter	Disk Processor	How to Migrate Volumes
		DKPF4-2W	17-x
		DKPF5-2W	18-x
		DKPF6-2W	
		DKPF7-2W	

* Migrate these volumes only when the disk adapters are in the 1-pair configuration (Standard Disk Access model).

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 section 6.9.6.

Notes:

- 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.
- Table 2.1 is not applicable if you are using the TagmaStore NSC subsystem, because disk adapters are in one-pair configuration. Check the usage rate of each parity group.

2.1.7 Viewing 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 a 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 subsystem.

For details on how to view usage statistics on DRRs, see section 6.9.7.

2.1.8 Viewing Write Pending Rate and Cache Memory Usage Statistics

The write pending rate indicates the ratio of write-pending data to cache memory capacity. The **Performance Monitor** 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 cache memory usage statistics, see section 6.9.8.

2.1.9 Viewing Access Path Usage Statistics

An access path is a path through which data and commands are transferred within a disk subsystem.

In a disk subsystem, channel adapters control data transfer between hosts and the cache memory. Disk adapters control data transfer between the cache memory and disk drives. Data transfer does not occur between channel adapters and disk adapters. Data is 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 section 6.9.9.

2.2 Viewing Disk Workload Statistics

If particular disks or data are heavily accessed, system performance may deteriorate. Performance Monitor lets you view statistics about parity groups and logical devices to help you detect bottlenecks in your system. If you mapped logical volumes in an external disk subsystem, 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 peak I/O access times.

Note: You cannot view workload statistics that expired after a specific storing period because such statistics are erased from the disk subsystem. The storing period of statistics is only short range (between 1 and 15 days) and that changes depending on the collecting interval specified by the user. For details on the storing period of statistics, see section 2.6.

Workload information displayed by Performance Monitor is as follows:

- **I/O rate:**

The *I/O rate* indicates how many I/Os are made to the disk 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 disk 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 disk workloads. For details on how to view workload statistics about disks, see section 6.10.

2.3 Viewing Port Traffic Statistics

Performance Monitor tracks host ports and disk subsystem 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 section 6.5.1.

Notes:

- You will be unable to view workload statistics that expired after specific storing period because such statistics are erased from the disk subsystem. The storing period of statistics is only short range (between 1 to 15 days) and that changes depending on the collecting interval specified by the user. For details on the storing period of statistics, see section 2.6.
- 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.

2.4 Viewing 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 section 6.5.2.

Note: You cannot view workload statistics that expired after specific storing period because such statistics are erased from the disk subsystem. The storing period of statistics is only short range (between 1 to 15 days) and that changes depending on the collecting interval specified by the user. For details on the storing period of statistics, see section 2.6.

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.

2.5 Viewing Traffic between HBAs and Disk Subsystem Ports

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

If Server Priority Manager is enabled, Performance Monitor lets you view statistics about traffic between HBAs and disk subsystem ports. The traffic statistics reveal the number of I/O requests that have been made from hosts and also reveals the size of data transferred between hosts and disk subsystem ports.

Note: You cannot view workload statistics that expired after a specific storing period because such statistics are erased from the disk subsystem. The storing period of statistics is only short range (between 1 to 15 days) and that changes depending on the collecting interval specified by the user. For details on the storing period of statistics, see section 2.6.

2.6 Understanding Statistical Storage Ranges

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

- **Storing in short range**

Performance Monitor collects statistics at a user-specified interval of between 1 and 15 minutes, and stores them between 1 and 15 days. If the user specifies 1 minute as the collecting interval, the storing period is 1 day, and if the user specifies 15 minutes as the collecting interval, the storing period is 15 days. If there are CUs numbered 40 to FF (i.e., if there are CU groups other than the CU group 0), the storing period is 8 hours if the user specifies 1 minute as the collecting interval or 5 days if the user specifies 15 minutes as the collecting interval.

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-minute intervals, and stores them for 3 months (i.e., 93 days). However, if there are CUs numbered 40 to FF (i.e., if there are CU groups other than the CU group 0), Performance Monitor collect statistics at fixed 15-minute intervals and store them for 1 month (i.e., 31 days).

The usage statistics about resources in the disk subsystem 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 Monitor windows can display statistics within the range of the storing periods above. You can specify a part of the storing period to display statistics.

All statistics, except some information related to Volume Migration, can be displayed in short range (for the storing period corresponding to the collecting interval setting) on Performance Monitor windows. In addition, usage statistics about resources in the disk subsystem can be displayed in both short range and long range (for the storing period corresponding to the CU group for use by the user) 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 Monitor windows in section 6.1. For more about the relationship between collection interval and the storing period of the statistics, see section 6.7.

2.7 Overview of Performance Monitor Operations

This section briefly describes how to use Performance Monitor to monitor the disk subsystem. For more detail, see Chapter 6.

2.7.1 Start Monitoring

To start monitoring the disk subsystem, you start Performance Monitor and display the Monitoring Options window (Figure 2.2)



Figure 2.2 Monitoring Options Window of Performance Monitor

To start monitoring, select **Enable** for the Current Status option in Monitoring Switch, and specify the gathering interval, that is interval to collect the information, for the Gathering Interval option. After specifying that interval, click **Apply**. Performance Monitor starts to obtain statistics about the disk subsystem and the saves the statistics at the specified interval.

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 that 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.

Moreover, the resource usage in the disk subsystem 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 section 6.7.

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 z/OS, Universal Replicator, and Universal Replicator for 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.

2.7.2 View the Monitoring Results

To view the results of monitoring, use the **Performance Monitor** window of Performance Monitor. When you select the Performance Management tab, the **Performance Monitor** 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 disk subsystem are displayed in the list on the right side of the tree (see Figure 2.3).




Figure 2.3 Performance Monitor 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 disk subsystem.
- **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 disk subsystem.

- **WWN tab**

Enables you to view statistics about traffic at the path between Host Bus Adapters (HBAs) and ports.

Two sliders are displayed on the upper right of the **Performance Monitor** 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 2004/02/27 16:00 and the date and time on the right is 2004/02/28 16:00, you can view usage statistics for the period of 2004/02/27 16:00 to 2004/02/28 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 Monitor** window. The information in the window is updated at the specified gathering interval (every 1 to 15 minutes). Note that you cannot view the monitoring results displayed in the **Physical** tab and the **LDEV** tab in real time.

2.7.3 Notes on Performance Monitor

Keep the following in mind while viewing Performance Monitor windows:

- **Powering off the disk subsystem**

If you power off the disk subsystem during monitoring, monitoring stops only while the disk subsystem is powered off. When you power on the disk subsystem again, Performance Monitor cannot display the monitoring data while the disk subsystem 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 their storing periods because such statistics are erased from the disk subsystem.
- In short range, if I/O workloads between hosts and the disk subsystem become heavy, the disk subsystem gives higher priority to I/O processing than to monitoring processing, therefore, part of the monitoring data might be missing. If monitoring data is frequently missing, use the **Gathering Interval** option in the **Monitoring Options** window to make the collection interval longer. For details, refer to section 2.7.1.

Note: Short-range and long-range monitoring data may have some margin of error.

- Viewing the **LDEV**, **Port-LUN**, and **WWN** tabs
 - Monitoring results are stored for the last 1 to 15 days depending on the specified gathering interval. If the storing period have passed since a monitoring result was obtained, the result is erased from the disk subsystem, and you will not be able to view that monitoring result.
 - If I/O workloads between hosts and the disk subsystem become heavy, the disk subsystem gives higher priority to I/O processing than to 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 make the collection interval longer. For details, refer to section 2.7.1.
 - (WWN tab) If you want to start monitoring traffic between host bus adapters and disk subsystem ports, you must make settings before starting monitoring. For details, see section 8.4.1.
- Displaying the monitoring data
 - In lists of the **Performance Monitor** window contents, a hyphen (-) might be displayed in monitoring data column to indicate 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 Monitor** 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.
- Replacing the micro program

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

Chapter 3 Volume Migration Overview

Volume Migration enables you to optimize your data storage and retrieval on the TagmaStore USP and NSC subsystem. Volume Migration tuning operations can be used to balance resource utilization and resolve bottlenecks of system activity.

This chapter provides an overview of Volume Migration as follows:

- Section 3.1 Using Volume Migration to Balance Work Loads
- Section 3.2 About Migrating Logical Volumes Manually
- Section 3.3 About Migrating Logical Volumes Automatically
- Section 3.4 Requirements and Restrictions

3.1 Using Volume Migration to Balance Work Loads

If I/O accesses from hosts converge on particular 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 disk drive or a fast disk drive. Volume Migration can balance workloads among 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.

Note: 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 section 2.1.3.

- Volume Migration operations are completely nondisruptive—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 subsystem-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 logical 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 subsystem 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 disk subsystem. For details on CCI, please refer to the *Command Control Interface (CCI) User and Reference Guide*.

Caution: When an error condition exists in the TagmaStore USP and NSC subsystem, resource usage can increase or become unbalanced. Do not use data collected during an error condition as the basis for planning Volume Migration operations.

3.2 About Migrating Logical Volumes Manually

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

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

Note: You can also migrate external volumes manually. However, you cannot examine the estimated usage rate after migration because the usage rate of an external volume cannot be collected.

3.2.1 Migrating Logical Volumes Manually

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

- A logical volume for migrating, and
- The migration destination for the specified logical volume.

To migrate volumes, you execute your migration plan.

Figure 3.1 gives an example of a migration plan as it would be displayed in the **Volume Migration** window. 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
-	-	2	00:00	OPEN-3	2.28 GB	5(3D+1P)	1-1	DKR1B-J047FC	00:CLPR0	02:06	5(3D+1P)	2-1	DKR1B-J047FC	00:CLPR0	USP

Figure 3.1 Sample Migration Plan in 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 logical volumes as shown in Figure 3.2.
2. Choose a volume to be migrated from among high-usage volumes.




Figure 3.2 Usage Rate for Logical Volume

3. After determining a migration destination, find fast parity groups (i.e., fast disk drives).

To find fast disk drives, display the **Volume Migration** window and check classes (Figure 3.3). Classes indicate disk access speed for parity groups. Fastest parity groups (or groups of the 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 Universal Storage Platform documentation, the terms *class* and *HDD class* mean the same thing.




Figure 3.3 Classes and Parity Groups in the Volume Migration Window

4. Choose a low-usage parity group from the fast parity groups. You can find low-usage volumes in the **Volume Migration** window (see Figure 3.4).




Figure 3.4 Parity Group Usage Example

5. Display the contents of the low-usage parity group and look for disk areas preceded by icons, as shown in Figure 3.5.

Note: When you use Volume Migration for the first time, you must reserve logical volumes for migration destinations in advance. 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 instructions on how to reserve logical volumes for migration destinations, see section 7.3.




Figure 3.5 Specifying Migration Destination

6. Choose the destination disk area from among those marked with the icon.

TagmaStore USP documents sometimes use the term reserve volumes to refer to disk areas reserved as migration destinations (that is, preceded with the icon).

When you specify the volume that you want to migrate and the destination disk area, Volume Migration calculates the estimated volume usage rate after migration, and displays the estimated usage rate in the window (see Figure 3.6). 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 3.6, if the logical volume 00:01 in parity group 1-1 is migrated to the reserved volume 02:01 in parity group 1-4, the parity group 1-1 can expect a 2 percent decrease in the average usage rate (from 13 percent to 11 percent). The destination parity group 1-4 can expect a 31 percent increase in the average usage rate (from 14 percent to 45 percent). If the estimated result is not satisfactory, you can change the migration destination.

Source LDEV					
00:01	PG: 1-1	Ave.(Total) :	13 → 11 %	Max. :	14 → 12 %
Target (Reserved) LDEV					
02:01	PG: 1-4	Ave.(Total) :	14 → 45 %	Max. :	15 → 46 %

Figure 3.6 Estimated Usage Rate for Parity Groups

You can create multiple migration plans. When volume migration is finished, a volume to be migrated and a volume designated the migration destination will have replaced each other. For example in Figure 3.5, the logical volume 02:01 will be displayed in the parity group 1-1 and the logical volume 00:01 will be displayed in the parity group 1-4 at the upper-left tree. For more detailed information about Volume Migration, see section 7.4.

3.2.2 About the Manual Volume Migration Mechanism

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

Note: TagmaStore USP and NSC documents (including this user's guide) 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.

To migrate a volume, do the following:

1. Copy the data on the source volume to the target volume (see Figure 3.7).
2. Transfer host access to the target volume to complete the migration (see Figure 3.8).

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 subsystem.

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 logical volume. If there are deferential data formed by the update, the deferential 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 deferential 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 disk subsystem. The indication of the workloads between hosts and the disk subsystem shall be 50 IOPS or less for updating I/O.




Figure 3.7 Data Flow during a Volume Migration Operation

For auto and manual migration operations, Volume Migration checks the current subsystem 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.

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.




Figure 3.8 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, the RAID level before the replacement might be displayed or an internal error might occur in the disk subsystem. In this case, select **Refresh** to update the display of the window.

3.3 Automating Volume Migrations

If you automate the creation and execution of migration plans, you do not need to analyze usage rates of subsystem resources and speed of 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.

Note: Auto migration does not support external volumes.

3.3.1 Migrating Logical Volumes Automatically

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 section 6.2.
2. If necessary, specify fixed parity groups. If you automate volume migrations, some logical volumes and parity groups might suffer degraded I/O speeds. For example, if a logical volume 00:01 is migrated from the parity group Grp1-1 of the class A to parity group Grp-2-1 of the class B, the volume 00:01 is migrated to a lower-speed 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, logical volumes in Grp1-1 will not be migrated to any other parity group. Also, no logical volumes will be migrated to Grp1-1. For details on how to specify fixed parity groups, see section 7.5.1.

Note: In manual migration, a fixed parity group operates the same as a normal parity group. In this case, you can specify a logical 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 logical volumes might be migrated to parity groups in that class. If the disk usage limit for one class is low, fewer logical 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 section 7.5.2.

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 the time period on which you want statistical analysis to focus. 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 section 7.5.3.
5. Specify the time when Volume Migration executes migration plans. Hitachi Data Systems recommends specifying a time when disk usage rate is low because volume migrations typically produce heavy workloads on disk drives.
Volume Migration executes migration plans every day at the specified time.
6. Specify parameters that affect the system performance. The **Volume Migration** window lets you specify parameters that affect the system performance after volumes are migrated. For details, see section 7.5.3.
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**. For details, see section 7.5.3.

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 sections 7.5.3 and 7.5.4.

After auto migration finishes, you can use Performance Monitor to check whether the workloads on 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.

3.3.2 About the Auto Migration Mechanism

The auto migration function allows you to specify the maximum usage rate for each HDD class in the TagmaStore USP and NSC subsystem. The maximum usage rate applies to all the parity groups in the HDD class. 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 disk drive activity while providing load balancing of disk drive usage. Volume Migration uses the estimated usage rates to verify each proposed auto migration, and will not perform a migration operation that 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 3.9 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 3.9 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 (see Figure 3.10).




Figure 3.10 Keeping Reserve Volumes in High HDD Class Available

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 five (5) percent difference in estimated disk usage between the two volumes (see Figure 3.11). If the anticipated difference in disk usage is less than 5%, auto migration is considered an ineffective strategy, and the volume is not moved.




Figure 3.11 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 (see Figure 3.12). If the anticipated difference in disk usage is less than 20%, this migration is considered an ineffective strategy, and the volume is not moved.




Figure 3.12 Migrating a Volume Automatically within Same HDD Class

3.4 Migration Requirements and Restrictions

The requirements and restrictions for performing Volume Migration operations are as follows:

- **Selection of CU group.** When you select one of the following CU groups, the **Volume Migration** button is deactivated (the button is grayed out) and cannot be selected:
 - CU group 1 (CU 40-7F)
 - CU group 2 (CU 80-BF)
 - CU group 3 (CU C0-FE)
- **User types.** If the user type of your user account is *storage partition administrator*, you cannot use Server Priority Manager. For details on the limitations when using Server Performance Manager logged in as a storage partition administrator, see Appendix B.
- **Combinations of logical volumes.** The source and target volumes must be managed by the same TagmaStore USP and NSC subsystem. (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 **CU:LDEV** number of volumes accept a value between 00:00 and 3F: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 *customized volumes* (CV), or both the volumes are normal volumes.
- If Data Retention Utility is used with Volume Migration, volume management areas (VMA) are set on both the source and target volumes, or VMAs are not set on both the volumes.

Note: You cannot set VMA on a volume already reserved by Volume Migration. If you want to set VMA on a volume, perform the VMA setting first, and then reserve that volume. For details, see Table 3.1.

Table 3.1 Combinations of Volumes That Can Be Migrated

Source Volume	Target Volume					
	Normal Volume (not OPEN-V)	Normal Volume (OPEN-V, without VMA setting)	Normal Volume (OPEN-V, with VMA setting)	CV (not OPEN-V)	CV (OPEN-V, without VMA setting)	CV (OPEN-V, with VMA setting)
Normal Volume (not OPEN-V)	✓					
Normal Volume (OPEN-V, without VMA setting)		✓			✓	
Normal Volume (OPEN-V, with			✓			✓

Source Volume	Target Volume					
	Normal Volume (not OPEN-V)	Normal Volume (OPEN-V, without VMA setting)	Normal Volume (OPEN-V, with VMA setting)	CV (not OPEN-V)	CV (OPEN-V, without VMA setting)	CV (OPEN-V, with VMA setting)
VMA setting)						
CV (not OPEN-V)				✓		
CV (OPEN-V, without VMA setting)		✓			✓	
CV (OPEN-V, with VMA setting)			✓			✓

Checkmarks in Table 3.1 identify combinations that can be migrated. Blank squares are associated with volume combinations that cannot be migrated.

Specify the logical 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.

■ **Source volumes.** The following volumes cannot be used as source volumes:

- Volumes set as command devices (reserved for use by the host).
- Volumes that are in an abnormal or inaccessible condition (e.g., pinned track, fenced).
- Volumes used by the 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 that form a Cross-System Copy pair.

Note: A volume that forms a Cross-System Copy pair in the TagmaStore USP and NSC disk subsystem is displayed on the **Volume Migration** window, but you cannot manipulate it. A volume that forms a Cross-System Copy pair in the external disk subsystem is not displayed on the Volume Migration window.

- Volumes that 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 z/OS as a journal volume.
- Volumes that 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.

- Volumes being shredded.
- **TrueCopy for z/OS.** If the status of a TrueCopy for 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 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 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.
- **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.
- **Universal Replicator for z/OS:** If the status of a Universal Replicator for z/OS volume is Pending duplex or Duplex, the volume cannot be used as a source volume.
Note: If you specify a Universal Replicator for z/OS volume as a source volume, you cannot specify an external volume as the target volume.
- **Universal Replicator:** If the status of a Universal Replicator volume is COPY or PAIR, the volume cannot be used as a source volume. Note that if you specify a Universal Replicator volume as a source volume, you cannot specify an external volume as the target volume.
- **ShadowImage for z/OS and ShadowImage.** Usage of ShadowImage for 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 z/OS volumes whose status is SP-Vend or V-Split.
 - ShadowImage volumes whose status is COPY(SP) or PSUS(SP).
 Table 3.2 specifies which non-cascade volumes can be used as source volumes.

Table 3.2 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 3.3 specifies which cascade volumes can be used as source volumes.

Table 3.3 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:

- XRC operation
- CC operation
- TrueCopy for 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 z/OS operation that changes the volume status to COPY.
- ShadowImage for 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 or Universal Replicator for z/OS operation.
- **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 Hitachi 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 which are used by the Extended Remote Copy (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 which are used by the IBM 3990 Concurrent Copy (CC) host software function.

- Volumes which are reserved for ShadowImage operations.
- Volumes that have Cache Residency Manager or Cache Residency Manager for IBM z/OS data stored in cache.
- Volumes that form a Cross-System Copy pair.

Note: A volume that forms a Cross-System Copy pair in the TagmaStore USP and NSC disk subsystem is displayed on the **Volume Migration** window, but you cannot manipulate it. A volume that forms a Cross-System Copy pair in the external disk subsystem is not displayed on the **Volume Migration** window.

- Volumes used by Universal Replicator Universal Replicator for z/OS as a primary data volume, secondary data volume, or either of two journal volume types (master journal and restore journal).
- Volumes in which Volume Retention Manager specified the Read Only or Protect attribute.
- Volumes that 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.
- Virtual volumes and pool volumes.
- Volumes that are in an abnormal or inaccessible condition (e.g., pinned track, fenced).

You can also reserve external volumes as target volumes. However, if you specify a TrueCopy for z/OS volume or TrueCopy volume as a source volume, you cannot specify an external volume 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 TagmaStore USP programs, and also depending on the emulation types and sizes of the migrated volumes.

Volume Migration uses *differential tables* to manage differentials incurred by the write operation while migrating volumes. Differential tables are the resources shared by Copy-on-Write Snapshot, Volume Migration, ShadowImage, and ShadowImage for z/OS, Compatible Mirroring for IBM FlashCopy®, Compatible Mirroring for IBM FlashCopy Version 2, and Cross-System Copy. Therefore, the maximum number of differential tables that Volume Migration can use is the number of total differential tables in the disk subsystem minus the number 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 section 3.4.1. For details on the number of differential table required for migrating a volume, see sections 3.4.2 and 3.4.3.

- **Auto migration planning.** Volume Migration will not plan an auto migration operation if any of the following conditions apply:
 - 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 if any of the following conditions apply:
 - 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 subsystem is 60% or higher.
- **Manual migration execution.** Volume Migration will not execute a manual migration operation when the current write pending rate for the subsystem is 60% or higher.
- **If heavy workloads are imposed on your disk subsystem.** Avoid volume migration operations when heavy workloads are imposed onto your disk subsystem due to I/O requests from hosts. Volume migration attempts under those conditions could fail. To ensure success on subsequent attempts, please reduce the workloads on your disk subsystem before retrying volume migration.

Copy Threshold Option: When you set this option you can minimize the degradation of host I/O performance by temporarily stopping Volume Migration operations when the subsystem load is heavy.

Caution: Do not use the Copy Threshold option while hosts are frequently updating the volumes being migrated. This option temporarily stops Volume Migration copy processing when the subsystem is overloaded, thus increasing the possibility of volume migration failure.

If Volume Migration is performing migration operations during heavy subsystem load, it is more likely that host server I/O performance will become degraded. The Copy Threshold option temporarily halts Volume Migration operations during periods of heavy subsystem load.

For information about the setting of the Copy Threshold option, please call the Support Center.

Note: Copy operations that are stopped by the Copy Threshold option will resume when the subsystem load lightens. If this option is enabled, not only Volume Migration operations, but also the copy operations of these program products will stop when the subsystem load is heavy.

- ShadowImage
- ShadowImage for z/OS
- Compatible Mirroring for IBM FlashCopy

- Compatible Mirroring for IBM FlashCopy Version 2
- Copy-on-Write Snapshot
- Cross-System Copy
- **Progress Value:** If a volume with heavy workloads for updating I/Os between hosts and the disk subsystem is migrated with a volume with a light workloads for updating I/Os between hosts and the disk subsystem, the migration of the volumes may not complete, even if the progress value indicates 100%. In this case, perform either of the following to migrate the volumes.
 - Reduce the workloads for update I/Os between hosts and the disk subsystem and complete the migration.
 - Stop the volume migration for the volume with heavy workloads for update I/Os between hosts and the disk subsystem, execute a migration of another volume, and retry the migration that stopped.
- **Effects on Other Program Products:** When the workloads for update I/Os between hosts and the disk subsystem are heavy, it will take longer to complete the volume migration (see Table 3.4 for copying times). Longer copying time may be required for the following program products:
 - Volume Migration
 - ShadowImage
 - ShadowImage for z/OS
 - Compatible Mirroring for IBM FlashCopy
 - Cross-System Copy
- The copying times shown depend on the number of pairs set to the program product and may be doubled to the copying times shown in Table 3.4.

