

Symmetrix
Open Systems Environment
Volume I

PRODUCT GUIDE

P/N 200-999-563
REV F

Copyright © 2001 EMC Corporation. All rights reserved.

Printed March 2001

No part of this publication may be reproduced or distributed in any form or by any means, or stored in a database or retrieval system, without the prior written consent of EMC Corporation.

The information contained in this document is subject to change without notice. EMC Corporation assumes no responsibility for any errors that may appear.

All computer software programs, including but not limited to microcode, described in this document are furnished under a license, and may be used or copied only in accordance with the terms of such license. EMC either owns or has the right to license the computer software programs described in this document. EMC Corporation retains all rights, title and interest in the computer software programs.

EMC Corporation makes no warranties, expressed or implied, by operation of law or otherwise, relating to this document, the products or the computer software programs described herein. EMC CORPORATION DISCLAIMS ALL IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE. In no event shall EMC Corporation be liable for (a) incidental, indirect, special, or consequential damages or (b) any damages whatsoever resulting from the loss of use, data or profits, arising out of this document, even if advised of the possibility of such damages.

Trademark Information

EMC², EMC, MOSAIC:2000, and Symmetrix are registered trademarks and EMC Enterprise Storage, The Enterprise Storage Company, The EMC Effect, Connectrix, EDM, SDMS, SRDF, TimeFinder, PowerPath, InfoMover, FarPoint, EMC Enterprise Storage Network, EMC Enterprise Storage Specialist, EMC Storage Logic, Universal Data Tone, E-Infostructure, and Celerra are trademarks of EMC Corporation.

Adaptec is a trademark of Adaptec, Inc.

AIX and AS/400 are registered trademarks of International Business Machines

Compaq is a registered trademark of Compaq Computer Corporation.

Fujitsu is a registered trademark of Fujitsu Limited.

Intel and Pentium are registered trademarks of Intel Corporation.

Hewlett-Packard, HP and HP-UX are registered trademarks of Hewlett-Packard Company.

Intel is a registered trademark of Intel Corporation.

Lucent is a registered trademark of Lucent Technologies, Inc.

Motorola is a registered trademark of Motorola, Inc.

NCR is a registered trademark of National Cash Register.

NEC is a trademark of NEC Corporation.

Network File System and NFS are trademarks of Sun Microsystems, Inc.

Novell is a registered trademark of Novell, Inc.

OpenVMS trademark of Digital Equipment Corporation.

Origin is a trademark and SGI and Challenge are registered trademarks of Silicon Graphics Inc.

OS/2 is a registered trademark of International Business Machines Corporation.

Pyramid is a registered trademark of Siemens Nixdorf.

SCO and UnixWare are registered trademarks of the Santa Cruz Operation, Inc.

Sequent is a registered trademark of Sequent Computer Systems, Inc.

SGI and Challenge are registered trademarks of Silicon Graphics Inc.

Siemens is a registered trademark of Siemens Aktiengesellschaft (or Siemens AG).

Solaris and SunOS are trademarks of Sun Microsystems, Inc.

Stratus is a registered trademark of Stratus Computer, Inc.

Tru64 is a trademark of Compaq Computer Corporation.

UNIX is a registered trademark in the United States and other countries and is licensed exclusively through X/Open Company Ltd.

VERITAS and VxFS are registered trademarks of VERITAS Software Corporation.

Windows NT is a registered trademark of Microsoft Corporation.

All other trademarks used herein are the property of their respective owners.



Chapter 1	Introducing Open Systems Environments	
	Overview	1-2
	Symmetrix Enhancement Options.....	1-13
	Hardware Option.....	1-13
	Software Options	1-13
	Information Protection Options.....	1-14
	Information Sharing Options	1-17
	Information Management Options	1-17
	EMC Enterprise Backup Solutions	1-23
Chapter 2	General Information	
	Optimizing Performance	2-2
	Open Systems (FBA) Data Format	2-3
	Fibre Channel and SCSI Interfaces	2-3
	Logical Volume Structure	2-4
	Hyper-Volume Extension	2-16
	Meta Volumes.....	2-18
Chapter 3	Performance Tools	
	UNIX Performance Tools	3-2
	sar	3-2
	vmstat	3-3
	iostat.....	3-3
Chapter 4	NCR Environment	
	Overview	4-2
	Useful NCR Utilities and Functions.....	4-4

Symmetrix Configuration.....	4-6
Incorporating Symmetrix	4-7
Displaying the New Devices.....	4-10
Configuring New Devices	4-11
Using sysadm.....	4-11
Configuring New Devices Using NCR Script	4-14
Configuring Devices from the Command Line.....	4-15
Adding Devices Online.....	4-18
System and Error Messages	4-19

Chapter 5 **IBM NUMA-Q Environment**

Overview.....	5-2
Useful IBM NUMA-Q Utilities and Functions	5-5
Symmetrix Configuration.....	5-6
Incorporating Symmetrix	5-7
Target ID Support	5-7
Wide, Synchronous, & Queueing	
Support 1.....	5-7
Wide, Synchronous, & Queueing	
Support 2.....	5-8
LUN Support below DYNIX/ptx 4.4.2	5-10
Fibre Channel Topology	5-12
Fibre Channel Class of Service	5-14
Configuring the IBM NUMA-Q System.....	5-15
Multipathing and Multiporting Support	5-15
Configuring /Deconfiguring Devices Online	5-18
Formatting and Partitioning Definition Files	5-19
New features in DYNIX/ptx 4.6.0 and above	5-19
Creating Physical Volumes.....	5-23
Create the Mount Directory	5-24
Create a New Filesystem	5-24
Mount the Directory.....	5-24
IBM NUMA-Q Volume Manager	5-26
Create Volumes Using ptx/SVM.....	5-27
Creating Mirrored Volumes using ptx/SVM	5-29
Dissociating Mirrored Volumes.....	5-31
Useful ptx/SVM Command Line Utilities.....	5-33
System and Error Messages	5-40
Using the Command Line	5-40
Using the Menu Utility	5-40

Chapter 6 HP 9000 Server Environment

Overview	6-2
Useful HP-UX Utilities and Functions	6-5
Symmetrix Configuration	6-8
Using Symmetrix as a Root Disk	6-9
Creating a Bootable Disk (10.x & 11.x)	6-9
Hardware Mapping 10.0 & 11.x (offline)	6-13
Recognizing New Devices and LUNs	6-15
HP-UX 10.x and 11.x	6-15
Configuring Logical Volumes	6-17
Creating Physical Volumes	6-17
Volume Groups and Logical Volumes	6-17
Creating the Mount Directory	6-18
Creating a New Filesystem	6-18
Mount the Directory	6-18
Using SAM	6-20
HP 9000 K- and D-class Servers	6-21
Channel Failover (10.x and 11.x)	6-22
Setting Up Alternate Links	6-22
Primary Link Failure	6-23
Viewing the Current Links	6-24
Common Serial Number Option	6-26
Adding Devices Online	6-27
System and Error Messages	6-28
Diagnostic Patches	6-28

Chapter 7 HP 9000 Workstation Environment

Overview	7-2
Useful HP-UX Utilities and Functions	7-5
HP-UX 10.x & 11.x	7-5
Symmetrix Configuration	7-8
Using Symmetrix as a Root Disk	7-9
Creating a Bootable Disk	7-9
Hardware Mapping 10.x & 11.x (offline)	7-12
Recognizing New Devices and LUNs	7-15
HP-UX 9.04	7-15
HP-UX 10.x and 11.x	7-15
Configuring Logical Volumes	7-17
Creating Physical Volumes	7-17
Volume Groups & Logical Volumes	7-17
Creating the Mount Directory	7-18
Creating a New Filesystem	7-18

	Mount the Directory	7-18
	Using SAM.....	7-19
	Adding Devices Online.....	7-20
	Performance Issues with HP-UX	7-21
	System and Error Messages	7-22
	Diagnostic Patches.....	7-22
Chapter 8	NEC Environment	
	NEC Servers and Workstations	8-2
Chapter 9	AIX Environment	
	Overview.....	9-2
	Terminology	9-2
	Useful AIX Utilities and Functions	9-5
	Useful EMC Commands and Functions.....	9-10
	Symmetrix Configuration.....	9-11
	Incorporating Symmetrix	9-12
	With AIX 4.2 or Higher	9-12
	Device Characteristics for Symmetrix	9-14
	Configuring Disks Using Commands	9-15
	Make a Volume Group	9-15
	Make the Volume Group Available.....	9-17
	Create Logical Volumes	9-18
	Creating a Mount Point	9-18
	Create and Mount Filesystems	9-18
	Configuring Disks Using SMIT	9-20
	Add a Volume Group.....	9-20
	Add Logical Volumes to the Volume Group	9-20
	Creating Filesystems	9-21
	Adding Devices Online.....	9-23
	System and Error Messages	9-24
	Viewing using SMIT.....	9-24
	Viewing using errpt.....	9-24
Chapter 10	SGI/CHALLENGE Environment	
	Overview	10-2
	Useful IRIX Utilities and Functions	10-3
	Symmetrix Configuration.....	10-5
	Configuring the SGI System.....	10-7
	Partitioning and Labeling the Devices	10-8
	Creating the Filesystem	10-12

Creating and Mounting an xfs Filesystem	10-12
Creating a Stripe Volume Set	10-14
Using XLV	10-14
Using the XVM Volume Manager	10-17
Example	10-18
Deleting a Stripe Volume Set	10-19
Channel Failover	10-20
Setting up Alternate Paths	10-20
Starting Channel Failover	10-21
Viewing the state of paths	10-21
Adding Devices Online	10-24
System and Error Messages	10-25

Chapter 11 SGI/ORIGIN Environment

Overview	11-2
Useful IRIX Utilities and Functions	11-3
Symmetrix Configuration	11-5
Configuring the SGI System	11-6
Partitioning and Labeling the Devices	11-7
Tag Command Queuing Enabled	11-8
Creating the Filesystem	11-10
Creating and Mounting an xfs Filesystem	11-10
Creating a Stripe Volume Set	11-12
Using XLV	11-12
Deleting a Stripe Volume Set	11-14
Using the XVM Volume Manager	11-15
Example	11-16
Channel Failover	11-18
Setting up Alternate Paths	11-18
Starting Channel Failover	11-19
Viewing the state of paths	11-19
Adding Devices Online	11-22
System and Error Messages	11-23

Chapter 12 Solaris Environment

Overview	12-2
Useful Solaris Utilities and Functions	12-4
Hardware Mapping	12-5
Symmetrix Configuration	12-6
SPARCstations and SPARCservers	12-7
UltraSPARC Series	12-8

Incorporating Symmetrix	12-9
To Recognize LUNs	12-9
Format, Partition, Label	12-12
Adding Devices Online	12-14
Solstice DiskSuite	12-15
Creating a Concatenated Volume	12-15
Concatenated Volumes for UFS Logging	12-16
Creating a Striped Volume	12-17
Creating a Striped Volume for UFS Logging	12-18
Creating a Mirrored Striped Volume	12-20
Mirrored Striped Volume for UFS Logging	12-21
DiskSuite State Database Replicas	12-23
Useful DiskSuite Commands	12-24
Veritas Volume Manager	12-25
Initializing Disks	12-25
Creating a Concatenated Volume	12-26
Creating a Striped Volume	12-30
Creating a Striped Mirrored Volume	12-33
Create and Mount a File System	12-40
Intent Logging	12-40
UFS File System on Raw Device	12-41
VxFS File System on Raw Device	12-42
UFS File System on SDS Device	12-43
VxFS File System on SDS Device	12-43
UFS File System on VxVM Device	12-44
VxFS File System on VxVM Device	12-45
File System Expansion	12-46
System and Error Messages	12-48

Chapter 13 Solaris x86 Environment

Overview	13-2
Host Setup Overview for New Symmetrix	13-3
Useful Solaris x86 Utilities and Functions	13-4
Host Bus Adapters	13-6
Symmetrix Configuration	13-9
Incorporating Symmetrix	13-11
To Recognize LUNs	13-11
System Specification File Modification	13-12
Partitioning and Labeling a New Device	13-13
Disk Device File Naming Conventions	13-13
Partitioning and Slicing Symmetrix Devices	13-13
Creating and Mounting a UFS File System	13-17

Create a New File System	13-17
Create a Mount Directory	13-17
Mount the Filesystem	13-17
Adding Devices Online	13-18
System and Error Messages	13-19

Chapter 14 Compaq Tru64 UNIX Environment

Overview	14-2
EMC Files, Compaq Patches & Online Docs	14-2
Compaq Tru64 UNIX Utilities & Functions	14-4
Symmetrix Configuration	14-6
Host Configuration	14-8
Installing the HBA	14-8
Console-Level Scan and Configuring a Boot Device	14-8
Rebuilding the Tru64 UNIX Kernel	14-9
Adding the Symmetrix Device Entry	14-9
Device Management	14-11
Device Naming Conventions	14-11
Disklabel and Device Partitions	14-12
Adding and Managing Devices Online	14-13
Using File Systems	14-15
Creating & Mounting UNIX File System	14-15
Advanced File System (AdvFS)	14-15
Creating & Mounting Advanced File System	14-16
Reconstructing an AdvFS Domain	14-17
LSM	14-18
Example 1: Setting Up LSM	14-18
Example 2: Creating mirrored Volume	14-19
Example 3: Creating 4-way striped volume	14-20
System and Error Messages	14-22

Chapter 15 Fujitsu Environment

Overview	15-2
Useful Fujitsu Utilities and Functions	15-5
Symmetrix Configuration	15-7
Incorporating Symmetrix	15-8
Fujitsu DS-90 or GP7000D Host	15-8
Fujitsu PRIMEPOWER (GP7000F) Host	15-9
Format, Partition, Label	15-12
Fujitsu DS-90 or GP7000D Host	15-12
Fujitsu PRIMEPOWER (GP7000F) Host	15-14

Format, Partition, Label using sysadm.....	15-15
Adding Devices Online.....	15-16
Fujitsu DS-90 or GP7000D Host	15-16
Fujitsu PRIMEPOWER (GP7000F) Host.....	15-16
Solstice DiskSuite.....	15-17
Creating a Concatenated Volume.....	15-17
Concatenated Volumes for UFS Logging.....	15-18
Creating a Striped Volume	15-19
Creating a Striped Volume for UFS Logging.....	15-20
Creating a Mirrored Striped Volume	15-22
Mirrored Striped Volume for UFS Logging	15-23
DiskSuite State Database Replicas	15-25
Useful DiskSuite Commands.....	15-26
Veritas Volume Manager.....	15-27
Initializing Disks.....	15-27
Creating a Concatenated Volume.....	15-28
Creating a Striped Volume	15-32
Creating a Striped Mirrored Volume	15-36
Create and Mount a File System.....	15-42
Intent Logging.....	15-42
UFS File System on Raw Device.....	15-43
VxFS File System on Raw Device.....	15-44
UFS File System on SDS Device	15-45
VxFS File System on SDS Device.....	15-46
UFS File System on VxVM Device	15-47
VxFS File System on VxVM Device	15-48
File System Expansion	15-49
System and Error Messages	15-50

Chapter 16 Fujitsu Siemens PRIMEPOWER (GP7000F) Environment

Overview.....	16-2
Useful Solaris Utilities and Functions	16-4
Hardware Mapping.....	16-5
Symmetrix Configuration.....	16-6
PRIMEPOWER (GP7000F) Servers	16-7
Incorporating Symmetrix	16-8
To Recognize LUNs.....	16-8
Format, Partition, Label	16-11
Adding Devices Online.....	16-12
Solstice DiskSuite.....	16-13
Creating a Concatenated Volume.....	16-13
Concatenated Volumes for UFS Logging	16-14

Creating a Striped Volume.....	16-15
Creating a Striped Volume for UFS Logging	16-16
Creating a Mirrored Striped Volume.....	16-17
Mirrored Striped Volume for UFS Logging	16-19
DiskSuite State Database Replicas.....	16-21
Useful DiskSuite Commands	16-22
Veritas Volume Manager	16-23
Initializing Disks	16-23
Creating a Concatenated Volume	16-24
Creating a Striped Volume.....	16-28
Creating a Striped Mirrored Volume.....	16-32
Create and Mount a File System	16-38
Intent Logging	16-38
UFS File System on Raw Device	16-39
VxFS File System on Raw Device	16-40
UFS File System on SDS Device.....	16-41
VxFS File System on SDS Device	16-41
UFS File System on VxVM Device.....	16-42
VxFS File System on VxVM Device.....	16-43
File System Expansion.....	16-44
System and Error Messages.....	16-46

Chapter 17 **Stratus Fault Tolerant Environment**

Overview	17-2
Useful Stratus Utilities and Functions	17-4
Symmetrix Configuration	17-7
Continuum 400 Series Host Cabling	17-9
Editing the /stand/conf File.....	17-11
Recognizing New Devices and LUNs.....	17-14
Configuring Logical Volumes.....	17-15
Creating Physical Volumes	17-15
Volume Groups & Logical Volumes	17-15
Creating the Mount Directory	17-16
Creating a New Filesystem.....	17-16
Mount the Directory	17-16
Using SAM	17-17
Adding Devices Online	17-18
System and Error Messages.....	17-19
Diagnostic Patches	17-19

Chapter 18	Lucent Environment	
	DNCP-400 Hardware System	18-2
Chapter 19	SCO UnixWare Environment	
	Overview.....	19-2
	Host Bus Adapters.....	19-3
	Useful Functions and Utilities	19-7
	Symmetrix Configuration.....	19-9
	Configuring a New Device.....	19-10
	Disk Device File Naming Conventions	19-10
	Format and Partition a Disk.....	19-10
	Setting Up a Raw Device.....	19-12
	Create and Mount a File System	19-13
	Adding Devices Online.....	19-15
	System and Console Error Messages	19-16
Chapter 20	Linux Environment	
	Host Connectivity	20-2
	Useful Linux Utilities and Functions	20-3
	Symmetrix Configuration.....	20-5
	Host Configuration with Adaptec AHA-2944UW HBA.....	20-7
	Installing the HBA.....	20-7
	Kernel Configuration & Compilation with Adaptec HBA	20-10
	Introduction.....	20-10
	Required Kernel Modifications	20-10
	Additional Notes	20-11
	Configuring and Compiling the Kernel	20-16
	New Installation	20-16
	Kernel Upgrade	20-19
	Partitioning and Labeling a New Device	20-24
	Partitioning Symmetrix Devices.....	20-25
	Creating and Mounting a File System	20-29
	System and Error Messages	20-30
Appendix A	Symmetrix Device Definition Files	
	Overview	A-2
	Using the World Wide Web	A-3
	Connecting to EMC's FTP Server	A-4
	Transferring the Device Definition Files	A-5
	The inq Command	A-6

Supported Hosts A-6

Appendix B SCSI Directors and Cabling

 SCSI Directors B-2

 SCSI Cabling B-3

 Multiple SCSI Hosts to Symmetrix..... B-3

 SCSI Interface Tests..... B-4

 Cable Tests..... B-5

Appendix C Host Bus Adapter Configuration

 Host Bus Adapters Tables C-2

Appendix D vi Commands

 Editing Tasks D-2

2-1	3:1 Logical Volume Mapping Example	2-18
2-2	Concatenated Volumes	2-19
2-3	Striped Data	2-20
5-1	Multiporting (only with DYNIX/ptx 4.4.2 and above)	5-16
5-2	ptx/SVM Mirrored Volume Structure in DYNIX/ptx 4.4.2	5-26
6-1	Bootable Disk - Sample Configuration	6-9
6-2	I/O Configuration (HP-UX 10.0)	6-14
6-3	Current Links (vgdisplay option)	6-24
6-4	Symmetrix Volume Configuration	6-25
7-1	Bootable Disk - Sample Configuration	7-9
7-2	I/O Configuration (HP-UX 10.x)	7-13
9-1	Detail Data Format Diagrammed	9-29
10-1	disk_config File	10-7
13-1	Solaris X86 fdisk Menu Example	13-15
14-1	Partitioning Scheme Using "unknown" Disk Type	14-12
17-1	Host Cabling Configuration (7100)	17-10
17-2	Host Cabling Configuration (8000)	17-11
20-1	Adaptec AHA-2944UW Banner	20-7
20-2	SCSI Support Menu	20-17
20-3	SCSI Low-Level Drivers Submenu	20-18
20-4	SCSI Support Menu	20-20
20-5	SCSI Low-Level Drivers Submenu	20-21
A-1	Sample Output of inq Command	A-6
B-1	Split Cable Configuration (2 SA Ports)	B-3
B-2	"Y" Cable Configuration (1 SA Port)	B-4

1-1	Open Systems Supported Hosts	1-3
1-2	Information Protection Options	1-14
1-3	Information Sharing Options	1-17
1-4	Information Management Options	1-18
1-4	Information Management Options	1-19
1-5	Host Platforms Supported by EMC ControlCenter	1-22
1-6	Enterprise Backup Solutions	1-23
2-1	Total Open Systems Cylinders	2-5
2-2	Maximum Cylinders per Logical Volume for Each Device	2-6
2-3	Cylinders per Logical Volume - 4 GB Physical Device	2-7
2-4	Cylinders per Logical Volume - 9 GB Physical Device (5.25")	2-8
2-5	Cylinders per Logical Volume - 9 GB Physical Device (3.5")	2-9
2-6	Cylinders per Logical Volume - 18 GB Physical Devices	2-10
2-7	Cylinders per Logical Volume - 23 GB Physical Devices	2-11
2-8	Cylinders per Logical Volume - 36 GB Physical Devices	2-12
2-9	Cylinders per Logical Volume - 47 GB Physical Devices	2-13
2-10	Cylinders per Logical Volume - 50 GB Physical Devices	2-14
2-11	Cylinders per Logical Volume - 73 GB Physical Devices	2-15
2-12	Disk Device Configuration	2-17
2-13	Meta Volume Stripe Sizes	2-20
3-1	sar Statistic Summary	3-2
4-1	NCR Operating System Features	4-2
4-2	Disk Device Configuration	4-6
5-1	DYNIX/ptx 4.2.x and below Host Environment	5-3
5-2	DYNIX/ptx 4.3.0 and above Host Environment	5-4
5-3	Disk Device Configuration	5-6
5-4	Symmetrix Device Definition Files	5-21
6-1	Qualified Operating Systems for HP 9000 Server Hosts	6-2
6-2	HP 9000 Server Operating System Features	6-3
6-3	Capacity Limits of Recent HP-UX Releases	6-4

6-4	Disk Device Configuration	6-8
6-5	K-class Server Configuration	6-21
6-6	HP-9000 D-Class Servers	6-21
6-7	Diagnostic Cumulative Patches	6-28
6-8	Logtool Decoding Patches	6-29
7-1	HP/9000 Workstation Matrix	7-2
7-2	HP 9000 Workstation OS Features	7-3
7-3	Capacity Limits of Recent HP-UX Releases	7-4
7-4	Disk Device Configuration	7-8
7-5	Diagnostic Cumulative Patches	7-22
8-1	NEC UNIX Systems Matched to HP	8-2
9-1	AIX Operating System Features Summary	9-4
9-2	Disk Device Configuration	9-11
9-3	Device Characteristics	9-14
9-4	Factor -t, Physical Partition Numbers, Disks per VG	9-16
9-5	Limitations of LVM in AIX Versions	9-17
9-6	Detail Data	9-27
9-7	Sense Data Format Explained	9-29
10-1	SGI/Challenge Operating System Features	10-2
10-2	Disk Device Configuration	10-5
11-1	SGI/Origin Operating System Features	11-2
11-2	Disk Device Configuration	11-5
12-1	Sun/Solaris Operating System Features	12-3
12-2	Disk Device Configuration	12-6
12-3	Bit Settings	12-11
13-1	Solaris x86 Operating System Features	13-2
13-2	Adaptec AHA2944UW	13-6
13-3	NCRpqs 2.0 Quad Card	13-7
13-4	NCRpqs 2.1 Quad Card	13-7
13-5	Qlogic QLA2100F (Fibre Channel) Card	13-8
13-6	QLogic QLA2200-EMC	13-8
13-7	Disk Device Configuration	13-9
13-8	Bit settings for SCSI Options	13-12
14-1	Disk Device Configuration	14-6
14-2	Tru64 UNIX SCSI Device Support	14-7
15-1	Fujitsu DS-90 Operating System	15-3
15-2	Fujitsu GP7000D Operating System	15-3
15-3	Fujitsu Siemens PRIMEPOWER Solaris Operating System Features ..	15-4
15-4	Disk Device Configuration	15-7
15-5	Bit Settings	15-11
16-1	Fujitsu Siemens PRIMEPOWER Solaris Operating System Features ..	16-3
16-2	Disk Device Configuration	16-6
16-3	Bit Settings	16-9

17-1	Stratus with HP-UX v10.20.1 OS Features	17-2
17-2	Disk Device Configuration	17-7
17-3	Director Bits to be Set	17-8
19-1	SCO UnixWare 7 Summary of Features	19-2
19-2	QLogic Host Bus Adapter QLA2200F-EMC	19-3
19-3	QLogic Host Bus Adapter QLA1240D	19-4
19-4	QLogic Host Bus Adapter QLA1041D	19-5
19-5	Adaptec Host Bus Adapter AHA2944UW	19-6
19-6	Disk Device Configuration	19-9
19-7	Partition Table	19-11
20-1	HBAs for Linux Hosts on SCSI	20-2
20-2	Disk Device Configuration	20-5
20-3	HBAs for Linux Hosts on SCSI	20-6
B-1	Symmetrix SCSI Directors	B-2
C-1	NCR	C-2
C-2	IBM NUMA-Q (Sequent)	C-3
C-3	HP 9000	C-3
C-4	NEC	C-6
C-5	IBM	C-7
C-6	Silicon Graphics	C-10
C-7	Sun	C-11
C-8	Sun Solaris x86	C-14
C-9	Host BIOS Chart	C-15
C-10	Compaq (DEC)	C-17
C-11	Fujitsu	C-18
C-12	Fujitsu Siemens/FSC	C-19
C-13	Stratus	C-20
C-14	Lucent	C-20
C-15	SCO UnixWare 7	C-21
C-16	Linux	C-23
D-1	Starting vi	D-2
D-2	Moving the Cursor	D-2
D-3	Inserting Text	D-2
D-4	Deleting Text	D-3
D-5	Searching	D-3
D-6	Saving and Exiting vi	D-3
D-7	Miscellaneous vi Commands	D-3

As part of a continuous effort improve and enhance the performance and capabilities of the Symmetrix product line, EMC from time to time releases new revisions of Symmetrix hardware and microcode. Therefore, some functions described in this manual may not be supported by all revisions of Symmetrix microcode or hardware presently in use. If your Symmetrix unit does not offer a function described in this manual, please contact your EMC representative for a hardware upgrade or microcode update.

This manual describes the features and operation of the Symmetrix products in the open systems environment. This manual is intended for the storage administrator, system programmer, or operator who will be involved in acquiring, managing, or operating the Symmetrix subsystem.

This manual is part of a two volume set. This volume covers UNIX-based host environments. Linux is also presented. Volume II covers other UNIX-based hosts, including DG/UX and Siemens and non-UNIX hosts, such as Windows 2000 PC-based servers.

Here is an overview of where information is located in this manual.

Chapter 1, “**Introducing Open Systems Environment**,” describes the Symmetrix Fast-Wide Differential or Ultra-Wide Differential SCSI interface, supported UNIX-based hosts and operating systems, and Symmetrix open systems environment hardware and software options.

Chapter 2, “**General Information**,” discusses performance in the open systems environment, open systems data formats, and the logical volume structure.

Chapter 3, "**Performance Tools**," covers UNIX performance tools and Symmetrix software tools.

Chapter 4, "**NCR Environment**," describes how to incorporate Symmetrix in the NCR environment.

Chapter 5, "**IBM NUMA-Q Environment**," describes how to incorporate Symmetrix in the IBM NUMA-Q (formerly Sequent) environment.

Chapter 6, "**HP 9000 Server Environment**," describes how to incorporate Symmetrix in the HP/9000 server environment.

Chapter 7, "**HP 9000 Workstation Environment**," describes how to incorporate Symmetrix in the HP/9000 workstation environment.

Chapter 8, "**NEC Environment**," lists NEC servers and workstations and maps them to corresponding HP models.

Chapter 9, "**AIX Environment**," describes how to incorporate Symmetrix in the IBM AIX environment.

Chapter 10, "**SGI/CHALLENGE Environment**," describes how to incorporate Symmetrix in the SGI/Challenge environment.