Table 3.4 Migrated Volume Capacity and Estimated Delay in Copying Speed

Capacity of Migrated Volume (MB)	Estimated Delay in Copying Speed (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 estimated times in Table 3.4 assume that the workload for update I/Os for the migrated volume is 50 IOPS for one volume.

- **Maintenance.** Do not perform migration Volume Migration operations during subsystem maintenance activities (e.g., cache or drive installation, replacement, uninstallation, or replacement of the microprogram).

- **Auto migration parameters.** Auto migration parameters are initialized when a change is made to the disk subsystem configuration, as with the installation of new drives or LDEVs. The parameters include the disk usage limit and the migration time. After a change is made to the subsystem configuration, you must specify auto migration parameters again.
- **Powering off the disk subsystem.** If you want to turn the power off at the disk subsystem, please confirm whether migration of volumes is finished. If migration is unfinished, we strongly recommend that you do not turn power off.

If you turn the power off when migration is in progress, data migration stops abruptly. If you turn power on later, Volume Migration resumes data migration. If the data remains in the shared memory (a volatile memory), Volume Migration attempts to copy only the data that has not been migrated to the migration destination. If data is no longer in shared memory, however, Volume Migration attempts to copy all available data to the migration destination, in effect re-copying even data that had been successfully migrated before power loss. This can add significant amounts of time to copy operations.

Note: TagmaStore USP and NSC documentation sometimes uses the abbreviation "PS ON" to mean turning on the power supply to the disk subsystem. Similarly, the abbreviation "PS OFF" is sometimes used to mean turning off the power supply to the disk subsystem.

3.4.1 Calculating the Number of Concurrent Migration Plans

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

- How much shared memory is available for differential tables:
 - You may use 13,652 differential tables if no additional shared memory for differential tables is installed.
 - You may use 30,718 differential tables if additional shared memory for differential tables is installed.

To install additional shared memory for differential tables, contact your HDS representative.

- The emulation type and capacity of each volume to be migrated:

The number of differential tables needed to migrate one volume differs depending on the emulation type and size of the volume. For the number of differential tables needed for migrating a mainframe volume, see section 3.4.2. For the number of differential tables needed for migrating an open-system volume, see section 3.4.3.

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

$$\Sigma(\alpha) \leq (B)$$

$\Sigma(\alpha)$ stands for the total number of differential tables needed for migrating all volumes.

(B) stands for the number of differential tables available in the TagmaStore USP and NSC subsystem.

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 the required differential tables can be found by the calculation described in section 3.4.2. When you apply this number to the equation, it will be as follows:

$$[(1 \times 20) = 20] \leq 13,652 \text{ or } 30,718$$

Since this equation is true, you can create all the migration plans that you want to create.

This example shows a method that can be used to calculate the maximum number of migration plans when only Volume Migration is running. However, the total number of differential tables used by Copy-on-Write Snapshot, ShadowImage, and ShadowImage for z/OS, Compatible Mirroring for IBM FlashCopy, Compatible Mirroring for IBM FlashCopy Version 2, Cross-System Copy, and Volume Migration should be within the value of (β). For details on how to calculate the number of differential tables used by the programs other than Volume Migration, refer to these manuals:

- *Hitachi TagmaStore® ShadowImage User's Guide (MK-94RD204)*
- *Hitachi TagmaStore® ShadowImage for z/OS® User's Guide (MK-94RD212)*
- *Hitachi TagmaStore® Universal Volume Manager User's Guide (MK-94RD220)*
- *Hitachi TagmaStore® Copy-on-Write Snapshot User's Guide (MK-95RD277)*

3.4.2 Determining the Number of 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 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.*

(Y): The number of the control cylinders. (See Table 3.7)

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

$$(1,916 \times 32)$$

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

Note that you should 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.

$$(3,339 + 6) \times 15 \div (1,916 \times 32) = 0.81836$$

When you round up 0.81836 to the nearest whole number, it becomes 1. Therefore, the total number of the differential table for one migration plan is 1 when emulation type is 3390-3.

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

Table 3.5 Number of 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

3.4.3 Determining the Number of Differential Tables for Open System Volume Migration

When you migrate open-system volumes, use the expression in Table 3.6 to calculate the total number of the required differential tables per migration plan.

Table 3.6 Total Number of 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) * (Y): The number of the control cylinders. (See Table 3.7) (Z): The number of the slots that can be managed by a differential table. $(1,916 \times 32)$ Note that you should 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. $(2,403,360 \div 48 + 8 \times 15) \div (1,916 \times 32) = 0.81860$ When you round up 0.81860 to the nearest whole number, it becomes 1. Therefore, the total number of the differential 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) * (Z): The number of the slots that can be managed by a differential table. $(1,916 \times 32)$ Note that you should 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. $(3,019,898,880 \div 256) \div (1,916 \times 32) = 192.40084$ When you round up 192.40084 to the nearest whole number, it becomes 193. Therefore, the total number of the differential tables for one migration plan is 193 when emulation type is OPEN-V.

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

Table 3.7 Number of Control Cylinders by Emulation Type

Emulation Type	Number of 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)

Chapter 4 Server Priority Manager Overview

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.

This chapter provides an overview of Server Priority Manager as follows:

- Section 4.1 Safeguarding the Performance of High-Priority Hosts
- Section 4.2 Disabling Upper Limit Control Automatically
- Section 4.3 Using Server Priority Manager
- Section 4.4 Notes on Using Server Priority Manager

4.1 Safeguarding Performance of High-Priority Hosts

In an SAN (storage area network) environment, the disk subsystem is usually connected to many host servers. Some host servers (such as production servers) require high performance but others often do not.

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. Any decline in development server performance, while bothersome to developers, would not be as likely to impact an entire organization negatively as a decline in production server performance. Accordingly, production servers typically have priority over development servers.

Server Priority Manager enables you to limit the number of I/O requests from development servers to the disk subsystem. Server Priority Manager also allows you to limit the size of data that should be transferred between the development servers and the disk subsystem. This reduces production server response time and enables better overall performance.

Throughout this chapter, 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.

4.2 Disabling Upper-Limit Control Automatically

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*. A threshold is a value that indicates the timing at which upper-limit control is disabled. If threshold control is used, upper-limit control is automatically disabled when traffic between production servers and the disk subsystem decreases to a certain level. For example, if a threshold of 500 IOPS (500 I/Os per second) is applied to the entire disk subsystem, 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 IOPS. If the number of I/Os from the production servers increases and exceeds 500, 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 disk subsystem, the I/O rate limit for development servers is disabled when the amount of data transferred between the disk subsystem and all the production servers is below 20 MB/s.

4.3 Using Server Priority Manager

Procedures for using Server Priority Manager differ depending on the connection between host bus adapters and disk subsystem 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.
- 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 cannot define the server priority by the port, because one port can connect to multiple host bus adapters, and 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 disk subsystem.
- 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 operating procedures associated with each tab.

Note: Host bus adapters (HBAs) are adapters contained in hosts. They serve as host ports for connecting the hosts and the disk subsystem.

4.3.1 If One-to-One Connections Link HBAs and Ports

Figure 4.1 shows an example of a network in which each host bus adapter is connected to only one port on the disk subsystem (Henceforth, this network is referred to as network A). Host bus adapters and the disk subsystem ports are connected directly, rather than through hubs and switches.




Figure 4.1 Network A (One-to-One connections between HBAs and Ports)

If one-to-one connections are established between HBAs and ports, take the following steps:

1. Set priority to ports on the disk subsystem 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 4.2 shows a portion of the **Server Priority Manager** window, where the Attribute column abbreviation **Prio.** indicates that the associated port is given high priority, and the abbreviation **Non-Prio.** indicates that the associated port is given low priority.

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 4.2 Priority Specified in the Server Priority Manager Window

2. Monitor traffic at ports. You must obtain statistics about traffic at each port on the disk subsystem. There are two types of traffic statistics:
 - The I/O rate is the number of I/Os per second (IOPS).
 - The transfer rate is the size of data transferred between a host and the disk subsystem.

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 Monitor** window lets you view a line graph that illustrates changes in traffic.

Figure 4.3 is a graph that illustrates 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 IOPS (I/Os per second) at first. The I/O rate for 2A was approximately 100 IOPS at first. However, as the I/O rate for 2A gradually increased from 100 to 200, the I/O rate for 1A and 1C decreased from 400 to 200. This 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 IOPS. To do that, you must set an upper limit to the I/O rate for the port 2A.

For detailed information about monitoring traffic, see sections 8.3.1 and 8.3.3.




Figure 4.3 Traffic at Ports (Example)

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 should be approximately 90 percent of the peak traffic level. In network A, the peak I/O rate for the non-prioritized port (2A) is 200 IOPS (I/Os per second). So, the recommended upper limit for 2A is 180 IOPS. For details on how to set an upper limit, see section 8.3.4.

4. Check the result of applying upper-limit values. After applying upper-limit values, you must measure traffic at ports. You 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. If the I/O rate reaches 400, 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 disk subsystem. In network A, if the upper limit is set to 180 but the I/O rate for 1A and 1C is still below 400, the administrator needs to change the upper limit until the I/O rate reaches 400.

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 for the entire disk subsystem:

In network A, if you set a threshold of 500 IO/s to the disk subsystem, 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 section 8.3.5.

4.3.2 If Many-to-Many Connections Link HBAs and Ports

Figure 4.4 gives an example of a network in which a production server and a development server are connected to the disk subsystem (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 4.4 Network B (Many-to-Many connections between HBAs and Ports)

If many-to-many connections are established between HBAs and ports, take the following steps to use the Server Priority Manager:

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 *TagmaStore USP and NSC LUN Manager User's Guide*.
2. Ensure that all host bus adapters connected to ports in the disk subsystem 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 4.5 Specifying Host Bus Adapters to be Monitored

Note: Server Priority Manager cannot 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 4.4, you must give high priority to wwn01 and low priority to wwn02 and wwn03.

To give priority to the three host bus adapters, take the following steps (see Figure 4.6):

- 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 4.6 Priority Specified in 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 disk subsystem. 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 4.7 illustrates 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 IOPS (I/Os per second) at first. The I/O rate at the path between 1A and the non-prioritized WWNs (wwn02 and wwn03) was approximately 100 at first. However, as the I/O rate for non-prioritized WWNs (wwn02 and wwn03) gradually increased from 100 to 200, the I/O rate for the prioritized WWN (wwn01) decreased from 400 to 200. This indicates that performance on the high-priority production server has degraded. If you were the network administrator, you would probably want to maintain the I/O rate for the prioritized WWN (wwn01) at 400 IOPS.

For more information about monitoring traffic, see sections 8.4.2 and 8.4.3.




Figure 4.7 Traffic at Ports (Example 2)

5. Set an upper limit to traffic between ports and the non-prioritized WWN to prevent any 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 I/Os per second (IOPS). The peak I/O rate at the paths between port 1C and the non-prioritized WWNs is also 200. The peak I/O rate at the paths between port 2A and the non-prioritized WWNs is also 200. The peak I/O rate at the paths between port 2C and the non-prioritized WWNs is also 200. So, the recommended upper limit for the non-prioritized WWNs is 720 IO/s ($= 200 \times 4 \times 0.90$).

If your network is configured as shown in Figure 4.4, you must do the following (see Figure 4.8):

- 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 4.8 is the result of setting the upper limit of 720 to the paths between 1A and the non-prioritized WWNs. For details on how to set an upper limit, see section 8.3.4.




Figure 4.8 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 has improved to a desirable level.

In network B, the desirable I/O rate for the prioritized WWN is 400 I/Os per second (IOPS). If production server performance has 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 disk subsystem. In network B, if the upper limit is set to 720 but the I/O rate for wwn01 is still below 400, the administrator needs to change the upper limit until the I/O rate reaches 400.

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 subsystem, regardless of the number of prioritized WWNs. For example, if there are three prioritized WWNs in the network and the threshold is 100, 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.

For details on how to set a threshold, see section 8.4.5.

4.4 Notes on Using Server Priority Manager

The following subsections discuss Server Priority Manager interaction with transfer rates and with remote copy functions.

Note that if your user ID is of the *storage partition administrator* type, you cannot use Server Priority Manager. For details on the limitations when using Performance Manager logged in as a storage partition administrator, see Appendix B.

4.4.1 Selecting CU Groups


- When you select one of the following CU groups, the **SPM** button is deactivated (the button is grayed out) and cannot be selected:
 - CU group 1 (CU 40-7F)
 - CU group 2 (CU 80-BF)
 - CU group 3 (CU C0-FE)

4.4.2 Calculating I/O 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 port 1A changes as illustrated in Graph 1 in Figure 4.9. 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 Monitor** window, 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, if as in Figure 4.9 with its gathering interval of 1 minute, you set an upper limit of 150 I/Os per second (IOPS) 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 IOPS. 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)



Graph 2: I/O rate displayed in Performance Monitor
(The Detail check box is not selected)



Graph 3: I/O rate displayed in Performance Monitor
(The Detail check box is selected)

Figure 4.9 Line Graphs indicating Port Traffic Changes

4.4.3 Using Server Priority Manager with Remote Copy Functions

When the remote copy functions (TrueCopy, TrueCopy for z/OS, Universal Replicator, and Universal Replicator for z/OS) are used in your environment, Server Priority Manager monitors write I/O requests issued from initiator ports of your disk subsystem.

- If you specify an RCU target port as a prioritized port, I/O requests from the initiator port are targeted by threshold control.

- If you specify an RCU target port as a non-prioritized port, I/O requests from the initiator port are not targeted by upper limit control.

Chapter 5 Preparing for Performance Manager Operations

Preparation for Performance Manager operations is discussed in the following sections:

- Section 5.1 System Requirements
- Section 5.2 Understanding the Role of Your Web Client Computer

5.1 System Requirements

To use Performance Manager, you need:

- The TagmaStore Universal Storage Platform subsystem, and
- Performance Manager software (Volume Migration and Server Priority Manager are optional, but Performance Monitor is required), and
- A Web client computer (intended for use as a Storage Navigator) connected to the TagmaStore USP and NSC 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™ applet, 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 user's guide for the TagmaStore USP and NSC Storage Navigator.

5.2 Understanding the Role of Your Web Client Computer

Performance Manager operations require the Storage Navigator program, which is downloaded to your Web client computer. Your Web client computer must be connected to the TagmaStore Universal Storage Platform via LAN. Performance Manager operations also depend on specific Web browser settings. For details about these and other configuration settings, see the user's guide for Storage Navigator.

Note: 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.

Chapter 6 Using Performance Monitor

Performance Monitor usage is explained in the following sections:

- Section 6.1 About the Performance Monitor Window
- Section 6.2 About Other Windows in the Interface
- Section 6.3 About the Physical Tab of the Performance Monitor Window
- Section 6.4 About the LDEV Tab of the Performance Monitor Window
- Section 6.5 About the Port-LUN Tab of the Performance Monitor Window
- Section 6.6 About the WWN Tab of the Performance Monitor Window
- Section 6.7 About the Monitoring Options Window
- Section 6.9 Monitoring Resources in the Disk Subsystem
- Section 6.10 Monitoring Disk Drives
- Section 6.11 Monitoring Ports
- Section 6.12 Monitoring LU Paths
- Section 6.13 Monitoring Paths Between Host Bus Adapters and Ports

Note: If the user type of your account is *storage partition administrator*, the functions you can use are limited. For details, see Appendix B.

6.1 About the Performance Monitor Window

The **Performance Monitor** window displays the monitoring results for disk subsystem performance as collected by Performance Monitor. The window includes four tabs: WWN, Port-LUN, LDEV, and Physical:

- **Physical tab**
Displays the usage statistics about resources in the disk subsystem (see section 6.3)
- **LDEV tab**
Displays the statistics about workload on disks (see section 6.4)
- **Port-LUN tab**
Displays the statistics about traffic at ports and LU paths in the disk subsystem (see section 6.5)
- **WWN tab**
Displays the statistics about traffic at the path between host bus adapters and ports (see section 6.6)

6.2 About Other Windows in the Interface

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

- **Monitoring Options** window

This window lets you make settings for monitoring and displaying the disk subsystem performance. You can start or stop collecting statistics by using this window (see section 6.7). Settings in this window are effective for all the monitoring windows of Performance Monitor.

- **TC Monitor** window

This window displays the information about remote copy operation of TrueCopy. The contents of this window are the same as those of the **Usage Monitor** window of TrueCopy. For more information about either window, please refer to the *TrueCopy User and Reference Guide*.

- **TCz (TrueCopy for z/OS) Monitor** window

This window displays the information about remote copy operation of TrueCopy for z/OS. The contents of this window are the same as those of the **Usage Monitor** window of Hitachi TrueCopy for z/OS. For more information about either window, please refer to the *TrueCopy for z/OS® User and Reference Guide*.

- **UR Monitor** window

This window displays the information about remote copy operation of Universal Replicator. The contents of this window are the same as those of the **Usage Monitor** window of Universal Replicator. For details, please refer to the *Universal Replicator User and Reference Guide*.

- **URz Monitor** window

This window displays the information about remote copy operation of Universal Replicator for z/OS. The contents of this window are the same as those of the **Usage Monitor** window of Universal Replicator for z/OS. For details, please refer to the *Universal Replicator for z/OS® User and Reference Guide*.

This chapter describes options, buttons and other components displayed in each tab of the **Performance Monitor** and **Monitoring Options** windows.

Note: If remote copy functions (TrueCopy, TrueCopy for z/OS, Universal Replicator, and Universal Replicator for z/OS) are not installed in your environment, their corresponding tabs (**TC Monitor** tab, **TCz Monitor** tab, **UR Monitor** tab, or **URz Monitor** tab) are inactivate and cannot be displayed.

6.3 About the Physical Tab of the Performance Monitor Window

When you select the  icon in the Storage Navigator main window, Performance Management starts and the **Performance Monitor** window is active.

The **Performance Monitor** window includes the **Physical** tab, which lets you view usage rates for parity groups, logical volumes, channel processors, disk processors, and array control processors. In addition, when you use Universal Volume Manager to map volumes in an external disk subsystem (disk subsystem other than the TagmaStore Universal Storage Platform) to the internal volumes, the **Physical** tab also lets you view usage conditions of volumes in the external disk subsystem (i.e., external volumes), and groups of external volumes.




Figure 6.1 The Physical Tab of the Performance Monitor Window

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

- When **Monitoring Switch** is set to **Enable**, Performance Monitor is monitoring the disk subsystem (a **Disable** setting indicates that the disk subsystem is not being monitored).
- **Gathering Interval** indicates the interval of collecting statistics in short-range monitoring. For example, if **1 min.** is displayed and *shortrange* is selected as the storing period of statistics, the list and graph in the **Physical** tab display the statistics obtained every minute.

Note: The gathering interval in long-range monitoring is fixed at 15 minutes. If you select *longrange* 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 more about the storage period for statistics, see section 2.6.

- In the drop-down list on the right of Monitoring Data, you can select the storing period of statistics from shortrange and longrange (see Figure 6.2).




Figure 6.2 Using the Drop-Down List to Select a Storage Period

The storing period of statistics is the range of monitoring data 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 Monitor** window.

Differences between shortrange and longrange are described in sections 6.3.1 and 6.3.2.

- **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 desired terms within this range as the target for display in lists and graphs.

For example, if you want to view usage statistics within the range of 10:30 July 1 2004 to 22:30 July 31 2004, you set 2004/07/01 10:30 to the **From** box, set 2004/07/31 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 result of the calculation. Length displays in days when you select **longrange**, and in minutes when you select **shortrange**.

Monitoring Term Notes:

- The **From** and **To** boxes are grayed out if the monitoring data (that is, usage statistics) is not stored in the disk subsystem.
- 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
	CHA-1F (NAS 4Port) *	a channel adapter for use in a NAS environment
	CHA-2U (iSCSI 4Port) *	a channel adapter for use in an iSCSI environment

* The channel adapter number and number of ports that displayed on the right side of the icon are examples.

- Icons displayed below the **Access Path Usage** 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 to 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.

Note: A volume existing in an external disk subsystem and mapped to a volume in the TagmaStore USP and NSC disk subsystem using Universal Volume Manager is called an *external volume*. An *external volume group* is a group of external volumes for managing them and that does not contain any parity information unlike a parity group. However, Performance Monitor windows treat external volume 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 concatenated. If two or more parity groups are concatenated, logical 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.

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

If you select the ACP (Array Control Processor) folder in the tree, the list displays a list of disk adapters (see Figure 6.3), 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 6.3 List of Disk Adapters

Notes:

- When you select the parity group icon, 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.

For details on the list contents, see section 6.9.

- The **Volume Migration** button starts Volume Migration if that program is enabled and **longrange** is specified for the display range.
- 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 shorrange, 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 6.4 shows the drop-down lists to select an item displayed in the graph and to select the highest value of the Y-axis.




Figure 6.4 Selecting an Item Displayed on the Graph

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 on the values of the selected item, Arrange the graph by changing the highest value of the Y-axis.

6.3.1 When You Select Long-Range Storage

The **Performance Monitor** 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 **longrange**, the **Volume Migration** button is activated. The system administrator can start Volume Migration to migrate logical volumes for balancing workloads based on the monitoring results displayed in the **Performance Monitor** window.

Notes for Longrange:

- When you select **longrange**, the value displayed at **Gathering Interval** is ineffective. The gathering interval is fixed to 15 minutes regardless of the displayed value.
- When **longrange** 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 **shortrange**. For details on viewing the usage statistics of external volumes and external volume groups, see sections 2.1.3 and 2.1.4.
- When **longrange** is selected, you cannot view usage statistics about cache memory.
- When **longrange** is selected, you cannot view the ratio of ShadowImage for z/OS and ShadowImage or other program processing to all processing. For details on the ratio of individual program processing to all processing, see section 6.9.2.

6.3.2 When You Select Short-Range Storage

The **Performance Monitor** 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, that is the range of monitoring data that can be displayed, changes between 1 and 15 days depending on the gathering interval.

When you select **shortrange**, the Volume Migration button is deactivated. In addition, estimated usage rates of logical volumes after migration by Volume Migration cannot be displayed.

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

For details on the relationship between collection interval and the storing period of the statistics, see section 2.6.

6.4 About the LDEV Tab of the Performance Monitor Window

When you select the  icon in the Storage Navigator main window, Performance Monitor starts and the **Performance Monitor** window is active. The **Performance Monitor** window includes the **LDEV** tab (Figure 6.5), 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 logical volumes.




Figure 6.5 LDEV tab of Performance Monitor Window

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

- When **Monitoring Switch** is set to **Enable**, Performance Monitor is monitoring the disk subsystem (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. For example, if **1 min.** is displayed, the information obtained every minute is displayed in the list and the graph.
- The drop-down list on the right of **Monitoring Data** indicates the storing period of statistics (monitoring data). The statistics displayed in the **LDEV** tab are stored only in short range. The range of monitoring data that can be displayed in the window is between 1 and 15 days depending on the gathering interval.

- **Monitoring Term** lets you narrow the range of statistics that are 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 desired term within this range as the target of display in lists and graphs.

For example, if you want to view statistics within the range of 10:30 July 1 2004 to 10:30 July 2 2004, you set 2004/07/01 10:30 to the **From** box, set 2004/07/02 10:30 to the **To** box, and then select 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 times 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.

Monitoring Term Notes:

- The **From** and **To** boxes are grayed out if the monitoring data (that is, obtained statistics) is not stored in the disk subsystem.
 - 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, and Box V1) are displayed below the subsystem 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). The RAID level is displayed to the right of each ID. When you select a parity group, the list on the right lists logical volumes in the parity group.

The IDs of external volume groups begin 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.

The IDs of V-VOL groups begin with the letter "V". A folder name that includes the letter "V" and a number, such as Box V1, 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 Monitor windows treat V-VOL groups the same as parity groups for convenience.

The parity group icon () can represent a single parity group, or two or more parity groups that are connected together. If two or more parity groups are connected together, logical 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 maintenance personnel.
 - When a ShadowImage or ShadowImage for z/OS quick restore operation is being performed, a **Storage Navigator** window may display old information (status before the quick restore operation) on logical volume (LDEV) configurations. In this case, wait until the quick restore operation completes, and then select Refresh () 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 response time).

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

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.

Note: 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 **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 selecting the **Draw** button, you should use the drop-down list at the right corner (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.




Figure 6.6 Selecting an Item to Display in the Graph (LDEV Tab)

The items that can be selected in the drop-down list change 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 statistical types that correspond to them are shown in Table 6.1.

Table 6.1 Cascade Volumes and Volume Migration Source Volumes

Item Selected in 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 disk drive.	✓	✓
Trans.	The transfer rate. The volume of data transferred per second.		✓
✓ marks selectable			

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 6.7 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 6.2. The information in the graph depends on the item selected in the drop-down list on the right of the list.

Table 6.2 Information Displayed in the LDEV Tab Graph

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 <i>Note:</i> 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.	- Number of data transfers from the cache memory to disk drives ("Cache to Drive"). - Number of data transfers from disk drives to the cache memory in sequential access mode ("Drive to Cache Sequential"). - Number of data transfers from disk drives to the cache memory in random access mode ("Drive to Cache Random").
Trans.	The graph does not display detailed information.

6.5 About the Port-LUN Tab of the Performance Monitor Window

When you select the  icon in the Storage Navigator main window, Performance Monitor starts and the **Performance Monitor** window is active. The **Performance Monitor** window includes the **Port-LUN** tab, which lets you view statistics about I/O rates, transfer rates, and average response time at disk subsystem 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.




Figure 6.8 Port-LUN Tab of the Performance Monitor Window

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

- When Monitoring Switch is set to **Enable**, Performance Monitor is monitoring the disk subsystem. (When Monitoring Switch is set to **Disable**, Performance Monitor is not monitoring the disk subsystem.)
- **Gathering Interval** indicates that the statistics are collected at the interval displayed here. 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 are displayed in the list and the graph.
- 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 1 and 15 days depending on the gathering interval.

- **Monitoring Term** lets 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 desired term within this range as the target of display in lists and graphs.

For example, if you want to view statistics within the range of 10:30 July 1 2004 to 10:30 July 2 2004, you set 2004/07/01 10:30 to the **From** box, set 2004/07/02 10:30 to the **To** box, and then select 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 disk subsystem.

- 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 that 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 port icon indicates either:

- a Fibre Channel port in **Standard** mode to which LUN security is applied, or
- an iSCSI port to which LUN security is applied.

Note: An iSCSI port is not supported in the present version of the TagmaStore USP and NSC.

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

Note: For detailed information about LUN security, please refer to the user's guide for LUN Manager.



This port icon indicates either:

- a Fibre Channel port in **Standard** mode to which LUN security is **NOT** applied, or
- an iSCSI port to which LUN security is NOT applied.

Note: An iSCSI port is not supported in the present version of TagmaStore USP.



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. For details of host groups, see the *USP and NSC LUN Manager User's Guide*.

When you double-click a host group, an item named LUN () is displayed. When you select LUN, the list on the right lists LU paths. For information about LU paths, see the *USP and NSC 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 sections 6.10 and 6.11.
- The **SPM** button starts the Server Priority Manager program if that program has been enabled.

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 conveniently for your purposes, as shown in Figure 6.9.




Figure 6.9 Port-LUN Tab Y Axis Detail

- 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, refer to Detail Statistics about Workload on a Port below.
- When you select LUN (📁) in the tree and select a LUN (an address of a logical 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 section 6.5.1.

6.5.1 Viewing Port Workload Statistics

Figure 6.10 shows an example of a graph displaying the workload of a port. In this example, the user selected the port CL1-A from the list and then selected **Draw**. In this example, the gathering interval is specified as 1 minute. The graph contents change 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 I/Os per second (IOPS). The average workload on CL1-A is 200 IOPS. The minimum workload on CL1-A is 100 IOPS (refer to the graph on the right).

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




Figure 6.10 Graphs Illustrating Changes Port Workloads

6.5.2 Viewing Statistics about Workload 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 6.3.

Table 6.3 Information Displayed in the Port-LUN Tab Graph

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 <i>Note:</i> 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 disk drives)	- Number of data transfers from the cache memory to disk drives ("Cache to Drive"). - Number of data transfers from disk drives to the cache memory in sequential access mode ("Drive to Cache Seq."). - Number of data transfers from disk drives to the cache memory in random access mode ("Drive to Cache Rnd.").

Note: The IO, Read, and Write items in the table above can only be selected when I/O rates are displayed.

6.6 About the WWN Tab of the Performance Monitor Window

When you select the  icon in the Storage Navigator main window, Performance Monitor starts and the **Performance Monitor** window is active. The **Performance Monitor** 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 disk subsystem.

Notes:

- The **WWN** tab is grayed out and cannot be used if Server Priority Manager is not enabled.
- 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 disk subsystem. One **WWN** represents one host bus adapter.




Figure 6.11 The WWN Tab of the Performance Monitor Window

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

- When Monitoring Switch is set to **Enable**, Performance Monitor is monitoring the disk subsystem (When Monitoring Switch is set to **Disable**, Performance Monitor is not monitoring the disk subsystem).
- **Gathering Interval** indicates that the statistics are collected at the interval displayed here. 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 are displayed in the list and the graph.

- 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, **shortrange** 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 1 and 15 days depending on the gathering interval.

- **Monitoring Term** lets 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 desired term within this range as the target of display in lists and graphs.

For example, if you want to view statistics within the range of 10:30 July 1 2004 to 10:30 July 2 2004, you set 2004/07/01 10:30 to the **From** box, set 2004/07/02 10:30 to the **To** box, and then select 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 disk subsystem.

- 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 that 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. 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.
- 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 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 displays 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. If many-to-many connections are established between HBAs and ports, make sure that all the traffics between HBAs and ports is monitored. See section 8.4.1 for how to monitor the traffic between the HBA and the all the connected ports.

- The list displays statistics (that is, I/O rates, transfer rates, or average response time).
- The **SPM** button starts the Server Priority Manager program if that program has been enabled.

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, 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 a 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 6.12 WWN Tab Y Axis Detail

6.7 About the Monitoring Options Window

When you select the  icon in the Storage Navigator main window, Performance Monitor starts. When you select the **Monitoring Options** tab, the **Monitoring Options** window is displayed. Use it to make settings for obtaining usage rates about disk drives, channel processors, disk processors, etc.



Figure 6.13 The Monitoring Options Window of Performance Monitor

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

- **Current Status**

To start obtaining statistics of the disk subsystem (that is, monitoring), select **Enable**. To stop monitoring, select **Disable**. The default setting is **Disable**.

- **Gathering Interval**

Specify the interval to obtain usage statistics about the disk subsystem for short range monitoring. This option is activated when you specify Enable for Current Status. 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.

This option is effective only for:

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

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

- The **Apply** button applies settings in the **Monitoring Options** window to the disk subsystem.
- 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 by multiplying the gathering interval by 1440. For example, if you specify one minute for the gathering interval, the statistics for one day can be stored at the maximum value derived 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 Monitor 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 **longrange** 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 are used in conjunction with the monitoring switch settings in the Usage Monitor windows of TrueCopy, TrueCopy for z/OS, Universal Replicator, and Universal Replicator for 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 change automatically.

6.8 Starting and Stopping Disk Subsystem Monitoring

To monitor the disk subsystem, start Performance Monitor 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:

1. Log onto Storage Navigator. The Storage Navigator main window is displayed.
2. Select the  icon. The **Performance Monitor** window is displayed.

To start monitoring the disk subsystem:

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 Monitor** 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 (see Figure 6.13).
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 (refer to section 6.7).
5. Select **Apply**. Performance Monitor starts monitoring the disk subsystem.

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 disk subsystem:

1. Start Performance Monitor and select the **Monitoring Options** tab. The **Monitoring Options** window is displayed (see Figure 6.13).
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 disk subsystem.

6.9 Monitoring Resources in the Disk Subsystem

This section describes how to view usage statistics on resources in the disk subsystem.

Before taking the following steps, you need to start monitoring in accordance with the procedure described in section 6.8 and obtain the usage statistics.

6.9.1 Viewing Usage Statistics about 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 Monitor** window is displayed.
2. In the tree, select the **Physical** tab.
3. In the drop-down list above the tree, select the storing period of statistics from **longrange** and **shortrange** for display.

For information on statistical storage periods, refer to section 2.6.

When you view usage statistics about parity groups, the items displayed in the list by selecting **longrange** and **shortrange** are the same.

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

The list on the right displays usage statistics about parity groups (Figure 6.14). The displayed statistics are the average and the maximum usage rates for the period specified in the **From** and **To** boxes.

5. To display a graph that illustrates changes in usage rate for parity groups, select the desired parity groups in the list and then select the **Draw** button.

Parity Group Statistics Notes:

- If you change the date and time in the **From** and **To** boxes and then select 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, refer to section 6.3).
- 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).
- 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.
- The range of monitoring and the gathering interval affects the time period represented by a graduation on the horizontal axis.

- The usage rate for a parity group may not be equal to the sum of the usage rate for each logical volume in that parity group. This is because the **Performance Monitor** window rounds off fractions below the decimal point to the nearest whole number when displaying the usage rate for each logical volume.

	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	12	13	01:CLPR1
	1-3	5(3D+1P)	DKR1B-J047FC	13	14	02:CLPR2
	1-4	5(3D+1P)	DKR1B-J047FC	14	15	03:CLPR3

Figure 6.14 Sample List of Parity Group Usage Rates

The list displays the following items:

- When the green check mark  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, please refer to the *Virtual Partition Manager User's Guide*.

6.9.2 Viewing Usage Statistics about Logical Volumes in Parity Groups

Performance Monitor monitors logical 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 logical volume usage within that period.

To view usage statistics about logical volumes in a parity group:

1. Ensure that the **Performance Monitor** window is displayed.
2. In the tree, select the **Physical** tab.
3. In the drop-down list above the tree, select the storing period of statistics from **longrange** and **shortrange** for display.

For details on the types of storing period of statistics, refer to section 2.6.

When you view usage statistics about logical volumes, some items displayed in the list differ depending on the selection of the storing period: **longrange** or **shortrange**. If you want to use Volume Migration to migrate logical volumes, select **longrange**. This enables you to confirm the estimated usage rate of parity groups after migration.

If you want to examine the ratio of ShadowImage for z/OS and ShadowImage or other individual program processing to all the processing in the physical drive, select **shortrange**.

Note: **Shortrange** and **longrange** usage statistics for the same term might be slightly different because the monitoring precision of these two interval types differs. Differences in read rates and write rates between the **shortrange** and **longrange** interval types are larger than ranges associated with other usage statistics.

4. In the tree, double-click the Parity Group folder. The folder opens. 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 logical volumes in the specified parity group (Figure 6.15). The displayed statistics are the average and the maximum usage rates for the period specified in the **From** and **To** boxes.
6. If you want to display a graph that illustrates changes in usage rate for logical volumes, select the desired logical volumes in the list and then select the **Draw** button.

Logical Volume Parity Group Statistics Notes:

- If you change the date and time in the **From** and **To** boxes and then select 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, refer to section 6.3).
- If an exclamation mark (!) is displayed before a usage rate, the reported logical 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).
- If a plus sign (+) is displayed before a usage rate 0, such as "+0", the usage rate is not completely 0 but less than 1.
- The range of monitoring and the gathering interval affects the time period represented by a graduation on the horizontal axis.
- The usage rate for a parity group may not be equal to the sum of the usage rate for each logical volume in that parity group. This is because the **Performance Monitor** window rounds off fractions below the decimal point to the nearest whole number when displaying the usage rate for each logical volume.

When selecting longrange 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.	
	00:01	OPEN-3	11	12	3	4	5	6	13 -> 7	14 -> 8	00:CLPR0
	00:02	OPEN-3	2	3	4	5	6	7	13 -> 8	14 -> 9	00:CLPR0
	00:03	OPEN-3	13	14	5	6	7	8	13 -> 9	14 -> 10	00:CLPR0
	00:04	OPEN-3	4	5	6	7	8	9	13 -> 10	14 -> 11	00:CLPR0
	00:05	OPEN-3	15	16	7	8	9	10	13 -> 11	14 -> 12	00:CLPR0

When selecting shorrange for storing period of statistics:

	LDEV	Emulation	Usage(%)		Read Rate(%)		Write Rate(%)		ShadowImage(%)		CLPR
			Ave.(Total)	Max.	Rnd.	Seq.	Rnd.	Seq.	Ave.(Total)	Max.	
	00:01	OPEN-3	11	12	13	14	15	16	17	18	00:CLPRO
	00:02	OPEN-3	12	13	14	15	16	17	18	19	00:CLPRO
	00:03	OPEN-3	13	14	15	16	17	18	19	20	00:CLPRO
	00:04	OPEN-3	14	15	16	17	18	19	20	21	00:CLPRO
	00:05	OPEN-3	15	16	17	18	19	20	21	22	00:CLPRO

Figure 6.15 Examples of Logical Volume Usage Rates Displayed

The list displays the following items:

-  : When the green check mark  icon is displayed to the left of a logical volume, the graph illustrates changes in usage rate for the logical volume.
- **LDEV:** This column indicates logical volumes (LDEVs). The number on the left of the colon is the CU image number. The number on the right of the colon is the LDEV number.
- **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 **longrange** 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 **shortrange** 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 logical 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 z/OS
- ShadowImage
- Cross-system Copy
- Compatible Mirroring for IBM FlashCopy
- Compatible Mirroring for IBM FlashCopy Version 2
- 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, please refer to their respective user guides.

- **CLPR:** This column indicates numbers and names of CLPRs corresponding to each parity group which the logical volume belongs to, in the format "CLPR-number:CLPR-name". For details on CLPRs, please refer to *Virtual Partition Manager User's Guide*.

6.9.3 Viewing Usage Statistics about 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 **shortrange** for the storing period of statistics.

To view usage statistics about external volume groups:

1. Ensure that the **Performance Monitor** window is displayed.
2. In the tree, select the **Physical** tab.
3. In the drop-down list above the tree, select **shortrange** for the storing period of statistics. If you select **longrange**, no statistics are displayed.
4. In the tree, select the External Group folder. The list on the right displays usage statistics about external volume groups (Figure 6.16). The displayed statistics are the usage statistics for the period specified in the **From** and **To** boxes.
5. If you want 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 select the **Draw** button.

External Volume Group Statistics Notes:

- If you change the date and time in the **From** and **To** boxes and then select 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, refer to section 6.3).
- 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 6.16 Sample External Volume Group Usage Statistics

The list displays the following items:

-  : When the green check mark  icon is displayed to 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 TagmaStore USP and NSC disk subsystem 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 disk subsystem 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, please refer to the *Virtual Partition Manager User's Guide*.

6.9.4 Viewing Usage Statistics about External Volumes in EV 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.

Notes:

- You can view the usage statistics about external volumes only when you select **shortrange** for the storing period of statistics.
- The usage statistics about external volumes that are set for use by Cross-System Copy are not displayed in the window.

To view usage statistics about external volumes in an external volume group:

1. Ensure that the **Performance Monitor** window is displayed.

2. In the tree, select the **Physical** tab.
3. In the drop-down list above the tree, select **shortrange** for the storing period of statistics.
If you select **longrange**, no statistics appears in the list.
4. In the tree, select the External Group folder. The folder opens. A list of external volume groups appears 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 (Figure 6.17). 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 select 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, refer to section 6.3).

6. If you want to display a graph that illustrates changes in usage rate for external volumes, select the desired external volumes in the list and then select 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
	06:04 #	OPEN-9	2000	2.930	02:CLPR2
	06:05 #	OPEN-9	2001	2.931	02:CLPR2
	06:06 #	OPEN-9	2002	2.932	02:CLPR2
	06:07 #	OPEN-9	2003	2.933	02:CLPR2

Figure 6.17 Sample External Volume Usage Rates Displayed

The list displays the following items:

-  : When the green check mark  icon is displayed to the left of an external volume, the graph illustrates changes in usage rate for the external volume.
- **ExLDEV:** This column indicates external volumes. The number on the left of the colon is the CU image number. The number of the right of the colon is the ExLDEV number. The symbol "#" at the end of numbers indicates the logical 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 the TagmaStore USP and NSC disk subsystem 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 disk subsystem 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, please refer to the *Virtual Partition Manager User's Guide*.

6.9.5 Viewing Usage Statistics about 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 Monitor** window is displayed.
2. In the tree, select the **Physical** tab.
3. In the drop-down list above the tree, select the storing period of statistics from **longrange** and **shortrange** for display.

For details on statistical storage periods, refer to section 2.6. When you view usage statistics about channel processors, the items displayed in the list by selecting longrange and shortrange are the same.

4. In the tree, do either of the following:
 - If you want to view usage statistics about all the channel processors in your disk subsystem, 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 (Figure 6.18). 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 select 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 section 6.3).

5. To display a graph that illustrates changes in usage rate for channel processors, select the desired channel processors in the list and then select 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 6.18 Sample Channel Processor Usage Rates Displayed

The list displays the following items:

-  : When the green check mark  icon is displayed to 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.

6.9.6 Viewing Usage Statistics about 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 Monitor** window is displayed.
2. In the tree, select the **Physical** tab.
3. In the drop-down list above the tree, select the storing period of statistics from **longrange** and **shortrange** for display. For details on the types of storing period of statistics, refer to section 2.6. When you view usage statistics about disk processors, the items displayed in the list by selecting **longrange** and **shortrange** are the same.
4. In the tree, double-click the **ACP** folder.
5. Select **DKP** from below the **ACP** folder. The list on the right displays usage statistics about disk processors (Figure 6.19). 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 select 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, refer to section 6.3).

6. 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 select 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 6.19 Sample Disk Processor Usage Rates

The list displays the following items:

- : When the green check mark  icon is displayed to 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.

6.9.7 Viewing Data Recovery and Reconstruction Processor Usage

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 Monitor also displays a graph that illustrates changes in DRR usage within that period.

Note: Usage statistics can only be viewed if you are monitoring the disk subsystem.

To view usage statistics about DRRs:

1. Ensure that the **Performance Monitor** window is displayed.
2. In the tree, select the **Physical** tab.
3. In the drop-down list above the tree, select the storing period of statistics from **longrange** and **shortrange** for display. For details on statistical storage periods, refer to section 2.6.

When you view usage statistics about DRRs, the items displayed in the list by selecting **longrange** and **shortrange** are the same.

4. In the tree, double-click the **ACP** (in some configurations, **DKA**) folder.
5. Select DRR from below the **ACP** or **DKA** folder. The list on the right displays usage statistics about DRRs (Figure 6.20). 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 select the **Apply** button, Performance Monitor recalculates usage rates and updates information in the table. To change the date and time in the **From** and **To** boxes, use the arrow buttons and the sliders (for details, refer to section 6.3).

- To display a graph that illustrates changes in usage rate for DRRs, select the desired DRRs in the list and then select 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 6.20 Example of DRR Usage Rates Displayed

The list displays the following items:

-  : When the check mark icon () is displayed in this column, 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.

6.9.8 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 **shortrange** 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 cache memory usage statistics:

- Ensure that the **Performance Monitor** window is displayed.
- In the tree, select the **Physical** tab.
- In the drop-down list above the tree, select the storing period of statistics from **longrange** and **shortrange** for display. For details on statistical storage period options, refer to section 2.6.

The items displayed in the list change depending on storing period selected. Write pending rates are displayed by both selections, but usage statistics of the cache memory are displayed only when you select **shortrange**.

- In the tree, select the **Cache** folder.

The list on the right displays the write pending rates. If you select **shortrange** for the storing period of statistics, the list also displays usage statistics about the cache memory (Figure 6.23). 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 select 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, refer to section 6.3).

5. 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 select 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 **shortrange** 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 illustrate the graph of the write pending rate, and select **Usage** to illustrate that of usage statistics of the 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.




Figure 6.21 Selecting an Item to be Displayed in the Graph (Shortrange)

Examples of write pending rate displayed in the lists are shown below:

When selecting longrange 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 shorrange 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 6.22 Examples of Write Pending and Cache Usage Rates Displayed

The list displays the following items:

-  : When the check mark icon () is displayed in this column, the graph illustrates changes in the write pending rate and usage statistics about 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, please refer to the *Virtual Partition Manager User's Guide*.
- **Usage:**

This item is displayed only when you select shorrange 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.

6.9.9 Viewing Usage Statistics about Access Paths

The channel adapters (CHAs) and the disk adapters (DKAs) transfer data to the cache switch (CWS) and the shared memory (SM) when I/O requests are issued from hosts to the disk subsystem. In some configurations, DKAs are called array control processors (ACPs). Also, the cache switch transfers data to the cache memory.

Performance Monitor monitors 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 that illustrates changes in path usage within that period.

To view usage statistics about paths:

1. Ensure that the **Performance Monitor** window is displayed.
2. In the tree, select the **Physical** tab.

3. In the drop-down list above the tree, select the storing period of statistics from **longrange** and **shortrange** for display.

For details on the types of storing period of statistics, refer to section 2.6.

When you view usage statistics about paths, the items displayed in the list by selecting **longrange** and **shortrange** 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**, below the Access Path Usage folder.
 - To check usage statistics about paths between adapters (CHAs and DKAs) and the shared memory, select **Adapter-SM**, below the Access Path Usage folder.
 - To check usage statistics about paths between cache switches and the cache memory, select **CSW-Cache**, below the Access Path Usage folder.

The list on the right displays the average and maximum usage rate for the specified paths (Figure 6.23). 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 select 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, refer to section 6.3).

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 select 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 6.23 Examples of Usage Statistics Displayed

The list displays the following items:

-  : When the check mark icon () is displayed in this column, the graph illustrates changes in the 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.

6.10 Monitoring Disk Drives

The **LDEV** tab of the **Performance Monitor** window lets you check workloads on physical disk drives (parity groups) or on logical 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 section 6.2 and obtain the usage statistics.

6.10.1 Viewing I/O Rates for Disks

Performance Monitor monitors disk drives and measures I/O rates (that is, the number of disk I/Os per second or IOPS).

To view I/O rates:

1. Ensure that the **Performance Monitor** window is displayed.
2. Select the **LDEV** tab (Figure 6.5). 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 either of the following:
 - 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 (Figure 6.24).

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-".

- To view the I/O rate for each logical volume, select the parity group or external volume group that contains the logical volumes.

The list on the right displays the I/O rate for each logical volume in the selected group (Figure 6.24).

The displayed statistics are the average and the maximum I/O 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 select 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, refer to section 6.3).

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 logical 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, refer to Table 6.1.
 - c) Select the **Draw** button. A graph is displayed below the list. The horizontal axis indicates the time.