Chapter 11, "**SGI/ORIGIN Environment**," describes how to incorporate Symmetrix in the SGI/Origin environment.

Chapter 12, "**Solaris Environment**," describes how to incorporate Symmetrix in the Solaris environment.

Chapter 13, "**Solaris x86 Environment**," describes how to incorporate Symmetrix in the Solaris x86 environment.

Chapter 14, "**Compaq Tru64 UNIX Environment**," describes how to incorporate Symmetrix in the Compaq (formerly Digital) Tru64 UNIX environment.

Chapter 15, "**Fujitsu Environment**," describes how to incorporate Symmetrix in the Fujitsu GP7000D and earlier hosts environment.

Chapter 16, "**Fujitsu Siemens PRIMEPOWER (GP7000F) Environment**," describes how to incorporate Symmetrix in the Fujitsu Siemens PRIMEPOWER (formerly GP7000F) Solaris OS environment.

Chapter 17, "**Stratus Fault Tolerant Environment**," describes how to incorporate Symmetrix in the Stratus environment.

Chapter 18, "**Lucent Environment**," lists Lucent (formerly Ascend) systems and maps them to corresponding Stratus models.

Chapter 19, "**SCO UnixWare Environment**," describes how to incorporate Symmetrix in the SCO UnixWare 7 environment.

Chapter 20, "**Linux Environment**," describes how to incorporate Symmetrix in the SCO UnixWare 7 environment.

Appendix A, "**Symmetrix Device Definition Files**," describes how to obtain special device definition files needed for use with open systems hosts from EMC.

Appendix B, "**SCSI Directors and Cabling**," describes the types of cables to use when connecting to a server, the channel attachments, and how to test the cables and connections.

Appendix C, "**Host Bus Adapter Configuration**," lists Host Bus Adapters for the systems in this volume. Appendix C also describes how to configure selected host adapters for those hosts requiring them when connecting to Symmetrix open systems environment.

Appendix D, "**vi Commands**," lists the most commonly used vi editor commands.

The **Glossary** defines terms used in this manual.

Conventions Used in this Manual

EMC uses the following conventions for notes, cautions, warnings, and danger notices.

A note presents information that is important, but not hazard-related.



CAUTION

A caution contains information essential to avoid damage to the system or equipment. The caution may apply to hardware or software.

EMC uses the following type style conventions in this guide:

Boldface	◆ Specific filenames or complete paths. ◆ Window names and menu items in text. ◆ Selections you can make from the user interface, including buttons, icons, options, and field names. ◆ Emphasis in cautions and warnings.
<i>Italic</i>	◆ New terms or unique word usage in text. ◆ Command line arguments when used in text.
Fixed space	Examples of specific command entries that you would type, displayed text, or program listings. For example: <pre>QUERY [CUU=cuu VOLSER=volser]</pre>
<i>Fixed italic</i>	Arguments used in examples of command line syntax.

Where to Get Help

Obtain technical support by calling your local sales office.

For service, call:

(800) 782-4362 (SVC-4EMC)

or

(800) 543-4782 (543-4SVC)

and ask for Customer Service.

If you are located outside the USA, call the nearest EMC office for technical assistance. These offices are listed at the back of this manual.

Your Comments

Your suggestions will help us continue to improve the accuracy, organization, and overall quality of the user publications. Please e-mail us at **techpub_comments@emc.com** to let us know your opinion or any errors concerning this manual.

Introducing Open Systems Environments

This chapter describes the Symmetrix Fast-Wide Differential or Ultra-Wide Differential SCSI interface, supported UNIX-based hosts and operating systems, and Symmetrix open systems environment hardware and software options.

- ◆ Overview 1-2
- ◆ Symmetrix Enhancement Options..... 1-13

Overview

Symmetrix is a high performance Integrated Cached Disk Array (ICDA®) designed for online data storage.

Symmetrix is available in several models. All models differ in the number of disk devices, the number of storage directors, and base cache size per model.

The Symmetrix Fast-Wide or Ultra-Wide Differential SCSI channel and fibre channel interfaces attach to a variety of open system hosts.

Table 1-1 lists the host systems and operating systems supported by Symmetrix in the open systems environment. You should consult your EMC sales representative for an updated list of supported hosts.

The EMC Customer Engineer (CE) should also contact the EMC Configuration Specialist for updated online information. This information is necessary to configure Symmetrix to support the customer's host environment.

Refer to Volume II of this product guide for information on using Symmetrix with PC servers including Windows 2000, DG/UX, Fujitsu Siemens, AS/400, OpenVMS systems hosts and others.

For a list of host controller cards (HBAs) used on each of the various open systems hosts supported by Symmetrix, refer to Appendix C, *Host Bus Adapter Configuration*.

Table 1-1 Open Systems Supported Hosts

System	Host Models or Series	Min. Operating System
Bull	Escala D, E, M, R, S (S100, S120), T Series EPC400, EPC800 Series	AIX 4.2.1
	Escala EPC Series (EPC400, RL470, EPC1200), Escala E Series (E), Escala T Series (T)	AIX 4.3.0
	Escala E Series (E230), T Series (T430), EPC Series (EPC430)	AIX 4.3.1
	Escala EPC Series (EPC440, EPC1200A)	AIX 4.3.2
	Escala E Series (E250), Escala T Series (T450, T610), Escala EPC Series (EPC450, EPC610, EPC810, EPC2400, EPC2450)	AIX 4.3.3
Compaq (formerly Digital Equipment Corp.)	AlphaStation and AlphaServer Series 3000 AXP Models 300-900, 1000, 2000, 2100, 4000/4100, 8200, 8400 DS10, DS20, ES40, GS60, GS80, GS140, GS160, GS320	Digital UNIX V3.2A (DEC OSF/1) Open VMS V6.2 or higher
Data General	AviiON 5500/88k-based, 4900/5900/8600 Intel-based	DG/UX R4.11MU03
	AViiON 3750, 25000 NUMA Server	DG/UX R4.20MU06
EMC Celerra File Server	Data Mover DM6 Series	NAS 1.2.41
	Data Mover DM7 Series	NAS 1.2.51
	Control Station CS6-E & CS-E Series	SCO UnixWare 2.1.3
Fujitsu	DS90 (Japan only)	SVR4 V4.20
	GP7000D	UXP/DS V20L10
	PRIMEPOWER (GP7000F) models M200, M400, M600, M800, M1000, M2000	Fujitsu Solaris 2.6
Fujitsu Siemens/FSC	RM400, RM600	SINIX V5.42
	RM400C, RM600E/20/30/60/70	Reliant UNIX V5.43C10
	RM600E40/E45/E80/E85, RM600CS42	Reliant UNIX V5.45A10
	SR2000 (RM600E/20/60)	Reliant UNIX 2000-Y v7.5 (5.43C20)
	PRIMEPOWER (GP 7000F) Models M200, M400, M600, M800, M1000, M2000	Solaris 2.6
	RM300E, RM400E, RM400C	Reliant UNIX V5.44B10
	DS2000 (RM600E/70), SR2000 (RM600E/30/70)	Reliant UNIX 2000-Y v8.0 (5.44C30)

Table 1-1 Open Systems Supported Hosts (*continued*)

System	Host Models or Series	Min. Operating System
HP	HP 3000: 900 Series	MPE/ix 5.0 SubSys C.50.02 PP C5007 PA
	HP 9000: 700 Series: 715/100XC, 725/100, 735/99, 735/125, 755/99, 755/125 800 Series: E, F, G, H, I, T500 D-class, K-class (SCSI), R-Class B-Series (all), C-Series: J-200, J-210, J-210CX, J-250 T520	HP-UX 10.01
	HP 9000: T600 (SCSI & Fibre Channel), K-class (Fibre Channel) HP 9000: A-Class (Staccato A180, A180C)	HP-UX 10.20 TFC
	HP 9000: V-Class, V2200, V2250, V2500, V2600	HP-UX 11.0
	HP 9000: N-Class, L1000, L2000	HP-UX 11.0 990P
	HP 9000: V2500 SCA, V2600 SCA	HP-UX 11.10
	HP 9000: A-Class A500/440MHz, A400/440MHz	HP-UX 11.ACE
IBM	MCA-based RS/6000 (7015-R30, 7015-R40, 7015-R50) & SP2 PCI-based RS/6000 (7025-F50, 7026-H50) & SP2	AIX 4.2.1
	PCI-based RS/6000: 7025-F50, 7026-H50 PCI-based SP2 nodes: 9076-2xx, 9076-3xx, 9076-4xx, 9076-500, 9076-550, 9076-T70	AIX 4.3.0
	7013-S70, 7015-S70, 7025-H70, 7026-H70, 7017-S70, 7017-S70 as SP2 node	AIX 4.3.1
	7013-S7A, 7015-S7A, 7017-S7A, 7017-S7A as SP2 node	AIX 4.3.2
	PCI-based RS/6000: 7044-170, 7044-270, 7025-F50, 7025-F80, 7025-H70, 7026-H50, 7026-H70, 7026-H80, 7026-M80, 7017-S70, 7017-S70 as SP2 node, 7017-S7A, 7017-S7A as SP2 node, 7017-S80, 7017-S80 as SP2 node, pSeries 640 (B80), 680 (S85) PCI-based SP2 nodes: 9076-2xx, 9076-3xx, 9076-4xx, 9076-500, 9076-550, 9076-T70	AIX 4.3.3
	AS/400 9406 systems: Models 3xx, 5xx, 6xx, 7xx	OS/400 V3R2

Table 1-1 Open Systems Supported Hosts (continued)

System	Host Models or Series	Min. Operating System
IBM NUMA-Q (Sequent)	Symmetry 2000 and 5000	DYNIX/ptx 4.2.1
	NUMA-Q	DYNIX/ptx 4.4.2
	NUMA-Q 2000	DYNIX/ptx 4.4.4
ICL	SmartController SDS	Open VME 2
	XtraSERVER P1000i	SCO UnixWare V2.1.2 and 7.0.1 Microsoft Windows NT 4.0 Server Enterprise Edition (OEM copy) plus Trimetra XtraSERVER NT Additions package BN553205
Linux	Compaq Intel-based platforms	RedHat v2.2.9
	NEC ExpressServer 5800 120Ra2	RedHat v2.2.10
	HP Netserver LH Series FSC Primergy 460, 470, 670, 870 Series, H400, K400, N400, N800 Dell PowerEdge 2300	RedHat v2.2.11
	HP Netserver LX Series	RedHat v2.2.15
	VALinux FullOn 2200 Series, 3500 NEC ExpressServer 5800 120Ra2 Compaq AP200/Proliant 800, 850, 1600, 1850, 8500 Dell Poweredge 2300 FSC Primergy 460, 470, 670, 870 Series, H400, K400, N400, N800 Data General AViiON 1400 HP Netserver LH, LX Series VA Linux FullOn 2200 Series VA Linux 3500	RedHat v6.2
Lucent	DNCP-400	HP-UX B.10.20.1

Table 1-1 Open Systems Supported Hosts (*continued*)

System	Host Models or Series	Min. Operating System
NCR	34xx, 35xx, 41xx, 45xx, 3600, 3600AP	UNIX SVR4 MPRAS 2.03.01
	4300, 4380	NT 4.0 SP3 UNIX SVR4 MPRAS 3.01
	4700 (8-node) and 5150 (128-node)	UNIX SVR4 3.01
	4400, 48XX, S50	NT 4.0 SP3 UNIX SVR4 3.02 Solaris 7 x86
	4300, 4400	Solaris 7 x86
NEC	HP9000 D-Class: NX7000/22, NX7000/23, NX7000/27, NX7000/D280, NX7000/32, NX7000/33, NX7000/37, NX7000/D380, NX7000/D390, TX7/D390, TX7/D280 HP9000 K-Class: NX7000/220, NX7000/260, NX7000/K370, NX7000/K380, NX7000/420, NX7000/460, NX7000/K570, NX7000/K580, TX7/K370, TX7/K380 NX7000/T600/V2250, 2500 , TX7/V2250, 2500	HP UX 11.0
	UP4800/760EX, 860, 860AD, 760REX, 860R, 860RAD, 780, 790, 880, 890, 880AD, 890AD	UX4800 R13.5, R13.6

Table 1-1 Open Systems Supported Hosts (continued)

System	Host Models or Series	Min. Operating System
PC: Intel-based Novell NetWare Servers (SCSI)	All Intel-based PC Servers with Microchannels	NetWare 3.12
	All Intel-based PC Servers with EISA bus	
	HP Netserver LH (LH-PRO), LHII, LHIII, LH-4, LX (LX-PRO), LXR (LXR-PRO), LXR-PRO8, LXR-8000, LXR-8500 IBM Netfinity 5000, 5500, 5500-M10, 5500-M20, 7000, 7000-M10, 8500 FSC Primergy 460, 470, 670, 700, 870 series, H400, K400, N400, N800 Compaq Proliant 850, 1600, 1850, 2500, 3000, 5000, 5500, 6000, 6400R, 6500, 7000, 8500 (Intel Pentium Pro and XEON models), DL360, DL380, DL580, ML350, ML370, ML530, ML570 NCR Worldmark 4300, 4380, 4400, 47XX, 48XX, 52XX, S50 Dell Poweredge 6100, 6300, 6350 HP Netserver LH (LH-PRO), LHII, LHIII, LH-4, LX (LX-PRO), LXR (LXR-PRO), LXR-PRO8, LXR-8000, LXR-8500	Novell NetWare 4.11 SP6
PC: Intel-based Novell NetWare Servers (Fibre Channel)	Compaq Proliant 850, 1600, 1850, 2500, 3000, 5000, 5500, 6000, 6400R, 6500, 7000, 8500 (Intel Pentium Pro and XEON models), DL360, DL380, DL580, ML350, ML370, ML530, ML570 IBM Netfinity 5000, 5500, 5500-M10, 5500-M20, 7000, 7000-M10, 8500 FSC Primergy 460, 470, 670, 870 series, H400, K400, N400, N800 Dell 6100, 6300, 6350 HP Netserver LH (LH-PRO), LHII, LHIII, LH-4, LX (LX-PRO), LXR (LXR-PRO), LXR-PRO8, LXR-8000, LXR-8500	Novell NetWare 5.00 SP1

Table 1-1 Open Systems Supported Hosts (*continued*)

System	Host Models or Series	Min. Operating System
PC: Intel-based NT Servers (SCSI)	All Intel-based PC Servers with Microchannels All Intel-based PC Servers with EISA Bus Compaq Proliant 850, 1600, 1850, 2500, 3000, 5000, 5500, 6000, 6400R, 6500, 7000, 8000, 8500 (Intel Pentium Pro and XEON models), DL360, DL380, DL580, ML350, ML370, ML530, ML570 FSC Primergy 460, 470, 670, 870 series, H400, K400, N400, N800 IBM Netfinity 5000, 5500, 5500-M10, 5500-M20, 6000R, 7000, 7000-M10, 7600, 8500(R) Dell Poweredge 2300, 2450, 4400, 6100, 6300, 6350, 6450, 8450 DG AV8600 Unisys AQUANTA QS/2V NEC Express 5800/140Ma, 5800/140Ha, 5800/180Ha	NT 3.5
	NCR Worldmark Server 4300, 4380, 4400, 47XX, 48XX, 52XX, S50 NCR Worldmark Server 3400, S50 HP Netserver LH (LH-PRO), LHII, LH-3, LH-4, LX (LX-PRO), LXR (LXR-PRO), LXR-PRO8, LXR-8000, LXR-8500, LT6000r, LH3000, LH6000 DG AV4900, AV5900 (Pentium)	NT 4.0 SP3
	Unisys SMP 5200 (ALR 2X), SMP5400 SPO Intel 4X, SMP6400 (QR6/QS6) Unisys QS6/QR6 ALR 4X, HS6/HR6 ALR 6X, QR2/QS2, DR2/DS2, ES5043, ES5045, ES5085, ES2043, ES2045, ES2085, ES2025, ES2023 Unisys ES7000 XS6/XR6 Unisys 10X	NT 4.0 SP5
	Network Engines P7000	NT 4.0 SP6A

Table 1-1 Open Systems Supported Hosts (*continued*)

System	Host Models or Series	Min. Operating System
PC: Intel-based NT Servers (Fibre Channel)	HP Netserver LH (LH-PRO), LHII, LHIII, LH-4, LX (LX-PRO), LXR (LXR-PRO), LXR-PRO8, LXR-8000, LXR-8500, LT6000r, LH3000, LH6000 Compaq Proliant 850, 850(R), 1600, 1850, 2500, 3000, 5000, 5500, 6000, 6400R, 6500, 7000, 8000, 8500 (Intel Pentium Pro and XEON models), DL360, DL380, DL580, ML350, ML370, ML530, ML570 FSC Primergy 460, 470, 670, 870 series, H400, K400, N400, N800 NCR Worldmark Server 4300, 4380, 4400, 47XX, 48XX, 52XX, S50 IBM Netfinity 5000, 5500, 5500-M10, 5500-M20, 6000R, 7000, 7000-M10, 7600, 8500(R) Dell Poweredge 2300, 2450, 4400, 6100, 6300, 6350, 6450, 8450 DG AV1400, AV3704R, AV3800, AV8600, AV8900, AV8950(R)	NT 4.0 SP3
	Unisys AQUANTA QS/2V NEC Express 5800/140Ma, 5800/140Ha, 5800/180Ha	NT 4.0 SP4
	Unisys ES5043, ES5045, ES5085, DR2/DS2, ES2043, ES2045, ES2085, QR2/QS2, ES2025, ES2023, ES7000 SP6, QR2/QS2	NT 4.0 SP5A
	ICL Trimetra P2000	NT 4.0 SP6A

Table 1-1 Open Systems Supported Hosts (continued)

System	Host Models or Series	Min. Operating System
PC: Intel-based Windows 2000 Servers (SCSI)	Compaq Proliant 1600, 1850, 2500, 3000, 5500, 6000, 6400R, 6500, 7000, 8000, 8500, DL360, DL380, DL580, ML350, ML370, ML530, ML570 IBM Netfinity 5500, 5500-M10, 5500-M20, 6000R, 7000, 7000-M10, 7600, 8500 Dell 2300, 2450, 4400, 6100, 6300, 6350, 6450, 8450 NCR 4300, 4380, 4400 NEC Express 5800/140Ha, 5800/140Ma, 5800/120Ra2, 5800/180Ha FSC Primergy 460, 470, 670, 870 series, H400, K400, N400, N800 Unisys QS/2V, ES7000 Network Engines P7000 ICL Trimetra P2000 HP LXR-8500, LT6000r, LH6000(R), LH3000, LH 3, LXR-8000, LH 4	Windows 2000 Advanced Server Build 2195
PC: Intel-based Windows 2000 Servers (Fibre Channel)	Compaq Proliant 1600, 1850, 2500, 3000, 5500, 6000, 6400R, 6500, 7000, 8000, 8500, DL360, DL380, DL580, ML350, ML370, ML530, ML570 IBM Netfinity 5500, 5500-M10, 5500-M20, 6000r, 7000, 7000-M10, 7600, 8500 Dell 2300, 2450, 4400, 6100, 6300, 6350, 6450, 8450 NCR 4300, 4380, 4400 NEC Express 5800/140Ha, 5800/140Ma, 5800/120Ra2, 5800/180Ha HP LXR-8500, LT6000r, LH6000, LH3000, LH 3, LXR-8000, LH 4 ICL Trimetra P2000 IBM Netfinity 5500, 5500-M10, 5500-M20, 6000r, 7000, 7000-M10, 7600, 8500 FSC Primergy 460, 470, 670, 870 series, H400, K400, N400, N800 Unisys QS/2V, ES2025, ES2043, ES2045, ES2085, ES5043, ES5045, ES5085, ES7000 DG AviiON 1400, 3704R, 3800, 8900, 8950, 8950R	Windows 2000 Advanced Server Build 2195

Table 1-1 Open Systems Supported Hosts (continued)

System	Host Models or Series	Min. Operating System
PC: Intel-based OS/2 Servers (SCSI)	All Intel-based PC Servers with Microchannels All Intel-based PC Servers with EISA bus HP Netserver LH (LH-PRO), LHII, LHIII, LH-4, LX (LX-PRO), LXR (LXR-PRO), LXR-PRO8, LXR-8000 IBM Netfinity 5000, 5500, 5500-M10, 5500-M20, 7000, 7000-M10 FSC Primergy 460, 670, 700, and 870 series Compaq Proliant 850, 1600, 1850, 2500, 3000, 5000, 5500, 6000, 6500, 7000 (Intel Pentium Pro and XEON models), DL360, DL380, DL580, ML350, ML370, ML530, ML570 NCR Worldmark 4300, 4380, 4400 Dell Poweredge 6100, 6300, 6350	OS/2 2.1
SCO UnixWare 7	FSC Primergy 460, 470, 670, 870 series, H400, K400, N400, N800 Dell Poweredge 2300, 2450, 4400, 6100, 6300, 6350, 6450, 8450 HP Netserver LH (LH-PRO), LHII, LHIII, LH-4, LX (LX-PRO), LXR (LXR-PRO), LXR-PRO8, LXR-8000 NCR Worldmark Server 4300, 4380, 4400 ICL XtraSERVER P1000	7.0.1
	FSC Primergy 460, 470, 670, 870 series, H400, K400, N400, N800 Dell Poweredge 2300, 2450, 4400, 6100, 6300, 6350, 6450, 8450 HP Netserver LH (LH-PRO), LHII, LHIII, LH-4, LX (LX-PRO), LXR (LXR-PRO), LXR-PRO8, LXR-8000 NCR Worldmark Server 4300, 4380, 4400 ICL P2000	7.1.0
Siemens-Pyramid	Nile Series	Data Center OSx 94D079.6
	RM 1000	Data Center OS DC/OSx 95m079 Reliant UNIX V5.43C0053
Silicon Graphics	Challenge S, L, DM, XL	IRIX 5.3
	Origin 200, 2000	IRIX 6.4.1
Stratus	Continuum 400 Series (Models AC POLO 7100, 8000)	HP-UX B.10.20.1

Table 1-1 Open Systems Supported Hosts (*continued*)

System	Host Models or Series	Min. Operating System
Sun	SPARCServer & SPARCcenter 10, 20, 1000, 2000	Solaris 2.3
	Sbus-based Servers: Ultra Workstation (Ultra 2) Ultra Enterprise (Ultra 3000, 3500, 4000, 4500, 5000, 5500, 6000, 6500, 10000) PCI-based Workstations & Servers: Netra 1120, 1125, 1400, 1405, T1; Ultra 5, 10, 30, 60, 80, 220(R), 250, 420(R), 450	Solaris 2.5.1
Sun Solaris x86	Siemens Primergy 460, 670, 700, and 870 series, H400, K400, N400, N800 NCR Worldmark 4300, 4380, 4400, 47XX, 48XX, 52XX, S50 Dell Poweredge 2300, 2450, 4400, 6100, 6300, 6350, 6450, 8450	Solaris 7 x86
Unisys	U6000/550 & 580 SMP6400 (QR6/QS6) SMP61000 Unisys 10X QR6/QS6 ALR4X HS6/HR6 ALR6X XS6/XR6 Unisys 10X DR2/DS2 QR2/QS2, ES2023, ES2025, ES2043, ES2045	UNIX SVR4 UNIXWare 2.1.3

Symmetrix Enhancement Options

EMC offers the following hardware and software options for enhancing Symmetrix performance and control in the open systems environment. Consult your EMC Systems Engineer for the latest information on these options.

Hardware Option

Phone Multiplexer

EMC offers a Phone Multiplexer for domestic customer sites with multiple Symmetrix units. The Symmetrix Phone Multiplexer consists of a (Private Branch Exchange) PBX with a Direct Inward Systems Access (DISA) board set that switches incoming calls to the correct Symmetrix unit. The Phone Multiplexer allows any mix of Symmetrix units (Models 3xxx, 4xxx, 5xxx, and 8xxx) to use a common telephone line to communicate with the EMC Customer Support Center. Each Symmetrix unit has its own extension in the Phone Multiplexer network.

Consult your EMC Systems Engineer for the latest information on this option.

Software Options

EMC offers several software options for the Symmetrix 8430 based on their functionality:

- ◆ Information Protection (Table 1-2 on page 1-14)
- ◆ Information Sharing (Table 1-3 on page 1-17)
- ◆ Information Management (Table 1-4 on page 1-18)
- ◆ EMC Control Center Supported Host Platforms (Table 1-5 on page 1-22)
- ◆ Enterprise Backup Solutions (Table 1-6 on page 1-23)

All of the software products are furnished under a license. Refer to the copyright page in this manual for the complete licensing statement. For the latest information on each of the EMC software products described in this section, contact your EMC Systems Engineer.

Information Protection Options

EMC software provides a variety of information protection/business continuance options. Table 1-2 describes the information protection software options.

Table 1-2 Information Protection Options

Option	Description	Prerequisite	Supported Hosts/Servers	Comments
Symmetrix Remote Data Facility (SRDF)	Business continuance solution that maintains a mirror image of data at a logical volume level in multiple Symmetrix units located up to hundreds of miles apart.	Symmetrix ESP for transfers between mainframe and open systems Symmetrix models.	All hosts and servers that connect to the Symmetrix system.	SRDF can be controlled by an EMC Customer Engineer from the service processor or by the customer using EMC ControlCenter host-based software. Refer to the SRDF Product Guide, P/N 200-999-554
• Campus Solution	Allows source and target units to be located up to 30 mi (60 km) apart.	None	Same as SRDF.	No host resources required.
• Extended Distance Solution	Allows source and target units to be located over 30 mi (60 km) apart.	None	Same as SRDF.	No host resources required.
• SRDF FarPoint	Extended distance solution option that reduces the number of communication lines and enhances synchronous mode performance. Can be used in Campus Solution environments with links greater than 9 mi (15 km).	None	Same as SRDF.	No host resources required.
• SRDF Multi-Hop	Extended distance solution option that enables a Business Continuance Volume BCV on the source side of an SRDF configuration to perform an incremental synchronization with a target volume on the other side, copying only changed tracks from source to target volume.	None	Same as SRDF.	Multi-hop mirroring uses EMC TimeFinder™ software to mirror to a third Symmetrix site only data that has changed.

Table 1-2 Information Protection Options (continued)

Option	Description	Prerequisite	Supported Hosts/Servers	Comments
Symmetrix Data Migration Service (SDMS)	Provides continuous business operations and data availability when data is automatically migrated from older mainframe storage devices to a new Symmetrix system.	Symmetrix ESP for transfers of mainframe and open systems data into the same Symmetrix system. Two Remote Link Directors during migration.	All supported CKD mainframe and open systems	Transparent to the host operating system and host applications. Must be Symmetrix-to-Symmetrix donor and recipient must be running the same code family. Symmetrix Data Migration Facility P/N 200-999-559
Nondisruptive Symmetrix Data Migration Service (ND-SDMS)	Software utility designed to perform data migration from one Symmetrix to another Symmetrix system without application or system interruption.	Symmetrix ESP for transfers between mainframe and open systems Symmetrix models. Two RLDs during migration.	All supported CKD mainframe and open systems	Nondisruptive to host and applications. ND-SDMS redirects the I/O from the donor controller to the Symmetrix controller. Migration takes place at the Symmetrix level.
EMC TimeFinder Open Systems	Business continuance solution that makes copies of data of standard Symmetrix devices available on special devices called Business Continuance Volumes (BCVs). BCVs are standard Symmetrix volumes that are preconfigured in the Symmetrix unit on which the business continuance operation will run, and have separate host addresses.	None	UNIX servers, Windows NT servers that connect to Symmetrix	Refer to EMC Solutions Enabler SYMCLI TimeFinder Component P/N 200-999-634 TimeFinder Command Line Interface for UNIX P/N 200-999-593 TimeFinder Command Line Interface for WIndows NT P/N 200-999-600
EMC TimeFinder for AS/400	Same as EMC TimeFinder	None	AS/400	AS/400 Control Software Package
EMC GeoCluster for HPMC (Hewlett-Packard MetroCluster)	An essential component of Hewlett-Packard's MC/ServiceGuard® with MetroCluster software. Provides open systems disaster tolerant solution enabling geographically dispersed HP-UX clusters with automatic failover of applications.	SRDF	HP 9000	HP MetroCluster with EMC SRDF

Table 1-2 Information Protection Options (*continued*)

Option	Description	Prerequisite	Supported Hosts/Servers	Comments
EMC GeoCluster for MSCS (Microsoft Cluster Services)	Provides automatic failover of geographically dispersed server clusters and EMC Enterprise Storage Systems	SRDF	Windows NT environments	EMC SRDF with Microsoft Cluster Services
EMC CopyPoint™ for AS/400 V1.0	Allows a secondary, offline copy of production data to be accessed by a secondary AS/400 system.	SRDF	AS/400	Operating system level V3R6 or less CISC support only, SRDF required
EMC CopyPoint for AS/400 V2.2.9	Allows a secondary, offline copy of production data to be accessed by a secondary AS/400 system.	SRDF or TimeFinder	AS/400	V3R7, V4R1, V4R2, V4R3 RISC support only or TimeFinder required

Information Sharing Options

Symmetrix provides centralized, sharable information storage that supports changing environments and mission-critical applications. Table 1-3 describes the information sharing software options.