Note: If the graph does not display changes in the I/O rate (for example, if the line in the graph runs vertically), 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 select the **Draw** button. The detailed graph contents change as described in Table 6.2 in section 6.4.

Notes:

- If more than one parity group or logical volume is selected in the list, you cannot select the Detail check box to view detailed information.
- The I/O rate for a parity group or an external volume group may not be equal to the sum of the I/O rate for each logical volume in that group. This is because the **Performance Monitor** window omits fractions below the decimal point when displaying the I/O rate for each logical volume.

I/O rate for parity groups and external volume groups (Subsystem folder 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 (Box 1 folder 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 logical volumes:

	LDEV	Emulation	IO Rate(IOPS)	Read(IOPS)	Write(IOPS)	Read Hit(%)	Write Hit(%)	Back Trans.(cou...)	Response Time(ms)	CLPR
	00:00	OPEN-3	1	2	3	4	5	6	5.008	00:CLPR0
	00:01	OPEN-3	2	3	4	5	6	7	5.009	00:CLPR0
	00:02	OPEN-3	3	4	5	6	7	8	5.010	00:CLPR0
	00:03	OPEN-3	4	5	6	7	8	9	5.011	00:CLPR0
	00:04	OPEN-3	5	6	7	8	9	10	5.012	00:CLPR0
	00:05	OPEN-3	6	7	8	9	10	11	5.013	00:CLPR0

Figure 6.24 I/O Rates for Disks

The list displays the following:

-  : When the green check mark  icon is displayed to the left of an item, the graph displays workload statistics about the parity group, the external volume group, or the logical 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", the group is a V-VOL group.
- **LDEV:** Indicates the logical volume ID. If the ID ends with the symbol "#", the logical volume is an external volume. If the ID ends with the symbol "V", the logical 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 logical volume per second.
- **Read:** Indicates the number of read requests to the parity group, the external volume group, or the logical volume per second.
- **Write:** Indicates the number of write requests to the parity group, the external volume group, or the logical 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 logical volume) and the cache memory.
- **Response Time:** This column indicates the time for replying from the parity group, the external volume group, or the logical 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: For the volumes whose emulation type is mainframe or the parity groups that contain mainframe volumes, the values of Response Time are always zero (0).

- **CLPR:** This column indicates numbers and names of CLPRs corresponding to each parity group, external volume group, or logical volume in the format "CLPR number:CLPR name". For details on CLPRs, please refer to the *Virtual Partition Manager User's Guide*.

6.10.2 Viewing Transfer Rates for Disks

To view transfer rates:

1. Ensure that the **Performance Monitor** window is displayed.
2. Select the **LDEV** tab (Figure 6.5). 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 either 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 (Figure 6.25).
 - To view the transfer rate for each logical volume, select the parity group or external volume group that contains the logical volumes.
The list on the right displays the transfer rate for each logical volume in the selected group (Figure 6.25).

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-".

The displayed statistics are the average and the maximum transfer 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 select 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, refer to section 6.4).

5. 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 logical 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, refer to Table 6.1 in section 6.4.
- c) Select the **Draw** button. A graph is displayed below the list. The horizontal axis indicates the time.

Note: If the graph does not display changes in the transfer rate (for example, if the line in the graph runs vertically), 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 select the **Draw** button. The detailed graph contents change as described in Table 6.2 in section 6.4.

Note: If more than one parity group or logical 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 (Subsystem folder 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 (Box 1 folder 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 logical volumes:

	LDEV	Emulation	Trans.(MB/s)	Read Hit(%)	Write Hit(%)	Back Trans.(count/sec)	Response Time(ms)	CLPR
	00:00	OPEN-3	1.960	4	5	6	5.008	00:CLPR0
	00:01	OPEN-3	1.961	5	6	7	5.009	00:CLPR0
	00:02	OPEN-3	1.962	6	7	8	5.010	00:CLPR0
	00:03	OPEN-3	1.963	7	8	9	5.011	00:CLPR0

Figure 6.25 Transfer Rates for Disks

Note: The transfer rate for a parity group or an external volume group may not be equal to the sum of the transfer rate for each logical volume in that group. This is because the **Performance Monitor** window omits fractions below the decimal point when displaying the transfer rate for each logical volume.

The list displays the following:

-  : When the green check mark  icon is displayed to the left of an item, the graph displays workload statistics about the parity group, the external volume group, or the logical 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 logical volume ID. If the ID ends with the symbol "#", the logical volume is 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 logical 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 logical volume) and the cache memory.
- **Response Time:** Indicates the time for replying from the parity group, the external volume group, or the logical 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.

Notes:

- Response times of volumes whose emulation type is mainframe and of parity groups that contain mainframe volumes are always 0.
- 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 logical volume in the format "CLPR number:CLPR name". For details on CLPRs, please refer to the *Virtual Partition Manager User's Guide*.

6.11 Monitoring Ports

Performance Monitor monitors ports on the disk subsystem 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 disk subsystem. Before taking the following steps, you need to start monitoring in accordance with the procedure described in section 6.8 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.

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

6.11.1 Viewing I/O Rates for Ports

Performance Monitor monitors ports on the disk subsystem 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 Monitor** window is displayed.
2. Select the **Port-LUN** tab (Figure 6.8). The tree displays a list of ports. If you double-click a port icon, the host groups corresponding to that port are displayed.
3. Select IOPS from the drop-down list on the right side of the window.
4. In **Monitoring Term**, do either of the following:
 - To view the I/O rate in real time, 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, select the **From** option, change the date and time in the **From** and **To** boxes, and then select **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** and **From** options, refer to section 6.5.

5. In the tree, select the Subsystem folder. The list on the right displays I/O rates for the ports on the subsystem (Figure 6.26).

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 (Figure 6.26).

- If you select a host group (🖨️) in the tree, the list displays I/O rates for host bus adapters in the host group (Figure 6.26).
6. 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.

Note: If the graph does not display changes in the I/O rate (for example, if the line in the graph runs vertically), 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	18	1.999	Prio.
	All Non-Prio.	-	13	19	1.998	Non-Prio.
	CL1-A(E8)	-	15	70	2.000	Prio.
	CL1-B(E4)	-	65	120	2.050	Non-Prio.
	CL1-X(E2)	-	115	170	2.100	Non-Prio.
	CL2-a(E1)	-	165	220	2.150	Prio.

All Prio. indicates all the prioritized ports.

All Non-Prio. indicates all the non-prioritized ports

Notes:

- iSCSI ports are always *prioritized ports*. Therefore, the **Attribute** column displays **Prio.** for each iSCSI port.
- When calculating I/O rates to be displayed in the **All Prio.** row, Performance Monitor uses I/O rates of iSCSI ports and other prioritized ports.

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 6.26 I/O Rates for Ports

The list displays the following:

- : When the green check mark  icon is displayed to the left of an item, the graph displays changes in workload statistics about the item.
- **Port:** Indicates ports on the disk subsystem.
- **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 adapter 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 for **Port-LUN** tab, two types of aliases are displayed for host bus adapters: SPM name and nickname. If you select a port icon in the tree, SPM names defined by Server Priority Manager are displayed. If you select a host group icon in the tree, nicknames defined by LUN Management are displayed. We recommend that you specify the same name for an SPM name and a nickname for convenience of host bus adapter management.

6.11.2 Viewing Transfer Rates for Ports

Performance Monitor monitors ports on the disk subsystem and measures transfer rates (that is, the volume of data transferred in one second).

To view transfer rates:

1. Ensure that the **Performance Monitor** window is displayed.
2. Select the **Port-LUN** tab (Figure 6.8). The tree displays a list of ports. If you double-click a port icon, the host groups corresponding to that port are displayed.

3. Select **MB/s** from the drop-down list on the right side of the window.
4. In Monitoring Term, do either of the following:
 - To view the transfer rate in real time, 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, select the **From** option, change the date and time in the **From** and **To** boxes, and then select **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** and **From** options, refer to section 6.5.

5. In the tree, select the **Subsystem** folder. The list on the right displays transfer rates for the ports on the subsystem.

Tips:

- If you select a port 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 (Figure 6.27).

Note: You cannot view information about host bus adapters if the host group is not registered in LUN Manager.

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 ports or host bus adapters (WWNs).
 - b) Select 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), 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 Subsystem folder is selected):

	Port	Current(MB/s)	Ave.(MB/s)(Total)	Peak(MB/s)	Response Time(ms)	Attribute
	All Non-Prio.	-	1.498	1.998	1.998	Non-Prio.
	All Prio.	-	1.499	1.999	1.999	Prio.
	CL1-A(E8)	-	1.500	2.000	2.000	Prio.
	CL1-B(E4)	-	1.549	2.049	2.050	Non-Prio.
	CL1-X(E2)	-	1.598	2.098	2.100	Non-Prio.
	CL2-a(E1)	-	1.646	2.146	2.150	Prio.

All Prio. indicates all the prioritized ports.

All Non-Prio. indicates all the non-prioritized ports.

Notes:

- iSCSI ports are always *prioritized ports*. Therefore, the **Attribute** column displays **Prio.** for each iSCSI port.
- When calculating transfer rates to be displayed in the **All Prio.** row, Performance Monitor uses transfer rates of iSCSI ports and other prioritized ports.

Transfer rate for host bus adapters connected to a specified port (when a port 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 6.27 Transfer Rates for Ports

The list displays the following:

-  : When the green check mark  icon is displayed to the left of an item, the graph displays changes in workload statistics about the item.
- **Port:** Indicates ports on the disk subsystem.
- **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:** 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 are displayed. If you select a host group icon in the tree, nicknames defined by LUN Manager are displayed. We recommend that you specify the same name for an SPM name and a nickname for convenience of host bus adapter management.

6.11.3 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. Check the **Detail** check box. The graph contents change, as described in Figure 6.10.

6.12 Monitoring LU Paths

Performance Monitor monitors LU paths and measures I/O rates (the number of I/Os per second) and transfer rates (the volume of data transferred per second). This section describes how to view I/O rates and transfer rates of LU paths on the disk subsystem. Before taking the following steps, you need to start monitoring in accordance with the procedure described in section 6.8 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.

6.12.1 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 Monitor** window is displayed.
 2. Select the **Port-LUN** tab (Figure 6.8). The tree displays a list of ports on the disk subsystem.
 3. Select IOPS from the drop-down list on the right side of the window.
 4. In **Monitoring Term**, do either 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, select the **From** option, change the date and time in the **From** and **To** boxes, and then select **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** and **From** options, refer to section 6.5.
5. In the tree, double-click a port and then select a host group. The LUN icon () is displayed.
 6. Select LUN. The list on the right displays a list of LU paths and I/O rates (Figure 6.28).
 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 LUNs.
 - b) 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), 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 6.2.

Note: If more than one row is selected in the list, you cannot select the Detail check box.

	LUN	CU:LDEV	Emulation	Paths	Current(IOPS)	Ave.(IOPS)(Total)	Max.(IOPS)	Response Time(ms)
	0000	00:00	OPEN-3 * 1	2	-	15	20	2.000
	0001	00:00	OPEN-3 * 1	2	-	65	70	2.050
	0002	00:01	OPEN-3 * 1	2	-	115	120	2.100

Figure 6.28 I/O Rates for LU Paths

The list displays the following:

-  : When the green check mark  icon is displayed to the left of an item, the graph displays changes in workload statistics about the item.
- **LUN:** Indicates LUNs (logical unit numbers).
- **CU:LDEV:** Indicates IDs of logical volumes. The number of the left of the colon is the CU image number. The number of the right of the colon is the LDEV number. The symbol "#" at the end of numbers indicates the logical volume is an external volume. The symbol "V" at the end of numbers indicates the logical volume is a virtual volume.
- **Emulation:** Indicates emulation types.
- **Paths:** Indicates the number of LU paths (i.e. paths to logical 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.

6.12.2 Viewing LU Paths Transfer Rates

Performance Monitor monitors LU paths and measures transfer rates (that is, the volume of data transferred in one second).

To view transfer rates:

1. Ensure that the **Performance Monitor** window is displayed.
2. Select the **Port-LUN** tab (Figure 6.8). The tree displays a list of ports on the disk subsystem.

3. Select MB/s from the drop-down list on the right side of the window.
4. In Monitoring Term, do either 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, select the **From** option, change the date and time in the **From** and **To** boxes, and then select **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** and **From** options, refer to section 6.5.

5. In the tree, double-click a port and then select a host group (🖥️). The LUN icon (📊) is displayed.
6. Select **LUN**. The list on the right displays a list of LU paths and transfer rates (Figure 6.29).
7. 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) Select 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), 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. 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 6.2.

Note: If more than one row is selected in the list, you cannot select the Detail check box.

	LUN	CU:LDEV	Emulation	Paths	Current(MB/s)	Ave.(MB/s)(Total)	Max.(MB/s)	Response Time(ms)
☐	0000	00:00	OPEN-3 * 1	2	-	1.500	2.000	2.000
☐	0001	00:00	OPEN-3 * 1	2	-	1.549	2.049	2.050
☐	0002	00:01	OPEN-3 * 1	2	-	1.598	2.098	2.100

Figure 6.29 Transfer Rates for LU Paths

The list displays the following:

-  : When the green check mark  icon is displayed to the left of an item, the graph displays changes in workload statistics about the item.
- **LUN:** Indicates LUNs (logical unit numbers).
- **CU:LDEV:** Indicates IDs of logical volumes. The number of the left of the colon is the CU image number. The number of the right of the colon is the LDEV number. The symbol "#" at the end of numbers indicates the logical volume is an external volume. The letter "V" with a number indicates the logical volume is a virtual volume.
- **Emulation:** Indicates emulation types.

- **Paths:** Indicates the number of LU paths (i.e. paths to logical 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 path when I/O accesses are made from the host to the LU path. 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).

6.13 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 disk subsystem. HBAs are contained in host servers and serve as ports for connecting the host servers to the disk subsystem. 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 perform following operations:

1. Perform a Server Priority Manager operation (see sections 8.4.1 and 8.4.2).
2. Start the monitoring by Performance Monitor (refer to section 6.8) to obtain the traffic data between host bus HBAs and ports.

6.13.1 Viewing I/O Rates between HBAs

Performance Monitor monitors traffic between HBAs in the hosts and ports on the disk subsystem, 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 Monitor** window is displayed.
2. Select the **WWN** tab (Figure 6.11). The tree displays a list of SPM groups . Below that list, an item named **Not Grouped** is displayed.
 - If you double-click an SPM group, host bus adapters  in the SPM group are displayed.
 - If you double-click **Not Grouped**, host bus adapters  that do not belong to any SPM group are displayed.
3. Select IOPS from the drop-down list on the right side of the window.
4. In Monitoring Term, do either of the following:
 - To view the I/O rate in real time, 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, select the **From** option, change the date and time in the **From** and **To** boxes, and then select **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** and **From** options, refer to section 6.5.

In the tree, do either 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 (Figure 6.30).
 - 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.
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 one or more SPM groups, WWNs, or ports.
 - b) Select the **Draw** button.

Notes:

- If the graph does not display changes in the I/O rate (for example, if the line in the graph runs vertically), 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 displays 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. If many-to-many connections are established between HBAs and ports, make sure that all the traffic between HBAs and ports is monitored. See section 8.4.1 for how to monitor the traffic between the HBA and the all the connected ports.

I/O rate at host bus adapters (SPM group or Not Grouped 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 (subsystem folder 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 (host bus adapter 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 6.30 I/O Rates for Host Bus Adapters (WWN Tab)

The list displays the following:

-  : When the green check mark  icon is displayed to 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 the 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 adapter in the **Server Priority Manager** windows.
- **Port:** Indicates ports on the disk subsystem.
- **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).

6.13.2 Viewing Transfer Rates between HBAs

Performance Monitor monitors traffic between HBAs in the hosts and ports on the disk subsystem, and measures transfer rates (that is, the size of data transferred in one second).

To view transfer rates:

1. Ensure that the **Performance Monitor** window is displayed.
2. Select the **WWN** tab (Figure 6.11). The tree displays a list of SPM groups. Below SPM groups, an item named Not Grouped is displayed.
 - If you double-click an SPM group, the tree displays host bus adapters in the SPM group.
 - If you double-click Not Grouped, the tree displays host bus adapters that do not belong to any SPM group.
3. Select MB/s from the drop-down list on the right side of the window.

4. In Monitoring Term, do either of the following:
 - To view the transfer rate in real time, 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, select the **From** option, change the date and time in the **From** and **To** boxes, and then select **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** and **From** options, refer to section 6.5.

5. In the tree, do either 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 (Figure 6.31).
 - 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 (Figure 6.31).

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. 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), 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 displays 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. If many-to-many connections are established between HBAs and ports, make sure that all the traffic between HBAs and ports is monitored. See section 8.4.1 for how to monitor the traffic between the HBA and the all the connected ports.

Transfer rate at host bus adapters (SPM group or Not Grouped 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 (subsystem folder 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 (host bus adapter 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 6.31 Transfer Rates for Host Bus Adapters (WWN Tab)

The list displays the following:

-  : When the green check mark  icon is displayed to 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 adapter in the **Server Priority Manager** windows.
- **Port:** Indicates ports on the disk subsystem.
- **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).

Chapter 7 Using Volume Migration

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

Volume Migration can also be used to migrate external volumes in external volume groups. External volumes cannot be manipulated manually by system administrators, but can be manipulated automatically through Volume Migration.

This chapter discusses the windows and operations for Volume Migration in the following sections:

- Section 7.1 About the Volume Migration Window
- Section 7.2 Starting Volume Migration
- Section 7.3 Reserving Migration Destinations for Logical Volumes
- Section 7.4 Migrating Logical Volumes Manually
- Section 7.5 Automating Volume Migrations
- Section 7.6 Viewing the Migration History Log

Notes:

- From an open-system host, users can also use ShadowImage commands of CCI to perform manual volume migration. For details on CCI, please refer to *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 account 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 Appendix B.

7.1 About the Volume Migration Window

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

7.1.1 About the Manual Plan Tab in the Volume Migration Window

When you select the **Volume Migration** button in the **Performance Monitor** window, **Volume Migration** starts and the **Volume Migration** window is displayed. The **Volume Migration** window includes the **Manual Plan** tab, which lets you create and execute manual migration plans.

Operations performed with this tab are discussed in section 7.4.




Figure 7.1 Manual Plan Tab in the Volume Migration Window

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

- **Monitoring Term:** Displays the monitoring period specified in the **Performance Monitor** 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 for 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 disk subsystem. The Parity group folder contains icons of parity groups, and the External Group folder contains icons of external volume groups. When you double-click an icon for a parity group or an external volume group, a list of logical volumes in the group appears.

The graphics of parity group icons and external volume group icons indicate the RAID levels, HDD classes, and presence of volumes that can be the target of manual migration.

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.
No volumes in this fixed parity group can be used as target volumes for manual migration.
-  : Indicates a fixed parity group when it appears below the Parity group folder.
This fixed parity group contains volumes that can be used as target volumes for manual migration; such volumes are marked by the  icon.
-  : Indicates a RAID-1 parity group. The HDD class for this parity group is A.
No volumes in this parity group can be used as target volumes for manual migration.
-  : Indicates a RAID-1 parity group. The HDD class for this parity group is A.
Volumes marked by  in this parity group can be used as target volumes for manual migration.
-  : Indicates a RAID-1 parity group. The HDD class for this parity group is B.
No volumes in this parity group can be used as target volumes for manual migration.
-  : Indicates a RAID-1 parity group. The HDD class for this parity group is B.
Volumes marked by  in this parity group can be used as target volumes for manual migration.
-  : Indicates a RAID-1 parity group. The HDD class for this parity group is C.
No volumes in this parity group can be used as target volumes for manual migration.
-  : Indicates a RAID-1 parity group. The HDD class for this parity group is C.
Volumes marked by  in this parity group can be used as target volumes for manual migration.
-  : Indicates a RAID-1 parity group. The HDD class for this parity group is D.
No volumes in this parity group can be used as target volumes for manual migration.
-  : Indicates a RAID-1 parity group. The HDD class for this parity group is D.
Volumes displayed by  in this parity group can be used as target volumes for manual migration.
-  : Indicates a RAID-5 parity group. The HDD class for this parity group is A.

No volumes in this parity group can be used as target volumes for manual migration.

 : Indicates a RAID-5 parity group. The HDD class for this parity group is A. Volumes marked by  in this parity group can be used as target volumes for manual migration.

 : Indicates a RAID-5 parity group. The HDD class for this parity group is B. No volumes in this parity group can be used as target volumes for manual migration.

 : Indicates a RAID-5 parity group. The HDD class for this parity group is B. Volumes marked by  in this parity group can be used as target volumes for manual migration.

 : Indicates a RAID-5 parity group. The HDD class for this parity group is C. No volumes in this parity group can be used as target volumes for manual migration.

 : Indicates a RAID-5 parity group. The HDD class for this parity group is C. Volumes marked by  in this parity group can be used as target volumes for manual migration.

 : Indicates a RAID-5 parity group. The HDD class for this parity group is D. No volumes in this parity group can be used as target volumes for manual migration.

 : Indicates a RAID-5 parity group. The HDD class for this parity group is D. Volumes marked by  in this parity group can be used as target volumes for manual migration.

 : Indicates a RAID-6 parity group. The HDD class for this parity group is A. No volumes in this parity group can be used as target volumes for manual migration.

 : Indicates a RAID-6 parity group. The HDD class for this parity group is A. Volumes displayed by  in this parity group can be used as target volumes for manual migration.

 : Indicates a RAID-6 parity group. The HDD class for this parity group is B. No volumes in this parity group can be used as target volumes for manual migration.

 : Indicates a RAID-6 parity group. The HDD class for this parity group is B. Volumes displayed by  in this parity group can be used as target volumes for manual migration.

 : Indicates a RAID-6 parity group. The HDD class for this parity group is C. No volumes in this parity group can be used as target volumes for manual migration.

 : Indicates a RAID-6 parity group. The HDD class for this parity group is C.

Volumes displayed by  in this parity group can be used as target volumes for manual migration.

 : Indicates a RAID-6 parity group. The HDD class for this parity group is D.
No volumes in this parity group can be used as target volumes for manual migration.

 : Indicates a RAID-6 parity group. The HDD class for this parity group is D.
Volumes displayed by  in this parity group can be used as target volumes for manual migration.

 1-4(!14/15 %, 03:CLPR3)
① ② ③

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.

 1-3[1-4](112/13 %, 01:CLPR1)
④

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

Figure 7.2 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.

 : Indicates an external volume group when it appears below the External Group folder.

This external volume group contains volumes that can be used as target volumes for manual migration; such volumes are marked by the  icon.



- 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 7.3 Information about External Volume Groups


The following are logical volume icons, contained in parity groups and external volume groups that can be displayed in the tree.

Note: The logical volumes contained in an external volume group are called external volumes. In the tree of **Manual Plan** tab, logical volumes in a parity group and external volumes are expressed by the same icons.

Logical volume icons:

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

When you specify a source volume in the **Manual Plan** tab, 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 a target volume, 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 operate it.
-  : Indicates a logical volume that can be used as a migration destination.
-  : Indicates a logical volume that is being migrated, or a normal logical volume that is registered in a migration plan. A logical volume that is being migrated by a program other than Volume Migration is also indicated by this icon.



- ① The logical volume ID. The number before the colon is the CU image number.
The number after the colon is the LDEV 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 symbol (!) appears before a usage rate, the usage rate might be incorrect.
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 emulation type
- ④ The size of the logical volume
- ⑤ The number and name of the CLPR.
The number before the colon is the CLPR number. The name after the colon is the CLPR name.