Table 1-3 Information Sharing Options

Option	Description	Prerequisite	Supported Hosts/Servers	Comments
Symmetrix Enterprise Storage Platform™ (ESP)	Allows simultaneous storage and access of mainframe data and open systems data on the same Symmetrix system, or between two or more open systems and mainframe Symmetrix models.	None	All hosts and servers that connect to the Symmetrix systems	Refer to the Enterprise Storage Platform Product Guide,. P/N 200-999-556.
EMC InfoMover	Provides high-speed transfer between heterogeneous and similar MVS and UNIX-based host systems through a shared Symmetrix Enterprise Storage system.	Symmetrix ESP for transfers between mainframe and open systems Symmetrix models.	Mainframe, Open Systems	InfoMover is installed on each host platform that will use the software for file transfer. Files can be transferred among all combinations of MVS and UNIX platforms including: • MVS to UNIX • UNIX to MVS • UNIX to UNIX Refer to the InfoMover Product Guide P/N 200-999-619 and the InfoMover Release Notes, P/N 200-999-620.

Information Management Options

Symmetrix systems improve information management by allowing users to consolidate storage capacity for multiple hosts and servers. EMC offers powerful graphical user interface (GUI)-based tools that simplify and enhance Symmetrix configurations, performance, and status information gathering and management. Table 1-4 describes the information management options. Table 1-5 describes the host platforms supported by the EMC ControlCenter.

Table 1-4 Information Management Options

Option	Description	Prerequisite	Supported Hosts/Servers	Comments
EMC ControlCenter	Obtains the status of and monitors the performance of a Symmetrix unit attached to a host. Used with MVS, UNIX, and Windows NT platforms.	License to use EMC ControlCenter	Refer to Table 1-5.	Two user interfaces: • SM-GUI • SM-CLI Monitors and controls multiple Symmetrix units from a single workstation with SymmConsole Provides mainframe support through an MVS agent. Refer to EMC ControlCenter 4.0 for Windows with Agents for Windows NT, UNIX, and MVS. P/N 200-999-640 Also obtain the Solution Enabler Kit (separate installation procedure from EMC ControlCenter)
EMC ControlCenter – Workload Analyzer	Allows a storage administrator to collect, analyze, and graph historical performance data to help identify performance trends or implementation deficiencies	EMC ControlCenter	Intel platforms Windows 95, Windows 98, and Windows NT 4.0	Refer to EMC ControlCenter – Workload Analyzer Product Guide, P/N 200-999-627
EMC ControlCenter – SRDF TimeFinder Manager	Allows users to both monitor and control SRDF disaster recovery operations from a remote location.	Solutions Enabler	Same as EMC ControlCenter. Supports UNIX, and Windows NT	Refer to EMC ControlCenter SRDF Manager for UNIX P/N 200-999-565 EMC ControlCenter SRDF Manager for Windows NT P/N 200-999-589 EMC ControlCenter TimeFinder Manager for UNIX P/N 200-999-592 EMC ControlCenter TimeFinder Manager for Windows NT P/N 200-999-591

Table 1-4 Information Management Options

Option	Description	Prerequisite	Supported Hosts/Servers	Comments
EMC ControlCenter	Obtains the status of and monitors the performance of a Symmetrix unit attached to a host. Used with MVS, UNIX, and Windows NT platforms.	License to use EMC ControlCenter	Refer to Table 1-5.	Two user interfaces: • SM-GUI • SM-CLI Monitors and controls multiple Symmetrix units from a single workstation with SymmConsole Provides mainframe support through an MVS agent. Refer to EMC ControlCenter 4.0 for Windows with Agents for Windows NT, UNIX, and MVS. P/N 200-999-640 Also obtain the Solution Enabler Kit (separate installation procedure from EMC ControlCenter)
EMC ControlCenter – Workload Analyzer	Allows a storage administrator to collect, analyze, and graph historical performance data to help identify performance trends or implementation deficiencies	EMC ControlCenter	Intel platforms Windows 95, Windows 98, and Windows NT 4.0	Refer to EMC ControlCenter – Workload Analyzer Product Guide, P/N 200-999-627
EMC ControlCenter – SRDF TimeFinder Manager	Allows users to both monitor and control SRDF disaster recovery operations from a remote location.	Solutions Enabler	Same as EMC ControlCenter. Supports UNIX, and Windows NT	Refer to EMC ControlCenter SRDF Manager for UNIX P/N 200-999-565 EMC ControlCenter SRDF Manager for Windows NT P/N 200-999-589 EMC ControlCenter TimeFinder Manager for UNIX P/N 200-999-592 EMC ControlCenter TimeFinder Manager for Windows NT P/N 200-999-591

Table 1-4 Information Management Options (continued)

Option	Description	Prerequisite	Supported Hosts/Servers	Comments
EMC ControlCenter – Symmetrix Database Tuner Option	Provides ORACLE® database performance monitoring and tuning software optimized for Symmetrix physical placement.	Requires EMC ControlCenter Console and Symmetrix Manager	GUI runs on ControlCenter Console (Windows NT based)	Refer to EMC ControlCenter Symmetrix Database Tuner P/N 300-999-079 and EMC ControlCenter Symmetrix Database View P/N 300-999-078
EMC ControlCenter – Resource View	Provides end-to-end mapping of open systems data structures (databases, filesystems, and volume groups) to Symmetrix logical and physical volumes.	Requires EMC ControlCenter Console and Symmetrix Manager	Selected UNIX and Windows NT filesystems, volume managers and Oracle databases.	Refer to Resource View for Windows NT P/N 300-999-033
EMC ControlCenter – Symmetrix Disk Reallocation	Allows you to reallocate Symmetrix logical volumes including FBA devices, BCVs, and meta devices.	EMC ControlCenter	Same as EMC ControlCenter. Includes Windows NT, and UNIX.	Refer to EMC ControlCenter – SDDR Component Product Guide Contains SM-GUI component EMC ControlCenter – SDDR Component for Windows NT P/N 200-999-594 EMC ControlCenter – SDDR Component for UNIX P/N 200-999-595
EMC PowerPath for Windows NT	Enhances Symmetrix capabilities by providing automatic detection of and recovery from server-to-Symmetrix connection failures.	None	Intel-based servers Windows NT V4.0 with Service Pack 3, 4, 5.	Establishes from 2 - 32 SCSI paths between a system host and a Symmetrix device configured for the server. Allows you to configure data paths, maintain the paths, and tune the system for performance online. Refer to EMC PowerPath Product Guide P/N 200-999-628 and EMC PowerPath Release Notes, P/N 200-999-629

Table 1-4 Information Management Options (continued)

Option	Description	Prerequisite	Supported Hosts/Servers	Comments
EMC PowerPath for UNIX	Enhances Symmetrix capabilities by providing automatic detection of and recovery from server-to-Symmetrix connection failures.	COMPAQ and Siemens hosts require an RPQ. Digital and Reliant UNIX require an RPQ.	UNIX-based servers.	Establishes from 2 - 32 SCSI paths between an open system host and a Symmetrix device configured for the server. Allows you to configure data paths, maintain the paths, and tune the system for performance online. Various restrictions apply depending upon the host/server selected. Refer to the EMC PowerPath Product Guide, P/N 200-999-598 and EMC PowerPath Release Notes, P/N 200-999-603
Volume Logix	Allows the user to query and define access control on Symmetrix volumes across an Enterprise Storage Network (ESN) from a single, central, administrator station.	None	Any server supported by Volume Logix, running Volume Manager Administrator, or running Volume Logix Host Utilities requires a Fibre Channel host bus adapter (HBA). Volume Logix Administrator GUI and CLI.	Refer to the Volume Logix Product Guide, P/N 300-999-024 and Volume Logix Release Notes, P/N 300-999-025

Table 1-5 Host Platforms Supported by EMC ControlCenter

Vendor	Hardware	Operating System	Host Agent ^a
Bull	Escala	AIX®	4.3.2, 4.3.3
Compaq	AlphaStation™, AlphaServer™	Compaq Tru64™ UNIX	4.0E, 4.0F, 5.0A
Fujitsu Siemens	PRIMEPOWER (GP7000F) and models M200, M400, M600, M800, M1000, M2000	Solaris®	2.6, 7, 8
HP	HP 9000 A, D, E, F, G, H, I, K, L, N, V, and T-500, T-520 and T600 Series	HP-UX®	10.20, 11.0
IBM	RS/6000®, SP2®	AIX	4.3.2, 4.3.3
Microsoft®	PC servers with PCI Bus	Windows NT® Windows® 2000	SP4, SP5
NCR	3xxx and 5100 Series	UNIX® System V, Release 4	3.01, 3.02
NEC	All except P590	HP-UX	10.20, 11.0
Sequent	Symmetry® 2000 and 5000	DYNIX/ptx™	4.2.4 to 4.4.5, 4.4.7, 4.5.0, 4.5.1
	NUMA-Q®	DYNIX	4.4.2 to 4.4.5, 4.4.7, 4.5.0, 4.5.1
Siemens	RM 400 and 600	RELIANT® UNIX	5.43, 5.44, 5.45
Silicon Graphics®	Challenge® S, L, DM, XL	IRIX®	6.2
	Origin	IRIX64	6.5.3, 6.5.5, 6.5.8
Sun	SPARCserve™r, Enterprise Server	Solaris	2.6, 7, 8
	Intel PCI Platforms	Solaris x86	

a. Host agents are required for at least one system connected to a Symmetrix.

EMC Enterprise Backup Solutions

Enterprise Backup Solutions is the integration of EMC SRDF and TimeFinder with backup solutions from EMC and other vendors. These partner applications enable nondisruptive, point-in-time backup of corporate databases and applications. Table 1-6 describes the Enterprise backup Solutions.

Table 1-6 Enterprise Backup Solutions

Option	Description	Prerequisite	Supported Hosts/Servers	Comments
EMC Data Manager (EDM)	Centralized open systems high-performance database backup system. Includes all hardware, software, and service support. It automates tasks such as scheduling, tape management, and offsite disaster media storage.	Dependent upon selected host system.	UNIX, OS/2, NetWare®, Windows NT	Supports database backup and restore, on line and off-line, for ORACLE, Sybase®, Informix, and Microsoft SQL. Used with either EMC TimeFinder or SRDF. Refer to the EDM Product Description Guide.
Fast Dump Restore Safeguard Open Storage (FDRSOS)	Storage management software option that provides high speed, full volume backup and restore of open systems (SCSI) volumes to mainframe (MVS) tape or disk storage.	Requires Symmetrix connection to mainframe devices and Symmetrix ESP software.	Mainframe and open systems. Runs under MVS/XA, MVS/ESA, and OS/390 operating systems	Product is sold, supported, and serviced by Innovation Data Processing.
Omniback II V3.0	Provides automated backup and recovery of SRDF and TimeFinder volumes while the user continues to have access to the data.	Requires TimeFinder and/or SRDF (depending upon the implementation) and ES-Backup	Open systems	Sold by Hewlett Packard Company. Integration with data base products, as well as OPENVIEW®, SAP, and file systems.
Legato NetWorker for EMC Symmetrix	Provides data protection and storage management services for all network sizes. Provides backup, recovery, archive and automated disaster recovery for file/print, database, and application servers and their networked clients - regardless of platform or operating system.	Requires TimeFinder and/or SRDF (depending upon the implementation) and ES-Backup	Oracle 7.2.X and 7.3.X	Sold by Legato. Integration of SRDF and TimeFinder allows users to create separately addressable copies of their production database. Also perform high-speed backup from a second host while the production application remains on line and fully available to users.

This chapter discusses performance in the open systems environment, useful UNIX performance tools, open systems data formats, and the logical volume structure.

- ◆ Optimizing Performance.....2-2
- ◆ Open Systems (FBA) Data Format.....2-3
- ◆ Logical Volume Structure.....2-4

Optimizing Performance

Achieving optimal performance in the open systems environment requires careful and detailed planning of the Symmetrix configuration according to the requirements of the host(s) you are connecting to Symmetrix and your performance needs.

Carefully review issues listed below with your EMC representative before the EMC Customer Engineer installs the Symmetrix system. To assist in this process, refer to the sample worksheets and checklists provided in the Planning and Installation appendix in the product manual for your particular Symmetrix model.

- ◆ When configuring multiple hosts to a Symmetrix unit, distribute workloads by ranking from most busy to least busy.
- ◆ Data storage capacity required for each host connected to Symmetrix.
- ◆ Number of channels available from each host.
- ◆ The nature of the applications executed on the host connected to Symmetrix.
- ◆ The availability of a Logical Volume Manager (LVM) on the host and the use of data striping.
- ◆ The use of hyper-volumes on Symmetrix, hyper-volume size, and the allocation of hyper-volumes between different hosts, different channels, and different applications.
- ◆ Maximum drive and filesystem sizes supported by each host connected to Symmetrix.
- ◆ Any requirements for device sharing.
- ◆ Any requirements for a mixture of different RAID levels in the same Symmetrix subsystem.
- ◆ Number of Fast-Wide or Ultra-Wide Differential SCSI channel directors used and the number of ports used on each director.
- ◆ Fibre Channel
- ◆ Host-level mirroring considerations for device distribution in Symmetrix.
- ◆ The possibility of upgrading Symmetrix with additional drives in the future and its effect on the configuration if the model installed is not at maximum capacity.

Open Systems (FBA) Data Format

When using a Fast-Wide or Ultra-Wide Differential SCSI or a Fibre Channel interface to an open system processor, the Symmetrix system appears as standard SCSI disk devices with data stored in fixed-block architecture (FBA) format.

FBA disk devices store data in fixed sized blocks (typically 512 bytes). A disk device using FBA format is viewed as a large array of blocks. The physical position of the block (cylinder and track) is usually not significant to the host. When requesting disk access for read or write, the host addresses a file by the logical block address (LBA) of the starting block and a count of the total blocks used by the file. Symmetrix Fast-Wide and Ultra-Wide Differential SCSI channel directors and disk directors control access to cache and disk devices, responding to host requests as a standard SCSI disk device.

Fibre Channel and SCSI Interfaces

Symmetrix uses the SCSI protocol on its Fibre Channel interface; as a result, Fibre Channel interfaces and standard SCSI interfaces look the same to host devices. Each Symmetrix is recognized and addressed by controller number and logical unit number (LUN).

Implementation of the Fibre Channel interface is handled at the I/O driver and hardware level, so configuration of devices at the host is the same for SCSI or Fibre Channel.

For more information on Fibre Channel, see *Symmetrix Fibre Channel Product Guide*, P/N 200-999-642, EMC Corporation.

Logical Volume Structure

The Fast-Wide or Ultra-Wide Differential SCSI channel or Fibre Channel directors connecting open systems hosts to Symmetrix interact with cache. Therefore, there is no mechanical meaning to cylinders, tracks, and heads on a Symmetrix logical volume from the front end (host) point of view.

Symmetrix, however, uses a *logical* geometry definition for its logical volume structure. This geometry is reflected in the SCSI mode sense data available to the host. Symmetrix uses this logical volume structure:

- ◆ Each logical volume has N cylinders
- ◆ Each cylinder has 15 tracks (heads)
- ◆ Each track has 64 blocks of 512 bytes

Therefore, a Symmetrix logical volume with N cylinders has a usable block capacity of:

$$N * 15 * 64$$

“N” for each volume is specified during AutoInstall Symmetrix configuration by the EMC Customer Engineer.

To calculate the size of the logical volume in bytes for SCSI hosts:

$$\begin{aligned} &\text{\# of Cylinders} * \text{Heads} * \text{Blocks} * 512 \\ &(N * 15 * 64 * 512) \end{aligned}$$

For the maximum number of cylinders per logical volume for various device sizes, refer to Table 2-2 on page 2-6.

Devices with low-level formatting for other than 512-byte blocks have a different number of available cylinders. For example, drives formatted for AS/400 hosts use a 520-byte block format resulting in fewer logical cylinders.

For information on the maximum number of cylinders per logical volume in an AS/400 environment, refer to Volume II.

Table 2-1 summarizes the Symmetrix disk device cylinder capacities for 512-byte block format used for all hosts and 520-byte block format used by the IBM AS/400.

Table 2-1 Total Open Systems Cylinders

Drive Size	Available Cylinders ^a		
	512-byte block	520-byte block ^b	520-byte block ^c
3 GB (5.25")	5,862	Unsupported	Unsupported
4 GB (3.5")	8,672	8,534	8,522
9 GB (5.25")	17,726	Unsupported	Unsupported
9 GB (3.5")	18,414	17,540 (Load source disk) 17,484 (Emulates 9337-590 disk)	n/a
18 GB (3.5")	36,830	17,540 (Load source disk) 17,484 (Emulates 9337-590 disk)	n/a
23 GB (5.25")	46,942	Unsupported	Unsupported
36 GB (3.5")	73,662	n/a	n/a
47 GB (5.25")	95,262	Unsupported	Unsupported
50 GB (3.5") ^d	101,264	Unsupported	Unsupported
73 GB (3.5") ^{d, e}	149,063	Unsupported	Unsupported

a. The maximum number of cylinders that can be configured for a single Symmetrix volume is 20,478 (microcode revision level 5264 or lower) and 32,510 (microcode revision level 5265 and higher). See Table 2-2.

b. Symmetrix attached to AS/400 hosts only.

c. Symmetrix attached to UNIX hosts (servers) when AS/400 is present:

1. With microcode revision level 5566, devices for servers are low-level formatted to 512-byte block size and devices for the AS/400 are low-level formatted to 520-byte block size.
2. With microcode revision level 5266 and lower, all devices for servers and the AS/400 must be low-level formatted to 520-byte block size.

d. 50 GB and 73 GB drives supported in microcode revision level 5x66 and higher.

e. The 73 GB drive does not support the RAID-S configuration on Symmetrix.

Individual hosts might require specific cylinder counts for their file systems. Refer to the chapter in this manual for your specific host to determine its cylinder requirements.

The number of cylinders per volume for each physical drive is summarized in the following tables.

Table 2-2 Maximum Cylinders per Logical Volume for Each Device

Physical Drive	Logical Volumes	Maximum Cylinders per Volume
Microcode Revision Level 5566 or higher:		
73 GB	5	32,510/32,510/32,510/32,510/19,023
50 GB	4	32,510/32,510/32,510/3,646
Microcode Revision Level 5265 or higher:		
47 GB	3	32,510/32,510/30,598
36 GB	3	32,510/32,510/8,638
23 GB	2	32,510/14,454
18 GB	2	32,510/4,318
Microcode Revision Level 5264 or lower:		
47 GB	5	20,478/20,478/20,478/20,478/13,342
36 GB	4	20,478/20,478/20,478/12,222
23 GB	3	20,478/20,478/6,006
18 GB	2	20,478/16,350

The cylinder counts shown in the tables are for Symmetrix units attached to UNIX and/or PC hosts only. These cylinder counts are different for Symmetrix units attached to AS/400 hosts or both AS/400 and UNIX hosts.

Tables 2-3 through 2-10 which follow show the number of cylinders for each ratio of logical volumes to physical device (1:1, 2:1, 3:1, etc. splits) for microcode revision level 5265 (5566 or higher for the 50 GB and 73 GB drives). This is depicting the maximum cylinder counts, for 'n' number of splits.

Some Symmetrix units presently in use are configured with logical volumes having more cylinders than those listed in the tables. When updating Symmetrix microcode or adding volumes to these Symmetrix units, the EMC Product Support Engineer performing the upgrade will preserve the existing volume sizes.

To utilize full disk capacity, the 18, 23, 36, 47, 50 and 73 GB drives should be configured with the minimum number of supported logical volumes shown in the tables.

Table 2-3 Cylinders per Logical Volume - 4 GB Physical Device

Number of Logical Volumes per Disk Device	Cylinders per Logical Volume	GB ₁₀ per Logical Volume ^a	GB ₂ per Logical Volume ^a	Number of Logical Volumes per Disk Device	Cylinders per Logical Volume	GB ₁₀ per Logical Volume ^a	GB ₂ per Logical Volume ^a
1	8,672	4.26	3.97	17	508	0.25	0.23
2	4,335	2.13	1.98	18	480	0.23	0.22
3	2,889	1.42	1.32	19	455	0.22	0.20
4	2,167	1.06	0.99	20	432	0.21	0.19
5	1,733	0.85	0.79	21	411	0.20	0.18
6	1,444	0.71	0.66	22	392	0.19	0.18
7	1,237	0.60	0.56	23	375	0.18	0.17
8	1,082	0.53	0.49	24	359	0.17	0.16
9	962	0.47	0.44	25	345	0.17	0.15
10	865	0.42	0.39	26	332	0.16	0.15
11	787	0.38	0.36	27	319	0.15	0.14
12	721	0.35	0.33	28	308	0.15	0.14
13	665	0.32	0.30	29	297	0.14	0.13
14	618	0.30	0.28	30	287	0.14	0.13
15	576	0.28	0.26	31	278	0.13	0.12
16	540	0.26	0.24	32	269	0.13	0.12

- a. GB₂ values are based on the convention 1 GB = 1024 × 1024 × 1024.
 GB₁₀ values are based on the convention 1 GB = 1000 × 1000 × 1000. Although these values are expressed differently, their values are equivalent. The operating systems use binary values for logical volume configuration.

Table 2-4 Cylinders per Logical Volume - 9 GB Physical Device (5.25")

Number of Logical Volumes per Disk Device	Cylinders per Logical Volume	GB ₁₀ per Logical Volume ^a	GB ₂ per Logical Volume ^a	Number of Logical Volumes per Disk Device	Cylinders per Logical Volume	GB ₁₀ per Logical Volume ^a	GB ₂ per Logical Volume ^a
1	17,726	8.71	8.11	17	1,041	0.51	0.47
2	8,862	4.35	4.05	18	983	0.48	0.45
3	5,907	2.90	2.70	19	931	0.45	0.42
4	4,430	2.17	2.02	20	884	0.43	0.40
5	3,544	1.74	1.62	21	842	0.41	0.38
6	2,953	1.45	1.35	22	804	0.39	0.36
7	2,531	1.24	1.15	23	769	0.37	0.35
8	2,214	1.08	1.01	24	737	0.36	0.33
9	1,968	0.96	0.90	25	707	0.34	0.32
10	1,771	0.87	0.81	26	680	0.33	0.31
11	1,610	0.79	0.73	27	655	0.32	0.30
12	1,475	0.72	0.67	28	631	0.31	0.28
13	1,362	0.66	0.62	29	609	0.29	0.27
14	1,264	0.62	0.57	30	589	0.28	0.27
15	1,180	0.58	0.54	31	570	0.28	0.26
16	1106	0.54	0.50	32	552	0.27	0.25

- a. GB₂ values are based on the convention 1 GB = 1024 x 1024 x 1024.
 GB₁₀ values are based on the convention 1 GB = 1000 x 1000 x 1000. Although these values are expressed differently, their values are equivalent. The operating systems use binary values for logical volume configuration.

Table 2-5 Cylinders per Logical Volume - 9 GB Physical Device (3.5")

Number of Logical Volumes per Disk Device	Cylinders per Logical Volume	GB ₁₀ per Logical Volume ^a	GB ₂ per Logical Volume ^a	Number of Logical Volumes per Disk Device	Cylinders per Logical Volume	GB ₁₀ per Logical Volume ^a	GB ₂ per Logical Volume ^a
1	18,414	9.05	8.42	17	1,081	0.53	0.49
2	9,206	4.52	4.21	18	1,021	0.50	0.46
3	6,137	3.01	2.80	19	967	0.47	0.44
4	4,602	2.26	2.10	20	919	0.45	0.42
5	3,681	1.80	1.68	21	875	0.43	0.40
6	3,067	1.50	1.40	22	835	0.41	0.38
7	2,629	1.29	1.20	23	799	0.39	0.36
8	2,300	1.13	1.05	24	765	0.37	0.35
9	2,044	1.00	0.93	25	735	0.36	0.33
10	1,840	0.90	0.84	26	706	0.34	0.32
11	1,672	0.82	0.76	27	680	0.33	0.31
12	1,533	0.75	0.70	28	656	0.32	0.30
13	1,415	0.69	0.64	29	633	0.31	0.29
14	1,313	0.64	0.60	30	612	0.30	0.28
15	1,226	0.60	0.56	31	592	0.29	0.27
16	1,149	0.56	0.52	32	574	0.28	0.26

- a. GB₂ values are based on the convention 1 GB = 1024 × 1024 × 1024.
 GB₁₀ values are based on the convention 1 GB = 1000 × 1000 × 1000. Although these values are expressed differently, their values are equivalent. The operating systems use binary values for logical volume configuration.

Table 2-6 Cylinders per Logical Volume - 18 GB Physical Devices

Number of Logical Volumes per Disk Device	Cylinders per Logical Volume	GB ₁₀ per Logical Volume ^a	GB ₂ per Logical Volume ^a	Number of Logical Volumes per Disk Device	Cylinders per Logical Volume	GB ₁₀ per Logical Volume ^a	GB ₂ per Logical Volume ^a
1 ^b	b	b	b	17	2,165	1.06	0.99
2	18,414	9.05	8.42	18	2,044	1.00	0.93
3	12,275	6.03	5.61	19	1,937	0.95	0.88
4	9,206	4.52	4.21	20	1,840	0.90	0.84
5	7,364	3.62	3.37	21	1,752	0.86	0.80
6	6,136	3.01	2.80	22	1,672	0.82	0.76
7	5,260	2.58	2.40	23	1,599	0.78	0.73
8	4,602	2.26	2.10	24	1,533	0.75	0.70
9	4,090	2.01	1.87	25	1,471	0.72	0.67
10	3,681	1.80	1.68	26	1,415	0.69	0.64
11	3,346	1.64	1.53	27	1,362	0.67	0.62
12	3,067	1.50	1.40	28	1,313	0.64	0.60
13	2,831	1.39	1.29	29	1,268	0.62	0.58
14	2,629	1.29	1.20	30	1,226	0.60	0.56
15	2,453	1.20	1.12	31	1,186	0.58	0.54
16	2,300	1.13	1.05	32	1,149	0.56	0.52

a. GB₂ values are based on the convention 1 GB = 1024 x 1024 x 1024.
 GB₁₀ values are based on the convention 1 GB = 1000 x 1000 x 1000. Although these values are expressed differently, their values are equivalent. The operating systems use binary values for logical volume configuration.

b. Unsupported. Maximum logical volume size is 32,510 cylinders.

Table 2-7 Cylinders per Logical Volume - 23 GB Physical Devices

Number of Logical Volumes per Disk Device	Cylinders per Logical Volume	GB ₁₀ per Logical Volume ^a	GB ₂ per Logical Volume ^a	Number of Logical Volumes per Disk Device	Cylinders per Logical Volume	GB ₁₀ per Logical Volume ^a	GB ₂ per Logical Volume ^a
1 ^b	b	b	b	17	2,761	1.35	1.26
2	23,482	11.54	10.74	18	2,607	1.28	1.19
3	15,654	7.69	7.16	19	2,470	1.21	1.13
4	11,740	5.77	5.37	20	2,346	1.15	1.07
5	9,392	4.61	4.29	21	2,235	1.09	1.02
6	7,826	3.84	3.58	22	2,133	1.04	0.97
7	6,708	3.29	3.07	23	2,040	1.00	0.93
8	5,869	2.88	2.68	24	1,955	0.96	0.89
9	5,217	2.56	2.38	25	1,877	0.92	0.85
10	4,695	2.30	2.14	26	1,804	0.88	0.82
11	4,268	2.09	1.95	27	1,738	0.85	0.79
12	3,912	1.92	1.79	28	1,675	0.82	0.76
13	3,611	1.77	1.65	29	1,618	0.79	0.74
14	3,353	1.64	1.53	30	1,564	0.76	0.71
15	3,129	1.53	1.43	31	1,513	0.74	0.69
16	2,934	1.44	1.34	32	1,466	0.72	0.67

- a. GB₂ values are based on the convention 1 GB = 1024 × 1024 × 1024.
 GB₁₀ values are based on the convention 1 GB = 1000 × 1000 × 1000. Although these values are expressed differently, their values are equivalent. The operating systems use binary values for logical volume configuration.
- b. Unsupported. Maximum logical volume size is 32,510 cylinders.