Figure 7.4 Information about Logical Volumes

Before you specify a volume to be migrated (that is, source volume), the icons of fixed parity groups and external volume groups are displayed with a red pin (📌). The icons of not-fixed parity groups appear without a blue frame. (For example, a parity group whose RAID level is RAID-1 and HDD class is A, is expressed by an icon .) Below these icons, four types of logical volume icons are displayed: (red), (blue), (green), and (yellow). Among these icons, you can specify as a source volume.

When you specify a source volume, Volume Migration analyzes all the parity groups and external volume groups to search for a volume that can be a target volume. If Volume Migration finds a volume that can be a target volume, the icon of a parity group or external volume group that contains the found volume changes as follows:

- Icons of a fixed parity group or external volume group:
The color of the pin on the icon changes from red (📌) to green (📌).
- Icons of a variable (not-fixed) parity group:
The icon is framed in by a blue line. For example, an icon of parity group whose RAID level is RAID-1 and HDD class is A, changes from to .

These changed icons (such as and) contain a logical volume icon (🔗) whose graphic is changed from to . You can specify a volume with icon as a target volume.

- The **S** button specifies the volume selected in the tree as the migration source.
- The **T** button specifies the volume selected in the tree as the migration target (destination).
- **Source LDEV:** The text box lets you specify the source volume (that is, the logical 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:** The text box 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 how the usage rates have changed for the source volume and the target volume.
- The **Set** button adds a new manual migration plan to the list. The new manual migration plan consists of the logical volume in Source LDEV and the logical volume in Target (Reserved) LDEV.
- **List:** Each row in the list indicates a migration plan, which is a pair of a source volume and its corresponding target volume created in the **Manual Plan** window. A migration plan before being applied to the disk subsystem is displayed in blue, and a migration plan after being applied (or during execution) to the disk subsystem 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 an error code if an error occurs during execution of a migration plan. To view detailed information about an error, select and right-click the error code and then select **Detail** from the pop-up menu.

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 select 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 logical volume ID. The number before the colon (:) is the CU image number. The number after the colon is the device ID. To the left of the device ID, the  icon indicates that the logical volume is being migrated or is registered in a migration plan.
- **Emulation** indicates the emulation type of the logical volume.
- **Capacity** indicates the capacity of the logical 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 **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.
- **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, please refer to the *Virtual Partition Manager User's Guide*.

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

- **LDEV** indicates the logical volume ID. The number before the colon (:) is the CU image number. The number after the colon is the device ID. To the left of the device ID, an icon which shows the type of the reserved volume is displayed. 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, please refer to the *Virtual Partition Manager User's Guide*.
- **Owner** indicates the program which uses this volume as a target volume. If Volume Migration uses this volume, "USP" is displayed. If another program (such as CCI) uses this volume, "Other [XX]" is displayed. [XX] shows an ID of the program.

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 using the **Delete** button.

- 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 disk subsystem 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 select the **Apply** button.

The **Delete** button has no effect when both blue and black migration plans are selected. Confirm that only blue or black migration plans are selected before you select Delete.

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

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

- The **Refresh** button updates information in the **Manual Plan** tab. The button updates the displaying of the tree after the execution of manual migration plan
- The **Apply** button applies manual migration plans in the list to the disk subsystem to start Volume Migrations.
- The **Reset** button: If you select this button after you select the **Apply** button, execution of a manual migration plan is canceled.
- The **Close** button closes the **Volume Migration** window.

7.1.2 About the Auto Plan Tab in the Volume Migration Window

When you select the Volume Migration button in the **Performance Monitor** window, Volume Migration starts and the **Volume Migration** window is displayed. This window includes the **Auto Plan** tab, which lets you make settings for automating volume migrations.

Manual Plan **Auto Plan** **Attribute** **History**

Volume Migration

Monitoring Term : 2004/05/09 00:00 - 2004/05/10 00:00
Gathering Interval : 15 Minute(s)

Auto Migration Plans

Auto Migration : Disable Migration Status : Not planned yet Auto Refresh
Plan Creation : 2004/12/06 01:45 Next Migration : 2004/12/06 01:45 ☐ Enable ☒ Disable

%	Source LDEV						Target LDEV			
	LDEV	Emulation	Capacity	RAID	PG	HDD	LDEV	RAID	PG	HDD
10	00:00	OPEN-3	2.28 GB	S(3D+1P)	1-1	DKR1B-J047FC	02:00	S(3D+1P)	2-1	DKR1B-J047FC
-	00:01	OPEN-3	2.28 GB	S(3D+1P)	1-1	DKR1B-J047FC	02:01	S(3D+1P)	2-1	DKR1B-J047FC

Make New Plan
Delete
Delete All

Auto Plan Parameters

Auto Migration Function ☒ Enable ☐ Disable **Default**

Sampling Term
Date: ☐ None ☒ Every day ☐ Once every X days: 2 ☐ Once a week: Sunday ☐ Once a month: 1
Time: From: 22 : 15 To: 01 : 15

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

Migration Time 01 : 15

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

Set
Reset
Close

Figure 7.5 Auto Plan Tab in the Volume Migration Window

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

The **Auto Plan** tab in the **Volume Migration** window displays the following:

- **Monitoring Term:** Displays the monitoring period specified in the **Performance Monitor** 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 set to **Enable**, auto migrations take place at the time displayed at Next Migration. If this option is set to **Disable**, auto migrations do 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 not yet executed them.
- 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 can 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 notification message is displayed 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 starting from the top of the list when you perform auto migration.

- The % column indicates the progress of the auto migration. A hyphen appears if Volume Migration has not executed the auto migration plan yet.
- The Source LDEV columns display information about source volumes.
 - LDEV indicates the logical volume ID. The number before the colon (:) is the CU image number. The number after the colon is the device ID.

- Emulation indicates the emulation type of the logical volume.
 - Capacity indicates the capacity of the logical 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 **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.

- The Target LDEV columns display information about target volumes:
 - LDEV indicates the logical volume ID. The number before the colon (:) is the CU image number. The number after the colon is the device ID.
 - 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 are the following buttons:

- **Make New Plan** discards the existing auto migration plans and then creates new auto migration plans.
- **Delete** deletes the auto migration plan selected in the list.
- **Delete All** 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. You can select either **Enable** or **Disable**.
 - Sampling Term:

The Volume Migration auto migration function analyses 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 set to **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 disk subsystem.
- The **Set** button applies each parameter in Auto Plan Parameters to the disk subsystem.
- The **Reset** button restores each parameter in Auto Plan Parameters to the value before change.
- The **Close** button closes the **Volume Migration** window.

7.1.3 About the Attribute Tab in the Volume Migration Window

When you select the **Volume Migration** button in the **Performance Monitor** window, Volume Migration starts and the **Volume Migration** window is displayed. The **Volume Migration** window includes the **Attribute** tab, which lets you find out which parity group belongs to which HDD class. The **Attribute** tab also lets you view detailed information about parity groups and external volume groups.

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




Figure 7.6 Attribute Tab in the Volume Migration Window

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

- **Monitoring Term** displays the monitoring period specified in the **Performance Monitor** 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 and external 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.

- The icons of parity groups and external 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.

Note that 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 7.7 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 **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.
- **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).

If a plus sign (+) 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).

- **Total** indicates the total number of logical volumes in the parity group.
- **Reserved** indicates the number of reserved logical 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, please refer to *Virtual Partition Manager User's Guide*.

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

	LDEV	Emulation	Capacity	Ave.(%)(Total)	Max.(%)	CLPR	Owner
	02:04	OPEN-3	2.28 GB	+0	1	00:CLPR0	USP
	02:05	OPEN-3	2.28 GB	!1	!2	00:CLPR0	Other [FF]
	02:06	OPEN-3	2.28 GB	2	3	00:CLPR0	-
	02:07	OPEN-3	2.28 GB	!3	!4	00:CLPR0	-
	03:00	OPEN-V	2.28 GB	4	5	00:CLPR0	-
	03:01	OPEN-V	2.28 GB	!5	!6	00:CLPR0	USP

Figure 7.8 Information about Logical Volumes in a Parity group

- **LDEV** indicates the logical volume ID. The number before the colon (:) is the CU image number. The number after the colon is the device ID.
- The  icon indicates a normal logical volume that can be migrated but cannot be used as a migration destination. You cannot change this logical volume to a reserved volume.
- The  icon also indicates a normal logical volume that can be migrated. If you change the attribute of this volume to a reserved volume () , this logical volume can be used as a migration destination (that is, target volume).
- The  (blue) icon indicates a volume reserved as a migration destination by Volume Migration. A reserved volume cannot be migrated. You can change this volume to a normal logical volume, and then use this volume as a source volume.
- The  (green) icon 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 operate it, and you cannot change the attribute of this volume.

Note: The list does not display logical volumes that are used as a LUSE.

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

If a plus sign (+) 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 logical 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).
- **CLPR** indicates the number and name of the CLPR corresponding to the parity group which the logical volume belongs to, in the format "CLPR-number:CLPR-name". For details on CLPRs, please refer to the *Virtual Partition Manager User's Guide*.

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

If Volume Migration reserved this volume, "USP" is displayed.

If one of other programs reserved this volume, "Other[XX]" is displayed. [XX] shows an ID of the program.

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

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 7.9 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 marked with the  icon. Note that 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 logical volumes in the external volume group.
- **Reserved** indicates the number of reserved logical 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, please refer to 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., logical volumes in the external volume groups).

	ExLDEV	Emulation	Capacity	CLPR	Owner
	06:04 #	OPEN-9	2.28 GB	02:CLPR2	-
	06:05 #	OPEN-9	2.28 GB	02:CLPR2	-
	06:06 #	OPEN-9	2.28 GB	02:CLPR2	-
	06:07 #	OPEN-9	2.28 GB	02:CLPR2	Other [FF]
	06:08 #	OPEN-V	2.28 GB	02:CLPR2	USP
	06:09 #	OPEN-V	2.28 GB	02:CLPR2	USP
	06:0A #	OPEN-V	2.28 GB	02:CLPR2	Other [FF]

Figure 7.10 Information about External Volumes

- **ExLDEV** indicates the ID of the external volume. The number before the colon (:) is the CU image number. The number after the colon is the device ID.
The icons displayed to the left of IDs have the same meaning here as the icons for logical volumes in parity groups.
Note: The list does not display logical volumes that are used as a LUSE.
- **Emulation** indicates the emulation type of the logical volume.

- **Capacity** indicates the capacity of the logical 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, please refer to 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 logical volumes in parity groups.
- The **Apply** button applies the settings in the **Attribute** tab to the disk subsystem.
- The **Reset** button discards changes in the **Attribute** tab and restores the original settings.
- The **Close** button closes the **Volume Migration** window.

7.1.4 About the History Tab in the Volume Migration Window

When you select the **Volume Migration** button in the **Performance Monitor** window, Volume Migration starts and the **Volume Migration** window is displayed. The **Volume Migration** window includes the **History** tab, 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.

Volume Migration

Monitoring Term : 2005/12/09 00:00 - 2005/12/15 00:00
Gathering Interval : 15 Minute(s)

Auto Migration History

2005/12/14 14:48 The parameters for creating auto migration plans were initialized.
2005/12/14 00:07 The parameters for creating auto migration plans were initialized.
2005/12/13 23:44 The parameters for creating auto migration plans were initialized.
2005/12/13 15:40 The parameters for creating auto migration plans were initialized.
2005/12/13 15:11 The parameters for creating auto migration plans were initialized.
2005/12/12 20:38 The parameters for creating auto migration plans were initialized.
2005/12/09 20:06 Failed to create the auto migration plans. Cause: Failed to get utilization statistical data
0 auto migration plan(s) were created.
2005/12/09 20:04 The creation of auto migration plans started.
2005/09/18 10:03 Source candidate = 00:00 in Group 1-1
Failed to specify proper target volumes.
2005/09/18 10:01 Source candidate = 00:00 in Group 1-1

Migration History

Date	Time	Action	Source[Parity Gr.]	Target[Parity Gr.]	Owner
2005/12/12	02:17	Migration Started	06:00[3-1]	00:00[2-1]	Other [FF]
2005/12/12	02:32	Migration Completed	06:01[3-1]	00:01[2-1]	Other [FF]
2005/12/12	02:47	Migration Canceled by User	06:02[3-1]	00:02[2-1]	Other [FF]
2005/12/12	03:02	Migration Failed	06:03[3-1]	00:03[2-1]	Other [FF]
2005/12/12	03:17	Migration Canceled by Controller	06:04[E1-1]	00:04[4-1]	USP
2005/12/12	03:32	Migration Canceled by ShadowImage	06:05[E1-1]	00:05[4-1]	USP
2005/12/08	03:47	Migration Canceled by TC	06:06[E1-1]	00:06[4-1]	USP
2005/12/08	04:02	Migration Canceled by Data Migration	06:00[XX-XX]	00:00[XX-XX]	USP

Refresh Close

Figure 7.11 History Tab in the Volume Migration Window

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

- **Monitoring Term** displays the monitoring period specified in the **Performance Monitor** 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. For details on the displayed messages, see Table 7.1:
 - The **Erase** button lets you erase all information about a particular plan.
- **Migration History** displays logs of all the events associated with 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 of the event.
- **Time** displays the time of the event.
- **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 Volume Migration executed this migration plan, "USP" is displayed. If another program (such as CCI) executed this migration plan, "Other [XX]" is displayed. [XX] shows an ID of the program.

Erase lets you erase all of the plan 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.
- **Refresh** updates information in the **History** tab.
 - **Close** closes the **Volume Migration** window.

7.2 Starting Volume Migration

To use 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 select 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 changes to red ().

2. Ensure that the **Physical** tab is active in the **Performance Monitor** window.
3. In the drop-down list at the right of Monitoring Data, select **longrange**. The **Volume Migration** button is activated.
4. Select the **Volume Migration** button. The **Volume Migration** window is displayed.

7.3 Reserving Migration Destinations for Logical Volumes

To be able to migrate volumes, you must reserve disk spaces (actually, logical 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 logical volume is reserved.

Throughout the TagmaStore USP and NSC 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. In addition, the term "normal volume" is sometimes used to refer to a volume that is not reserved as a migration destination.

To reserve a logical volume as a migration destination:

1. Start Volume Migration and select the **Attribute** tab in the **Volume Migration** window (Figure 7.6).
2. In the tree, locate a logical 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, select a parity group (or an external volume group). The list on the right displays a list of logical volumes in that group.
3. In the list to the right of the tree, locate logical volumes preceded by the  icon. From those logical volumes, select the desired logical volume and right-click it.
Note: If you want to reserve two or more logical volumes, you can select the desired logical volumes and right-click the selection.
4. Select **Reserved LDEV** from the pop-up menu. The icon changes a blue  icon.
5. Select the **Apply** button. The settings in the window are applied to the disk subsystem. The specified logical volume is reserved as a migration destination.

Notes:

- When you click the **Apply** button one time, up to 64 logical volumes can be reserved as migration destinations. If you click the **Apply** button to reserve more than 64 logical volumes, an error occurs. If you want to reserve more than 64 logical volumes, you must follow the above procedure at least twice. For example, if you want to reserve 100 logical volumes, you need to reserve 64 or fewer logical volumes and then reserve the remaining logical volumes.
- If you right-click a logical 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.
- A logical volume preceded by the green  icon is reserved by a program other than Volume Migration and you cannot change the attribute of the logical volume by using the pop-up menu.

7.4 Migrating Logical Volumes Manually

To migrate logical volumes manually, you must create migration plans. A migration plan identifies a source volume and its corresponding target (destination) volume. You prepare a migration plan for each logical 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 logical volumes manually, stop the auto migration function first by selecting the **Auto Plan** tab 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 disk subsystem.

Caution: When performing a series of manual migration operations, remember that the configuration changes after each migration is complete.

To migrate logical volumes manually:

1. Start Volume Migration (refer to section 7.2) and select the **Manual Plan** tab in the **Volume Migration** window (Figure 7.1).
2. In the tree, double-click a parity group (or an external volume group). A list of logical volumes in that parity group (or external volume group) is displayed. Logical volumes that can be migrated are indicated by the  icon.
3. From the logical volumes with the  icon, select a logical volume you want to migrate, which is called *source volume*, and then select the **S** button on the left side of **Source LDEV**.

The selected volume is defined as the source volume, and it is displayed in the **Source LDEV** text box.

Tip: If you are using Windows®, you can drag the source volume from the tree and drop it to the **Source LDEV** text box.

When you define a source volume, some of parity group icons and external volume group icons change as follows:

- Icons of a fixed parity group or external volume group:

The color of the pin on the icon changes from red () to green ()

- Icons of a not-fixed parity group:

The icon is framed in by a blue line. For example, an icon of parity group whose RAID level is RAID-1 and HDD class is A, changes from  to .

These changed icons (such as  and ) contain a logical volume that you can specify as a target volume.

4. In the tree, find a parity group (or external volume group) to which you want to migrate the source volume and then double-click it. Select a parity group (or external volume group) from icons with a green pin () or icons framed in by a blue line (such as )

A list of logical volumes in the selected parity group (or external volume group) is displayed. Logical volumes that can be a migration target are indicated by the  icon.

5. From the logical volumes with the  icon, select a logical volume you want to define as the migration destination, which is called *target volume*, and then select the T button on the left side of **Target LDEV**.

The selected volume is defined as the target volume, and it is displayed in the **Target LDEV** text box.

Tip: If you are using Windows, you can drag the target volume from the tree and drop it to the **Target LDEV** text box.

Note: Only the volumes marked by the  (reserved volume) icon can be used as migration destinations.

If you specify a logical volume in a parity group as the source and target volumes, Volume Migration calculates an estimated usage rate of the source and target parity groups after migration. The estimated results are displayed on the right of the Source LDEV and the Target LDEV text box. The line graph indicates how the usage rate has changed for the source volume and the target volume.

If you specify an external volume (logical volume in an external volume group) as the source and target volumes, Volume Migration cannot calculate an estimated usage rate of the source and target external volume groups after migration. In this case, a hyphen (-) appears instead of the estimated usage rate.

6. Check the estimated results in **Source LDEV**, **Target LDEV** and the graph:
 - If you anticipate getting the desired disk performance, go to the next step.
 - If not, return to step 2 or 6 to change the source and target volumes.
7. Select **Set**. 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 TagmaStore USP and NSC documentation set, the terms "migration plan" and "Volume Migration pair" 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. Newly-added migration plans are displayed below existing migration plans.
9. Select the **Apply** button. The migration plans are applied to the disk subsystem and multiple migrations start concurrently. The font color of the registered migration plans changes from blue to black.

When a logical volume starts migration, the % column displays the progress of the migration. When the migration is finished, the migration plan of that volume automatically disappears from the list.

If an error occurs during migration, an error code is displayed in the **Stat** column. To view detailed information about the error, select and right-click the error code and then select **Detail** from the pop-up menu.

Cautions:

- If you delete a migration plan that is being executed, data integrity in the target volume cannot be guaranteed.
- When an error condition exists in the TagmaStore USP and NSC subsystem, 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.

Notes:

- Up to eight migration plans can be applied with each click of the **Apply** button. To create and execute more than eight migration plans, use the **Apply** button several times. While executing migration plans, you can apply additional migration plans.
- The maximum number of the migration plans that can be executed concurrently might be restricted depending on the usage of other TagmaStore USP and NSC programs, and also depending on the emulation types and sizes of the migrated volumes. For details, refer to section 3.4.1.

Related topics:

- To delete a migration plan, select a migration plan in the list and then select the **Delete** button.
- If you select a migration plan in black (a migration plan that is waiting for execution or being executed) and then select the **Delete** button, the character "D" appears in the DEL column, and the migration plan will not be deleted until you select 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.

7.5 Automating Volume Migrations

You can set Volume Migrations to take place every day at a 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 programs other than Volume Migration.

7.5.1 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 logical 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 (refer to section 7.2) and select the **Attribute** tab in the **Volume Migration** window (Figure 7.6).
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 flagged with 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 **Volume Migration** window.

Tip: If you want to enable logical 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 logical volumes in the fixed parity groups as the source or target volumes.

7.5.2 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 (see section 7.2) and select the **Attribute** tab in the **Volume Migration** window (Figure 7.6).

2. In the tree, right-click an HDD class folder and then select **Change Class Threshold** from the pop-up menu. The **Change Threshold** dialog box is displayed (Figure 7.12).
3. From the Threshold drop-down list, select the value that you want to use as the limit and then select **OK**. In the tree, the selected value is displayed at the right of the specified class.
4. Select the **Apply** button in the **Volume Migration** window. The specified limit is applied to the disk subsystem.




Figure 7.12 The Change Threshold Dialog Box

7.5.3 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 (see section 7.2) and select the **Auto Plan** tab in the **Volume Migration** window (Figure 7.5).
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.

4. In Number of Sampling Points, do either 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:

- 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.

Notes:

- Volume Migrations might impose heavy workloads on disks. For this reason, you are recommended to 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.

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 **Set**. The settings in this window are applied to the disk subsystem. Migration plans are 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 **Volume Migration** 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, select the **Delete All** button.
- To delete an auto migration plan, select the plan in the list and then select **Delete**.

7.5.4 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, refer to section 7.5.2.
 - Change parameters in Auto Plan Parameters. For details, refer to section 7.5.3.
2. In the **Auto Plan** tab of the **Volume Migration** window, 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 **Volume Migration** window.

7.6 Viewing the Migration History Log

The **History** tab (Figure 7.11) in the **Volume Migration** window displays logs of auto migration operations and manual migration operations.

To view migration logs:

1. Start Volume Migration to open the **Volume Migration** window.
2. Select the **History** tab in the **Volume Migration** window.
 - To view logs of auto migration operations, look at Auto Migration History. For messages that might appear in Auto Migration History, refer to Table 7.1.
 - To view logs of manual migration operations, look at Migration History. For messages that might appear in Migration History, refer to

- Table 7.6. In Migration History, logs of auto migration operations are displayed together with those of manual migration operations.

Note: Messages in the **History** tab may not be arranged in the order that the events to which they correspond actually take place.

Table 7.1 Messages in the Auto Migration Logs (1 of 5)

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.
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	

Table 7.2 Messages in the Auto Migration Logs (2 of 5)

B. Messages about Executing Individual Auto Migration Plans	
yyyy/mm/dd hh:min The auto migration plan (CU:LDEV->CU:LDEV) was executed normally. (a) (b) (c) The execution of the auto migration plan (CU:LDEV->CU:LDEV) was finished normally. (a) Execution end time (b) Source volume (c) Target volume	
yyyy/mm/dd hh:min The auto migration plan (CU:LDEV->CU:LDEV) was canceled because of timeout. (a) (b) (c) The auto migration plan (CU:LDEV->CU:LDEV) was canceled because the duration of the migration exceeded the time specified in Max. Migration duration of the Auto Plan tab. (a) Canceled time (b) Source volume (c) Target volume	

B. Messages about Executing Individual Auto Migration Plans	
yyyy/mm/dd hh:min The auto migration plan (CU:LDEV (X-X)->CU:LDEV (X-X)) was canceled because of invalid parity group information. (a) (b) (c) (d) (e) The auto migration plan (CU:LDEV (X-X)->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 (CU:LDEV->CU:LDEV) was stopped. Cause: XXXXXXXXXXXX (a) (b) (c) (d) The execution of the auto migration plan (CU:LDEV->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:	
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 logical 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 tab.
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, please refer to the Storage Navigator™ Error Codes manual.
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.

Table 7.3 Messages in the Auto Migration Logs (3 of 5)

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 = CU:LDEV in Group X-X, Target = CU:LDEV in Group X-X (a) (b) (c) (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	
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: XXXXXXXXXX (a) (b) Error location: Class X Group X-X (c) (d) The creation of auto migration plans failed due to the cause shown in XXXXXXXXXX. 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..
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.

C. Messages about Creating Auto Migration Plans (output when the creation started or ended)	
Failed to obtain XXXXXXXXXX.	Volume Migration failed to obtain the specified information or collected statistics data which are required for creating auto migration plans. XXXXXXXXXX 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 tab. If Volume Migration failed to obtain the statistics data, check the status of the collected statistics data.
Failed to write in XXXXXXXXXX.	Volume Migration failed to write in the information of XXXXXXXXXX. XXXXXXXXXX shows the type of the information specified for migration-plan creation. For remedial measures, contact your Hitachi Data Systems service engineer.
Invalid XXXXXXXXXX.	Volume Migration could not use XXXXXXXXXX because it was invalid. XXXXXXXXXX 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 tab. 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 remedial measures, contact your Hitachi Data Systems 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.ini file does not exist. The creation of auto migration plans failed because there was not a hihsmatm.ini 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 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 logical volume CU:LDEV of the parity group X-X. Therefore, Volume Migration created auto migration plans without using this volume. (a) Parity group (b) Logical volume	

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 has already been started. (a) Volume Migration has already started the creation of auto migration plans. (a) Creation start time

Table 7.4 Messages in the Auto Migration Logs (4 of 5)

D. Messages about Creating each Auto Migration Plan (output during the creation)	
yyyy/mm/dd hh:min Source candidate = CU:LDEV in Group X-X (a) (b) (c) Target candidate = CU: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 = 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 tab, and change the value, if required.
Insufficient reserved volumes.	There are insufficient reserved volumes. Check the number of reserved volumes and their locations.

Table 7.5 Messages in the Auto Migration Logs (5 of 5)

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
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) All the parity groups in the disk subsystem 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: ▪ All the volumes in the disk subsystem are external volumes ▪ All the parity groups in the disk subsystem 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 7.6 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 (Data Migration (*2)) Migration Canceled by Controller (UR (*3)) Migration Canceled by Controller (XRC(*4)) Migration Canceled by Controller (CC(*5))	Because each program product started processing, the migration operation was canceled by Volume Migration. 1 = indicates True Copy. *2 = indicates Hitachi Online Data Migration. *3 = indicates Universal Replicator. *4 = indicates Compatible Replication for IBM XRC, which is abbreviated as XRC Replication. *5 := indicates IBM 3990 Concurrent Copy.

Chapter 8 Using Server Priority Manager

Server Priority Manager enables you tune system performance by setting threshold values for I/O and data transfer rates.

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

This chapter discusses the windows and operations for Server Priority Manager in the following sections:

- Section 8.1 About the Server Priority Manager Window
- Section 8.2 Starting Server Priority Manager
- Section 8.3 If One-to-One Connections Link HBAs and Ports
- Section 8.4 If Many-to-Many Connections Link HBAs and Ports
- Section 8.5 Grouping Host Bus Adapters

8.1 About the Server Priority Manager Window

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.

8.1.1 About the 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 section 8.3.




Figure 8.1 Port Tab of 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, 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 select the **Apply** button.

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

Note: If you change the 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 disk subsystem.
- 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 sections 6.11.1 and 6.11.2.
- 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 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 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 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 disk subsystem. 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 deletes the settings on ports in channel adapters that have been removed from SVP.

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 the previous 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 activated only when any settings for ports in a removed channel adapter remain on SVP.

- The **Apply** button applies the settings in this window to the disk subsystem.
- The **Reset** button restores the last applied settings in the window. When you select 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 disk subsystem:
 - 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** check box is checked, the check mark disappears.
- The **Close** button closes the **Server Priority Manager** window.

8.1.2 About the 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 sections 8.4 and 8.5.




Figure 8.2 WWN Tab of 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 select 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 select 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 disk subsystem:

- Ports on the disk subsystem are shown below the **Subsystem** folder. The ports are indicated by icons such as  and .

When you double-click on a port, the tree expands 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 list of host bus adapters () whose traffic with the specified port is monitored is displayed.
- If you double-click **Non-Monitor**, the list of host bus adapters whose traffic with the specified port is not monitored is displayed.

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 Names) 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.

On the right of the host bus adapter icon () below **Non-Monitor**, only the WWN is displayed.

Note: When many-to-many connections are established between host bus adapters (HBAs) and ports, ensure that all the traffic between HBAs and ports is monitored. In such cases, all the connected HBAs should be displayed below **Monitor**. For details on how to move an HBA displayed below **Non-Monitor** to below **Monitor**, see section 8.4.1.

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 information about 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 of SPM groups, see section 8.5.
 - If you double-click an SPM group, host bus adapters in that group are displayed in the tree. The WWN and SPM names are displayed to the right of the host bus adapter icon () .

Note: If the WWN of a host bus adapter displays 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, make all the traffic between HBAs and ports monitored. See section 8.4.1 for how to monitor the traffic between the HBA and all the connected 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 WWN** 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 using the drop-down list at the upper left corner of the list. The displayed items 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.

The total number of WWNs, the number of prioritized WWNs, and the number of non-prioritized WWNs are displayed on the right side of the list.

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 is displayed when a port is selected in the upper-left tree and does not appear when a 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 same average value as the 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 same maximum value as the 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.

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 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 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.

- The Upper columns will not display if an SPM group or 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 disk subsystem. 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.

- 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 deletes from SVP the settings 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 previous settings for Server Priority Manager to be applied to a different channel adapter than the one newly-installed to that location.

The **Delete ports if CHA is removed** check box is activated only when any settings on ports in a removed channel adapter remain on SVP.

- The **Apply** button applies the settings in this window to the disk subsystem.
- 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 disk subsystem:
 - 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.

8.2 Starting Server Priority Manager

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 changes to red (.
2. Ensure that the **WWN** tab or the **Port-LUN** tab is active in the Performance Monitor 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 (Figure 8.1) if one-to-one connection is established between host bus adapters and disk subsystem ports. For details on operations in the **Port** tab, see section 8.3.
 - Use the **WWN** tab (Figure 8.2) if many-to-many connection is established between host bus adapters and disk subsystem ports. For details on operations in the **WWN** tab, see section 8.4.

8.3 If One-to-One Connections Link HBAs and Ports

If one-to-one connections are established between host bus adapters (HBAs) and disk subsystem ports, use the **Port** tab in the **Server Priority Manager** window to do the following:

- Set priority to ports on the disk subsystem.
- Measure traffic between host bus adapters and disk subsystem ports (see section 8.3.1).
- Analyze traffic statistics (see section 8.3.3).
- Set an upper limit to traffic at each non-prioritized port (see section 8.3.4).
- Set a threshold to the disk subsystem or to each prioritized port, if necessary (see section 8.3.5).

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 disk subsystem.

For the system configuration of one-to-one connections between host bus adapters and ports, refer to section 4.3.1.

8.3.1 Setting Priority for Ports on the Disk Subsystem

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 disk subsystem:

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**.

6. Click the **Apply** button.

The settings on the window are applied to the disk subsystem.

After priority has been set, you can implement the procedure for measuring traffic (I/O rates and transfer rates).

8.3.2 Measuring Traffic between HBAs and Ports

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 measure traffic between HBAs and ports:

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. 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.
6. Select the **Apply** button. The settings on the window are applied to the disk subsystem.
7. Follow the instructions in section 6.8 to measure traffic (i.e., I/O rates and transfer rates).

8.3.3 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 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. Note that you can also use the **Performance Monitor** window to analyze traffic statistics. Performance Monitor can display a line graph that indicates changes in traffic.

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 either 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 and threshold values, refer to section 4.3.1.

8.3.4 Setting Upper-Limit Values for Traffic at Non-Prioritized Ports

After you analyze traffic statistics about prioritized ports and non-prioritized ports, 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 disk subsystem 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 either of the following:
 - If you want 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.
 - If you want 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 either 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 disk subsystem. The upper-limit value that you entered turns black.

8.3.5 Setting a Threshold

If threshold control is used, upper-limit control is automatically disabled when traffic between production servers and the disk subsystem is reduced to a specified level. For details, refer to sections 4.2 and 4.3.1.

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, or
- Set one threshold value for the entire disk subsystem.

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. (If you want to set one threshold value for the entire disk subsystem, this step is unnecessary). 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**.
4. Do either 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.
 - To set one threshold to the entire disk subsystem, select the **All Thresholds** check box. Next, select **IOPS** or **MB/s** from the drop-down list 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.

Note: If you enter 0 (zero) in a cell, 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 disk subsystem and the upper limits of the non-prioritized ports are disabled.

5. Select the **Apply** button. The settings in the window are applied to the disk subsystem.

8.4 If Many-to-Many Connections Link HBAs and Ports

If many-to-many connections are established between host bus adapters (HBAs) and disk subsystem ports, you use the **WWN** tab in the **Server Priority Manager** window to do the following:

- Make all the traffic between host bus adapters and ports monitored.
- Set priority to host bus adapters.
- Measure traffic between host bus adapters and disk subsystem ports (see section 6.2).
- Analyze traffic statistics.
- 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 host bus adapters. If necessary, you can set one threshold value applied for the entire disk subsystem. When many-to-many connections are established between host bus adapters and ports, you cannot set individual thresholds for prioritized host bus adapters.

For the system configuration of many-to-many connections between host bus adapters and ports, refer to section 4.3.2.

8.4.1 Monitoring All Traffic between HBAs and Ports

When many-to-many connections are established between host bus adapters (HBAs) and ports, make all the traffics between HBAs and ports monitored, basically. **Server Priority Manager** can perform upper limit control and threshold control only for the monitored traffics.

To make all the traffics between host bus adapters and ports monitored:

1. Start **Server Priority Manager** (refer to section 8.2).
The **Server Priority Manager** window is displayed.
2. Ensure that the **WWN** tab (Figure 8.2) is displayed.
The two trees are displayed in the left side of the **WWN** tab. The upper-left tree lists ports in the disk subsystem.
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 traffic with the specified port is 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** dialog box (Figure 8.3) is displayed. This window lets you add a WWN of a host bus adapter to **Monitor**.

7. In the **Add WWN** dialog box, 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 those 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 using the same aliases, the management of 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 steps 6 to 8 to move all the host bus adapters displayed below **Non-Monitor** to below **Monitor**.

10. Click the **Apply** button in the **Server Priority Manager** window.

The settings on the window are applied to the disk subsystem.

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** dialog box in which you can specify the SPM name only is displayed.




Figure 8.3 Add WWN Dialog Box

If you add a port or host bus adapter to the disk subsystem 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 (Figure 8.2) 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**.

Notes:

- 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 section 8.5.2.
6. Select the **OK** button for the confirmation message that asks you whether to exclude the host bus adapter from the monitoring target.

The selected 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 disk subsystem.

8.4.2 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 traffic.

Prioritized WWNs are WWNs on which the processing has high priority and *non-prioritized WWNs* are WWNs on which the processing has low priority. Specify a host bus adapter existed in a server with high priority as a prioritized WWNs. Specify a host bus adapter existed in a server with low priority as a non-prioritized WWNs.

To set priority for 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 appears 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 are displayed below **Monitor**.
If some of the WWNs are missing, use the procedure in section 8.4.1 to move all WWNs to below **Monitor**.
7. Click **Monitor** to display information about the monitored host bus adapters in the list on the right side 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. If you want to specify more than one prioritized WWN, repeat this operation. The Attribute column displays **Prio**.
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 section 8.5.3.
9. Right-click a non-prioritized WWN in the list and then select **Prio ->> Non-Prio** from the pop-up menu. If there is more than one non-prioritized WWN, repeat this operation. The Attribute column displays **Non-Prio**.
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 section 8.5.3.
10. Repeat steps 4 to 9 for ports (except for the port selected in step 4). 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 disk subsystem.
12. Follow the instructions in section 6.8 to measure traffic (i.e., I/O and transfer rates).

8.4.3 Analyzing Traffic Statistics

Traffic statistics reveal the number of I/Os that have been made via ports from HBAs. They also reveal the amount of data that has been transferred between ports and HBAs. You must analyze 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. Note that you can also use the **Performance Monitor** window to analyze traffic statistics. Performance Monitor can display a line graph that indicates changes in traffic (For details, refer to sections 6.10.1 and 6.10.2).

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 either 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., I/O rates and transfer rates).

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. Each traffic type contains 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:

- Only the traffic statistics about the host bus adapters below *Monitor* are displayed in the list.
 - The **WWN Total** traffic statistics are also displayed in the list when you select an icon in the lower-left tree. If you select 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 section 8.5.
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, refer to section 4.3.2.

8.4.4 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 disk subsystem and the low priority ports and thus provide 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.

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. Select the **WWN** tab.
3. Do either of the following:
 - If you want 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.
 - If you want 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 either of the following:
 - If you want 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.
 - If you want 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.

The upper-limit value that you entered is displayed in blue.

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, refer to section 8.4.4.
- 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 disk subsystem. The upper-limit value that you entered turns black.

8.4.5 Setting a Threshold

If threshold control is used, upper-limit control is automatically disabled when traffic between production servers and the disk subsystem is reduced to a specified level. For details, refer to sections 4.2 and 4.3.2.

If many-to-many connections are established between host bus adapters and disk subsystem ports, you can set one threshold value for the entire disk subsystem. 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 below the **All Thresholds** check box, and do either of the following:
 - If you want to disable the upper limit on I/O rates, select **IOPS** from the drop-down list at the upper left corner of the list.
 - If you want to disable the upper limit on transfer rates, select **MB/s** from the drop-down list at the upper left corner of the list.

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.
6. Select the **Apply** button. The settings in the window are applied to the disk subsystem.

8.4.6 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 can identify HBAs more easily if you assign SPM names. WWNs are 16-digit hexadecimal numbers and cannot be changed. However, SPM names can be changed, and should not necessarily be 16-digit hexadecimal numbers.

The following is the procedure for changing an already-assigned SPM name.

To change an SPM name:

1. Start Server Priority Manager. The **Server Priority Manager** window is displayed.
2. Select the **WWN** tab.
3. In the upper-left tree, select a host bus adapter from below **Monitor** and then right-click the selection.

4. From the pop-up menu, select **Change SPM Name**. The **Change WWN and SPM Name** dialog box is displayed.
5. Enter a new SPM name in the SPM Name box and then select **OK**.
6. In the Server Priority Manager window, select the **Apply** button. The settings in the window are applied to the disk subsystem.




Figure 8.4 Change WWN and SPM Name Dialog Box

8.4.7 Replacing a Host Bus Adapter

If a host bus adapter fails, you need to replace the adapter with a new one. After you finish replacement, you 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 (🖨️) below Monitor, and then right-click the selection.
4. From the pop-up menu, select **Change SPM Name**. The **Change WWN and SPM Name** dialog box (Figure 8.4) 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, using up to 16 characters.

7. Select **OK** to close the **Change WWN and SPM Name** dialog box.
8. In the **Server Priority Manager** window, select the **Apply** button. The settings in the window are applied to the disk subsystem.

8.5 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 the 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.

8.5.1 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** dialog box (Figure 8.5), enter the name of the SPM group and then select **OK**. An SPM group is created, and an SPM group icon () is added to the lower-left tree.
6. Select an HBA from the upper-left tree and select a SPM group from the lower-left tree.
7. 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, 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.

8. Select the **Apply** button. The settings in the window are applied to the disk subsystem.




Figure 8.5 Add New SPM Group Dialog Box

8.5.2 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 icon () that contains the HBA to be deleted.
4. Below the SPM icon, right-click on the icon for the HBA () you want to delete.
5. Select **Delete WWN** from the pop-up menu. The selected HBA icon is deleted from the tree.
6. Select the **Apply** button. Settings on the window are applied to the disk subsystem.

8.5.3 Switching the 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 a 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 either 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 disk subsystem.

8.5.4 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 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.
 - c) To assign an upper limit to the transfer rate, enter the upper-limit value in the MB/s column of the list.

Tip: 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:
- d) Right-click the selected SPM group and then select **Change Upper Limit** on the pop-up menu. The **Change Upper Limit** dialog box (Figure 8.6) is displayed.
 - e) 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**.
 - f) 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 disk subsystem.




Figure 8.6 Change Upper Limit Dialog Box

Note: To confirm an upper limit value specified for each SPM group, select the Subsystem folder in the lower-left tree of the **WWN** window. The SPM groups are displayed in the list and you can confirm each upper-limit value.

8.5.5 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** dialog box 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 disk subsystem.




Figure 8.7 Rename SPM Group Dialog Box

8.5.6 Deleting an SPM Group

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 disk subsystem.

Chapter 9 Troubleshooting

This chapter gives troubleshooting information on Performance Monitor, Volume Migration and Server Priority Manager. For troubleshooting information on Storage Navigator, please refer to its *User's Guide* and *Error Codes* documents.

9.1 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, make sure that all the traffics between HBAs and ports monitored. For details on the measures, refer to section 8.4.1.
- When a part of monitoring data is missing:
While displaying Performance Monitor in short range, if I/O workloads between hosts and the disk subsystem become heavy, the disk subsystem gives higher priority to I/O processing than to monitoring processing, therefore, some monitoring data might be missing. If you find that monitoring data is frequently missing, use the **Gathering Interval** option in the **Monitoring Options** window to make the collection interval longer.

Note that the display of the **LDEV**, **Port-LUN**, and **WWN** tabs is fixed in short range.

9.2 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, please refer to the *Error Codes Guide* for Storage Navigator.

9.3 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 monitored. For details on the measures, refer to section 8.4.1.

9.4 Calling the Hitachi Data Systems Technical Support Center

If you need to call the Hitachi Data Systems Technical Support Center, be sure to provide as much information about the problem as possible. Include the circumstances surrounding the error or failure and the exact content of any error codes and/or messages displayed on the Storage Navigator. The worldwide Hitachi Data Systems Technical Support Centers are:

- Hitachi Data Systems North America/Latin America
San Diego, California, USA
1-800-446-0744
- Hitachi Data Systems Europe
Contact Hitachi Data Systems Local Support
- Hitachi Data Systems Asia Pacific
North Ryde, Australia
61-2-9325-3300
- Hitachi, Ltd., Disk Array Systems Division
322-2 Nakazato, Odawara-shi, Kanagawa-ken, 250-0872 Japan
+81-465-48-1111

Appendix A Exporting Monitoring Data to Files

The Export Tool enables you to export monitoring data (i.e., statistics) that can display in the Performance Monitor window to text files. The Export Tool also enables you to export monitoring data on remote copy operations performed by TrueCopy, TrueCopy for z/OS, Universal Replicator, and Universal Replicator for 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.

Note: You cannot use the Export Tool to export monitoring data on remote copy operations performed by Universal Replicator and Universal Replicator for z/OS, displayed in the **UR Monitor** and **URz Monitor** windows.

Figure A.1 is an example of a text file imported into spreadsheet software:

Transfer rates for LUs (Filename: LU_IOPS.csv)

LU_IOPS.csv					
Serial number : 28528(RAID500)		Indicates that the subsystem serial number is 28528. The word in parentheses is a code that indicates the subsystem type.			
From : 2004/11/25 10:48		Indicates that the data was obtained from 10:48 to 10:52 on Nov. 25, 2004. Sampling rate: 1 indicates that monitoring was performed every minute (at a one-minute interval).			
To : 2004/11/25 10:52					
sampling rate : 1					
No.	time	CL1-B.00(1B-G00).0000	CL1-B.00(1B-G00).0001	CL1-B.01(WinNT).0000	CL1-B.01(WinNT).0001
1	2004/11/25 10:49	10	20	30	40
2	2004/11/25 10:50	12	22	30	35
3	2004/11/25 10:51	12	21	29	35
4	2004/11/25 10:52	10	19	32	38
		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 **0000** and **0001**) indicate an LUN. For example, the heading **CL1-B.00(1B-G00).0001** indicates the port **CL1-B**, the host group ID **00**, the host group name **1B-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 A.1 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 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.

A.1 Files to be Exported

The Export Tool enables you to save monitoring data in the **Performance Monitor** window into files (see Tables A.2 through A.9). The Export Tool also enables you to save monitoring data about remote copy operations displayed in the **TC Monitor** and **TCz Monitor** windows (see Table A.10). 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 A.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, see the tables indicated in the "See" column.

Table A.1 Performance Monitor Windows and Monitoring Data Saved by Export Tool

Window	Monitoring Data	See
Physical tab in the Performance Monitor window	Statistics about resource usage and write pending rates	Table A.2
LDEV tab in the Performance Monitor window	Statistics about parity groups, external volume groups, or V-VOL groups	Table A.3
	Statistics about logical volumes in parity groups, in external volume groups, or V-VOL groups	Table A.4
Port-LUN tab in the Performance Monitor window	Statistics about ports	Table A.5
	Statistics about host bus adapters connected to ports	Table A.6
	Statistics about LUs	Table A.7
WWN tab in the Performance Monitor window	Statistics about SPM groups	Table A.8
	Statistics about host bus adapters belonging to SPM groups	Table A.9
TC Monitor window and TCz Monitor window	Statistics about remote copy operations by TrueCopy and TrueCopy for z/OS (in the whole disk subsystem)	Table A.10
	Statistics about remote copy operations by TrueCopy and TrueCopy for z/OS (at LUs)	Table A.11
	Statistics about remote copy operations by TrueCopy and TrueCopy for z/OS (at logical volumes controlled by a particular CU)	Table A.12
	Statistics about remote copy operations by TrueCopy and TrueCopy for z/OS (At CLPR)	Table A.13
UR Monitor window and URz Monitor window	Statistics about remote copy operations by Universal Replicator and Universal Replicator for z/OS (in the whole disk subsystem)	Table A.14
	Statistics about remote copy operations by Universal Replicator and Universal Replicator for z/OS (at journal groups)	Table A.15
	Statistics about remote copy operations by Universal Replicator and Universal Replicator for z/OS (at LUs)	* 1 millisecond is one-thousandth of 1 second. Table A.16

Window	Monitoring Data	See
UR Monitor window and URz Monitor window	Statistics about remote copy operations by Universal Replicator and Universal Replicator for z/OS (at logical volumes controlled by a particular CU)	Table A.17

The tables below describe ZIP files and CSV files to be saved by the Export Tool.

Table A.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 logical volumes in a parity group in long range
	PHY_Short_LDEV_x-y.csv	Usage rates for logical volumes in a parity group in short range
PhyExG_dat.ZIP	PHY_ExG_Response.csv	Average response time for external volume groups (in milliseconds*)
	PHY_ExG_Trans.csv	Amount of transferred data for external volume groups (in KB/sec)
PhyExLDEV_dat.ZIP	PHY_ExLDEV_Response_x-y.csv	Average response time for logical volumes in an external volume group (in milliseconds*)
	PHY_ExLDEV_Trans_x-y.csv	Amount of data transferred for logical volumes in an external volume group (in 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

ZIP File	CSV File	Data Saved in the File
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
	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

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 A.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 (in 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
	PG_CFW_Write_IOPS.csv	The number of write operations in "cache-fast-write" mode per second

ZIP File	CSV File	Data Saved in the File
	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 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 disk drives (i.e., parity groups, external volume groups, or V-VOL groups)
	PG_D2CS_Trans.csv	The number of data transfer operations from 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 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 (in microseconds)

Table A.4 Files with Statistics about Logical Volumes in Parity/External Volume Groups, or 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 (in 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
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

ZIP File	CSV File	Data Saved in the File
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 disk drives (i.e., logical volumes) per second
LDEV_C2D_Trans.ZIP	LDEV_C2D_Trans_x-y.csv	The number of data transfer operations from cache memories and disk drives (i.e., logical volumes)
LDEV_D2CS_Trans.ZIP	LDEV_D2CS_Trans_x-y.csv	The number of data transfer operations from disk drives (i.e., logical 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 disk drives (i.e., logical volumes) to cache memories in random access mode
LDEV_Response.ZIP	LDEV_Response_x-y.csv	The average response time at logical volumes (in microseconds)

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 logical volume in the parity group 1-2.