Table 2-8 Cylinders per Logical Volume - 36 GB Physical Devices

Number of Logical Volumes per Disk Device	Cylinders per Logical Volume	GB ₁₀ per Logical Volume ^a	GB ₂ per Logical Volume ^a	Number of Logical Volumes per Disk Device	Cylinders per Logical Volume	GB ₁₀ per Logical Volume ^a	GB ₂ per Logical Volume ^a
1 ^b	b	b	b	17	4,331	2.12	1.98
2 ^b	b	b	b	18	4,090	2.01	1.87
3	24,553	12.06	11.23	19	3,875	1.90	1.77
4	18,414	9.05	8.42	20	3,681	1.80	1.68
5	14,731	7.24	6.74	21	3,506	1.72	1.60
6	12,275	6.03	5.61	22	3,346	1.64	1.53
7	10,521	5.17	4.81	23	3,201	1.57	1.46
8	9,206	4.52	4.21	24	3,067	1.50	1.40
9	8,183	4.02	3.74	25	2,945	1.44	1.34
10	7,364	3.62	3.37	26	2,831	1.39	1.29
11	6,695	3.29	3.06	27	2,726	1.34	1.24
12	6,137	3.01	2.80	28	2,629	1.29	1.20
13	5,664	2.78	2.59	29	2,538	1.24	1.16
14	5,260	2.58	2.40	30	2,453	1.20	1.12
15	4,909	2.41	2.24	31	2,374	1.16	1.08
16	4,602	2.26	2.10	32	2,300	1.13	1.05

a. GB₂ values are based on the convention 1 GB = 1024 x 1024 x 1024.

GB₁₀ values are based on the convention 1 GB = 1000 x 1000 x 1000. Although these values are expressed differently, their values are equivalent. The operating systems use binary values for logical volume configuration.

b. Unsupported. Maximum logical volume size is 32,510 cylinders.

Table 2-9 Cylinders per Logical Volume - 47 GB Physical Devices

Number of Logical Volumes per Disk Device	Cylinders per Logical Volume	GB ₁₀ per Logical Volume ^a	GB ₂ per Logical Volume ^a	Number of Logical Volumes per Disk Device	Cylinders per Logical Volume	GB ₁₀ per Logical Volume ^a	GB ₂ per Logical Volume ^a
1 ^b	b	b	b	17	5,602	2.75	2.56
2 ^b	b	b	b	18	5290	2.60	2.42
3	31,753	15.60	14.53	19	5012	2.46	2.29
4	23,814	11.70	10.90	20	4761	2.34	2.18
5	19,050	9.36	8.72	21	4534	2.22	2.07
6	15,875	7.80	7.26	22	4328	2.12	1.98
7	13,607	6.68	6.22	23	4140	2.03	1.89
8	11,906	5.85	5.45	24	3967	1.95	1.81
9	10,583	5.20	4.84	25	3809	1.87	1.74
10	9,524	4.68	4.36	26	3662	1.80	1.67
11	8,658	4.25	3.96	27	3526	1.73	1.61
12	7,937	3.90	3.63	28	3400	1.67	1.55
13	7,326	3.60	3.35	29	3283	1.61	1.50
14	6,803	3.34	3.11	30	3173	1.56	1.45
15	6,349	3.12	2.90	31	3071	1.50	1.40
16	5,952	2.92	2.72	32	2975	1.46	1.36

a. GB₂ values are based on the convention 1 GB = 1024 × 1024 × 1024.
 GB₁₀ values are based on the convention 1 GB = 1000 × 1000 × 1000. Although these values are expressed differently, their values are equivalent. The operating systems use binary values for logical volume configuration.

b. Unsupported. Maximum logical volume size is 32,510 cylinders.

Table 2-10 Cylinders per Logical Volume - 50 GB Physical Devices^a

Number of Logical Volumes per Disk Device	Cylinders per Logical Volume	GB ₁₀ per Logical Volume ^b	GB ₂ per Logical Volume ^b	Number of Logical Volumes per Disk Device	Cylinders per Logical Volume	GB ₁₀ per Logical Volume ^b	GB ₂ per Logical Volume ^b
1 ^c	c	c	c	17	5,954	2.92	2.72
2 ^c	c	c	c	18	5,624	2.76	2.57
3 ^c	c	c	c	19	5,328	2.61	2.43
4	25,314	12.44	11.58	20	5,060	2.48	2.31
5	20,250	9.95	9.27	21	4,820	2.36	2.20
6	16,876	8.29	7.72	22	4,600	2.26	2.10
7	14,464	7.11	6.62	23	4,400	2.16	2.01
8	12,656	6.22	5.79	24	4,216	2.07	1.93
9	11,250	5.52	5.15	25	4,048	1.99	1.85
10	10,124	4.97	4.63	26	3,892	1.91	1.78
11	9,204	4.52	4.21	27	3,748	1.84	1.71
12	8,436	4.14	3.86	28	3,614	1.77	1.65
13	7,788	3.82	3.56	29	3,490	1.71	1.59
14	7,230	3.55	3.31	30	3,374	1.65	1.54
15	6,748	3.31	3.08	31	3,264	1.60	1.49
16	6,326	3.11	2.89	32	3,162	1.55	1.44

a. Supported by microcode rev level 5x66 or higher.

b. GB₂ values are based on the convention 1 GB = 1024 x 1024 x 1024.

GB₁₀ values are based on the convention 1 GB = 1000 x 1000 x 1000. Although these values are expressed differently, their values are equivalent. The operating systems use binary values for logical volume configuration.

c. Unsupported. Maximum logical volume size is 32,510 cylinders.

Table 2-11 Cylinders per Logical Volume - 73 GB^a Physical Devices^b

Number of Logical Volumes per Disk Device	Cylinders per Logical Volume	GB ₁₀ per Logical Volume ^c	GB ₂ per Logical Volume ^b	Number of Logical Volumes per Disk Device	Cylinders per Logical Volume	GB ₁₀ per Logical Volume ^b	GB ₂ per Logical Volume ^b
1 ^d	d	d	d	17	8,747	4.004	4,299
2 ^d	d	d	d	18	8,251	3.782	4,060
3 ^d	d	d	d	19	7,826	3.582	3,847
4 ^d	d	d	d	20	7,435	3.403	3,654
5	29,746	13.617	14,621	21	7,080	3.241	3,480
6	24,788	11.347	12,184	22	6,758	3.094	3,322
7	21,246	9.726	10,443	23	6,465	2.959	3,178
8	18,590	8.510	9,137	24	6,195	2.836	3,045
9	16,524	7.564	8,122	25	5,947	2.722	2,923
10	14,872	6.808	7,310	26	5,719	2.618	2,811
11	13,520	6.189	6,645	27	5,507	2.521	2,707
12	12,393	5.673	6,091	28	5,310	2.431	2,610
13	11,439	5.236	5,622	29	5,127	2.347	2,520
14	10,622	4.862	5,221	30	4,956	2,269	2,436
15	9,914	4.538	4,873	31	4,796	2,195	2,357
16	9,294	4.254	4,566	32	4,646	2,127	2,284

- a. The 73 GB drive does not support the RAID-S configuration on Symmetrix.
- b. Supported by microcode rev level 5x66 or higher.
- c. GB₂ values are based on the convention 1 GB = 1024 x 1024 x 1024.
GB₁₀ values are based on the convention 1 GB = 1000 x 1000 x 1000. Although these values are expressed differently, their values are equivalent. The operating systems use binary values for logical volume configuration.
- d. Unsupported. Maximum logical volume size is 32,510 cylinders.

Hyper-Volume Extension

The Hyper-Volume Extension feature provides configuration flexibility by removing physical device boundary limitations. The Hyper-Volume Extension feature allows you to split one physical device into two or more logical volumes.

Symmetrix models with 5.25-inch disks use 4 GB, 9 GB, 18 GB, 23 GB, or 47 GB disk devices. See Table 2-1 for maximum number of cylinders for each drive.

Disk device capacity will vary from manufacturer to manufacturer. Please refer to Table 2-1 through Table 2-9 for specific device capacities supported in this environment.

Symmetrix models with 3.5-inch disks use 4 GB, 9 GB, 18 GB, 36 GB, 50 GB or 73 GB disk devices. See Table 2-1 for maximum number of cylinders for each drive.

Each cylinder on a Symmetrix unit has 15 tracks. Each track contains 64 blocks, each of which contains 512 bytes.

Multiple Logical Volumes per Single Device

When splitting a single physical device into multiple logical volumes, Symmetrix allows logical-to- physical relationships indicated in the following table. EMC, however, recommends specific logical volume splits for each host environment as listed in each chapter of this product guide.

Table 2-12 Disk Device Configuration

		Symmetrix Logical Volume to Physical Splits								
		1:1	2:1	3:1	4:1	5:1	6:1	7:1	⇒	32:1 ^a
Physical Devices	4 GB	x	x							
	9 GB	x	x	x	x					
	18 GB		x	x	x	x	x	x	⇒	x
	23 GB		* ^b	x	x	x	x	x	⇒	x
	36 GB			*	x	x	x	x	⇒	x
	47 GB			*	*	x	x	x	⇒	x
	50 GB ^c				x	x	x	x	⇒	x
	73 GB ^{c, d}					x	x	x	⇒	x

- a. With microcode revision level 5263, maximum is 8:1. With 5264, maximum is 16:1. With 5265, possible maximum is 32:1. EMC recommends a maximum of 8:1 for best performance.
- b. * indicates ratios available only in microcode revision level 5265 or higher.
- c. Supported by microcode revision level 5x66 or higher.
- d. The 73 GB drive does not support the RAID-S configuration on Symmetrix.

At microcode level 5265 and later, the maximum logical volume size on any physical drive is 32,510 cylinders. To utilize full disk capacity, the 18, 23, 36, 47, 50 and 73 GB disk devices should be configured with the minimum number of supported logical volumes shown in the table above.

Symmetrix supports up to 4,096 logical volumes (address range 000 to 7FF) per unit. The logical-to-physical relationship chosen can automatically apply to all devices in the unit. It is also possible to customize the logical-to-physical relationship on each device as well as the size of each logical volume.

Although a Symmetrix system can support up to 4,096 logical volumes, the allocation per SCSI director is limited to 960 logical volumes.

In the example illustrated in Figure 2-1, the logical-to-physical ratio chosen is 8:1. The logical volume mapping that occurs appears similar to that shown in Figure 2-1.




Figure 2-1 3:1 Logical Volume Mapping Example

Meta Volumes

Several operating systems, such as Windows NT, some applications software, and some open systems environments, require larger volumes than are provided by standard Symmetrix physical disk devices.

A meta volume is a logical volume set created from individual physical disks to define volumes larger than the current Symmetrix maximum Hyper-Volume size of approximately 16 GB. Meta volumes are functionally the same as logical volume sets implemented with host volume manager software.

Physically, a meta volume is two or more Symmetrix Hyper-Volumes presented to a host as a single addressable device. It consists of a "head device," some number of "member devices" (optional), and a "tail device."

Meta Volume Size Requirements

Symmetrix meta volumes may contain as many as 255 devices and up to 3.825 terabytes in size. Meta volumes can be composed of non-sequential and non-adjacent volumes.

Meta Volume Performance

By allowing individual physical disk devices to be grouped together into a meta volume and the capability of meta volume addressing, Symmetrix enhances disk system functionality. To increase throughput and further improve performance, Symmetrix provides multiple I/O drive queues for meta volumes.

Accessing Data in a Meta Volume

Addressing of data contained in a meta volume can be organized in two different ways:

- ◆ Concatenated volumes
- ◆ Striped data

Concatenated Volumes

Concatenated volumes are volume sets that are organized with the first byte of data at the beginning of the first volume (Figure 2-2). Addressing continues to the end of the first volume before any data on the next volume is referenced. When writing to a concatenated volume, the first slice of a physical disk device is filled, then the second and so on, and on to subsequent physical disk devices.




Figure 2-2 Concatenated Volumes

Striped Data

Meta volume addressing by striping also joins multiple slices to form a single volume. However, instead of sequential address space, addresses are interleaved between slices (Figure 2-3). When writing to a striped volume, equal size stripes of data from each participating drive are written alternately to each member of the set.

Configuring Meta Volumes

Although members of a striped set do not have to all be the same size, the effective size of each member is the actual size of the smallest member rounded down to an even cylinder count. Therefore, care should be taken when configuring such a Meta volume, to minimize wasted space.




Figure 2-3 Striped Data

Striping data across the multiple drives in definable cylinder stripes was designed to benefit “random reads” by avoiding stacking multiple reads to a single spindle and disk director. This scheme creates a large volume, but additionally balances the I/O activity between the disk devices and the Symmetrix disk directors.

Table 2-13 defines the data stripe sizes and capacities for Symmetrix Meta Volumes.

Table 2-13 Meta Volume Stripe Sizes

Stripe Size	Stripe Capacity (KB)
2 Cylinders	960
4 Cylinders	1,920
8 Cylinders	3,840
16 Cylinders	7,680
32 Cylinders	15,360
64 Cylinders	30,720

Performance Tools

This chapter covers host-based performance features and tools.

- ◆ UNIX Performance Tools 3-2

UNIX Performance Tools

This section describes tools and utilities used to monitor performance from a UNIX system console.

sar

sar is a useful UNIX utility to monitor I/O statistics such as IOs per second, blocks per second, average queue length, calculated service time, etc.

For example:

```
sar -d 5 10
```

collects disk statistics in 5 second intervals for 10 sampling intervals. sar produces a statistic summary similar to the one shown on the following page for the entire sampling period.

Table 3-1 sar Statistic Summary

12:23:45	device	%busy	avque	r+w/s	blks/s	await	avserv
12:23:50	qd30	44	1.4	94	1974	1.8	4.6
	qd31	33	1.2	95	1977	0.7	3.5
	qd36	43	1.2	94	1980	1.1	4.5
	qd37	32	1.2	94	1961	0.6	3.4
12:23:55	qd0	0	1.0	0	2	0.0	12.1
	qd30	48	1.4	113	1798	1.7	4.3
	qd31	34	1.2	93	1721	0.8	3.7
	qd36	49	1.3	113	1782	1.4	4.3
	qd37	34	1.2	93	1734	0.8	3.7
	qd48	0	1.0	0	0	0.0	11.1
	qd60	0	1.0	0	7	0.0	5.5
Average	qd0	0	1.0	0	1	0.0	12.1
	qd30	46	1.4	103	1886	1.7	4.4
	qd31	34	1.2	94	1849	0.7	3.6
	qd36	46	1.3	104	1881	1.3	4.4
	qd37	33	1.2	94	1848	0.7	3.6
	qd48	0	1.0	0	0	0.0	11.1
	qd60	0	1.0	0	3	0.0	5.5

For more information on this utility, refer to the man pages for sar or to the operating system manuals for your particular host.

vmstat

The UNIX utility, `vmstat` collects and reports system and I/O statistics.

For more information on this utility, refer to the man pages for `vmstat` or to the operating system manuals for your particular host.

iostat

The UNIX utility, `iostat` interactively reports I/O statistics for each active disk on the system. Disk data is arranged in a four-column format.

Column Heading	Interpretation
device	Device name
bps	Kilobytes transferred per second
sps	Number of seeks per second
msps	Milliseconds per average seek

If two or more disks are present, data is presented on successive lines for each disk.

For example, to view the I/O statistics for all disks every 5 seconds until INTERRUPT or QUIT is pressed, type:

```
iostat 5
```

To view the I/O statistics for all disks every 5 seconds and terminate after 10 readings, type:

```
iostat 5 10
```

For more information on this utility, refer to the man pages for `iostat` or to the operating system manuals for your particular host.

This chapter describes how to incorporate Symmetrix in the NCR environment.

◆ Overview	4-2
◆ Useful NCR Utilities and Functions	4-4
◆ Symmetrix Configuration	4-6
◆ Incorporating Symmetrix.....	4-7
◆ Displaying the New Devices	4-10
◆ Configuring New Devices	4-11
◆ Adding Devices Online	4-18
◆ System and Error Messages.....	4-19

Overview

When using Symmetrix with an NCR host, Table 4-1 below summarizes operating system features in the NCR environment.

The EMC Customer Engineer (CE) should contact the EMC Configuration Specialist for updated online information.

Table 4-1 NCR Operating System Features

O/S	NCR UNIX SVR4 below MP_RAS 3.0.1	NCR UNIX SVR4 MP_RAS 3.2
Largest recognizable Symmetrix volume size	4 GB (8,388,608 ₁₀ blks)	Up to 1 TB
Largest defined slice	2 GB (4,194,304 ₁₀ blks)	Up to 1 TB
Largest filesystem	2 GB (4,194,304 ₁₀ blks)	Up to 1 TB
Supports multiple target IDs	Yes	Yes
Supports multiple LUNs per target ID ^{a, b}	Yes	Yes
Partition representing the entire disk	s0	s0
To see the drive configuration, use this utility:	prtconf	prtconf
To configure drives use this utility:	sysadm	sysadm
Special drive definition files needed	yes	yes

a. When specifying multiple LUNs, the first LUN must be LUN 0. However, the LUNs specified need not be consecutive.

b. If needed, an NCR host can support up to 16 LUNs per target ID.

A Symmetrix installation on an NCR host requires a kernel rebuild with the EMC Symmetrix definition.

- ◆ The NCR default SCSI disk definition does not support wide, queuing, and synchronous communication between the host and the disk. It also does not support LUNs. However, with minimal modifications to NCR structures and a kernel rebuild, an NCR host will recognize an EMC Symmetrix device and will support all these features.

- ◆ If a power failure occurs that takes down both the host and the Symmetrix unit, please note that it will take Symmetrix longer to become ready after the host is available. This time can range from a few minutes, up to 25 minutes depending on the Symmetrix configuration.
- ◆ For host controller cards which support attachment to Symmetrix units, see Appendix C.
- ◆ For an exact volume size of 4 GB, specify 8,738 Symmetrix cylinders. For an exact volume size of 2 GB, specify 4,369 Symmetrix cylinders.

UNIX MP-RAS level 3.0.1 and later supports the full 9GB limit of 9,206 Symmetrix cylinders

Useful NCR Utilities and Functions

This section lists NCR functions and utilities you may use to define and manage Symmetrix devices. The use of these functions and utilities is optional. They are listed for reference only.

- ◆ `sysadm`

A menu-driven utility through which devices can be configured and maintained.

- ◆ `prtconf`

Displays the current system configuration.

- ◆ `fdisk`

Partitions a disk and creates/updates the master boot block (block 0). A disk can be partitioned into DOS, UNIX, and other partitions.

- ◆ `dklayout`

Initializes a UNIX partition with boot code, physical device information structure, and volume table of contents (VTOC) structure.

- ◆ `prtvto`

Displays the VTOC structure of a device. The VTOC structure contains the slicing information for the UNIX partition. This partition is divided into 16 slices that are used to describe the total partition capacity.

- ◆ `edvtoc`

Enables you to update the volume table of contents (VTOC). This structure contains the slicing information for the UNIX partition. Use this utility if you need to change only the VTOC.



CAUTION

Do NOT use edvtoc on a drive with live data!

- ◆ `disksetup`

An interactive utility that enables the user to setup the UNIX partition. This partition *must* be defined using `fdisk` and `dklayout`.

- ◆ `shutdown -g0 -i6 -y`

Command that properly reboots the operating system.

Symmetrix Configuration

In the NCR host environment, you can configure the Symmetrix disk devices into logical volumes. EMC recommends the following logical volume ratios as indicated in the table below:

Table 4-2 Disk Device Configuration

		Symmetrix Logical Volume to Physical Splits ^a							
		1:1	2:1	3:1	4:1	5:1	6:1	7:1	8:1 ^b
Physical Devices	4 GB	x	x						
	9 GB	x	x	x	x				
	18 GB		x	x	x	x	x	x	x
	23 GB		*	x	x	x	x	x	x
	36 GB			*	x	x	x	x	x
	47 GB			*	*	x	x	x	x
	50 GB ^c				x	x	x	x	x

- a. * indicates ratios available only in microcode level 5265 or higher.
- b. With microcode revision level 5263, maximum is 8:1. With 5264, maximum is 16:1. With 5265, maximum is 32:1. With 5266, maximum is 32:1
- c. The 50 GB drive is supported only by microcode revision level 5x66.

For specific cylinder counts and logical volume sizes, refer to Table 2-1 on page 2-5 and the following tables in that chapter.

Due to NCR restrictions in early revisions of MP-RAS NCR, EMC Symmetrix devices should not be configured to have more than 8,738 cylinders. MP-RAS NCR Version 3.0.1 and later support the full 9,206 cylinders available in Symmetrix 9 GB devices (with a 2:1 split on a 3.5-inch drive).

The EMC Customer Engineer (CE) should contact the EMC Configuration Specialist for updated online information. This information is necessary to configure Symmetrix to support the customer's host environment.

Incorporating Symmetrix

If you are using an NCR host running NCR UNIX SVR4 MP_RAS 3.0.1 or later, the default kernel contains the correct EMC Symmetrix device definitions. Please skip to the next section

To enable the NCR host to support the Symmetrix system, EMC provides a modified **space.c** file. Alternatively, should you wish to modify your existing **space.c** file to support Symmetrix devices, EMC also provides a device configuration file that can be inserted into your **space.c** file.

FTP the Symmetrix device definition files: **space.README** and **space.c.emc** to your NCR host. These files are available on EMC's FTP server, ftp.emc.com, in **/pub/symm3000/ncr**. Make the modifications listed below.

See *Symmetrix Device Definition Files* on page A-1 for instructions on how to transfer these files to your host.

1. Power on the NCR host and wait for it to perform its self-tests.
2. When the self-tests complete and the login prompt appears, log in as **root**.
3. If you installed new host controller cards into your host prior to installing and connecting Symmetrix to your host, make sure the host controller cards are configured to the system. Use the **sysadm** utility to perform this check.

Pay attention to any host message that appears regarding these new host controller cards at boot time. Failing to properly install new host controller cards will prevent the host from accessing any devices attached to these host controller cards.

4. Change directory to **/etc/conf/pack.d/SCSI**.
5. Make a copy of the original **space.c** file and name it **space.old**.



CAUTION

If the current **space.c file on your host has been modified, do NOT perform step 6.**

6. Copy the EMC files to your current directory. If you are going to use the modified `space.c` file provided by EMC, rename **space.c.emc** to **space.c**. Otherwise, continue with the steps below.
7. Modify the **space.c** file as follows:

The necessary Symmetrix device definitions are included in the file **space.README**.

- Make sure the following definition is FALSE:

```
int ParallelBusSynchronousSupport = FALSE
```

- Search for the default device configuration,

```
SCSI_Device_Configuration_t Seagate_Configuration[]...
```

- Insert the EMC device configuration definition immediately following the “;” of this entry. The EMC device configuration section has the following format:

```
SCSI_Device_Configuration_t EMC_Configuration[]=
{
  TRUE, 0xff, 0xff, 0x10, FALSE, 0x08, &cscan_sort, 0x0f, FALSE /* EMC - SYMM*/
};
```

Make sure this parameter is set to TRUE

Both parameters must be set to 0xff to support synchronous transfers

Data Width (use "10" with Quad Adapter (Q720))

Queue size

16 Supported LUNs

- Add the EMC device identification table to your **space.c** file following the definition for Seagate:

Search for

```
DeviceId_t DeviceIdentificationTable
```

Scroll down to the end of the Seagate structure. Insert the EMC structure (provided in **space.README**) immediately following the “;” of the Seagate structure.

8. Rebuild the kernel for these changes to take effect.
 - a. Rebuild the kernel by executing the command below. `idbuild` will make a copy of the old kernel and put it under **/stand/unix.old**.

```
/etc/conf/bin/idbuild
```

Consult your System Administration manuals for additional information on the kernel rebuild procedure for your system if necessary.

- b. Reboot the system.

Displaying the New Devices

Use the following procedure to verify that the host can see all the Symmetrix devices assigned to that host.

1. Log in as **root**.
2. Type `prtconf > sys_config`. This creates a file called **sys_config** that contains the system configuration information.
3. Type `more sys_config` to view the system configuration or print this file for a hardcopy listing of all devices seen by the host.

Configuring New Devices

In order to use the Symmetrix devices, you must define partitions and slices. To define partitions and slices, you can use one of the following procedures:

- ◆ Use the `sysadm` utility described below
- ◆ Use the EMC-supplied script, `diskconfig.NCR`
- ◆ Use standalone commands

Using `sysadm`

Follow the procedure in this section to define partitions and slices using the `sysadm` utility.

Partitioning the Devices

To partition Symmetrix devices follow the procedure described below. You may also define partitions using individual commands (described later).

1. At the root prompt, type `sysadm` to access the **System Administration** menu (window 1).

The softkey blocks on the bottom of the screen display correspond with the F1 through F8 key actions.

2. Press the TAB key until the **storage_devices** option highlights in the **UNIX System V Administration** window. Press <CR>.
3. Choose the **device_config** option in the **Storage Device Operations** window (window 2).
4. Choose the **partition** option in the **Partition Devices** window (window 3).
5. Press <F2> **Choices** to display the devices the host sees on its SCSI bus.
6. Using the TAB key, highlight the external device to partition and press <CR>. For example, choose **c1t0d0s0**.
7. Press <F3> **Save**.
8. At the warning that you are modifying existing partitions, etc., press <F3> **Continue**.
9. If the device has never been partitioned, a prompt appears asking you to choose to use the whole disk or further partition the drive. Type `Y` <CR> to use the whole drive.

10. Press <F3> **Continue** to partition additional Symmetrix devices.
11. Repeat steps 5 through 10 for each Symmetrix device.

To configure a large number of Symmetrix devices, consider using the `dskconfig.NCR` script. See *Configuring New Devices Using NCR Script* on page 4-14.

Slicing the Symmetrix Devices

After you partition the Symmetrix devices, you need to specify the slices on the volume.

1. Return to the **Device Configuration and Maintenance** window (window 3) by pressing <F6> **Cancel** until that window is active.
2. Choose the **slice** option from the **Device Configuration and Maintenance** window.
3. Press <F2> **Choices** to display the devices the host sees on its SCSI bus.
4. Using the TAB key, highlight the external device to slice and press <CR>. The Slice screen of the selected device appears. For example, choose **c1t0d0s0**.
5. When prompted for the slice to modify, choose the logical disk device you want to use and the slice starting point. For example, choose **"slice 1"**.

A filesystem cannot be created larger than 2 GB (4,194,304 blocks). Refer to the recommended slicing tables for each Symmetrix disk size located in EMC's ftp server, <ftp.emc.com>, in `/pub/symm3000/ncr`.

Do not modify slice 0 or slice 7 on each LUN device.

6. Enter the size of the slice. For example, type **4194303**. This creates a 2 GB slice on the device.
7. Enter the starting block number.

The starting block number for a slice must be a block not used by any other slice on that physical device.

EMC recommends that you start each slice on 64-block multiples. (This is different than the default NCR slicing.) For example, if there are two 2 GB slices on a device, slice 1 starts at block 64 and is 4,194,240 blocks in size; slice 2 starts at block 4,194,305.

8. Press **x** to exit and update the device.

If you want to create additional slices on the Symmetrix device, repeat steps 5 through 7 specifying the appropriate slice to edit.

9. At the “**To re-edit?**” prompt, press **N**.
10. Repeat steps 2 through 9 for each additional Symmetrix device to partition.
11. After the host completes the slicing operation (this is indicated by an “**Update of table entries complete**” message) on the last Symmetrix device, press **<F6> Cancel** to return to the **System Administration** window.

The Symmetrix devices are now available to the host as raw devices. To create filesystem, continue with the next section “Creating Filesystems”.

Creating Filesystems

Now that the Symmetrix raw devices are available to the operating system, a new filesystem can be created for each Symmetrix device. To do this, use the “file_systems” option in the **System Administration** window.