Table A.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
	Port_KBPS.csv	The size of data transferred per second (in KB/sec)
	Port_Response.csv	The average response time at ports (in microseconds)

Table A.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 (in KB/sec)
	PortWWN_xx_Response.csv	The average response time between a port and the HBAs connected to that port (in microseconds)

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 A.7 Files with Statistics about 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 (in 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 disk drives (i.e., LUs).
	LU_D2CS_Trans.csv	The number of data transfer operations from disk drives (i.e., LUs) to cache memories in sequential access mode.
	LU_D2CR_Trans.csv	The number of data transfer operations from disk drives (i.e., LUs) to cache memories in random access mode.
	LU_Response.csv	The average response time at LUs (in microseconds)

Table A.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 (in KB/sec)
	PPCG_Response.csv	The average response time at SPM groups (in microseconds)

Table A.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 for HBAs belonging to an SPM group
	PPCGWWN_xx_KBPS.csv	The transfer rate for HBAs belonging to an SPM group (in KB/sec)
	PPCGWWN_xx_Response.csv	The average response time for HBAs belonging to an SPM group (in microseconds)
	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 (in KB/sec)
	PPCGWWN_NotGrouped_Response.csv	The average response time for HBAs that do not belong to any SPM group (in microseconds)

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 A.10 Files with Statistics about RC Operations by TC/TCz (Disk Subsystem)

ZIP File	CSV File	Data Saved in the File
RemoteCopy_dat.ZIP	RemoteCopy.csv	The following data in the whole disk subsystem are saved: ▪ The usage rate for sidefile cache ▪ The total number of remote I/Os (read and write operations) ▪ The total number of remote read I/Os ▪ 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 number of hits of initial copy remote I/Os ▪ The average transfer rate (kB/sec) for initial copy remote I/Os ▪ The average response time (in milliseconds) for initial copy ▪ The number of migration copy remote I/Os ▪ The number of hits of migration copy remote I/Os ▪ The number of update copy remote I/Os ▪ The number of hits of update copy remote I/Os ▪ The average transfer rate (kB/sec) for update copy remote I/Os ▪ The average response time (in milliseconds) 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

ZIP File	CSV File	Data Saved in the File
		I/Os - The average response time (in milliseconds) 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

Table A.11 Files with Statistics about RC Operations by TC/TCz (LUs)

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 (in milliseconds) for initial copy at LUs
	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 (in milliseconds) for update copy at LUs
	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 (in milliseconds) for asynchronous update copy at LUs
	RCLU_Scheduling_Recordset.csv	The number of scheduled recordsets

ZIP File	CSV File	Data Saved in the File
	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)
	RCLU_Out_of_Tracks.csv	The number of tracks that have not yet been copied by the initial copy or resync copy operation

Table A.12 Files with Statistics about RC Operations by TC/TCz (LUs controlled by 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 (in milliseconds) for initial copy at logical 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 (in milliseconds) for update copy at logical volumes
RCLDEV_Restore_Copy_RIO.ZIP	RCLDEV_Restore_Copy_RIO_xx.csv	The number of restore copy remote I/Os

ZIP File	CSV File	Data Saved in the File
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
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 (in milliseconds) for asynchronous update copy at logical 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

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 logical volumes controlled by the CU whose image number is 10.

Table A.13 Files Containing Statistics about Remote Copy Operations by TrueCopy and TrueCopy for z/OS (At CLPR)

ZIP File	CSV File	Data Saved in the File
RCCLPR_dat.ZIP	RCCLPR_SideFile.csv	Sidefile usage rate per CLPR (as a percentage)

Table A.14 Files with Statistics about RC Operations by UR/URz (Disk Subsystem)

ZIP File	CSV File	Data Saved in the File
UniversalReplicator.ZIP	UniversalReplicator.csv	The following data in the whole disk subsystem are saved: - The number of read I/Os per second - The number of read hit records per second - The number of write I/Os per second

ZIP File	CSV File	Data Saved in the File
		▪ The number of write hit records per second ▪ The amount of data that are read per second (in KB/sec) ▪ The amount of data that are written per second (in KB/sec) ▪ The initial copy hit rate (as a percentage) ▪ The average transfer rate for initial copy operations (in KB/sec) ▪ The number of asynchronous remote I/Os per second at the primary subsystem ▪ The number of journals at the primary subsystem ▪ The average transfer rate for journals in the primary subsystem (in KB/sec) ▪ The remote I/O average response time on the primary subsystem (in milliseconds) ▪ The number of asynchronous remote I/Os per second at the secondary subsystem ▪ The number of journals at the secondary subsystem ▪ The average transfer rate for journals in the secondary subsystem (in KB/sec) ▪ The remote I/O average response time on the secondary subsystem (in milliseconds)

Table A.15 Files with Statistics about RC Operations by UR/URz (Journal Groups)

ZIP File	CSV File	Data Saved in the File
URJNL_dat.ZIP	URJNL_Read_Record.csv	The number of read I/Os per second
	URJNL_Read_Hit.csv	The number of read hit records per second
	URJNL_Write_Record.csv	The number of write I/Os per second
	URJNL_Write_Hit.csv	The number of write hit records per second
	URJNL_Read_Transfer.csv	The amount of data that are read per second (in KB/sec)
	URJNL_Write_Transfer.csv	The amount of data that are written per second (in KB/sec)
	URJNL_Initial_Copy_Hit.csv	The initial copy hit rate (as a percentage)
	URJNL_Initial_Copy_Transfer.csv	The average transfer rate for initial copy operations (in KB/sec)
	URJNL_M-JNL_Asynchronous_RIO.csv	The number of asynchronous remote I/Os per second at the primary subsystem
	URJNL_M-JNL_Asynchronous_Journal.csv	The number of journals at the primary subsystem
	URJNL_M-JNL_Asynchronous_Copy_Transfer.csv	The average transfer rate for journals in the primary subsystem (in KB/sec)
	URJNL_M-JNL_Asynchronous_Copy_Response.csv	The remote I/O average response time on the primary subsystem (in milliseconds)
	URJNL_R-JNL_Asynchronous_RIO.csv	The number of asynchronous remote I/Os per second at the secondary subsystem
	URJNL_R-JNL_Asynchronous_Journal.csv	The number of journals at the secondary subsystem

ZIP File	CSV File	Data Saved in the File
	URJNL_R-JNL_Asynchronous_Copy_Transfer.csv	The average transfer rate for journals in the secondary subsystem (in KB/sec)
	URJNL_R-JNL_Asynchronous_Copy_Response.csv	The remote I/O average response time on the secondary subsystem (in milliseconds)
	URJNL_M-JNL_Data_Used_Rate.csv	Data usage rate for master journals (as a percentage)
	URJNL_M-JNL_Meta_Data_Used_Rate.csv	Metadata usage rate for master journals (as a percentage)
	URJNL_R-JNL_Data_Used_Rate.csv	Data usage rate for restore journals (as a percentage)
	URJNL_R-JNL_Meta_Data_Used_Rate.csv	Metadata usage rate for restore journals (as a percentage)

* 1 millisecond is one-thousandth of 1 second.

Table A.16 Files with Statistics about RC Operations by UR/URz (LUs)

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 (in KB/sec)
	URLU_Write_Transfer.csv	The amount of data that are written per second (in KB/sec)
	URLU_Initial_Copy_Hit.csv	The initial copy hit rate (as a percentage)
	URLU_Initial_Copy_Transfer.csv	The average transfer rate for initial copy operations (in KB/sec)

Table A.17 Files with Statistics about RC Operations by UR/URz (LUs controlled by 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 (in KB/sec)
URLDEV_Write_Transfer.ZIP	URLDEV_Write_Transfer_xx.csv	The amount of data that are written per second (in KB/sec)

ZIP File	CSV File	Data Saved in the File
URLDEV_Initial_Copy_Hit.ZIP	URLDEV_Initial_Copy_Hit_xx.csv	The initial copy hit rate (as a percentage)
URLDEV_Initial_Copy_Transfer.ZIP	URLDEV_Initial_Copy_Transfer_xx.csv	The average transfer rate for initial copy operations (in 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 logical volumes controlled by the CU whose image number is 10.

A.2 Preparing to Use the Export Tool

This section explains how to prepare to use the Export Tool.

A.2.1 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, please refer to the *Storage Navigator User's Guide* (MK-94RD206).

- **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, please refer to the *Storage Navigator User's Guide* (MK-94RD206).

- **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.

For detailed information about how to create a user ID, see the *Storage Navigator User's Guide*. When you create the user ID, note the following:

- Permissions of TagmaStore USP and NSC 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 subsystem 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, and TrueCopy for 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 can be saved into files.

If you specify "*storage partition administrator*" for the user ID, monitoring data on remote copy operations cannot be saved into files, and functions of the Export Tool are limited. See Appendix B for details.

■ The Export Tool program

The Export Tool is provided with the Performance Monitor product (installation media). For detailed information about how to install the Export Tool on a Windows computer, see section A.2.2. For detailed information about how to install the Export Tool on a UNIX computer, see section A.2.3.

A.2.2 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 appropriate 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.

A.2.3 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 either of the following:
 - If you are using Solaris™, please 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®, please 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.

A.3 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 (see section A.3.1)
- Preparing a batch file (see section A.3.2)
- Running the Export Tool (see section A.3.3)

A.3.1 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 A.2 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
shortrange 200410010850:200410010910
; Specifies term of data to be exported for data stored
; in short range
longrange 200409301430:200410011430
; 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 A.2 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.
- `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.

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 section A.2.1 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: 2004/10/01 01:00 - To: 2004/10/01 15:00 Interval : 1min.
Long Range	From: 2004/09/01 00:00 - To: 2004/10/01 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, 2004. Also, the monitoring data in long range is stored every 15 minutes in the term of Sep. 1, 2004, 0:00 through Oct. 1, 2004, 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.

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 by checking the log file for information issued by the show subcommand.

- group PhyPG Long and group RemoteCopy

The group subcommand specifies the type of data that you want to export. Specify a 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 A.2 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 z/OS in short range. You can describe multiple lines of the group subcommand to export multiple monitoring items at the same time.

The group subcommand specifies the type of data that you want to export.

- `shortrange 200410010850:200410010910`
`longrange 200409301430:200410011430`

The `shortrange` and `longrange` subcommands specify the term of monitoring data that you want to export. Use these subcommands when you want to narrow the export-target term (range) within the stored data. You can specify both the `shortrange` and `longrange` subcommands at the same time. The difference between these subcommands is as follows:

- The `shortrange` 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 `longrange` 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 Monitor** window with selecting **longrange**.)
Specify a term within "Long Range From XXX To XXX" which is output by the `show` subcommand.

In Figure A.2, the script `shortrange 200410010850:200410010910` specifies the term 8:50-9:10 on Oct. 1, 2004. 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 z/OS in the term specified by the `shortrange` subcommand.

Also, in Figure A.2, the script `longrange 200409301430:200410011430` specifies the term from Sep. 30, 2004, 14:30 to Oct. 1, 2004, 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 `longrange` subcommand.

If you run the Export Tool without specifying the `shortrange` or `longrange` 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 `shortrange` subcommand, see section A.4.6.

For detailed information about the `longrange` subcommand, see section A.4.7.

- `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 section A.4.8.

- `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 section A.4.9.

- apply

The apply subcommand saves monitoring data in files.

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.

A.3.2 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 A.3 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/JSanRmiServerEx.jar" -Xmx268435456
-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/JSanRmiServerEx.jar" -Xmx268435456 -
-Dmd.command=command.txt -Dmd.logpath=log sanproject.getmondat.RJMdMain ↵
```

Figure A.3 Scripts in Batch Files (runWin.bat and runUNIX.bat)

Note: The "↵" symbol indicates the end of a command line.

If the computer that runs 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 on 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 is as shown in Figure A.4:

Batch file for Windows computers (runWin.bat):

```
java -classpath "./lib/JSanExport.jar;./lib/JSanRmiServerEx.jar" -Dhttp.proxyHost=Jupiter
-Dhttp.proxyPort=8080 -Xmx268435456 -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/JSanRmiServerEx.jar" -Dhttp.proxyHost=Jupiter
-Dhttp.proxyPort=8080 -Xmx268435456 -Dmd.command=command.txt -Dmd.logpath=log
sanproject.getmondar.RJMdMain ↵
```

Figure A.4 Scripts in Batch Files (when specifying 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/JSanRmiServerEx.jar" -Dhttp.proxyHost=158.211.122.124
-Dhttp.proxyPort=8080 -Xmx268435456 -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/JSanRmiServerEx.jar" -Dhttp.proxyHost=158.211.122.124
-Dhttp.proxyPort=8080 -Xmx268435456 -Dmd.command=command.txt -Dmd.logpath=log
sanproject.getmondar.RJMdMain ↵
```

Figure A.5 Scripts in Batch Files (when specifying IP address of a proxy host)

Note: The "↵" symbol indicates the end of a command line.

A.3.3 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 (see Figure A.7 for reference). If you are using a Windows computer, you can double-click the batch file to execute the batch file.

```
c:\WINDOWS> c:\export> runWin.bat      ← Execute runWin.bat
```

Figure A.6 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 (e.g., the command prompt). The dots increment as export processing continues (see Figure A.8). 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 A.7 Example of Command Prompt Outputs when Export Tool Runs

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, refer to section A.1.

Notes:

- When an internal error occurs during export processing, the exclamation mark (!) appears and then retries export processing. Usually, the Export Tool can retry processing up to three times. If export processing does not finish through three retries, you need to quit the command prompt and then run the Export Tool again. Note, however, that you can change the maximum number of retries by using the **retry** subcommand. For detailed information about the **retry** subcommand, refer to section A.4.2.
- If you specify the **nocompress** operand for the **option** subcommand, the Export Tool saves files in CSV format instead of ZIP format.

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. For this reason, the exporting process might take a lot of time (see Table A.17). If you want to export statistics spanning a long period of time, it is recommended that you use the Export Tool more than one time, rather than using the tool only one time to export all the statistics. For example, if you want to export statistics spanning 24 hours, you could run the Export Tool eight times; each time you run the Export Tool, the tool should export statistics spanning three hours.

Table A.18 Estimates 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 logical 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, it is strongly recommended that you delete log files regularly.
- For information about how to solve errors with the Export Tool, see section A.5.
- The Export Tool returns a termination code when the Export Tool finishes (see Table A.18).

Table A.19 Termination Codes Returned by Export Tool

Termination Code	Meaning
0	The Export Tool finished successfully.
1	An error occurred when the set subcommand (see section A.4.1) 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 section A.4.1) executed. Some other user might have been logged on in Modify mode.
4	The user ID has none of permissions for Performance Monitor, TrueCopy, and TrueCopy for 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\JSanRmiServerEx.jar" -Xmx268435456 -
Dmd.command=command.txt -Dmd.logpath=log sanproject.getmondar.RJMdMain ↵
if %errorlevel%==1 echo THE SET SUBCOMMAND FAILED ↵
if %errorlevel%==3 echo THE SET SUBCOMMAND FAILED ↵
pause ↵
```

Figure A.8 Example of a Batch File, including a Reference to Termination Code

Note: The "↵" symbol indicates the end of a command line.

A.4 Command Reference

This section explains the syntax of subcommands that you can write in your command file.

Table A.20 Subcommand List

Subcommand	Function
svpip	Specifies the IP address of the SVP to be logged in.
retry	Makes settings on retries of export processing.
login	Logs the specified user into the SVP.
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.
group	Specifies the type of data that you want to export.
shortrange	Specifies the term of monitoring data to be exported for short-range monitoring data.
longrange	Specifies the term of monitoring data to be exported for long-range monitoring data.
outpath	Specifies the directory in which files should be saved.
option	Specifies whether to save files in ZIP format or in CSV format.
apply	Saves monitoring data in files.
set	Starts or ends monitoring the disk subsystem, and also specifies the gathering interval in short range monitoring.
help	Displays the on-line help for subcommands.

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 this square brackets and are delimited by vertical bars (|), you can select one of the operands. For detailed information, refer to Table A.13.

{ } 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 (|). For detailed information, refer to Table A.13.

... Indicates that a previously used operand can be repeated. For detailed information, refer to Table A.13.

Table A.21 Syntax Descriptions

Syntax	The Syntax Indicates 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:**

Note the following when you write script in your command file:

- Ensure that only one subcommand is used in one line.
- Note that 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 A.9 Examples of Comments

■ **Viewing the Online Help for subcommands:**

You can display the on-line Help to view the syntax of subcommands when you are working at the command prompt. To be able to view the on-line Help, you must use the **help** subcommand of the Export Tool. For detailed information about how to use the **help** subcommand, see section A.4.1 next.

A.4.1 svpip Subcommand

A.4.1.1 Syntax

`svpip△{ip-address|host-name}`

A.4.1.2 Description

The **svpip** subcommand

specifies the IP address or the host name of the SVP.

A.4.1.3 Operands

ip-address

Specifies the IP address of the SVP.

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 ("").

A.4.1.4 Example

The following example specifies the IP address of the SVP as 158.214.127.170.

```
svpip 158.214.127.170
```

A.4.2 Retry Subcommand

A.4.2.1 Syntax

`retry△[time=m] △[count=n]`

A.4.2.2 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.

A.4.2.3 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.

A.4.2.4 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
shortrange 200404010850:200404010910
outpath out
option compress
apply
```

A.4.3 Login Subcommand

A.4.3.1 Syntax

`login△userid△password`

A.4.3.2 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.

A.4.3.3 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 section A.2.1.

`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 ("").

A.4.3.4 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
```

A.4.4 Show Subcommand

A.4.4.1 Syntax

show

A.4.4.2 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 1 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, refer to sections 2.6, 6.3.1, and 6.3.2.

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: 2004/10/01 01:00 - To: 2004/10/01 15:00 Interval : 1min.
Long Range	From: 2004/09/01 00:00 - To: 2004/10/01 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 **shortrange** or **longrange** subcommand additionally, you can narrow the term of data to be exported.

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 Monitor** window.

The **show** subcommand outputs the period from April 28, 2004, 11:04 to May 3, 2004, 11:04.

The screenshot shows a web interface for monitoring data. 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 "2004 /05 /03 09 :34" and a "To" field with the value "2004 /05 /03 11 :04". Below these is a "Real Time" radio button and a dropdown menu showing the number "5". The second set on the right has a "From" field with the value "2004/04/28 11:04" and a "To" field with the value "2004/05/03 11:04". Below these is a "90 min." label and an "Apply" button. A red line connects the text box above to the "From" field of the second set of pickers. Red boxes highlight the date and time values in both the "From" and "To" fields of the second set.

Figure A.10 Information output by Show Subcommand

The **login** command must execute before the **show** subcommand executes.

A.4.5 Group Subcommand

A.4.5.1 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]][△...]|
}
```

A.4.5.2 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 A.21 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, see the tables indicated in the "See" column.

Table A.22 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 Monitor window	Usage statistics about parity groups	PhyPG_dat.ZIP (Note 1)	Table A.2
PhyLDEV		Usage statistics about logical volumes	PhyLDEV_dat.ZIP (Note 1)	
PhyExG		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 (Note 1)	
PhyCSW		Usage statistics about access paths and write pending rate	PhyCSW_dat.ZIP (Note 1)	
PG	LDEV tab in the Performance Monitor window	Statistics about parity groups, external volume groups, or V-VOL groups	PG_dat.ZIP	Table A.3
LDEV		Statistics about logical volumes in parity groups, in external volume groups, or in V-VOL groups	LDEV_XXXXX.ZIP (Note 2)	Table A.4
Port	Port-LUN tab in the Performance Monitor window	Statistics about ports	Port_dat.ZIP	Table A.5
PortWWN		Statistics about host bus adapters connected to ports	PortWWN_dat.ZIP	Table A.6
LU		Statistics about LUs	LU_dat.ZIP	Table A.7
PPCG	WWN tab in the Performance Monitor window	Statistics about SPM groups	PPCG_dat.ZIP	Table A.8
PPCGWWN		Statistics about host bus adapters belonging to SPM groups	PPCGWWN_dat.ZIP	Table A.9
Operand	Window of Performance Monitor	Monitoring Data Saved in the File	Saved ZIP File	See
RemoteCopy	TC Monitor window and TCz Monitor window	Statistics about remote copy operations by TrueCopy and TrueCopy for z/OS (in the whole disk subsystem)	RemoteCopy_dat.ZIP	Table A.10
RCLU		Statistics about remote copy operations by TrueCopy and TrueCopy for z/OS (at LUs)	RCLU_dat.ZIP	Table A.11
RCLDEV		Statistics about remote copy operations by TrueCopy and TrueCopy for z/OS (at logical volumes controlled by a particular CU)	RCLDEV_XXXXX.ZIP (Note 3)	Table A.12
RCCLPR		Statistics about remote copy operations by TrueCopy and TrueCopy for z/OS (at CLPR)	RCCLPR_dat.ZIP	Table A.13

UniversalReplicator	UR Monitor window and URz Monitor window	Statistics about remote copy operations by Universal Replicator and Universal Replicator for z/OS (in the whole disk subsystem)	UniversalReplicator.ZIP	Table A.14
URJNL		Statistics about remote copy operations by Universal Replicator and Universal Replicator for z/OS (at journal groups)	URJNL_dat.ZIP	Table A.15
URLU		Statistics about remote copy operations by Universal Replicator and Universal Replicator for z/OS (at LUs)	URLU_dat.ZIP	Table A.16
URLDEV		Statistics about remote copy operations by Universal Replicator and Universal Replicator for z/OS (at logical volumes controlled by a particular CU)	URLDEV_XXXXX.ZIP (<i>Note 4</i>)	Table A.17

Note 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.

Note 2: A ZIP file whose name begins with LDEV.

Note 3: A ZIP file whose name begins with RCLDEV.