1. Choose the **file_systems** option from the **System Administration** window.
2. Choose the **make** option from the **Manage File Systems** window.
3. Choose the **create** option from the **Create, Copy, or Move** Filesystems window.
4. Press **<F2> Choices** to display the devices the host sees on its SCSI bus.
5. Using the TAB key, highlight the external device on which to create a filesystem and press **<CR>**. For example, choose **c1t0d0s0**.
6. At the “Area that will contain the filesystem:” prompt, press **<F2> Choices** to display the devices. Choose the appropriate value and press **<CR>**.
7. At the “Filesystem type:” prompt, press **<F2> Choices** to display the various filesystem types available. Choose an appropriate type. Use **ufs** if you are unsure about which type to use.
8. At the “Label for filesystem:” prompt, press **<CR>** or enter a label.

9. At the “Once created, should the new filesystem be mounted?” prompt, press <CR> to accept the default value of **yes**.
10. At the “Filesystem name when mounted:” prompt, enter a name for the mount directory. For example, type **/syemm1**.
11. Press <F3> **Save** to save this information.
12. The host displays the information for your particular filesystem. For an s5 filesystem, it gives you the name of a prototype file (typically blank), the block size in bytes (default 1024), the number of blocks (slice size), and the number of files (in inodes). For a ufs filesystem, it gives you the number of blocks in filesystem, block size in bytes (default 4096), and fragment size in bytes (default 1024).
13. Press <F3> **Save**. The host creates the specified filesystem.
14. Repeat steps 4 through 13 for each additional filesystem to create.

Configuring New Devices Using NCR Script

The `dskconfig.NCR` script partitions, slices, and makes optimized filesystems on the NCR host once Symmetrix has been configured.

The script provides a number of options that allow you to perform various tasks. Refer to the syntax and examples listed below for more detail on these options.

```
dskconfig.NCR [-n] [-e] [-c] [-S] [-f file] [-s size] [-F type] <device name>
```

- | | |
|----|---|
| -n | Do not make filesystems. |
| -e | Print out syntax examples. |
| -c | Set this parameter if you want to create a slice file; you must specify -S <type> as well or the slice file is not created. |
| -f | Specify a slice file to be used. |
| -s | Specified the preferred slice size. (This is meaningful only when this option is used with -S.) |
| -S | For automatic slicing (uses a default slice of 2 GB). |
| -F | Specify filesystem type. |

Unless -S is used, **dskconfig.NCR** assumes there is a “slices” files in your current directory. Make sure this “slices” file has been edited appropriately to match your disk size, or use the -S option if you are unsure.

Example 1 will create ufs filesystems, mount points, and mount the filesystems for all slices (defined in a pre-existing “slices” file) on c1t0d0s0, except slice 0 and 7:

```
dskconfig.NCR -F ufs c1t0d0s0
```

Example 2 will partition and slice (but make no filesystems) the device based on the file, slices.new, in the current directory:

```
dskconfig.NCR -f slices.new -n c1t0d0s0
```

Example 3 will autosize the device and create a default “slices” table for it. It will ONLY create this “slices” file:

```
dskconfig.NCR -c -S c1t0d0s0
```

Example 4 will autosize the device and use a slice size of 2000000 blocks to create a “slices” file only:

```
dskconfig.NCR -c -S 2000000 -S c1t0d0s0
```

Example 5 will autosize, label, partition, and create ufs filesystems, mount points, and mount the filesystems using the slice defaults:

```
dskconfig.NCR -S c1t0d0s0
```

Configuring Devices from the Command Line

This section describes how to configure Symmetrix devices on the NCR host using the command line interface.

All procedures described in this section use c3t0d0s0 as the device and c3t0d0s1 as the slice for its examples.

Creating the Device File

At system boot, the `mktable` command automatically executes. This command will create device files (e.g., `/dev/rdsk/c1t0d0s0`) for all existing slices. For any new device it creates default slices 0 and 7. If you choose to create your own device files using “nodes”, be sure to run `mktable` at some time to update the system’s knowledge of the new slices.

To create a device file (e.g., `/dev/rdsk/c3t0d0s1`), use the following steps.

It is not necessary to create device files for every device. You may follow steps 1 through 6 in the section *Configuring a Filesystem Step-by-Step* for each new device, then run `/usr/sadm/sysadm/bin/mktable`. This will update system information as well as create the device files for your newly configured slices. (`mktable` may take several minutes to execute.)

1. Enter the following at the command prompt:

```
/usr/sadm/sysadm/bin/nodes -d /dev/rdisk/c3t0d0s1
```

The system will respond with two numbers:

`device_major` (i.e., 35) and `device_minor` (i.e., 49154).

2. Use these two numbers in the following commands:

```
/sbin/mknod /dev/rdisk/c3t0d0s1 c device_major device_minor
creates a character device file.
```

```
/sbin/mknod /dev/dsk/c3t0d0s1 b device_major device_minor
creates a block device file.
```

Configuring a Filesystem Step-by-Step

To create filesystem from a raw device, follow the steps below.

1. Create a 100% UNIX partition by typing:

```
fdisk -c -a -u 100 /dev/rdisk/c3t0d0s0
```

2. Initialize and VTOC the disk by typing:

```
dklayout -b -p -v -d /dev/rdisk/c3t0d0s0
```

3. Print the existing partition table from the disk on the screen by typing:

```
prtvtoc /dev/rdisk/c3t0d0s0
```

4. Create a slice file called "slices" containing the disk's current configuration by typing:

```
prtvtoc -f slices /dev/rdisk/c3t0d0s0
```

5. Edit the “slices” file to suit your needs and then apply it as follows:

```
edvtoc -f slices /dev/rdisk/c3t0d0s0
```

You can use an EMC-supplied slicing table available in EMC’s ftp server, **ftp.emc.com**, in **/pub/symm3000/ncr**. Please note that EMC recommends that each slice start on a 64-block boundary.

6. Create a “ufs” filesystem using the parameters below.

```
mkfs -F <FStype> -o <nsect=64>,<ntrack=15>,<rps=120>,<gap=0> <device> <Filesize>
```

For example:

```
mkfs -F ufs -o nsect=64,ntrack=15,rps=120,gap=0 /dev/dsk/c3t0d0s1 4194240
```

The given parameters are optimal for a Symmetrix configuration. Make sure they are set to these values for Symmetrix configurations.

Refer to the **mkfs.README** file available in EMC’s ftp server in **/pub/symm3000/ncr** for more information on the mkfs command.

7. Mount the filesystem by typing:

```
mkdir /mount_point
mount -F ufs /dev/dsk/c3t0d0s1 /mount_point
```

To mount this filesystem automatically at system boot, you must add an entry into **/etc/vfstab**. Please see the man page for a description of the fields shown in the sample entry below.

```
/dev/dsk/c3t0d0s1 /dev/rdisk/c3t0d0s1 /mount_point ufs 1 yes rw,suid
```

The filesystem is now ready for use.

Adding Devices Online

Whenever devices are added online to the Symmetrix unit or device channel addresses are changed, you will need to perform the actions described below in order to introduce the new devices to the system.

When configuring new devices on a live system, it is best to minimize the amount of disk activity the host is performing.

To inform the host about new devices, run the `mktable` command as follows:

```
% /usr/sadm/sysadm/mktable
```

This command may take several minutes to execute (depending on the number of resources on the host). This command updates the host's table of resources. If not executed, `sysadm` or `dskconfig.NCR` may not be able to work with a new device.

Once `mktable` has completed, determine which devices are the new ones by comparing the old list of devices using the `prtconf` command. Continue with the instructions on page 4-11 to introduce the new devices to the host.

System and Error Messages

The NCR operating system logs SCSI error messages in the `/etc/osm` file. It later moves the errors in this log file to the `/etc/osm.old` file. It moves the log file daily to `/var/adm/streams/error.xx.yy` where `xx` is the month and `yy` is the day.

`/etc/init.d/osm` is a script that manages the error logging. It has an environment variable, `GPLOGGEROPTS`, assigned in `/etc/default/osm`, that affects how the log file entries are saved.

`GPLOGGEROPTS` has the following options:

- ◆ `GPLOGGEROPTS="-ao"` appends errors to the log rather than truncate, create an `.old` file daily, and doesn't limit the size of the log.
- ◆ `GPLOGGEROPTS="-ac"` appends, ages, and compresses the log file daily, but doesn't limit the size of the log file.
- ◆ `GPLOGGEROPTS="-f 4000 -ao"` limits the log file to the first 4000 blocks, appends errors to the log, and creates an `.old` file daily.
- ◆ `GPLOGGEROPTS="-l 1000 -aw"` limits the log file to the last 1000 blocks, appends errors to the log, and ages the log daily with a day of week suffix `.[0-6]`.
- ◆ `GPLOGGEROPTS="-d all=3 -ac"` compacts log entries, appends errors to the log, as well as ages and compresses the log daily.

To display a list of the error message files, type

```
cd /var/adm/streams
```

to go to the directory containing the daily log files and then `ls` for a listing of the log files.

The following is an example of an error log file list from November 9 through December 27:

error.11-09	error.11-25	error.12-04	error.12-13	error.12-22
error.11-10	error.11-26	error.12-05	error.12-14	error.12-23
error.11-11	error.11-27	error.12-06	error.12-15	error.12-24
error.11-12	error.11-28	error.12-07	error.12-16	error.12-25
error.11-19	error.11-29	error.12-08	error.12-17	error.12-26
error.11-20	error.11-30	error.12-09	error.12-18	error.12-27
error.11-21	error.12-01	error.12-10	error.12-19	
error.11-22	error.12-02	error.12-11	error.12-20	
error.11-23	error.12-03	error.12-12	error.12-21	

This chapter describes how to incorporate Symmetrix in the IBM NUMA-Q (formerly Sequent) environment.

◆ Overview	5-2
◆ Useful IBM NUMA-Q Utilities and Functions	5-5
◆ Symmetrix Configuration	5-6
◆ Incorporating Symmetrix.....	5-7
◆ Configuring the IBM NUMA-Q System	5-15
◆ Configuring /Deconfiguring Devices Online.....	5-18
◆ Formatting and Partitioning Definition Files.....	5-19
◆ Creating Physical Volumes	5-23
◆ IBM NUMA-Q Volume Manager.....	5-26
◆ System and Error Messages.....	5-40

Overview

This chapter describes the IBM NUMA-Q environment, but you may still find some of the documentation references under the former Sequent name.

When using Symmetrix with an IBM NUMA-Q UNIX server, please note the following:

- ◆ For Symmetrix models 8430 and 8730, EMC recommends that you use the Ultra-SCSI director P/N 201-330-907.
- ◆ For Symmetrix models 55/3500, 52/32xx, and 51/31xx, EMC recommends that you use the Fast-Wide Differential SCSI director P/N 200-881-903.
- ◆ For Symmetrix models 57/3700, 54/34xx, and 53/33xx, EMC recommends that you use either the Fast-Wide Differential SCSI director P/N 201-207-917 (60 MHz processor) or P/N 201-207-927 (66 MHz processor) or Ultra-Wide Differential SCSI director P/N 201-277-917 (60 MHz processor) or P/N 201-277-927 (66 MHz processor).



CAUTION

With DYNIX/ptx V4.4.2 and above, you can reconfigure LUNs for Symmetrix. EMC strongly recommends, however, that for all DYNIX/ptx operating systems below 4.4.2, you do not reconfigure your host to support LUNs.

- ◆ Symmetrix attaches to an IBM Wide SCSI-2 Controller (Model QICW-CTRL-01) (Symmetry Series) or to an FC to SCSI bridge with firmware V1.5.3 (NUMA-Q system) or higher.
- ◆ DYNIX/ptx 4.2.x and below are valid operating system levels for the Symmetry Series only. DYNIX/ptx 4.3.0, 4.3.1, 4.4.0 and 4.6.x are valid operating system levels for NUMA-Q only. DYNIX/ptx 4.4.1 up to 4.5.x is valid for both the Symmetry Series and NUMA-Q systems.
- ◆ When configuring Symmetrix with SCSI, do not use target IDs 6, 7, 14, or 15 for IBM Symmetry. For NUMA-Q SCSI, do not use target IDs 6 and 7. These target IDs are used for host initiators in the cluster environment.

During installation, the EMC Customer Engineer configures the volume sizes of the Symmetrix devices based on the number of available cylinders specified in Table 2-1 on page 2-5 and the following tables in that chapter.

Table 5-1 below summarizes the IBM NUMA-Q environment when it is used as a UNIX server running DYNIX/ptx 4.2.x and below:

Table 5-1 DYNIX/ptx 4.2.x and below Host Environment

OS	DYNIX/ptx 4.2.x and below
Largest recognizable drive size	4 GB (8,388,608 ₁₀ blocks)
Largest recognizable drive with IBM NUMA-Q support tape	8 GB (17,016,960 ₁₀ blocks)
Largest defined partition	4 GB (8,388,608 ₁₀ blocks)
Largest filesystem	2 GB (4,194,304 ₁₀ blocks)
Supports multiple target IDs	yes
Supports multiple LUNs per target ID	not recommended
Partition representing the entire disk	see vtoc files
To see the drive configuration, use this utility	diskid
To configure drives use this utility	menu
Special drive definition files needed	yes ^a

a. Requires changes to SCSI **space** files and kernel rebuild.

Table 5-2 below summarizes the IBM NUMA-Q environment when it is used as a UNIX server running DYNIX/ptx 4.3.0 and above:

Table 5-2 DYNIX/ptx 4.3.0 and above Host Environment

O/S	DYNIX/ptx 4.3.0 and above
Largest recognizable drive size	1 TB (2,147,483,648 ₁₀ blocks)
Largest defined partition	1 TB (2,147,483,648 ₁₀ blocks)
Largest filesystem	1 TB (2,147,483,648 ₁₀ blocks)
Supports multiple target IDs	yes
Supports multiple LUNs per target ID	Yes (for 4.4.2 and above) No (for below 4.4.2)
Partition representing the entire disk	see vtoc files
To see the drive configuration, use this utility	diskid
To configure drives use this utility	devctl
Special drive definition files needed	yes ^a

a. Requires changes to SCSI **space** files and kernel rebuild for DYNIX/ptx 4.3.0 and 4.3.1 only.

Useful IBM NUMA-Q Utilities and Functions

You may optionally use these functions and utilities to define and manage Symmetrix devices. They are listed for reference only.

- ◆ `diskid`
Displays list of all disk devices available to the IBM system.
- ◆ `menu`
Invokes the IBM service and system management utility.
- ◆ `newfs /dev/pathname`
Creates a filesystem.
- ◆ `showcfg`
Shows the hardware configuration of the host system.
- ◆ `dmesg`
or
- ◆ `cat /usr/adm/ktlog/YYYY/MM/DD`
(or `more /usr/adm/ktlog/YYYY/MM/DD`)
Collects system diagnostic messages to form the error log.
- ◆ `dumpconf`
Dumps the kernel auto-configuration device table.
- ◆ `/etc/monitor -f`
Monitors system activity.

Symmetrix Configuration

In the IBM NUMA-Q host environment, you can configure the Symmetrix disk devices into logical volumes. EMC recommends the following logical volume ratios as indicated in the table below:

Table 5-3 Disk Device Configuration

		Symmetrix Logical Volume to Physical Splits ^a							
		1:1	2:1	3:1	4:1	5:1	6:1	7:1	8:1 ^b
Physical Devices	4 GB	x	x						
	9 GB	x	x	x	x				
	18 GB		x	x	x	x	x	x	x
	23 GB		*	x	x	x	x	x	x
	36 GB			*	x	x	x	x	x
	47 GB			*	*	x	x	x	x
	50 GB ^c				x	x	x	x	x
	73 GB ^{c,d}					x	x	x	x

- a. * indicates ratios available only in microcode level 5265 or higher.
- b. With microcode revision level 5263, maximum is 8:1. With 5264, maximum is 16:1. With 5265, maximum is 32:1.
- c. The 50 and 73 GB drives are supported only by microcode revision levels 5x66 and higher.
- d. The 73 GB drive does not support the RAID-S configuration on Symmetrix.

For specific cylinder counts and logical volume sizes, refer to Table 2-1 on page 2-5 and the following tables in that chapter.

The EMC Customer Engineer (CE) should contact the EMC Configuration Specialist for updated online information. This information is necessary to configure Symmetrix to support the customer's host environment.

Incorporating Symmetrix

The remainder of this section outlines the steps needed to incorporate a Symmetrix system in the IBM NUMA-Q environment.

The DYNIX/ptx V2.1.5, V4.0.3, V4.0.2, and V4.1.2 for the Symmetry series and V4.3.0 and V4.3.1 for NUMA-Q default SCSI disk definition does not support wide, synchronous SCSI, and command tag queueing between the host and the disk. However, with minimal modifications to DYNIX/ptx structures and a kernel rebuild, a IBM NUMA-Q UNIX server can support some or all of these features.

Target ID Support

DYNIX/ptx operating systems 4.4.2 and above support LUNs by default. Fibre Channel fabric connections to the Symmetrix (without a Fibre Channel-to-SCSI bridge) are supported only in DYNIX/ptx 4.4.4 and above.

Below 4.4.2, the IBM NUMA-Q environment supports target IDs only. In order to support LUNs, you must make modifications to your kernel tables and rebuild the kernel (not recommended).

EMC provides a Fast-Wide SCSI director and an Ultra-Wide SCSI director that support multiple target IDs. EMC recommends that you install either director in your Symmetrix unit for all connections with a IBM NUMA-Q host.

Wide, Synchronous, & Queueing Support 1

For DYNIX/ptx 4.2.x and below: If you are using DYNIX/ptx 4.3.0 or 4.3.1, please refer to the next section on how to add the EMC Symmetrix device definition description to the appropriate files. **If you are using a IBM NUMA-Q host running DYNIX/ptx 4.4.0 or above, the default kernel contains the correct EMC Symmetrix device definitions.**

To enable IBM NUMA-Q to support wide, synchronous, and queueing SCSI protocols, the EMC Symmetrix device definition description needs to be added to the `scsidisk_space.c` file.

You can choose to download the Symmetrix device definition, **scsidisk_space.c**, to the **/usr/conf/uts/io/scsidisk** directory on your IBM NUMA-Q host or modify your existing file as follows:

Before you make any changes to the structure, make a backup copy of your **scsidisk_space.c** file.

```
{
    "EMC"                /* Wide SCSI spare #1 */
    "SYMMETRIX",         /* vendor */
    1,                   /* product */
    0,                   /* 1=drive can do wide mode, 0=no */
    0,                   /* model */
    8,                   /* max requests queued */
    0,                   /* don't allow unordered */
    0,                   /* don't halt on error */
    1,                   /* DO negotiate synchronous */
    1,                   /* DO negotiate wide bus */
},
```

You may need to reapply this change to the **scsidisk_space.c** file after you upgrade your operating system, for example, from 4.1.3 to 4.2, as the new operating system may overwrite this file.

Rebuild the kernel for these changes to take effect.

Wide, Synchronous, & Queueing Support 2

For NUMA-Q at DYNIX/ptx 4.3.0 or 4.3.1: If you are using a Symmetrix system running DYNIX/ptx 4.2.x and below, please refer to the previous section on how to add the EMC Symmetrix device definition description to the appropriate files. **If you are using an IBM host running DYNIX/ptx 4.4.0 or above, the default kernel contains the correct EMC Symmetrix device definitions.**

To enable IBM NUMA-Q to support wide, synchronous, and queuing SCSI protocols, the EMC Symmetrix device definition description needs to be added to the **sd_space.c** file located in the **/usr/conf/uts/io/sd** directory and the **ids_space.c** file located in the **/usr/conf/uts/kernel/sd** directory on your NUMA-Q host.

Follow the steps below to make the appropriate file modifications.

Before you make any changes to the structure, make a backup copy of your **sd_space.c** file.

1. Modify your existing **sd_space.c** file as follows:

```
{
    "EMC"                /* Wide SCSI spare #1 */
    "SYMMETRIX",         /* vendor */
    1,                   /* product */
    1,                   /* 1=drive can do wide mode, 0=no */
    0,                   /* model */
    {
        8,               /* max requests queued */
        0,               /* don't allow unordered */
        0,               /* don't halt on error */
        1,               /* DO negotiate synchronous */
        1,               /* DO negotiate wide bus */
    }
},
```

Before you make any changes to the structure, make a backup copy of your **ids_space.c** file.

2. Modify the **ids_space.c** file located in the **/usr/conf/uts/kernel/ids** directory as follows:

Any lines preceded with "+" are to be inserted into the **ids_space.c** file.

```
.
.
.
IDS_VN SCSI CARTTAPE=11, /* Storage Tech, Fujitsu, Odetics */
IDS_VN SCSI QICTAPE=12, /* Tanberg, Archive */
IDS_VN SCSI UNISYS=13, /* Unisys, Kennedy */
+ IDS_VN SCSI EMC=14, /* EMC Corporation */

IDS_VN SCSI SIS=99, /* IBM NUMA-Q Integration Services */
IDS_VN SCSI SPARE=100 /* Spare table entry */
.
.
.
    ids_scsi_unisys_models,
};
+ /*
+ *EMC SCSI Model Registry
+ *    Device Class: SCSI
+ *    Vendor:      "EMC"
+ *
+ *NOTE: These are registered values, and must not change or slip.
+ */
```

```

+ ids_mod_rec_t ids_scsi_emc_models[] = {
+ { "IDS_MT_INVALID", 0 }, /* invalid sentinel entry */
+
+ { "SYMMETRIX", 1 }, /* EMC RAID Disk Subsystem */
+ };
+
+ ids_mod_tab_t ids_scsi_emc_table = {
+ /* --- table header --- */
+ {
+ 1, /* revision */
+ IDS_MK_MT)SIG(IDS_VN_SCSI_EMC), /* signature */
+ 1, /* byte cnt, max entry (255) */
+ IDS_MT_CNT(ids_scsi_emc_models), /* actual record count */
+ },
+
+ /* --- table data --- */
+ ids_scsi_emc_models,
+ };
+
+ /*
+ * SIS SCSI Model Registry
+ * Device Class:      SCSI
+ *
+ *
+ *
+ { "UNISYS",IDS_VN_SCSI_UNISYS, &ids_scsi_unisys_table, },
+ { "KENNEDY",IDS_VN_SCSI_UNISYS, &ids_scsi_unisys_table, },
+
+ { "EMC",IDS_VN_SCSI_EMC, &ids_scsi_emc_table, },
+
+ /* IBM NUMA-Q Integration Services */
+ { "SIS", IDS_VN_SCSI_SIS, &ids_scsi_sis_table, },

```

3. Rebuild the kernel for these changes to take effect.

LUN Support *below* DYNIX/ptx 4.4.2

NOT Supported by EMC: To enable IBM NUMA-Q or Symmetry servers to support LUNs in operating systems below 4.4.2, (its default environment supports target IDs only), EMC provides a **scsi_space.c**, **scsidisk_space.c**, and **qcic_space.c** file.

Before you make any changes to the structures, make a backup copy of your disk definition files, **qcic_space.c** and **scsidisk_space.c**.

The default system configuration file, located in `/etc/conf/symmetry/sb8000`, contains a list of hardware devices defining assigned ports and devices used during kernel rebuild. The standard file `system.std` may have been modified or replaced to customize the configuration. This file is maintained by the system administrator and is not modified by EMC during Symmetrix configuration.

FTP the Symmetrix device definition files from the EMC ftp server, `ftp.emc.com`, in directory `/pub/symm3000/sequent` to your IBM host. Download the file `scsi_space.c` to the `/usr/conf/uts/io/scsi` directory. Also download the file `qcic_space.c` to the `/usr/conf/uts/io/qcic` directory.

When kernel tables are changed to support LUNs, IBM may not properly recognize its internal `qd` devices when the kernel is rebuilt. To ensure that all internal devices are properly recognized, perform the following steps to have the internal devices poll for their target ID.

1. Power on the IBM server and wait for it to perform its self-tests.
2. When the self-tests complete and the login prompt appears, log in as **root**.

To verify that the host has identified the connected devices, type the following:

```
diskid > disklist
```

A listing of all devices found will be directed to **disklist** file.

3. Locate the `qd` number of the high range of internal storage devices currently in the system.

```
qd24: qcic 0 bus 0 target 8 lun 0 initiator 6
Vendor ...
qd25: qcic 0 bus 0 target 9 lun 0 initiator 6
Vendor ...
```

4. Make a hardcopy of the **disklist** file for future use.
5. Go to the `/etc/conf/uts/symmetry/sb8000/system.std` file.

This file is a generic listing of the hardware devices in the system. Do not edit the `system.std` file; make a copy to edit and use during kernel rebuild. Use a filename that begins with `system` and leave it in the `sb8000` directory.

6. Locate the entry for each internal drive. This entry appears similar to the following:

```
qd on qcic_scsi0 8
```

7. Change the target ID ("8" in the above example) to a question mark (?).

This question mark causes the system to poll that device for its target ID number during a kernel rebuild. For more information on polling for target ID, refer to the *DYNIX/ptx System Configuration and Performance Guide*, P/N 1003-51297-04.

8. Repeat steps 6 and 7 for each additional internal device.

Fibre Channel Topology

With the introduction of DYNIX/ptx 4.5.0, IBM NUMA-Q systems have the added capability to connect to the Symmetrix via Fibre Channel arbitrated loop. (Prior to DYNIX/ptx 4.5.0, IBM NUMA-Q could only connect to a Symmetrix via FC switches or SCSI bridges.)

The supported arbitrated loop configuration allows only a direct (hubless) connection from each HBA to a Symmetrix port via a supported FC cable. The FC topology must be either entirely

- a. arbitrated loop (FC-AL) using only direct connections (no hubs allowed)

or

- b. fabric switch (FC-SW), and not a mix of both.

In order to change the configuration of the NUMA-Q from fabric (the default) to arbitrated loop, complete the steps outlined below. These changes will affect all HBAs in the system. After making these changes, the system will be configured to operate only with direct arbitrated loop connections.



CAUTION

Before making changes to any kernel files, make a backup copy of each file prior to making the changes in case you need to go back to the original configuration.

1. For FC-AL support , modify your existing **ff_space.c**, located in **/usr/conf/uts/io/ff**, as follows :

```
.
.
.
/*
 * Topology to be configured -
 *   FC_TOPOLOGY_AL - Arbitrated Loop
 *   FC_TOPOLOGY_FABRIC - Fabric (switch)
 *   FC_TOPOLOGY_PT2PT- Point to point
 */
int    ff_topology = FC_TOPOLOGY_AL;

/*
 * If ff_class is set to -1, then the default class for the given
 * topology is used. The default class for FC_TOPOLOGY_AL is CLASS3
 * and for FC_TOPOLOGY_FABRIC or FC_TOPOLOGY_PT2PT is CLASS2.
 * If a specific class is needed other than the default class, then
 * the ff_class should be set to the specific class.
 */
int    ff_class = -1;
.
.
.
```

2. Recompile the kernel.
3. Reboot the host.

Fibre Channel Class of Service

With the introduction of DYNIX/ptx 4.5.0, IBM NUMA-Q systems which are connected directly to a Symmetrix via a switch are configured by default to communicate using Fibre Channel Class 2, but the kernel can be reconfigured to use Class 3. Fibre Channel arbitrated loop connections can only be configured to support Class 3.

1. To enable Class 3 support for fabric connections, modify your existing **ff_space.c**, located in **/usr/conf/uts/io/ff**, as follows :

```
.
.
.
/*
 * Topology to be configured -
 *   FC_TOPOLOGY_AL - Arbitrated Loop
 *   FC_TOPOLOGY_FABRIC - Fabric (switch)
 *   FC_TOPOLOGY_PT2PT- Point to point
 */
int    ff_topology = FC_TOPOLOGY_FABRIC;

/*
 * If ff_class is set to -1, then the default class for the given
 * topology is used. The default class for FC_TOPOLOGY_AL is CLASS3
 * and for FC_TOPOLOGY_FABRIC or FC_TOPOLOGY_PT2PT is CLASS2.
 * If a specific class is needed other than the default class, then
 * the ff_class should be set to the specific class.
 */
int    ff_class = FF_CLASS3;
.
.
.
```

2. Recompile the kernel.
3. Reboot the host.

Configuring the IBM NUMA-Q System

After the EMC Customer Engineer has assigned target IDs to Symmetrix devices and the Symmetrix channel and disk directors have been switched online, use the procedure below to verify that the host can see all the target Symmetrix devices assigned to that host.

1. Power on the IBM server and wait for it to perform its self-tests.
2. When the self-tests complete and the login prompt appears, log in as `root`.

To verify that the host has identified the connected devices, type the following:

```
diskid > diskfile
```

A listing of all devices found will be directed to **diskfile**. For example, entries such as the ones shown on the following page denote Symmetrix devices on IBM systems running DYNIX/ptx 4.4.4, 4.4.2, and 4.2.3.

Devices are named `qd` by default for 4.2.3 and below.

Devices are named `sd` by default for 4.3.0 and above, but you can change the `sd` name.

For 4.4.2 and above, verify that `diskid` reports multiple **sd's** for each device that is multiported.

Multipathing and Multiporting Support

DYNIX/ptx 4.4.2 and above support multipathing to the same EMC device. For each device, up to 8 initiators from a single host connected to 8 Symmetrix ports are supported. See the following figure for an example of multiporting.

Figure 5-1 shows multiporting with IBM hosts. The host knows that multiple references to sd1, one for each port, are referring to the same logical volume.




Figure 5-1 Multiporting (only with DYNIX/ptx 4.4.2 and above)

Here are diskid entries from a 4.4.4 system configured with Fibre Channel:

```
sd506: fabric1 cport 120 fc id 0x200d13, fcp lun 40:58:00:00:00:00:00
No Pbay on fabric1
Vendor EMC, product SYMMETRIX, revision 5265, serial # 5031390F6000
capacity 2157 Mbytes
sd506: fabric2 cport 128 fc id 0x200a13, fcp lun 40:58:00:00:00:00:00
No Pbay on fabric2
Vendor EMC, product SYMMETRIX, revision 5265, serial # 5031390F6000
capacity 2157 Mbytes
```

Here are diskid entries from a 4.5.0 system configured with Fibre Channel:

```
sd506: fabric1 cport 120 fc id 0x200d13, fcp lun 0X5840
No Pbay on fabric1
Vendor EMC, product SYMMETRIX, revision 5265, serial # 5031390F6000
capacity 2157 Mbytes
sd506: fabric2 cport 128 fc id 0x200a13, fcp lun 0X5840
No Pbay on fabric2
Vendor EMC, product SYMMETRIX, revision 5265, serial # 5031390F6000
capacity 2157 Mbytes
```

Here are `diskid` entries from a 4.4.2 system with SCSI. (Fibre is not supported on 4.4.2):

```
sd92: scsibus8 cport 16 target 0 lun 1, No Pbay on scsibus8
Vendor EMC, product SYMMETRIX, revision 5264, serial # 600249000000
capacity 2076 Mbytes
sd92: scsibus13 cport 1 target 0 lun 1, No Pbay on scsibus13
Vendor EMC, product SYMMETRIX, revision 5265, serial # 600249000000
capacity 2076 Mbytes
```

Here is the output from a 4.2.3 system. Notice that there is only one entry -- multiporting is only supported in 4.4.2 and above:

```
qd134: qcic 1 bus 2 target 10 lun 0 initiator 6
Vendor EMC, product SYMMETRIX, revision 5264, capacity 2177 Mbytes
No Pbay on qcic 1 bus 2
```

The capacity as reported using the `diskid` utility will differ when run on system using DYNIX/ptx 4.2.x and below vs. DYNIX/ptx 4.3.0 and above. DYNIX/ptx 4.2.x and below reports 1 MB as 1000 x 1000 bytes. DYNIX/ptx 4.3.0 and above reports 1 MB as 1024 x 1024 bytes.

Check the contents of `diskfile` to make sure all Symmetrix devices attached to the IBM system are recognized. Make a hardcopy of this file for future use. Also compare the contents of this file with `disklist`. Verify that all internal devices are correctly identified.

Configuring /Deconfiguring Devices Online



CAUTION

With DYNIX/ptx 4.4.x and above only.

(If you are below 4.4.x you need to reboot to configure or deconfigure devices.)

You must consult the man pages on the command `devctl` to use it for configuring and deconfiguring. Only examples are presented here.

`devctl` enables you to configure or deconfigure devices, to assign device names and to manage storing name assignments.

Use `devctl -c parent_name` to find devices connected to `parent_name`, as in these examples:

```
devctl -c scsibus8
or
devctl -c fabric0
```

New devices are assigned temporary names, for example `+sd10`.

To deconfigure devices, removing OS resources from `device_name`, use `devctl -d device_name`, as in the following example:

```
devctl -d sd10
```

To delete name assignments for deconfigured devices, use the command `devctl -D device_name`, as in the example below:

```
devctl -D sd10
```

You can generate permanent names for devices with temporary names. You can also assign permanent names to devices, using the option `-n`, as in the following example:

```
devctl -n +sd10 sd1
```

where

`+sd10` is the currently unnamed device

`sd1` is the new name

You can use `devctl -N` to assign permanent names to all devices with temporary names.

Formatting and Partitioning Definition Files

After verifying that the Symmetrix devices have been recognized by the host, download the files containing Symmetrix configuration information to the IBM system. These files define the drive geometry and default partitioning. There are different files to represent the various Symmetrix drives sizes compatible in this environment. At this time, EMC provides the files listed in Table 5-4.

If your Symmetrix unit has 3 GB disk devices, please contact your EMC Technical Support representative for the appropriate disk device definitions.

The **SYMMxG.geom** and **SYMMxG.scsi** files, downloaded to the **/etc/diskinfo** file, contain the recommended device naming conventions for Symmetrix devices. The **SYMMxG** files, downloaded to the **/etc/vtoc** file, contains the default partitioning information for Symmetrix devices.

The **SYMM41G** (default), **SYMM41G.share** (cluster environment), **SYMM41GF** (filesystems) are samples of VTOC files to accommodate the needs of clusters and multiple filesystem sections on a 4.35 GB device.

New features in DYNIX/ptx 4.6.0 and above

Boot Support on EMC Symmetrix

- ◆ Sharable device boot (boot on EMC devices) is supported on Fibre Channel fabric (switch) connected devices. FC-AL and Fibre Channel/SCSI bridge are not supported. Consult the man page for more information.
- ◆ A boot disk can be mirrored by means of SVM to any other EMC devices or local hard drives.

Useful Utilities

- ◆ `crvtoc` utility

Example:

```
crvtoc -n 2 -d sd20 -w -g
```

Puts a VTOC on sd20, partitioning it into 2 usable slices: sd20s0 and sd20s1, each about half the size of sd20. See the man page by running `man crvtoc` for details.

- ◆ `mapdev` utility

Shows the device path or physical specifier of a device for use in a boot string. See the man page by running `man mapdev` for details.



CAUTION

If you use any EMC Symmetrix device as the boot disk, never take the Symmetrix offline when the host is up. Doing so makes the host panic, resulting in down time and possible data loss.

- ◆ `diskown` utility

All disks in DYNIX/ptx 4.6.0 and above require a disk label before `system/SVM` can access them.

Example:

```
diskown -w -n 0xe -g sd20
```

Puts a node-owned disk label on `sd20`. See the man page by running `man diskown` for details.

See the following notes:

1. You can find a node ID and cluster/master ID by running the `showcfg` command.
2. The host constantly pings any disks that are configured as node-owned. You must label any BCV or RDF (R2 side) device as cluster-owned or other-cluster-owned even if the host is not in a cluster. Do this to prevent the host from pinging a BCV or RDF device during a BCV/RDF establish.
3. Raw writes directly to a whole device are not supported under DYNIX/ptx 4.6.0 and above. Raw writes may destroy the disk label on the device, rendering it unusable. If you need to do raw writes, encapsulate the device in an SVM volume or create a VTOC and `devbuild` the device.

Table 5-4 Symmetrix Device Definition Files

Description	Symmetrix Cylinders	Files	Download to directory
Elite-23 drive 8:1	5,869 cylinders	SYMM2.88G SYMM2.88G.scsi SYMM2.88G.geom	/etc/vtoc /etc/diskinfo /etc/diskinfo
Elite-23 drive 7:1	6,708 cylinders	SYMM3.29G SYMM3.29G.scsi SYMM3.29G.geom	/etc/vtoc /etc/diskinfo /etc/diskinfo
Elite-23 drive 6:1	7,826 cylinders	SYMM3.84G SYMM3.84G.scsi SYMM3.84G.geom	/etc/vtoc /etc/diskinfo /etc/diskinfo
Elite-23 drive 5:1	9,392 cylinders	SYMM4.61G SYMM4.61G.scsi SYMM4.61G.geom	/etc/vtoc /etc/diskinfo /etc/diskinfo
Elite-23 drive 4:1	11,740 cylinders	SYMM5.77G SYMM5.77G.scsi SYMM5.77G.geom	/etc/vtoc /etc/diskinfo /etc/diskinfo
Elite-23 drive 3:1	15,654 cylinders	SYMM7.69G SYMM7.69G.scsi SYMM7.69G.geom	/etc/vtoc /etc/diskinfo /etc/diskinfo
Elite-9 drive 4:1	4,430 cylinders	SYMM21G SYMM21G.scsi SYMM21G.geom	/etc/vtoc /etc/diskinfo /etc/diskinfo
Elite-9 drive 2:1 ^a	8,738 cylinders	SYMM4G SYMM4G.scsi SYMM4G.geom	/etc/vtoc /etc/diskinfo /etc/diskinfo
Elite-9 drive 2:1	8,862 cylinders	SYMM41G SYMM41G.scsi SYMM41G.geom	/etc/vtoc /etc/diskinfo /etc/diskinfo
Elite-9 drive 1:1	17,726 cylinders	SYMM83G SYMM83G.scsi SYMM83G.geom	/etc/vtoc /etc/diskinfo /etc/diskinfo

Table 5-4 Symmetrix Device Definition Files (*continued*)

Description	Symmetrix Cylinders	Files	Download to directory
Cuda-9 drive 4:1	4,602 cylinders	SYMM2.26G SYMM2.26G.scsi SYMM2.26G.geom	/etc/vtoc /etc/diskinfo /etc/diskinfo
Cuda-9 drive 3:1	6,137 cylinders	SYMM3.01G SYMM3.01G.scsi SYMM3.01G.geom	/etc/vtoc /etc/diskinfo /etc/diskinfo
Cuda-9 drive 2:1	9,206 cylinders	SYMM4.52G SYMM4.52G.scsi SYMM4.52G.geom	/etc/vtoc /etc/diskinfo /etc/diskinfo
Cuda-9 drive 1:1	18,414 cylinders	SYMM9.05G SYMM9.05G.scsi SYMM9.05G.geom	/etc/vtoc /etc/diskinfo /etc/diskinfo
4 GB drive 2:1	4,335 cylinders	SYMM2BG SYMM2BG.scsi SYMM2BG.geom	/etc/vtoc /etc/diskinfo /etc/diskinfo
4 GB drive 1:1	8,672 cylinders	SYMM4BG SYMM4BG.scsi SYM4BG.geom	/etc/vtoc /etc/diskinfo /etc/diskinfo

- a. Use this definition when configuring 4 GB disk drives when running DYNIX/ptx 4.2.1 and below.

Creating Physical Volumes

Once the host recognizes the Symmetrix devices, create a physical volume for each device when running DYNIX/ptx 4.2.x and below. When using DYNIX/ptx 4.3.0 and above, you may omit the procedure in this section for any Symmetrix logical volumes to be used as raw disks under the control of IBM NUMA-Q Volume Manager (ptx/SVM).

To create a physical volume, you need to use the DYNIX menus.

1. Type menu at the prompt.
2. Choose **A** (System Administration) from the menu that appears on the screen.
3. Choose **D** (Disk Drive Management) from the System Administration menu.
4. Choose **A** (Create a VTOC Device) from the Disk Drive Management menu.
5. Enter the Symmetrix device (its **sd** or **qd** number) as identified by the IBM NUMA-Q system on the first line.
 - The **sdxx** naming scheme is used by operating systems 4.3.0 and higher; you can also change the name **sd** to other names, including to the name **qd**.
 - The **qdx** naming scheme is used by operating systems below 4.3.0.

Then enter the name of the appropriate configuration file (**SYMMxG**) on the second line. Press the <F1> key.

If you are not using the default VTOC file, for example, creating two filesystems on a 4 GB drive, enter the Symmetrix device (its **sd** or **qd** number) as identified by the IBM NUMA-Q system on the first line, enter **SYMM4G** on the second line, and enter **/etc/vtoc/SYMM4GF** on the third line (non-standard VTOC). Press the <F1> key.

6. Repeat steps 4 and 5 for each additional Symmetrix device to create the VTOC for each device.

A VTOC is now defined for all Symmetrix devices. These are raw drives. To create a filesystem for each Symmetrix device, continue with the following sections.

Create the Mount Directory

If you want to use filesystems on the Symmetrix devices, you must create a mount directory for each of these devices. To do this, type a statement similar to the following for each physical volume:

```
mkdir /mntx
```

where “x” is an arbitrary number that distinguishes the mount directory

This example creates a mount directory called “**mntx**”.

Create a New Filesystem

Now that the Symmetrix raw devices and their respective mount directories are available to the operating system, a new filesystem can be created for each Symmetrix device. To do this, use the **newfs** command in a statement similar to the following:

```
newfs /dev/rdisk/qd48s0
```

If you are using ptx/SVM in OS 4.2.x and below, enter a statement similar to the following:

```
newfs /dev/rvol/volx
```

where volx is the name of the volume.

If you are using ptx/SVM in OS 4.4.x and above, enter a statement similar to the following:

```
newfs /dev/vx/rdsk/diskgroup/volumename
```

where diskgroup is the name of the diskgroup that volumename belongs to.

Mount the Directory

Mount the filesystem by typing a statement similar to the following:

```
mount /dev/dsk/qd48s0 /mntx
```

If you are using ptx/SVM with operating systems 4.4.x and above, mount the filesystem by typing a statement similar to the following:

```
mount /dev/vx/dsk/diskgroup/volumename /mntx
```

If you are using ptx/SVM with operating systems 4.2.x and below, mount the filesystem by typing a statement similar to the following:

```
mount /dev/vol/volx /mntx
```

The default limit for the number of mountpoints is 128 for operating systems 4.4.x. The default is 32 for operating systems 4.2.x and below.

For 4.4.x and below, please note that more than 128 mountpoints is not recommended. To increase the number of mountpoints beyond this limit, modify the `NMOUNT` argument in `/usr/conf/uts/kernel/i386_space/config`. Before making this change, copy the config file to `config.orig`. After making the changes, you must re-compile the kernel and reboot the operating system in order for the changes to take effect.

For 4.5.x and higher, none of the above applies. The maximum number of mount points is dynamic. No kernel modification or recompilation is required.

IBM NUMA-Q Volume Manager

The IBM NUMA-Q Volume Manager (ptx/SVM) is a tool for managing disk usage. This tool can be used to create logical disks, filesystems and mirrored filesystems, and provides other utilities for optimization and maintenance of disk devices. ptx/SVM is a required software component in a cluster environment, and may also be used to create and maintain non-shared systems. Refer to the *ptx/SVM Administration Guide* for detailed discussions and procedures for using ptx/SVM.

ptx/SVM uses a hierarchical system comprised of subdisks, plexes and volumes to define logical devices (Figure 5-2). A subdisk consists of all or part of a disk partition and is mapped to a physical device.

A plex consists of from one to 256 subdisks associated with it. A volume consists of 1 to 32 plexes in operating systems 4.4.x and above. In operating systems 4.2.x and below, a volume consists of 1 to 3 plexes.

When more than one plex is attached to a volume, the volume is mirrored. Figure 5-2 shows an example of a mirrored volume with two plexes.




Figure 5-2 ptx/SVM Mirrored Volume Structure in DYNIX/ptx 4.4.2

ptx/SVM may be used to define shared subdisks, plexes and volumes for use in a IBM NUMA-Q cluster system.

The ptx/SVM commands listed in this section are different for DYNIX/ptx 4.4.0 and above. This section will note the different syntax in all examples.

In operating systems 4.3.0 and higher, disk groups are added organization building blocks. A disk group is a collection of ptx/SVM disks which enable you to group disks into logical collections for administrative convenience. For example, you could include all disks containing database applications in a disk group. The figure below shows an example ptx/SVM Mirrored Volume Structure in DYNIX/ptx 4.4.2 -- a disk group in relation to subdisks, plexes and a mirrored volume.

ptx/SVM can be accessed through either a command line or the ptx/SVM menu system. Most ptx/SVM procedures can be accomplished using either interface. Procedures presented below use the command line method. To activate the ptx/SVM menu system, log in as **root** and type `svmadmin`

The `svmadmin` command is not available in DYNIX/ptx 4.4.0 and above.

Refer to the *ptx/SVM Administration Guide* for detailed information on using the ptx/SVM menu to initialize and create the **voladm** configuration database.

Create Volumes Using ptx/SVM

The following steps are an outline of the steps to create SVM volumes. For more complete information on creating volumes, see IBM NUMA-Q's *ptx/SVM Administration Guide*.

To create a volume using ptx/SVM, first create subdisks, then plex, and then volume using the `volmake` (`vxmake` in DYNIX/ptx 4.4.0 and later) command. For example, to create a volume named **vol40** on disk **/dev/dsk/qd30s2** use the procedure which follows. This reference, however, shows a device with a VTOC written to disk. A VTOC splits the disk into logical slices. In 4.4.x and above, a VTOC is not required to create a ptx/SVM volume.

1. Log in as **root**.

2. Create a subdisk named **qd30s2-0** at the beginning (offset=0) of partition **/dev/dsk/qd30s2** with a length of **4,104,192** sectors (2.1 GB).

```
volmake sd qd30s2-0 /dev/dsk/qd30s2,0,4104192
```

or, in DYNIX/ptx 4.4.0 and later

```
vxmake sd sd30s2-0 /dev/dsk/sd30s2,0,4104192
```

3. Create a plex and associate subdisk to it.

```
volmake plex vol40-0 sd=qd30s2-0
```

or, in DYNIX/ptx 4.4.0 and later

```
vxmake plex vol40-0 sd=sd30s2-0
```

4. Create a volume specifying a usage type of generic filesystem (**fs~~g~~en**) or raw partition (**g~~e~~n**), and attach plex to the volume.

```
volmake vol vol40 use_type=fsgen plex=vol40-0
```

or, in DYNIX/ptx 4.4.0 and later

```
vxmake vol vol40 use_type=gen plex=vol40-0
```

A plex containing more than one subdisk can be organized as either concatenated or striped. The default method is concatenated. Striped plexes can be created using the **volmake** (**vxmake** in **DYNIX/ptx 4.4.0 and later**) command. For example, to create two volumes with striped plexes from two 1 GB disk partitions:

1. Create two 500 MB subdisks on each partition.

```
volmake sd qd30s2-0 /dev/dsk/qd30s2,0,500M
volmake sd qd30s2-1 /dev/dsk/qd30s2,500M,500M
volmake sd qd32s2-0 /dev/dsk/qd30s2,0,500M
volmake sd qd32s2-1 /dev/dsk/qd30s2,500M,500M
```

or, in DYNIX/ptx 4.4.0 and later (sd, the default, or qd can be used)

```
vxmake sd qd30s2-0 /dev/dsk/sd30s2,0,500M
vxmake sd qd30s2-1 /dev/dsk/sd30s2,500M,500M
vxmake sd qd32s2-0 /dev/dsk/sd30s2,0,500M
vxmake sd qd32s2-1 /dev/dsk/sd30s2,500M,500M
```

2. Create plexes, associate subdisks, and specify stripe width.

```
volmake plex vol40-0 sd=qd30s2-0 layout=stripe st_width=32K
volmake plex vol40-1 sd=qd30s2-1 layout=stripe st_width=32K
volmake plex vol42-0 sd=qd32s2-0 layout=stripe st_width=32K
volmake plex vol42-1 sd=qd32s2-1 layout=stripe st_width=32K
```

or, in DYNIX/ptx 4.4.0 and later

```
vxmake plex vol40-0 sd=sd30s2-0 layout=stripe st_width=32K
vxmake plex vol40-1 sd=sd30s2-1 layout=stripe st_width=32K
vxmake plex vol42-0 sd=sd32s2-0 layout=stripe st_width=32K
vxmake plex vol42-1 sd=sd32s2-1 layout=stripe st_width=32K
```

3. Create two volumes specifying a usage type of generic filesystem (**fsngen**), and attach two plexes to each volume.

```
volmake vol vol40 use_type=fsngen plex=vol40-0,vol42-0
volmake vol vol42 use_type=fsngen plex=vol40-1,vol42-1
```

or, in DYNIX/ptx 4.4.0 and later

```
vxmake vol vol40 use_type=fsngen plex=vol40-0,vol42-0
vxmake vol vol42 use_type=fsngen plex=vol40-1,vol42-1
```

Creating Mirrored Volumes using ptx/SVM

To create a triple mirrored volume using ptx/SVM, create the subdisks, the plexes, and then the mirrored volume using the `volmake` (`vxmake` in DYNIX/ptx 4.4.0 and later) **command**. For example, to create a mirrored volume named **volmir1** using the devices **/dev/dsk/sd40s2**, **/dev/dsk/sd42s2** and **/dev/dsk/sd44s2**, do the following:

1. Log in as **root**.

2. Create subdisks **qd40s2-0**, **qd42s2-0** and **qd44s2-0**.

```
volmake sd qd40s2-0 path=/dev/dsk/qd40s2 offset=0 len=4093724
volmake sd qd42s2-0 path=/dev/dsk/qd42s2 offset=0 len=4093724
volmake sd qd44s2-0 path=/dev/dsk/qd44s2 offset=0 len=4093724
```

or, in DYNIX/ptx 4.4.0 and later

```
vxmake sd sd40s2-0 path=/dev/dsk/sd40s2 offset=0 len=4093724
vxmake sd sd42s2-0 path=/dev/dsk/sd42s2 offset=0 len=4093724
vxmake sd sd44s2-0 path=/dev/dsk/sd44s2 offset=0 len=4093724
```

3. Create the plexes **volmir1-0**, **-1**, and **-2**, and associate the subdisks (single subdisk association per plex).

```
volmake plex volmir1-0 sd=qd40s2-0
volmake plex volmir1-1 sd=qd42s2-0
volmake plex volmir1-2 sd=qd44s2-0
```

or, in DYNIX/ptx 4.4.0 and later

```
vxmake plex volmir1-0 sd=sd40s2-0
vxmake plex volmir1-1 sd=sd42s2-0
vxmake plex volmir1-2 sd=sd44s2-0
```

4. Create the mirrored volume specifying the usage type, raw (**use_type=gen**) or filesystem (**usr_type=fsgen**), and attach the plexes.

```
volmake vol volmir1 use_type=gen \
plex=volmir1-0, volmir1-1, volmir1-2
```

or, in DYNIX/ptx 4.4.0 and later

```
vxmake vol volmir1 use_type=gen \
plex=volmir1-0, volmir1-1, volmir1-2
```

5. Optional: To attach additional mirrored plexes to **volmir1**, use the **volplex att** (**vxplex att** in DYNIX/ptx 4.4.0 and above) command as follows:

```
volplex att volmir1 volmir1-3
```

or, in DYNIX/ptx 4.4.0 and later

```
vxplex att volmir1 volmir1-3
```

where **volmir1-3** is a previously created plex.



CAUTION

The plex changed to the **CLEAN** state is the source of volume data for resynchronization when the mirror volume is started. Be sure to specify a plex containing valid volume data when adding mirroring to existing plexes. For information on plex state change dependencies and effects, refer to the *IBM NUMA-Q ptx/SVM Administration Guide*.

6. Initialize the mirrored set by bringing the initial plex to the CLEAN state.

```
volume init clean volmir1 volmir1-0
```

or, in DYNIX/ptx 4.4.0 and later

```
vxvol init clean volmir1 volmir1-0
```

7. Start the volume, bringing it into the ENABLED state.

```
volume -o bg start volmir1
```

or, in DYNIX/ptx 4.4.0 and later

```
vxvol -o bg start volmir1
```

Dissociating Mirrored Volumes

To dissociate a plex from a mirrored volume (usually done to enable online backup), perform the following steps:

1. Log in as **root**.
2. Stop databases and unmount filesystems that use the parent volume.

If a database resides on the volume, stop it.

If a filesystem is mounted on the volume, unmount it.

If a database resides within a filesystem on the volume, stop the database and unmount the filesystem.

3. Dissociate one plex from the other plex(es) on a mirrored volume:

```
volplex dis plex_name
```

or, in DYNIX/ptx 4.4.0 and later

```
vxplex dis plex_name
```

4. Remount filesystem and restart the database if applicable.

5. Create and enable a temporary volume for the dissociated plex:

```
volmake vol temporary_vol_name use_type=fsgen plex=plex_name  
volume init enable temporary_vol_name
```

or, in DYNIX/ptx 4.4.0 and later

```
vxmake vol temporary_vol_name use_type=fsgen plex=plex_name  
vxvol init enable temporary_vol_name
```

The plex is now dissociated from its original volume and can be used for other operations. The original volume is available to the database. Perform all required operations (backup) then restore the original volume as follows:

1. Stop the temporary volume:

```
volume stop temporary_vol_name
```

or, in DYNIX/ptx 4.4.0 and later

```
vxvol stop temporary_vol_name
```

2. Dissociate the plex:

```
volplex dis plex_name
```

or, in DYNIX/ptx 4.4.0 and later

```
vxplex dis plex_name
```

3. Remove the temporary volume:

```
voledit rm temporary_vol_name
```

or, in DYNIX/ptx 4.4.0 and later

```
vxedit rm temporary_vol_name
```

4. Stop databases and unmount filesystems that use the original volume.

If a database resides on the volume, stop it.

If a filesystem is mounted on the volume, unmount it.

If a database resides within a filesystem on the volume, stop the database and unmount the filesystem.

5. Reattach the dissociated plex to its original volume:

```
volplex att volume_name plex_name
```

or, in DYNIX/ptx 4.4.0 and later

```
vxplex att volume_name plex_name
```

6. Remount the filesystem and/or restart the database.

Useful ptx/SVM Command Line Utilities

This section lists ptx/SVM commands you may use to define and manage the disk system under ptx/SVM.

The DYNIX/ptx 4.4.0 and later equivalent commands are given in *italics* immediately below the command for earlier versions of DYNIX/ptx.

Subdisk Operations

Command line entries associated with subdisk management.

- ◆ `volsd assoc plex subdisk`
`vxsd assoc plex subdisk`

Associate an existing subdisk with an existing plex.

- ◆ `volsd dis subdisk`
`vxsd dis subdisk`

Dissociate any subdisk from a plex.

- `volsd -s size split subdisk new_subdisk_1 new_subdisk_2`
`vxsd -s size split subdisk new_subdisk_1 new_subdisk_2`

Split a subdisk and create two new subdisks the length of **size**.

- ◆ `volsd join sub_disk1 sub_disk2 subdisk_new`
`vxsd join sub_disk1 sub_disk2 subdisk_new`

Join two subdisks into one new subdisk.

- ◆ `volsd mv subdisk_old subdisk_new`
`vxsd mv subdisk_old subdisk_new`

Move data and change associations from one subdisk to another existing subdisk that is the same size.

Plex Operations

Command line entries associated with plex management.

- ◆ `volplex att volume plex`
`vxplex att volume plex`

Attach an existing plex to an existing volume.

- ◆ `volplex det plex`
`vxplex det plex`

Detach a plex from a volume. Plex remains associated with the volume, but is not used for I/O.

- ◆ `volplex dis plex`
`vxplex dis plex`

Dissociate a plex from its volume.

- ◆ `volplex cp volume plex`
`vxplex cp volume plex`

Copy the contents of a volume to a plex.

- ◆ `volplex mv plex_old plex_new`
`vvxplex mv plex_old plex_new`

Move data and change associations from one plex to another existing plex that is the same size.

Volume Operations

Command line entries associated with volume management.

- ◆ `volume init clean volume [plex]`
`vxvol init clean volume [plex]`

Initialize a volume and put plex in the CLEAN state.

- ◆ `volume init zero volume`
`vxvol init zero volume`

Write zero blocks to all plexes in the volume.

- ◆ `volume init active volume`
`vxvol init active volume`

Set the ACTIVE state for all plexes associated with the volume.

- ◆ `volume [-o options] start volume`
`vxvol [-o options] start volume`

Starts (enables) a volume.

Some -o options are listed below:

`iosize=size`

Size of I/O requests used during plex resynchronization.

`bg`

Resynchronize in the background.

`slow`

Reduce performance impact of resynchronization.

- ◆ `volume startall`
`vxvol startall`

Start all volumes.