Note 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
```

A.4.5.3 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 Monitor** 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, refer to Table A.2 in section A.1.

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 both in the short range and long range will be 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 logical volume usage rates, which are displayed in the **Physical** tab of the **Performance Monitor** 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, refer to Table A.2 in section A.1.

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 both in the short range and long range will be 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 logical 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 Monitor** 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, refer to Table A.2 in section A.1.

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 logical volumes in external volume groups, which are displayed in the **Physical** tab of the **Performance Monitor** 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, refer to Table A.2 in section A.1.

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 Monitor** 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, refer to Table A.2 in section A.1.

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 both in the short range and long range will be 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 Monitor** 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 A.2 in section A.1.

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 both in the short range and long range will be 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 or external volume groups, or V-VOL groups, which are displayed in the **LDEV** tab of the **Performance Monitor** 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, refer to Table A.3 in section A.1.

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.

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 logical volumes, which are displayed in the **LDEV** tab of the **Performance Monitor** window. When statistics are exported to a ZIP file, multiple ZIP files whose names begin with LDEV_ will be output. For details on the statistics exported by this operand, refer to Table A.4 in section A.1.

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.

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 logical 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 Monitor** 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 A.5 in section A.1.

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. Note that 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:

CL1-A < CL1-B < ... < CL2-A < CL2-B < ... < CL3-a < CL3-b < ... < CL4-a < ... < 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 Monitor** 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, refer to Table A.6 in section A.1.

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. Note that 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:

CL1-A < CL1-B < ... < CL2-A < CL2-B < ... < CL3-a < CL3-b < ... < CL4-a < ... < 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 Monitor** 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, refer to Table A.7 in section A.1.

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. Note that 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 Monitor** 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, refer to Table A.8 in section A.1.

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. Note that 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 Monitor** 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 A.9 in section A.1.

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. Note that 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 z/OS in the whole disk subsystem. 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, refer to Table A.10 in section A.1.

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 z/OS at each 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, refer to Table A.10 in section A.1.

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. Note that 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 LU paths will be exported.

RCLDEV△[[*CU-id*]:[*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 z/OS at logical 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, refer to Table A.11 in section A.1.

When you specify variables *CU-id*, you can narrow the range of CUs that control the logical volumes whose monitoring data are to be exported. *CU-id* is an ID of a CU. The colon (:) indicates a range. For example, 00:05 indicates CUs from 00 to 05.

Ensure that the *CU-id* value on the left of the colon is smaller than the *CU-id* value on the right of the colon. For example, you can specify RCLDEV 00:05, but you cannot specify RCLDEV 05:00.

If *CU-id* is not specified, the monitoring data of all the logical 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 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, refer to Table A.13 in section A.1.

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 z/OS in the whole disk subsystem. When statistics are exported to a ZIP file, the file name will be UniversalReplicator.ZIP. For details on the statistics exported by this operand, refer to Table A.13 in section A.1.

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 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, refer to Table A.14 in section A.1.

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 z/OS at each 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, refer to Table A.15 in section A.1.

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. Note that 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 LU paths will be exported.

URLDEV△[[CU-id]:[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 z/OS at logical 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, refer to Table A.16 in section A.1.

When you specify variables *CU-id*, you can narrow the range of CUs that control the logical volumes whose monitoring data are to be exported. *CU-id* is an ID of a CU. The colon (:) indicates a range. For example, 00:05 indicates CUs from 00 to 05.

Ensure that the *CU-id* value on the left of the colon is smaller than the *CU-id* value on the right of the colon. For example, you can specify URLDEV 00:05, but you cannot specify URLDEV 05:00.

If *CU-id* is not specified, the monitoring data of all the logical volumes will be exported.

A.4.5.4 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 LU groups E1-1 to E1-5:

```
group PG E-1:E-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
```

A.4.6 Shortrange Subcommand

A.4.6.1 Syntax

shortrange△[[yyyyMMddhhmm][{+|-}hhmm]:[yyyyMMddhhmm][{+|-}hhmm]]

A.4.6.2 Description

The **shortrange** 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 **shortrange** subcommand is valid for monitoring data in short range. Monitoring data in short range is the contents of the following windows:

- The **Performance Monitor** window with selecting **shortrange** 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 **shortrange** subcommand whichever operand you specify to the **group** subcommand. If you run the Export Tool without specifying the **shortrange** subcommand, the data stored in the whole monitoring term will be exported.

The **login** subcommand must execute before the **shortrange** subcommand executes.

A.4.6.3 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 Monitor window (see Figure A.11).

The screenshot shows a 'Monitoring Term' configuration window. It has a blue sidebar on the left with the text 'Monitoring Term'. The main area is light blue and contains two sets of time pickers. The first set, on the left, has a 'From' field with '2004 /05 /03 09 :34' and a 'To' field with '2004 /05 /03 11 :04'. The second set, on the right, has a 'From' field with '2004/04/28 11:04' and a 'To' field with '2004/05/03 11:04'. Red boxes highlight the '2004/04/28 11:04' and '2004/05/03 11:04' values. Above the first set, a box labeled 'Starting time for collecting statistics' points to the 'From' field. Above the second set, a box labeled 'Ending time for collecting statistics' points to the 'To' field. Below the time pickers, there is a 'Real Time' section with a dropdown menu showing '5' and a '90 min.' label. An 'Apply' button is at the bottom right.

Figure A.11 Starting and Ending Times 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, **200401230000+0130** indicates Jan. 23, 2004. 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, **200401230000-0130** indicates Jan. 22, 2004. 22:30.

Subtracts time from the ending time for collecting monitoring data, if yyyyMMddhhmm is omitted.

Note: If the last two digits 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.

A.4.6.4 Examples

The examples below assume that:

- The starting time for collecting monitoring data is January 1, 2004, 00:00.
- The ending time for collecting monitoring data is January 2, 2004, 00:00.

shortrange 200401010930:200401011730

The Export Tool saves monitoring data within the range of Jan. 1, 9:30-17:30.

shortrange 200401010930:

The Export Tool saves monitoring data within the range of Jan. 1, 9:30 to Jan. 2, 00:00.

shortrange :200401011730

The Export Tool saves monitoring data within the range of Jan. 1, 0:00-17:30.

shortrange +0001:

The Export Tool saves monitoring data within the range of Jan. 1, 0:01 to Jan. 2, 00:00.

shortrange -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.

shortrange +0101:-0101

The Export Tool saves monitoring data within the range of Jan. 1, 1:01-22:59.

shortrange 200401010900+0130:200401011700-0130

The Export Tool saves monitoring data within the range of Jan. 1, 10:30-15:30.

shortrange 200401010900-0130:200401011700+0130

The Export Tool saves monitoring data within the range of Jan. 1, 7:30-18:30.

shortrange 200401010900-0130:

The Export Tool saves monitoring data within the range of Jan. 1, 7:30 to Jan. 2, 00:00.

A.4.7 Longrange Subcommand

A.4.7.1 Syntax

longrange△[[yyyyMMddhhmm][{+|-}ddhhmm]:[yyyyMMddhhmm][{+|-}ddhhmm]]

A.4.7.2 Description

The **longrange** 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 **longrange** 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 Monitor** window with selecting **longrange** as the storing period. The monitoring items whose data can be stored in long range are limited. Table A.23 shows the monitoring items to which the **longrange** subcommand can be applied, and also shows the operands to export those monitoring items.

Table A.23 Monitoring Items to which Longrange Subcommand can be Applied

Monitoring Data	Operands of Group Subcommand
Usage statistics about parity groups.	PhyPG Long
Usage statistics about logical 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 **longrange** subcommand, the data stored in the whole monitoring term will be exported.

The **login** subcommand must execute before the **longrange** subcommand executes.

A.4.7.3 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 Monitor window (see Figure A.12).




Figure A.12 Starting and Ending Times 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, **200401120000+010130** indicates Jan. 13, 2004, 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, **200401120000-010130** indicates Jan. 10, 2004. 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 up to one of the four values. Also, the value on the right of the colon will be rounded down to one of the four values. For example, if you specify 200401010013:200401010048, the specified value is regarded as 200401010015:200401010045.

A.4.7.4 Examples

The examples below assume that:

- The starting time for collecting monitoring data is Jan. 1, 2004, 00:00.
- The ending time for collecting monitoring data is Jan. 2, 2004, 00:00.

longrange 200401010930:200401011730

The Export Tool saves monitoring data within the range of Jan. 1, 9:30-17:30.

longrange 200401010930:

The Export Tool saves monitoring data within the range of Jan. 1, 9:30 to Jan. 2, 00:00.

longrange :200401011730

The Export Tool saves monitoring data within the range of Jan. 1, 0:00-17:30.

longrange +000015:

The Export Tool saves monitoring data within the range of Jan. 1, 0:15 to Jan. 2, 00:00.

longrange -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.

longrange +000115:-000115

The Export Tool saves monitoring data within the range of Jan. 1, 1:15-22:45.

longrange 200401010900+000130:200401011700-000130

The Export Tool saves monitoring data within the range of Jan. 1, 10:30-15:30.

longrange 200401010900-000130:200401011700+000130

The Export Tool saves monitoring data within the range of Jan. 1, 7:30-18:30.

longrange 200401010900-000130:

The Export Tool saves monitoring data within the range of Jan. 1, 7:30 to Jan. 2, 00:00.

A.4.8 Outpath Subcommand

A.4.8.1 Syntax

`outpath△[path]`

A.4.8.2 Description

The **outpath** subcommand specifies the directory to which monitoring data will be exported.

A.4.8.3 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.

A.4.8.4 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
```

A.4.9 Option Subcommand

A.4.9.1 Syntax

`option△[compress|nocompress]△[ask|clear|noclear]`

A.4.9.2 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.

A.4.9.3 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.

A.4.9.4 Example

The following example saves monitoring data in CSV files, not in ZIP files:

```
option nocompress
```

A.4.10 Apply Subcommand

A.4.10.1 Syntax

apply

A.4.10.2 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.

A.4.11 Set Subcommand

A.4.11.1 Syntax

`set△[switch={m | off}]`

A.4.11.2 Description

The **set** subcommand starts or ends monitoring the disk subsystem (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 section A.4.3) to log onto 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 section A.3.3.
- When the **set** subcommand starts or ends the monitoring or changes the gathering interval after the **Performance Monitor** window is started, the contents displayed in the **Performance Monitor** window does not change automatically in conjunction with the **set** subcommand operation. To display the current monitoring status in the Performance Monitor window, select the **Refresh** button.

A.4.11.3 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. Note that 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.

A.4.11.4 Examples

The following command file saves port statistics and then ends monitoring ports:

```
svpip 158.214.135.57
login expusr passwd
show
group Port
shortrange 200404010850:200404010910
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
```

A.4.12 Help Subcommand

A.4.12.1 Syntax

help

A.4.12.2 Description

The **help** subcommand displays the on-line help for subcommands.

If you want to view the on-line help, it is recommended that you create a batch file and a command file that are exclusively used for displaying the on-line help. For detailed information, see the *Example* below.

A.4.12.3 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/JSanRmiServerEx.jar" -Xmx268435456 -  
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 either of the following to view the on-line 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.

A.4.13 Java Command for Exporting Data in Files

A.4.13.1 Syntax

```
java△-classpath△class-path△property-parameters△sanproject.getmondar.RJMdMain
```

A.4.13.2 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.

A.4.13.3 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. You must, at least, 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.

- **-Xmx***memory-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 268435456, 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.

- **-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.

- **-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 on Jan. 1, 09:10:10.

A.4.13.4 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/JSanRmiServerEx.jar" -Dhttp.proxyHost=Jupiter
-Dhttp.proxyPort=8080 -Xmx268435456 -Dmd.command=command.txt -Dmd.logpath=log
sanproject.getmondat.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/JSanRmiServerEx.jar" -Xmx268435456
-Dmd.command=command.txt -Dmd.logfile=export.log
-Dmd.logpath=log sanproject.getmondat.RJMdMain ↵
```

Note: The "↵" symbol indicates the end of a command line.

A.5 Troubleshooting the Export Tool

Table A.24 explains possible problems with the Export Tool and probable causes and recommended solutions to the problems.

Table A.24 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, please 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, please install the correct version of JRE.
The Export Tool stops and the processing does not continue. The command prompt window was displaying progress of the export processing, but the window stopped displaying progress before the processing stops. The progress information does not seem to be updated anymore.	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. Please confirm whether the specified memory size is correct or not.
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. Please confirm whether you are using a computer that is not supported, or whether the computer is slow. Please 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 A.14, below.
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.
When the Export Tool terminates abnormally, the row of Check License is shown as UnmarshalException in the log file.	It could be caused by a combination of DKCMAIN/SVP program version and export tool version. Please confirm whether versions of these programs are correct.

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 A.24 lists the Export Tool messages and recommended actions against errors.

Table A.25 Messages to be Issued by 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. Please remove errors.
Illegal character: "character"	An illegal character is used. Please use legal characters.
Invalid length: token	The length is invalid. Specify a value that has a correct length.
Invalid range: range	The specified range is invalid. Specify the correct range.
Invalid value: "value"	The specified value is invalid. Specify a correct value.
Login failed	An attempt to log into the SVP failed. Probable causes are: 1. An incorrect operand is used for the svpip subcommand. 2. An incorrect operand is used for the login subcommand. 3. The specified user ID is used by another person, and the person is being logged in. 4. Two users are currently displaying the Performance Monitor window. 5. Two users are currently executing the Export Tool. Note: If the error cause is 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. Please specify a host name.
Missing output directory	No directory is specified for saving files. Please 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.
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: range	The value is outside the range. If the shorrange subcommand or the longrange 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 switchoperand, 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 z/OS, Universal Replicator, and Universal Replicator for

Possible Problems	Probable Causes and Recommended Action
	z/OS..
RMI server error (part-code, error-number)	An error occurs at the RMI server. For detailed information, please refer to the Storage Navigator Error Codes manual.
Unable to display help message	The Export Tool cannot display the on-line 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: file	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: host	The Export Tool cannot resolve the host name. Please specify the correct host name.
Unsupported command: <i>command</i>	The Export Tool does not support the specified command. Please specify a correct command.
Unsupported operand: operand	The specified operand is not supported. Please correct the specified operand.
Unsupported option: option	The specified option is not supported. Please correct the specified option.
Some file exists in path. What do you do? clear(c)/update(u)/stop(p) You selected "action". Is it OK? (y/n)	Files exist in path. - 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.
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 group subcommand. The number of targets is xxx and the total number is yyy (refer to <i>Note</i> below). The Export Tool ends collecting data. <i>Note:</i> For example, suppose that the disk subsystem contains 100 parity groups and the command file contains the following command line: group PG 1-1:1-2 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 subsystem 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 that some operands must be enclosed by double quotation marks ("). Please check the command line to find whether double quotation marks are missing.

Appendix B Limitations for Storage Partition Administrators

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

B.1 Limitations of Performance Monitor

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

Table B.1 Limitations for Storage Partition Administrators (Performance Monitor)

Window	Limited Function
Physical tab in the Performance Monitor 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 (see Figure B.1). Therefore, the user cannot start Volume Migration.
LDEV tab in the Performance Monitor window	The tree and list display only the information in the SLPR allocated to the user ID.
Port-LUN tab in the Performance Monitor window	▪ The tree and list display only the information in the SLPR allocated to the user ID. ▪ The SPM button is not displayed (see Figure B.1). Therefore, the user cannot start Server Priority Manager.
WWN tab in the Performance Monitor window	The WWN tab is not displayed (see Figure B.1). 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 (see Figure B.1). The user cannot view the information about remote copy operations performed by TrueCopy, TrueCopy for z/OS, Universal Replicator, and Universal Replicator for z/OS.
Monitoring Options window	The Monitoring Options window is not displayed (see Figure B.1). The user cannot start or stop monitoring, or change the gathering interval.

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




Figure B.1 Example of a Text File

B.2 Limitations of the Export Tool

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 z/OS, Universal Replicator, and Universal Replicator for z/OS (Table A.9 to Table A.16) cannot be exported.
- A storage partition administrator cannot use the **set** subcommand to start or stop monitoring, or change the gathering interval.

Appendix C Manual Volume Migration Using CCI

In the environment where Volume Migration is installed on the disk subsystem, users can use CCI to perform manual volume migration from an open-system host. This appendix explains the procedure and notes for using CCI to perform manual volume migration.

C.1 Migrating Volumes Using CCI

This section illustrates with an example the procedure when performing manual volume migration by using commands of CCI. In this example, It is assumed that `group1` means the group name in the CCI configuration definition file and `pair1` means the name of the volume which pair is target of the operation. For details on CCI operation, please refer to the *Command Control Interface (CCI) User and Reference Guide*.

To perform manual volume migration 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 1 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 C.1 Status Transition Chart of a Pair during Volume Migration by CCI

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 select the **All Refresh** button.
- Migration cannot be performed when either of the specified volumes is a LUSE volume.
- 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 C.1 Errors during Volume Migration or Migration Cancellation 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.
2056	The migration source volume is being used with Universal Replicator.	Release the volume from Universal Replicator, and then perform the volume migration.

Error Code	Cause	Measure
2057	The setting of VMA is different between the migration target volume and the migration source volume.	This volume migration cannot be executed. Check the configuration definition file.
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 disk subsystem.	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 disk subsystem is exceeded.	Reduce the number of pairs set in the disk subsystem, 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.
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	Release the TrueCopy pair, and then perform the

Error Code	Cause	Measure
	volume of TrueCopy.	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 with referring to Table 3.1.
2356	The migration target volume is being used with Cross-system Copy.	Release the Cross-system Copy pair, and then perform the volume migration.
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.
237E	The specified pair is being used with Cross-system Copy.	Release the Cross-system Copy 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.

Error Code	Cause	Measure
239D	The migration target volume is a system volume of ENAS.	This volume migration cannot be executed. Check the configuration definition file.

Acronyms and Abbreviations

ACP	Array control processor
CHA	Channel adapter (also called backend director)
CHIP	Client-host interface processor (HDS equivalent of CHA)
CHP	Channel processor
CSV	Comma-separated value format. Also known as comma-delimited format.
CSW	Cache switch
CU	Control unit (logical control unit)
CV	Customized volume
DASD	Direct-access storage device
DFD	DASD fast write
DKA	Disk adapter (also called backend director)
DKP	Disk processor
DRR	Data Recovery and Reconstruction Processor
ESCON	Enterprise System Connection (IBM trademark for optical channels)
HDD	Hard disk drive
LAN	Local-area network
LDEV	Logical device
LU	Logical unit
LUN	Logical unit number
NSC	Network Storage Controller
NVS	Nonvolatile storage
PDEV	Physical device
P-VOL	Primary volume
RAID	Redundant array of independent disks
SM	Shared memory
SMPL	Simplex
S-VOL	Secondary volume
SVP	Service processor
TC	TrueCopy
TCz	Hitachi TrueCopy for z/OS (formerly TC390)
USP	Universal Storage Platform
ZIP	A common data compression format.

Index

A

- apply subcommand, 244
- Attribute tab, 137
- auto migration mechanism, 30
- auto migration parameters, setting, 150
- auto migration plans, remaking, 152
- automating volume migrations, 149

C

- connections, HBA
 - many-to-many, 51
 - one-to-one, 48
- CSV files, 191

D

- data recovery and reconstruction processor
 - viewing statistics, 11
- disk drives, monitoring, 103
 - I/O rates, viewing disk, 103
 - transfer rates, viewing, 106
- disk usage rate, changing maximum, 149
- disk workload statistics, viewing, 13
- DRR. Refer to data recovery and reconstruction processor

E

- Export Tool
 - installing
 - in UNIX, 206
 - in Windows, 206
 - preparing a batch file, 211
 - requirements, 205
 - starting with java command, 248
 - subcommands, 216
 - to run, 212
 - troubleshooting, 250
- exporting monitoring data, 191
- external volume group statistics, 7

H

- HBA, deleting from SPM group, 186
- History tab, 142
- Host Bus Adapter, replacing, 183
- host bus adapter, traffic from, 15
- Host Bus Adapters, grouping, 185

I

- I/O rate, 13
- I/O rate, port, 14
- I/O rates between HBAs, viewing, 119
- installing
 - Export Tool
 - in UNIX, 206
 - in Windows, 206

J

- java command, 248
- JRE
 - requirements, 205

L

- logging in specified user to SVP, 220
- logical volumes, migrating
 - automatically, 29
 - manually, 26
- logical volumes, migrating manually, 146
- longrange storage, 70
- LU path transfer rates, 116
- LU paths, monitoring, 115

M

- mainframe volume migration, differential tables, 41
- manual migration, 146
- Manual Plan tab, 126
- migrating logical volumes
 - manually, 23
- migration
 - and workload balancing, 22
- migration destinations, reserving, 145
- migration plans, concurrent, 40
- migration requirements, 33
- migration, manual for logical volumes, 146
- monitoring
 - how to start, 17
 - viewing results of, 18
- monitoring data
 - determining when collected, 221
 - outputting, 221
 - specifying the type of, 225
- Monitoring Options panel, 85
- monitoring performance, procedure for, 87

N

- non- prioritized port, 171
- non-prioritized WWNs, setting upper limits for, 180
- normal volume, 145

O

- on-line help
 - displaying for subcommands, 247
 - viewing subcommands, 217
- open system volume migration, differential tables, 43

P

- parity group, 7
- parity group, setting a fixed, 149
- paths between HBAs and ports, monitoring, 119
- Performance Manager
 - system requirements, 60
- Performance Monitor
 - definition, 1
 - notes on operation, 19
 - overview, 5
 - overview of operations, 17
- Performance Monitor panel, 64
 - LDEV tab, 72
 - Physical tab, 66
 - Port-LUN tab, 77
 - WWN tab, 82
- Port tab in SPM, using the, 171
- Port tab, SPM panel, 162
- ports, viewing I/O rates for, 109
- ports, viewing transfer rates for, 111
- preparing a batch file
 - for Export Tool, 211
- prioritized port, 171
- Priority Access
 - definition, 1

R

- read hit ratio, 13
- referencing a termination code, 215
- requirements
 - Export Tool, 205
 - JRE, 205
- reserved volume, 145

S

- saving monitoring data, 212
- Server Priority Manager
 - overview, 46
 - starting, 170
 - using, 48
- Server Priority Manager panel, 162

- shortrange storage, 71
- SMP group, switching priority of, 186
- specifying
 - to compress monitoring data to ZIP files, 243
 - to overwrite or delete when saving monitoring data to files, 243
 - type of monitoring data, 225
- specifying IP address or host name of SVP, 218
- SPM group, containing multiple HBAs in an, 185
- SPM group, deleting, 188
- SPM group, renaming an, 188
- SPM name for HBA, changing, 182
- starting
 - Export Tool, 248
- storage ranges, statistical, 16
- subcommand
 - apply, 244
 - svip, 218
- subcommands
 - displaying on-line help for, 247
 - Export Tool, 216
 - viewing on-line help, 217
- svip subcommand, 218
- SVP
 - logging in specified user to, 220
 - specifying IP address or host name of, 218

T

- termination code, 214
- threshold, setting, 174
- threshold, setting a, 182
- to run
 - Export Tool, 212
- traffic statistics, analyzing, 179
- traffic, measuring between HBAs and ports, 172
- traffic, monitoring all between HBAs and ports, 176
- transfer rates, 13
- transfer rates between HBAs, viewing, 121
- transfer rates, port, 14
- troubleshooting
 - Export Tool, 250
 - technical support, 190

U

- upper-limit control
 - disabling SPM, 47
- upper-limit value, setting for HBAs in SPM group, 187
- upper-limit values, setting, 173
- usage statistics
 - cache memory, 98
 - channel processor, 95

- data recovery and reconstruction processor,
 - 11
- data recovery/reconstruction, 97
- disk LU path traffic, 15
- disk processor, 96
- external volume groups, 92
- for access paths, 100
- for disk subsystem resources, 6
- for external volumes in external volume groups, 93
- for logical volumes in parity groups, 89
- host bus adapter and port path traffic, 15
- parity group, 88
- viewing access path, 12
- viewing cache memory, 11
- viewing channel processor, 8
- viewing disk processor, 9
- viewing disk workload, 13
- viewing external volume, 8
- viewing external volume group, 7
- viewing logical volume, 7
- viewing LU path traffic, 15
- viewing parity group, 7
- viewing port traffic, 14
- write pending rate, 11, 98

V

- Volume Migration Panel
 - Attribute tab, 137
 - History tab, 142
 - Manual Plan tab, 126
- Volume Migration, starting, 144
- volume migrations, automating, 29, 149

W

- Web client computer, Performance Manager
 - operations using, 61
- write hit rate, 13
- write pending rate, 11
- write permission, 205
- WWN tab, SPM panel, 164
- WWN tab, using in many-to-many environment, 176

Z

- ZIP format, 191