- ◆ `volume stop volume`
`vxvol stop volume`

Stop (disable) a volume.

- ◆ `volume stopall`
`vxvol stopall`

Stop all volumes.

- ◆ `volume set len=length volume`
`vxvol set len=length volume`

Set the length of a volume.

- ◆ `voldisk disk|partition_name`
`vxdisk disk|partition_name`

List all of the ptx/SVM objects associated with a particular disk device or partition.

Other Object Control Operations

Command line entries associated with all object management.

- ◆ `voledit rm object`
`vxedit rm object`

Remove the object's record from the configuration database.

- ◆ `voledit rename old_name new_name`
`vxedit rename old_name new_name`

Change an object's name.

- ◆ `voledit set options object`
`vxedit set options object`

Change attributes of specified object.

Options for this command are listed below:

`use_type=gen|fsgen`

Change usage to raw (**gen**) or filesystem (**fsgen**) (volume must be disabled).

`len=length`

Change length of associated subdisk.

`writecopy=on|off`

Set writecopy on or off for a volume.

`user=name`

Change object owner.

`group=name`

Change object group id.

`mode=mode`

Change object mode.

- ◆ `volmend clear all object`
`vxmend clear all object`

Clear all object utility fields.

- ◆ `volmend fix clean|active|stale plex`
`vxmend fix clean|active|stale plex`

Set plex to CLEAN, ACTIVE, or STALE state.

- ◆ `volmend fix empty volume`
`vxmend fix empty volume`

Set all plexes on the specified volume to the EMPTY state.

- ◆ `volmend on object`
`vxmend on object`

Set specified object to the ONLINE state.

- ◆ `volmend off object`
`vxmend off object`

Set specified object to the OFFLINE state.

- ◆ `volmend -o force off object`
`vxmend -o force off object`

Set specified object, which is the last complete enabled plex in the volume, to the OFFLINE state.

Volume Statistics Reporting

Command line entries associated with statistics reporting.

- ◆ `volstat options`
`vxstat options`

Print statistics information on one or more volumes.

Options for this command are listed below:

`-i interval`

Print summary and then print changes occurring after specified interval.

`-c count`

Print summary the number of times specified by count.

`-s, -p, -v [object]`

Print statistics for subdisk, plex or volume.

`-fa volume`

Print statistics on atomic copies performed on volume.

-fp plex

Print statistics on plex read and write ioctls.

-fv volume

Print statistics on verified volume reads and writes.

-fc volume

Print statistics on corrected read operations.

-ff volume

Print number of failed read and write operations.

Display Volume Manager Records

Command line entries associated with ptx/SVM records.

- ◆ volprint -htv
- vxprint -htv

Print volume manager records.

System and Error Messages

The DYNIX/ptx operating system logs system and error messages including SCSI error messages and kernel trace messages into two log files called `/usr/adm/messages` and `/usr/adm/ktlog/YYYY/MM/DD`, where YYYY, MM, and DD are the year, month, and date, respectively.

You can read these logs using two methods:

- ◆ the command line
- ◆ the Menu utility

Using the Command Line

The DYNIX/ptx operating system logs system and error messages into a buffer log file called `/usr/adm/msgbuf`. When this buffer fills, the operating system moves the contents of the buffer log to a log file called `/usr/adm/messages`. The operating system stores old messages in the `/usr/adm/OLDmessages` file.

To read the buffer log file (`/usr/adm/msgbuf`), type:

```
dmesg
```

Read the `/usr/adm/messages` file using the `cat` or `more` command.

The IBM NUMA-Q operating system stores kernel related messages in the kernel trace log file, `/usr/adm/ktlog/YYYY/MM/DD` where YYYY, MM, and DD are the year, month, and date, respectively. Use the `cat` or `more` command to read the kernel trace log files.

Using the Menu Utility

The IBM NUMA-Q Menu Utility allows you to read the system and kernel trace log files. To invoke the Menu utility, type the following at the command line prompt:

```
menu
```

Perform the following steps to access and view the log files:

1. Choose **System Administration** from the Menu utility menu.
2. Choose **Log File Management** from the System Administration menu.
3. Choose **Kernel Trace Log Management** from the Log File Management menu.

4. Choose **View Kernel Management Trace Log** from the Kernel Trace Log Management menu.
5. Enter the desired duration of the trace you wish to view, such as starting date and ending date.

The operating system lists all error messages that have been captured during that time duration on the screen display.

This chapter describes how to incorporate Symmetrix in the HP-UX environment when working with HP 9000 servers.

◆ Overview	6-2
◆ Useful HP-UX Utilities and Functions.....	6-5
◆ Symmetrix Configuration	6-8
◆ Hardware Mapping 10.0 & 11.x (offline)	6-13
◆ Recognizing New Devices and LUNs.....	6-15
◆ Configuring Logical Volumes	6-17
◆ Using SAM.....	6-20
◆ HP 9000 K-class Servers	6-21
◆ Channel Failover (10.x and 11.x).....	6-22
◆ Adding Devices Online	6-27
◆ System and Error Messages.....	6-28

Overview

When using Symmetrix in an HP-UX environment on HP 9000 servers, note the following:

- ◆ See Table 1-1, Chapter 1 for a summary of the HP 9000 Series that is supported. See Table C-1, Appendix C, for a summary of host controller adapter cards for each series. The table below lists the qualified operating systems at time of this publication, including HP-UX 11.0. See your EMC Sales Representative for updates.

Table 6-1 Qualified Operating Systems for HP 9000 Server Hosts

Host	Earliest OS Qualified	Latest OS Qualified
800 Series: E, F, G, H, I, T500	HP-UX 9.04	HP-UX 10.01 HP-UX 10.20
T520	HP-UX 10.01	HP-UX 10.20
K-Class (SCSI)	HP-UX 9.04	HP-UX 11.0
T600 (SCSI)	HP-UX 10.20 TFC	
D-Class, R-Class	HP-UX 10.01	
	HP-UX 10.20 TFC	
K-Class (Fibre)	HP-UX 10.20 TFC	
T600 (Fibre)	HP-UX 10.20 TFC	
HP 9000 V-Class, V2200, V2250, V2500 ^a	HP-UX 11.0	
HP 9000 N-Class	HP-UX 11.0 990P	HP-UX 11.0 990P

a. Minimum OS revision is HP-UX 11.0 Rev 9812.

- ◆ The following notes list more detailed information related to the tables referred to above.

HP-UX version 10.00 does not support multiple LUNs when Symmetrix is attached to the HP-HSC 16-bit Differential Fast-Wide SCSI card (Model 2969A) *only*.

With HP-UX version 10.01 and later, multiple LUNs are supported with HP-UX version 10 and higher when Symmetrix is attached to the Model 2969A.

HP-UX 10.01 does not support mounted filesystems greater than 4 GB in size.

When defining Logical Unit Numbers (LUNs) in an HP-UX environment, do not skip numbers. Always start with LUN 0 and walk your way up numerically.

The table below summarizes the operating system features in the HP-UX environment.

Table 6-2 HP 9000 Server Operating System Features

O/S	HP-UX
Supports multiple target IDs	Yes ^a
Supports multiple LUNs per target ID	Yes
Partition representing the entire disk	s2
To see the drive configuration, use this utility:	<code>ioscan</code> or <code>SAM</code>
To configure drives use this utility:	<code>SAM</code> or line commands
Special drive definition files needed	When below HP-UX 10

a. LUNs 0-7, targets 0-F.

The table below summarizes the limits of recent HP-UX releases.

Table 6-3 Capacity Limits of Recent HP-UX Releases

Attribute	10.01	10.10	10.20	64-bit 11.00
File System	4 GB	128 GB	128 GB	128 GB
File Size	2 GB	2 GB	128 GB Local, 2 GB Network	128 GB local and network
Physical RAM	2 GB	3.75 GB	8 GB	4 TB
Shared Memory	1.75 GB	1.75 GB	2.75 GB	8 TB
Process Data Space	0.9 GB	1.9 GB	1.9 GB	4 TB
Number of File Descriptors	2 K	60 K	60 K	60 K
Number of User IDs	60 K	60 K	~2,000 K	~2,000 K

Useful HP-UX Utilities and Functions

This section lists HP-UX functions and utilities (with HP-UX 10.x and 11.x) you may use to define and manage Symmetrix devices. The use of these functions and utilities is optional. They are listed for reference only.

◆ `disksecn`

Manually creates a device ID.

◆ `ioscan`

An online hardware mapping program that provides a listing of devices connected to the host. It scans the system hardware for usable I/O systems devices or kernel I/O system data structures as appropriate and lists the results. Useful options for this function are:

<code>-C disk</code>	Restrict the output to all devices belonging to the disk class.
<code>-f</code>	Generate a full listing.
<code>-F</code>	Generate a compact listing.
<code>-n</code>	List device file names in the output.
<code>-M -H <hardware path></code>	Force the software driver specified hardware path) into the kernel.

◆ `insf`

Installs special files and assigns logical unit numbers to devices in the system. Used for introducing newly attached devices to the system without rebooting.

◆ `mknod`

Creates a directory, special, or ordinary file.

◆ `pvccreate`

Creates a physical volume.

- ◆ `pvchange`
Changes characteristics of a physical volume in a volume group. It also redirects I/O to different physical volume paths.
- ◆ `pvdisk`
Displays information about physical drives within a volume group.
- ◆ `vgscan`
Scans physical volumes for volume groups.
- ◆ `vgcreate`
Create volume group on a set of physical volumes.
- ◆ `vgchange`
Sets volume group availability.
- ◆ `vgdisplay`
Displays information about volume groups.
- ◆ `vgexport`
Exports a volume group and its associated logical volumes.
- ◆ `vgimport`
Imports a volume group onto a system.
- ◆ `vgremove`
Removes a volume group definition from the system.
- ◆ `lvcreate`
Stripes; creates logical volume in volume group.

- ◆ `lvchange`
Changes logical volume characteristics.
- ◆ `lvdisplay`
Displays information about logical volumes.
- ◆ `lvextend`
Increases space for logical volumes, but not size of filesystems.
- ◆ `lvreduce`
Decreases number of physical extents allocated to logical volume.
- ◆ `lvremove`
Removes one or more logical volumes.
- ◆ `newfs`
Creates a new filesystem.
- ◆ `bdf`
Shows all mounted filesystems and available free space.
- ◆ `lsdev`
Lists device drivers in the system.

Symmetrix Configuration

In the HP 9000 Server host environment, you can configure the Symmetrix disk devices into logical volumes. EMC recommends the following logical volume ratios as indicated in the table below:

Table 6-4 Disk Device Configuration

		Symmetrix Logical Volume to Physical Splits ^a							
		1:1	2:1	3:1	4:1	5:1	6:1	7:1	8:1 ^b
Physical Devices	4 GB	x	x						
	9 GB	x	x	x	x				
	18 GB		x	x	x	x	x	x	x
	23 GB		*	x	x	x	x	x	x
	36 GB			*	x	x	x	x	x
	47 GB			*	*	x	x	x	x
	50 GB ^c				x	x	x	x	x
	73 GB ^{c,d}					x	x	x	x

- a. * indicates ratios available only in microcode level 5265 or higher.
- b. With microcode revision level 5263, maximum is 8:1. With 5264, maximum is 16:1. With 5265, maximum is 32:1.
- c. The 50 and 73 GB drives are supported only by microcode revision levels 5x66 and higher.
- d. The 73 GB drive does not support the RAID-S configuration on Symmetrix.

For specific cylinder counts and logical volume sizes, refer to Table 2-1 on page 2-6 and the following tables in that chapter.

The EMC Customer Engineer (CE) should contact the EMC Configuration Specialist for updated online information. This information is necessary to configure Symmetrix to support the customer's host environment.

Using Symmetrix as a Root Disk

The high availability and data protection features of Symmetrix make its disks attractive candidates for selection as the system root disk. A Symmetrix device can be used as root disk; however, the following should be considered in the decision:

- ◆ Root, stand, and primary swap must reside within the first 2 GB of the physical boot disk.
- ◆ The time required to boot Symmetrix following a power outage may be significantly longer than the host boot time.

Creating a Bootable Disk (10.x & 11.x)

This section describes how to create an external boot disk other than the current root disk. This is an unusual configuration and you should consider other options for ensuring system integrity such as Mirror/UX. For this example, the following configuration will be used:




Figure 6-1 Bootable Disk - Sample Configuration

To create an external boot disk, follow the steps below:

1. Create a new volume group containing the disk c2t0d0 by typing statements similar to the following:

```
pvccreate -B /dev/rdisk/c2t0d0
mkboot /dev/rdisk/c2t0d0
mkboot -a "hpux (8/8.0.0)/stand/vmunix" /dev/rdisk/c2t0d0
mkdir /dev/vgalt
mknod /dev/vgalt/group c 64 0x010000
vgcreate /dev/vgalt /devdisk/c2t0d0
```

2. Create the logical volumes that are required to boot the system (root, swap, and usr) in the new volume group by typing statements similar to the following:

```
lvcreate -L 200 -n altroot -r n -s y -C y /dev/vgalt
lvcreate -L 500 -n altswap -r n -s y -C y /dev/vgalt
lvcreate -l 50 -n altdump -r n -s y -C y /dev/vgalt
lvcreate -L 100 -n altusr /dev/vgalt
```

The sizes here are arbitrary. Create the sizes you require.

3. Create filesystems and copy data from the original filesystems to the newly created filesystems using statements similar to the ones below:

```
newfs -F hfs /dev/vgalt/raltroot
newfs -F hfs /dev/vgalt/raltusr

mkdir /altroot
mkdir /altusr
mount /dev/vgalt/altroot /altroot
mount /dev/vgalt/altusr /altusr

cd /
find . -xdev | cpio -pdmux /altroot
cd /usr
find . -xdev | cpio -pdmux /altusr
```

4. Update BDRA for the new volume group to ensure that the system can boot from the alternative disk using statements similar to the ones below:

```
lvlnboot -r /dev/vgalt/altroot /dev/vgalt
lvlnboot -s /dev/vgalt/altswap /dev/vgalt
lvlnboot -d /dev/vgalt/altdump /dev/vgalt
```

5. Check the boot configurations by entering the following statement:

```
lvlnboot -v
```

The boot configuration should appear similar to the following:

```
Boot Definitions for Volume Group /dev/vg00:
Physical Volumes belonging in Root Volume Group:
    /dev/dsk/c0t0d0 -- Boot Disk
Root:   lvollon:/dev/dsk/c0t0d0
Swap:   lvoll2on:/dev/dsk/c0t0d0
Dump:   lvoll8on:/dev/dsk/c0t0d0
```

```
Boot Definitions for Volume Group /dev/vgalt:
Physical Volumes belongs in Root Volume Group:
    /dev/dsk/c2t0d0 -- Boot Disk
Root:   altrooton:/dev/dsk/c2t0d0
Swap:   altswapon:/dev/dsk/c2t0d0
Dump:   altdumpon:/dev/dsk/c2t0d0
```

6. Update the **fstab** file on the alternate boot disk to reflect its 'root', 'swap', and 'usr', and any other filesystems by entering the following statement:

```
vi /altroot/etc/fstab
```

Mount /altroot first.

Change the vg name and the lvol names to reflect their correct assignments.

7. Try rebooting from the alternate disk by booting from the hardware path of the new disk as follows:
 - a. Reboot the system by typing:

```
shutdown -r now
```

- b. Press any key and interrupt the boot process when the interact with boot message appears.
 - c. Change the primary boot path to the new boot disk by typing a statement similar to the following:

```
pa pri 8/8.0
```

- d. Verify that the change took place by typing:

```
pa pri Displays the new boot path
```

e. Boot from the new disk by typing:

```
bo pri Starts boot from primary boot path
```

f. Answer `no` to interact with the boot process.

Hardware Mapping 10.0 & 11.x (offline)

After the EMC Customer Engineer has assigned target IDs and LUNs to the Symmetrix devices and the Symmetrix channel and disk directors have been switched online, use the following procedure to verify that the host can see all the target Symmetrix devices assigned to that host. (Not for V series.)

1. Power on the HP-UX host system and wait for it to perform its self-tests.
2. When "Autoboot primary path enabled. To override, press any key within 10 seconds." appears on the system console, press any key.
3. At the "Boot from primary boot path (Y or N)?>" prompt appears, type `Y`
4. At the "Interact with IPL (Y or N)?>" prompt, type `Y`.
5. At the "ISL>" prompt, type `ls`
6. At the "ISL>" prompt, type `ODE`

Support CD Diagnostics must be loaded to use the ODE command (Offline Diagnostics Environment). For instructions on installing the Lif-Load product, refer to the *Support Media Users Manual*, HP P/N 92453-90010.

7. At the "ODE>" prompt, type `ls`
8. At the "ODE>" prompt, type `mapper`
9. At the "MAPPER>" prompt, type `run`

A screen similar to the one shown below appears.

Processor Identification:

```
Hardware Model: 280Hm Revision:      0
Software Model:4H, Revision:0
Hardware ID:0, Software ID: 324676221 (unsigned decimal)
Software Cap: 0x100000f0
HPUX Model String:9000/827
Processor Board Revisions:
    CPU - CPU Chip:      5
    System Controller:   3
    PDC - Processor Dependent Code:1.5
```

Cache and TLB Sizes:

```
Instruction Cache: 64K bytes, Instruction TLB: 96 Entries
Data Cache:      64K bytes, Data TLB:96 Entries
```

Coprocessors:None Installed

I/O Configuration:

```
Path  Component Name Type HW SW  Revisions
      ID  ModModHardwareFirmware
=====
8      HP-PB Fast Wide SCSI Card
8.0.0 Unknown SCSI device type (SYMMETRIX)
8.0.1 Unknown SCSI device type (SYMMETRIX)
8.0.2 Unknown SCSI device type (SYMMETRIX)
8.0.3 Unknown SCSI device type (SYMMETRIX)
.
.
.
```

Figure 6-2 I/O Configuration (HP-UX 10.0)

10. To exit from MAPPER, type `exit`
11. To exit from ODE, type `exit`
12. At the “ISL>” prompt, type `hpux`

The host operating system boots to the console login prompt.

13. Log in as **root** to perform the remaining functions described in this section.

Recognizing New Devices and LUNs

Now that you have a hardware map of the Symmetrix devices recognized by the host system, the host is now able to identify the connected devices and create the raw special device and the character files in **/dev/dsk**.

To verify that the host has identified the connected devices, type the following:

```
ioscan -fn [> filename]
```

where *filename* is the name of the file that will store this information. A listing of all devices found will be stored in the given filename.

Then type

```
insf -e
```

Or you can reboot the host.

HP-UX 10.x and 11.x

In **HP-UX 10.x and 11.x**, your screen display will appear similar to the following:

```
/etc/ioscan -fn > filename
more filename
.
.
.
disk 2    3.0/20.0.0disc3 CLAIMED DEVICE EMC SYMMETRIX
/dev/dsk/c3t0d0 /dev/rdisk/c3t0d0
/dev/floppy/c3t0d0/dev/rdisk/c3t0d0

disk2     4.0/20.0.1disk3 CLAIMED DEVICE EMC SYMMETRIX
/dev/dsk/c3t0d1 /dev/rdisk/c3t0d1
/dev/floppy/c3t0d1/dev/rdisk/c3t0d1
```

Note the card instance number next to the H/W Path for the disk devices.

Also note that the H/W Path will differ from one machine model to another.

Slice (partition) is not supported by HP 10.

For Fibre Channel

```
disk 603 8/0.8.0.10.3.1.1 sdiskCLAIMED DEVICE EMC SYMMETRIX  
                               /dev/dsk/c28t1d1           /dev/rdisk/c28t1d1
```

where in 10.3.1.1

10 is the loop ID.

3 is the bus.

1 is the target ID.

1 (rightmost) is the LUN.

where in c28t1d1

28 is the controller.

1 is the target ID.

1 (rightmost) is the LUN.

Configuring Logical Volumes

A logical volume is a flexible model for storage management. One or more physical devices are introduced into a volume group. Logical volumes are created from storage available for the volume group. Additional functionality such as striping, volumes that span multiple physical devices, etc. are achieved by volume management. You need to understand LVMs before implementing logical volumes on your host. For more information on logical volumes, refer to Chapter 9 in the *HP Concepts for the System Administrator* (P/N B2355-90029).

You must configure logical volumes to store data when using HP 10.x and 11.x.

Creating Physical Volumes

After you have verified that the host has recognized each Symmetrix device, you need to create a physical volume for each device.

To create a physical volume for each Symmetrix disk device (logical volume), enter a statement similar to the following for each physical volume to create:

For **HP-UX 10.x and HP-UX 11.x**:

```
pvccreate /dev/rdisk/c3t0d0
```

This statement creates a physical volume, `"/dev/rdisk/c3t0d0"`.

Physical volumes can also be created using the SAM utility.

Volume Groups and Logical Volumes

At this time you can use physical volumes or create volume groups and logical volumes. Volume groups and logical volumes can be created using HP-UX stand-alone commands or SAM. To create volume groups and logical volumes, refer to the HP documentation set.

Creating the Mount Directory

If you want to use filesystems on the volumes, you must create a mount directory for each of these volumes. To do this, type a statement similar to the following for each physical volume:

```
mkdir /symm
```

This example creates a mount directory called "symm".

Creating a New Filesystem

Now that the physical volumes and their respective mount directories are available to the operating system, a new filesystem must be created for each Symmetrix disk or LVM. To do this, use the **newfs** command in a statement similar to the following when using vxfs:

```
newfs /dev/rdisk/c27d0s2 symm2.9gb
```

Specify the disk name matching the volume size defined in your Symmetrix unit for this statement for hfs when volume groups are used in a statement similar to the following:

```
newfs /dev/vg-group/volx
```

Two file system types are used in HP-UX: vxfs and hfs. The default file system of the particular operating system utilized would be chosen unless the **-F hfs** or **-F vxfs** option is also included in the command line.

Mount the Directory

Before you can mount the directory in HP-UX 9.04, you need to add the path to **/etc/checklist**. To do this, type the following statement:

```
vi /etc/checklist
```

Specify each device in a statement similar to the following for HP-UX 9.04:

```
/dev/dsk/c27d0s2 /syemm hfs defaults 0 3 0
```

For HP-UX 10.x and volume groups, mount the directory by adding the path to **/etc/fstab**. To do this, type the following statement:

```
vi /etc/fstab
```

Specify each device in a statement similar to the following:

```
/dev/vg-group/volx
```

Mount the filesystem by typing a statement similar to the following in either operating system:

```
/etc/mount /syemm
```

Use the **bdf** command to show all mounted filesystems and the available free space. For example,

```
bdf
```

Filesystem	kbytes	used	available	%used	mounted on
/dev/vg00/lvol1	104299	52936	40933	56%	/
/dev/vg00/lvol5	349530	125692	188884	40%	/users
/dev/vg00/lvol4	124507	836	111220	1%	/tmp
/dev/vg00/lvol3	397867	137494	220586	38%	/usr
/dev/dsk/c10d0s22741142		104041	2362986	4%	/syemm5
/dev/dsk/c11d0s22741142		0	2741142	0%	/syemm
/dev/dsk/c12d0s22741142		104041	2362986	4%	/syemm3
/dev/dsk/c13d0s22741142		104041	2362986	4%	/syemm1
/dev/dsk/c29d0s22741142		104041	2362986	4%	/syemm4
/dev/dsk/c30d0s22741142		104041	2362986	4%	/syemm2

shows **/dev/dsk/c11d0s2** mounted.

Using SAM

HP-UX has a menu-driven utility for managing system resources, specifically disk drives. This utility is called SAM and users familiar with it might find it more useful to use SAM for configuring new Symmetrix disks.

HP 9000 K- and D-class Servers

The HP-9000 K-class servers are configured according to the options listed in the table below. Both of these HP controllers are supported by Symmetrix.

Table 6-5 K-class Server Configuration

Model	K100	K200	K400	K5xx	
				4-way	6-way
Max HP-PB I/O slots ^a	4	4	8	2	2
Max HP-HSC slots ^b	1	1	5	9	5

a. The HP-PB SCSI Controller model number is 28696A.

b. The HP-HSC SCSI Controller model number is 2969A.

The HP-9000 D-class servers are configured according to the options listed in the table below.

Table 6-6 HP-9000 D-class Servers

Model	D200/210	D250	D260	D270	D310	D350	D360	D370
Maximum number of HP-HSC slots ^a	5	5	5	5	8	8	8	8
Total HSC or EISA slots 0, 2, 3	3	3	3	3	2	3	3	3
Internal Disk (HSC) Slot 1	1	1	1	1	1	1	1	1

a. The HP-HSC SCSI controller model number is A4107A.

Channel Failover (10.x and 11.x)

Your particular host may provide the capability for dual cabling (using dual host controllers) to the same physical volume. If so, the Logical Volume Manager (LVM) can be configured with multiple paths to the same physical volume. In this configuration, if the primary path to the volume fails, the host automatically switches the path to the volume to an alternate connection or link without affecting the application.

If you are able to connect to EMC's anonymous FTP server, refer to the Channel Failover example in the **/pub/symm3000/hpux** directory. Otherwise, refer to the HP 10.x documentation for more information on this feature.

Setting Up Alternate Links

To set up alternate links, you must first create a physical volume for the device to become a primary link. For example,

```
pvcreate /dev/rdisk/c3t0d0
```



CAUTION

Do not issue `pvcreate` to the alternate path. This action will corrupt the device if data is on it from the primary path.

Then you create a volume group using the **vgcreate** command specifying *ONLY* the primary link device file name. You then use the **vgextend** command specifying BOTH the primary link and alternate link device file names add more devices or more paths.

Both device file names must represent two physical paths to the same Symmetrix logical volume. LVM configures the first device name as the primary link and the second device name as the alternate link.

When using the **vgextend** command, the primary path must be listed first.

For example,

```
vgcreate /dev/vg01 /dev/dsk/c3t0d0
```

creates the volume group **vg01** with **/dev/dsk/c3t0d0** as the path device.

To add the device, **/dev/dsk/c5t0d0**, as the alternate link to volume group **vg01**, type a command similar to the following:

```
vgextend /dev/vg01 /dev/dsk/c5t0d0
```

In this example, SCSI channels **c3** and **c5** are attached to different Symmetrix directors (see Figure 6-4). For load balancing and faster failovers, combine primary and alternate devices on the same channel. In Figure 6-4, the first four devices of **c3** are attached to FW SCSI channel director **SA-15**, port **C**. The last four devices of **c3** are attached to FW SCSI channel director **SA-16**, port **C**.

To add an alternate link to a physical volume that is already part of a volume group, use the **vgextend** command.

For example, if **/dev/dsk/c2t0d0** is already part of your volume group, but you wish to add **/dev/dsk/c4t0d0** to the physical volume, type:

```
vgextend /dev/vg02 /dev/dsk/c4t0d0
```

Primary Link Failure

If the primary link fails, LVM will automatically switch from the primary path to the device to the alternate link path on the second controller. While in this failure mode, LVM can also switch to a different controller at any time by entering a command similar to the following:

```
pvchange -s /dev/dsk/c2t1d0
```

When the primary link recovers, LVM automatically switches from the alternate link path to the primary link path. If you do not wish this automatic switch to occur, use the **pvchange** command in a statement similar to the following:

```
pvchange -S n /dev/dsk/c2t2d0
```

Viewing the Current Links

You can view the current links to a physical volume using the **vgdisplay** command with the **-v** option (Figure 6-3).

```

--- Volume groups ---
VG Name                /dev/vg_dd_1
VG Write Access         read/write
VG Status               available
Max LV                 255
Cur LV                 1
Open LV                 1
Max PV                 16
Cur PV                 1
Act PV                 1
Max PE per PV          1038
VGDA                   2
PE Size (Mbytes)        4
Total PE               1038
Alloc PE               512
Free PE                526
Total PVG              0

--- Logical volumes ---
LV Name                /dev/vg_dd_1/lvol1
LV Status               available/syncd
LV Size (Mbytes)        2048
Current LE              512
Allocated PE            512
Used PV                 1

--- Physical volumes ---
PV Name                /dev/dsk/c0t0d2
PV Name                /dev/dsk/c3t0d2      Alternate Link
PV Status               available
Total PE               1038
Free PE                526

```

Figure 6-3 Current Links (vgdisplay option)

Figure 6-4 shows the internal Symmetrix links to the volumes.




Figure 6-4 Symmetrix Volume Configuration

When assigning devices for primary and alternate paths, EMC recommends using different target IDs to identify the paths to reduce the possibility of mistakes when configuring LVM especially in cluster configurations.

For example, in Figure 6-4 above, device on the primary link path would be assigned TIDs of 1. Devices on the alternate link path would be assigned TIDs of 2.

Common Serial Number Option

Symmetrix typically generates its serial number information using the last two digits of the cabinet serial number, the three-digit Symmetrix device number, the director number to which that device is attached and the port number to which that device is connected.

The EMC Customer Engineer can set the Common Serial Number option on a port-by-port basis during Symmetrix configuration. This option causes the last three digits of the serial number information generated by Symmetrix (director number and port) to be set to zero, on all Symmetrix ports with the option enabled for the same Symmetrix device.

When using HP's alternate channel failover feature (PVL), the Common Serial Number option is not required, BUT, it is required for SAM to recognize alternate paths and allow configuration of alternate links using this utility. Therefore, when using SAM, be sure that the Common Serial Number option is set.

You can, however, configure alternate links from the command line without the Common Serial Number option enabled. Please be aware that there may also be I/O hardware scanning scripts that may rely on common serial numbers to determine multiple paths to a device.

Adding Devices Online

Whenever devices are added online to the Symmetrix unit or device channel addresses are changed, you will need to perform the actions described below in order to introduce the new devices to the system.

You can define newly connected physical volumes to the system without rebooting the host. This requires use of the **insf** command in a statement similar to the following:

```
insf -H 2/4.0.0
```

where **2/4.0.0** is a hardware path to which the new device has been connected.

Alternatively, you can use the following command sequence to look for all new hardware devices:

```
ioscan -fn  
insf -e
```

System and Error Messages

HP-UX 10.x and 11.x log system and error messages including SCSI error messages to a file called `/var/adm/syslog/syslog.log`.

Diagnostic Patches

Table 6-7 lists HP's Diagnostics Cumulative Patches. Installing the appropriate patches for your operating system level causes Symmetrix drives to be recognized properly by HP hosts.

Table 6-7 Diagnostic Cumulative Patches

Operating System	Patch Number
10.01	PHSS_14401
10.10	PHSS_17884
10.20	PHSS_17886 Symmetrix definition should already exist for STM and Sysdiag

These patches have an affect on the output of the STM and sysdiag utilities when run on Symmetrix devices. The affect of these patches are as follows:

- ◆ **SCSIDSK2** - HP's SCSI disk diagnostic utility (which includes commands like **INQUIRY**, **CAPACITY**, etc.) now works with Symmetrix devices, but reports the following message:

`SCSIDSK2 is not designed to test the selected device (which is "EMCSYMMETRIX")`

- ◆ The **IOMAP** utility recognizes Symmetrix devices as **EMCSYMMETRIX DISK**. (If the patch is not installed, Symmetrix devices are reported as "Device Unknown".)
- ◆ In the System Error Log Report, any entries in the log for Symmetrix devices now report the product name as "EMCSYMMETRIX" and the device class as "Hard Disk".

`logtool` decodes the hex bytes presented in the environmental data to text. Table 6-8 lists HP's patches for your operating system level for `logtool.s`

Table 6-8 Logtool Decoding Patches

Operating System	Patch Number
10.01	PHSS_14401
10.10	PHSS_17884
10.20	PHSS_17886

HP 9000 Workstation Environment

This chapter describes how to incorporate Symmetrix in the HP-UX environment on HP 9000 workstations.

◆ Overview	7-2
◆ Useful HP-UX Utilities and Functions.....	7-5
◆ Symmetrix Configuration	7-8
◆ Hardware Mapping 10.x & 11.x (offline)	7-12
◆ Recognizing New Devices and LUNs.....	7-15
◆ Configuring Logical Volumes	7-17
◆ Using SAM.....	7-19
◆ Adding Devices Online	7-20
◆ Performance Issues with HP-UX	7-21
◆ System and Error Messages.....	7-22

Overview

When using Symmetrix with HP 9000 workstations, please note the following:

- ◆ Symmetrix attaches to the HP 9000 workstations listed in Table 7-1 below.

Table 7-1 HP/9000 Workstation Matrix

HP Workstation	FW-SCSI Differential Channels	External Slots	Bus ^a	Host Adapter Card Product	External EISA	Operating System
715/100XC	1	1	EXT GSC	A4107A	n/a	11.x
725/100	4	4	EXT GSC	A4107A	see note ^b	10.20/11.x
735/99	1	0	internal	n/a	see note ^b	10.20
735/125	1	0	internal	n/a	see note ^b	10.20
755/99	1	0	internal	n/a	see note ^b	10.20
755/125	1	0	internal	n/a	see note ^b	10.20
C Series	3	4	EXT GSC	A4107A	n/a	11.x
J-200	3	5	EXT GSC	A4107A	n/a	11.x
J-210	3	5	EXT GSC	A4107A	n/a	11.x
J210CX	3	5	EXT GSC	A4107A	n/a	11.x

- The internal bus is a bus brought out to the rear of the host for connection to external devices. Determine the target ID of the internal bus using the `d.f` command before attaching Symmetrix to this bus. In general, the target ID of an internal bus is either 5 or 6.
 - EISA cards are BusLogic cards supplied by SBE Corporation or MDL Corporation. These cards require third party drivers. These cards are not supported by either Hewlett Packard or EMC at this time.
- ◆ For Symmetrix models 55/3500, 52/32xx, and 51/31xx, EMC recommends that you use the Fast-Wide Differential SCSI director P/N 200-881-903.

- ◆ For Symmetrix models 57/3700, 54/34xx, and 53/33xx, EMC recommends that you use either the Fast-Wide Differential SCSI director P/N 201-207-917 (60 MHz processor) or P/N 201-207-927 (66 MHz processor) or Ultra-Wide Differential SCSI director P/N 201-277-917 (60 MHz processor) or P/N 201-277-927 (66 MHz processor).
- ◆ For Symmetrix support you need HP-UX version 9.05 and higher.
- ◆ When using LUNs, do not skip LUN numbers and always start from LUN0.
- ◆ Avoid using target ID 7 which is reserved for the host adapter and target IDs 5 and 6 which are typically used by the internal drives.

The table below summarizes the operating system features in the HP-UX environment.

Table 7-2 HP 9000 Workstation OS Features

O/S	HP-UX
Supports multiple target IDs	Yes ^a
Supports multiple LUNs per target ID	Yes
Partition representing the entire disk	s2
To see the drive configuration, use this utility:	ioscan or SAM
To configure drives use this utility:	SAM or line commands
Special drive definition files needed	When below HP-UX 10

a. LUNs 0-7, targets 0-F.

The table below summarizes the limits of recent HP-UX releases (but not the limits of the workstations).

Table 7-3 Capacity Limits of Recent HP-UX Releases

Attribute	10.01	10.10	10.20	64-bit 11.00
File System	4 GB	128 GB	128 GB	128 GB
File Size	2 GB	2 GB	128 GB Local, 2 GB Network	128 GB local and network
Physical RAM	2 GB	3.75 GB	8 GB	4 TB
Shared Memory	1.75 GB	1.75 GB	2.75 GB	8 TB
Process Data Space	0.9 GB	1.9 GB	1.9 GB	4 TB
Number of File Descriptors	2 K	60 K	60 K	60 K
Number of User IDs	60 K	60 K	~2,000 K	~2,000 K

Useful HP-UX Utilities and Functions

This section lists HP-UX functions and utilities you may use to define and manage Symmetrix devices. The use of these functions and utilities is optional. They are listed for reference only.

HP-UX 10.x & 11.x

◆ ioscan

An online hardware mapping program that provides a listing of devices connected to the host. It scans the system hardware for usable I/O systems devices or kernel I/O system data structures as appropriate and lists the results. Useful options for this function are:

<code>-C disk</code>	Restrict the output to all devices belonging to the disk class.
<code>-f</code>	Generate a full listing.
<code>-F</code>	Generate a compact listing.
<code>-n</code>	List device file names in the output.
<code>-M -H <hardware path></code>	Force the software driver specified hardware path) into the kernel.

◆ insf

Installs special files and assigns logical unit numbers to devices in the system. Used for introducing newly attached devices to the system without rebooting.

◆ mknod

Creates a directory, special, or ordinary file.

◆ pvcreate

Creates a physical volume.

◆ pvchange

Changes characteristics of a physical volume in a volume group.

- ◆ `pvdisplay`
Displays information about physical drives within a volume group.
- ◆ `vgscan`
Scans physical volumes for volume groups.
- ◆ `vgcreate`
Create volume group on a set of physical volumes.
- ◆ `vgchange`
Sets volume group availability.
- ◆ `vgdisplay`
Displays information about volume groups.
- ◆ `vgexport`
Exports a volume group and its associated logical volumes.
- ◆ `vgimport`
Imports a volume group onto a system.
- ◆ `vgremove`
Removes a volume group definition from the system.
- ◆ `lvcreate`
Stripes; creates logical volume in volume group.
- ◆ `lvchange`
Changes logical volume characteristics.

- ◆ `lvdisplay`
Displays information about logical volumes.
- ◆ `lvextend`
Increases space for logical volumes.
- ◆ `lvreduce`
Decreases number of physical extents allocated to logical volume.
- ◆ `lvremove`
Removes one or more logical volumes.
- ◆ `newfs`
Creates a new filesystem.
- ◆ `bdf`
Shows all mounted filesystems and available free space.
- ◆ `lsdev`
Lists device drivers in the system.

Symmetrix Configuration

In the HP 9000 Workstation host environment, you can configure the Symmetrix disk devices into logical volumes. EMC recommends the following logical volume ratios as indicated in the table below:

Table 7-4 Disk Device Configuration

		Symmetrix Logical Volume to Physical Splits ^a							
		1:1	2:1	3:1	4:1	5:1	6:1	7:1	8:1 ^b
Physical Devices	4 GB	x	x						
	9 GB	x	x	x	x				
	18 GB		x	x	x	x	x	x	x
	23 GB		*	x	x	x	x	x	x
	36 GB			*	x	x	x	x	x
	47 GB			*	*	x	x	x	x
	50 GB ^c				x	x	x	x	x
	73 GB ^{c, d}					x	x	x	x

- a. * indicates ratios available only in microcode level 5265 or higher.
- b. With microcode revision level 5263, maximum is 8:1. With 5264, maximum is 16:1. With 5265, maximum is 32:1.
- c. The 50 and 73 GB drives are supported only by microcode revision levels 5x66 and higher.
- d. The 73 GB drive does not support the RAID-S configuration on Symmetrix.

For specific cylinder counts and logical volume sizes, refer to Table 2-1 on page 2-5 and the following tables in that chapter.

The EMC Customer Engineer (CE) should contact the EMC Configuration Specialist for updated online information. This information is necessary to configure Symmetrix to support the customer's host environment.

Using Symmetrix as a Root Disk

The high availability and data protection features of Symmetrix make its disks attractive candidates for selection as the system root disk. A Symmetrix device can be used as root disk; however, the following should be considered in the decision:

- ◆ EMC recommends that a Symmetrix device should not be used as the only root disk. Another root disk should be maintained on a local bus to aid in fault isolation and to ensure that the system can be booted by itself if Symmetrix is offline.
- ◆ The time required to boot Symmetrix following a power outage may be significantly longer than the host boot time.

Creating a Bootable Disk

This section describes how to create an external boot disk other than the current root disk. This is an unusual configuration and you should consider other options for ensuring system integrity such as, for example, Mirror/UX.

For this example, the following configuration will be used:




Figure 7-1 Bootable Disk - Sample Configuration

To create an external boot disk, follow the steps below:

1. Create a new volume group containing the disk c2t0d0 by typing statements similar to the following:

```
pvccreate -B /dev/rdisk/c2t0d0
mkboot /dev/rdisk/c2t0d0
mkboot -a "hpux (8/8.0.0)/stand/vmunix" /dev/rdisk/c2t0d0
mkdir /dev/vgalt
mknod /dev/vgalt/group c 64 0x010000
vgcreate /dev/vgalt /devdisk/c2t0d0
```

2. Create the logical volumes that are required to boot the system (root, swap, and usr) in the new volume group by typing statements similar to the following:

```
lvcreate -L 200 -n altroot -r n -s y -C y /dev/vgalt
lvcreate -L 500 -n altswap -r n -s y -C y /dev/vgalt
lvcreate -l 50 -n altdump -r n -s y -C y /dev/vgalt
lvcreate -L 100 -n altusr /dev/vgalt
```

The sizes here are arbitrary. Create the sizes you require.

3. Create filesystems and copy data from the original filesystems to the newly created filesystems using statements similar to the ones below:

```
newfs -F hfs /dev/vgalt/raltroot
newfs -F hfs /dev/vgalt/raltusr

mkdir /altroot
mkdir /altusr
mount /dev/vgalt/altroot /altroot
mount /dev/vgalt/altusr /altusr

cd /
find . -xdev | cpio -pdmux /altroot
cd /usr
find . -xdev | cpio -pdmux /altusr
```

4. Update BDRA for the new volume group to ensure that the system can boot from the alternative disk using statements similar to the ones below:

```
lvlnboot -r /dev/vgalt/altroot /dev/vgalt
lvlnboot -s /dev/vgalt/altswap /dev/vgalt
lvlnboot -d /dev/vgalt/altdump /dev/vgalt
```

5. Check the boot configurations by entering the following statement:

```
lvlnboot -v
```

The boot configuration should appear similar to the following:

```
Boot Definitions for Volume Group /dev/vg00:
Physical Volumes belonging in Root Volume Group:
  /dev/dsk/c0t0d0 -- Boot Disk
Root: lvollon:    /dev/dsk/c0t0d0
Swap: lvoll2on:  /dev/dsk/c0t0d0
```



```
Dump: lvol8on: /dev/dsk/c0t0d0
```

```
Boot Definitions for Volume Group /dev/vgalt:
Physical Volumes belongs in Root Volume Group:
/dev/dsk/c2t0d0 -- Boot Disk
Root: altrooton: /dev/dsk/c2t0d0
Swap: altswapon: /dev/dsk/c2t0d0
Dump: altdumpon: /dev/dsk/c2t0d0
```

6. Update the **fstab** file on the alternate boot disk to reflect its 'root', 'swap', and 'usr', and any other filesystems by entering the following statement:

```
vi /altroot/etc/fstab
```

```
Mount /altroot first.
```

Change vg name and the lvol names to reflect their correct assignments.

7. Try rebooting from the alternate disk by booting from the hardware path of the new disk as follows:

- a. Reboot the system by typing:

```
shutdown -h now
```

- b. Press any key and interrupt the boot process when the interact with boot message appears.
- c. Change the primary boot path to the new boot disk by typing a statement similar to the following:

```
pa pri 8/8.0
```

- d. Verify that the change took place by typing:

```
pa pri (Displays the new boot path)
```

- e. Boot from the new disk by typing:

```
bo pri (Starts boot from primary boot path)
```

- f. Answer no to interact with the boot process.

Hardware Mapping 10.x & 11.x (offline)

After the EMC Customer Engineer has assigned target IDs and LUNs to the Symmetrix devices and the Symmetrix channel and disk directors have been switched online, use the following procedure to verify that the host can see all the target Symmetrix devices assigned to that host.

1. Power on the HP-UX host system and wait for it to perform its self-tests.
2. Press and hold down the ESC key for approximately 5 seconds.

If the "BOOT_ADMIN>" prompt appears, follow these steps.

3. At the "BOOT_ADMIN>" prompt, type
`boot pri isl`
4. At the "ISL>" prompt, type `ls`
5. At the "ISL>" prompt, type `ODE`

Support CD Diagnostics must be loaded to use the ODE command (Offline Diagnostics Environment). For instructions on installing the Lif-Load product, refer to the *Support Media Users Manual*, HP P/N 92453-90010.

Continue with step 11 below.

If you don't see the "BOOT_ADMIN>" prompt, follow these steps

6. At the "Selecting a system to boot. To stop selection process, press and hold ESCAPE key.", press the ESC key.

When the Boot selection process stops, a menu appears.

7. Choose **a** (Enter Boot Administration Mode) from the menu.
8. At the "BOOT_ADMIN>" prompt, type
`boot pri isl`
9. At the "ISL>" prompt, type `ls`
10. At the "ISL>" prompt, type `ODE`

Support CD Diagnostics must be loaded to use the ODE command (Offline Diagnostics Environment). For instructions on installing the Lif-Load product, refer to the *Support Media Users Manual*, HP P/N 92453-90010.

11. At the "ODE>" prompt, type `ls`
12. At the "ODE>" prompt, type `mapper`
13. At the "MAPPER>" prompt, type `run`

A screen similar to the one shown in Figure 7-2 for a 735/125 Workstation appears.

Processor Identification:

```

Hardware Model:      206H Revision:      0
Software Model:      4H, Revision:      0
Hardware ID: 0, Software ID: 2015770545 (unsigned decimal)
Software Cap:      0xFFFFFFFF
HPUX Model String:   9000/735
Processor Board Revisions:
    CPU - CPU Chip:          20
    System Controller:       3
    PDC - Processor Dependent Code:  2.7

```

Cache and TLB Sizes:

```

Instruction Cache:  256K bytes, Instruction TLB:120 Entries
Data Cache:        256K bytes, Data TLB:      120 Entries

```

Coprocessors: Floating Point Co-processor is installed

I/O Configuration:

Path	Component Name	Type ID	HW Mod	SW Mod	Hardware	Revisions Firmware
1	Built-in Graphics					
2	Comm Bus Adapter					
2/0/1	Core SCSI					
2/0/2	Core LAN					
2/0/3	Core HiL					
2/0/7.0.0	Unknown SCSI device type (SYMMETRIX)					
.						
.						
.						

Fast-Wide Internal Bus Controller

Figure 7-2 I/O Configuration (HP-UX 10.x)

14. To exit from MAPPER, type `exit`
15. To exit from ODE, type `exit`
16. At the “ISL>” prompt, type `hpux`

The host operating system boots to the console login prompt.

17. Log in as **root** to perform the remaining functions described in this section.

Recognizing New Devices and LUNs

The host is now able to identify the connected devices and create the raw special device and the character files in **/dev/dsk**.

HP-UX 9.04

To verify that the host has identified the connected devices, type the following:

```
/etc/ioscan -fn [> filename]
```

where **filename** is the name of file in which to store this information.

A listing of all devices found will be stored in the given filename.

```
/etc/ioscan -fn > filename
more filename
Class      H/W Path  Driver      H/W Status   S/W Status
=====
graphics  1.0.0      graph3      ok (0x577)   ok
scsi       2.0.1      c700        ok (0x7071)  ok
disk       2.0.1.0.0  scsi        ok (0x202)   ok
lan        2.0.2      lan01       ok (0x7072)  ok
hil        2.0.3      hil         ok (0x7073)  ok
serial     2.0.4      asio0       ok (0x7075)  ok
serial     2.0.5      asio0       ok (0x7075)  ok
parallel   2.0.6      parallel    ok (0x7074)  ok
scsi       2.0.7      c700        ok (0x707c)  ok
disk       2.0.7.0.0  scsi        ok (0x202)   ok
disk       2.0.7.1.0  scsi        ok (0x202)   ok
disk       2.0.7.5.0  scsi        ok (0x202)   ok
disk       2.0.7.6.0  scsi        ok (0x202)   ok
audio      2.0.8      audio       ok (0x707f)  ok
```

Internal bus controller

Internal bus controllers
(target IDs 5 and 6)

Symmetrix supports the use of s2 only.

HP-UX 10.x and 11.x

To verify that the host has identified the connected devices, type the following:

```
/etc/ioscan -fnC disk [> filename]
```

where **filename** is the name of file in which to store this information.

A listing of all devices found will be stored in the given filename.

For example, in **HP-UX 10.x**, your screen display will appear similar to the following:

```
/etc/ioscan -fnC > filename
more filename
.
.
.
disk      23.0/20.0.0   disc3      CLAIMED DEVICE EMC SYMMETRIX
           /dev/dsk/c3t0d0      /dev/rdisk/c3t0d0
           /dev/floppy/c3t0d0   /dev/rdisk/c3t0d0

disk      24.0/20.0.1   disk3      CLAIMED DEVICE EMC SYMMETRIX
           /dev/dsk/c3t0d1      /dev/rdisk/c3t0d1
           /dev/floppy/c3t0d1   /dev/rdisk/c3t0d1
```

Note the card instance number next to the H/W Path for the disk devices.

Also note that the H/W Path will differ from one machine model to another.

Slice (partition) is not supported by HP 10.

Configuring Logical Volumes

This section applies to HP-UX 10.x and 11.x

A logical volume is a flexible model for storage management. One or more physical devices are introduced into a volume group. Logical volumes are created from storage available for the volume group. Additional functionality such as striping, volumes that span multiple physical devices, etc. are achieved by volume management. You need to understand LVMs before implementing logical volumes on your host. For more information on logical volumes, refer to HP documentation for system administrators.

You must configure logical volumes to store data when using HP-UX 10.x and 11.x.

Creating Physical Volumes

After you have verified that the host has recognized each Symmetrix device, you need to create a physical volume for each device.

To create a physical volume for each Symmetrix disk device (logical volume), enter a statement similar to the following for each physical volume to create:

```
pvccreate -f /dev/rdisk/c3t0d0
```

This statement creates a physical volume, `"/dev/rdisk/c3t0d0"`.

Physical volumes can also be created using the SAM utility.

Volume Groups & Logical Volumes

At this time you can use physical volumes or create volume groups and logical volumes. Volume groups and logical volumes can be created using HP-UX stand-alone commands or SAM. To create volume groups and logical volumes, refer to the HP documentation set.

Creating the Mount Directory

If you want to use filesystems on the volumes, you must create a mount directory for each of these volumes. To do this, type a statement similar to the following for each physical volume:

```
mkdir /symm
```

This example creates a mount directory called “symm”.

Creating a New Filesystem

Now that the physical volumes and their respective mount directories are available to the operating system, a new filesystem must be created for each Symmetrix disk. Specify the disk name matching the volume size defined in your Symmetrix unit for this statement for HP-UX 10.x when volume groups are used in a statement similar to the following:

```
newfs /dev/vg-group/volx
```

Mount the Directory

For HP-UX 10.x and volume groups, before you can mount the directory you need to add the path to **/etc/fstab**. To do this, type the following statement:

```
vi /etc/fstab
```

Specify each device in a statement similar to the following:

```
/dev/vg-group/volx
```

Mount the filesystem by typing a statement similar to the following:

```
/etc/mount /symm
```

Use the **bdf** command to show all mounted filesystems and the available free space. For example,

```
bdf
```

Filesystem	kbytes	used	available	%used	mounted on
/dev/vg00/lvol1104299	529364093356%				/
/dev/vg00/lvol15349530	125692188884	40%			/users
/dev/vg00/lvol14124507	836	111220	1%		/tmp
/dev/vg00/lvol13397867	13749422058638%				/usr
/dev/dsk/c1t0d02741142	10404123629864%				/symm5
/dev/dsk/c2t0d02741142	0	27411420%			/symm
/dev/dsk/c3t0d02741142	10404123629864%				/symm3
/dev/dsk/c4t0d02741142	10404123629864%				/symm1

shows **/dev/dsk/c2t0d0** mounted.

Using SAM

HP-UX has a menu-driven utility for managing system resources, specifically disk drives. This utility is called SAM and users familiar with it might find it more useful to use SAM for configuring new Symmetrix disks.

Adding Devices Online

Whenever devices are added online to the Symmetrix unit or device channel addresses are changed, you will need to perform the actions described below in order to introduce the new devices to the system.

You can define newly connected physical volumes to the system without rebooting the host. This requires use of the **insf** command in a statement similar to the following:

```
insf -H 2/4.0.0
```

where **2/4.0.0** is a hardware path to which the new device has been connected.

Performance Issues with HP-UX

This section covers some performance issues when operating in the HP-UX environment.

In HP-UX 10, the default value of **rotdelay** is 0. Therefore, it is not necessary to change the value of this parameter.

System and Error Messages

HP-UX 10.x and HP-UX log system and error messages including SCSI error messages to a file called `/var/adm/syslog/syslog.log`.

Diagnostic Patches

Table 7-5 lists HP's Diagnostics Cumulative Patches. Installing the appropriate patches for your operating system level causes Symmetrix drives to be recognized properly by HP hosts.

Table 7-5 Diagnostic Cumulative Patches

Operating System	Patch Number
10.01	PHSS_7583
10.10	PHSS_7585
10.20 and 11.0	Symmetrix definition should already exist for STM and Sysdiag

These patches have an affect on the output of the STM and sysdiag utilities when run on Symmetrix devices. The affect of these patches are as follows:

- ◆ SCSIDSK2 - HP's SCSI disk diagnostic utility (which includes commands like INQUIRY, CAPACITY, etc.) now works with Symmetrix devices, but reports the following message:

`SCSIDSK2 is not designed to test the selected device (which is "EMCSYMMETRIX")`

- ◆ The IOMAP utility recognizes Symmetrix devices as EMCSYMMETRIX DISK. (If the patch is not installed, Symmetrix devices are reported as "Device Unknown".)
- ◆ In the System Error Log Report, any entries in the log for Symmetrix devices now report the product name as "EMCSYMMETRIX" and the device class as "Hard Disk".

This chapter lists NEC UNIX system models and maps them to corresponding HP models.

- ◆ NEC Servers and Workstations..... 8-2

NEC Servers and Workstations

Each of the NEC systems in the table below corresponds to the HP model on the same line. For more information about HP see the HP chapters and other references to HP in the Index of this manual.

Table 8-1 NEC UNIX Systems Matched to HP

NEC NX Model with HP-UX 11.0	HP Model
PowerPath License Level: Workgroup	
NX7000/22	HP9000/D220
NX7000/23	HP9000/D230
NX7000/27	HP9000/D270
NX7000/D280	HP9000/D280
TX7/D280	HP Technical D280
NX7000/32	HP9000/D320
NX7000/33	HP9000/D330
NX7000/37	HP9000/D370
NX7000/D380	HP9000/D380
NX7000/D390	HP9000/D390
TX7/D390	HP Technical D390
NX7000/220	HP9000/K220
NX7000/260	HP9000/K260
NX7000/K370	HP9000/K370
TX7/K370	HP Technical K370
NX7000/K380	HP9000/K380
TX7/K380	HP Technical K380
PowerPath License Level: Department	
NX7000/420	HP9000/K420
NX7000/460	HP9000/K460

Table 8-1 NEC UNIX Systems Matched to HP (*continued*)

NEC NX Model with HP-UX 11.0	HP Model
NX7000/K570	HP9000/K570
NX7000/K580	HP9000/K580
	HP N-Series
PowerPath License Level: Enterprise	
NX7000/P590	None: Special NEC model
TX7/P590	Special NEC model
NX7000/T600	HP9000/T600
NX7000/V2250	HP9000/V2250
TX7/V2250	HP Technical V2250
PowerPath License Level: Global	
NX7000/V2500	HP9000/V2500
TX7/V2500	HP Technical V2500

This chapter describes how to incorporate Symmetrix in the IBM AIX environment.

◆ Overview	9-2
◆ Useful AIX Utilities and Functions	9-5
◆ Useful EMC Commands and Functions	9-10
◆ Symmetrix Configuration	9-11
◆ Incorporating Symmetrix	9-12
◆ Configuring Disks Using Commands	9-15
◆ Configuring Disks Using SMIT	9-20
◆ Adding Devices Online	9-23
◆ System and Error Messages	9-24

Overview

When using Symmetrix in an AIX environment, please read the bulleted items below before beginning for the terminology used in this section:

- ◆ Symmetrix requires the RISC System/6000™ to be running AIX version 4.2.1 or higher.
- ◆ Symmetrix attaches to IBM 2416 (MCA), 2412 (MCA), 6209 (PCI) Fast-Wide Differential SCSI Controller cards, 6207 (PCI) Ultra-wide SCSI Differential Controller card, 6204 (PCI) Universal Differential Ultra SCSI Controller card, or to the 6205 (PCI) Dual Channel Ultra 2 SCSI Controller card.
- ◆ The Fast-Wide SCSI Controller card has a default target ID of 7. Unless you change the default target ID of this controller card, choose target IDs 0 through 6 and 8 through 15 with AIX 4.2.1 or greater.
- ◆ **When initially configuring the Symmetrix devices, you must use ODM definitions for Symmetrix supplied by EMC.** These are available on the EMC ftp server, `ftp.emc.com` in `pub/symm3000/aix/ODM_DEFINITIONS`.
- ◆ In the Bull® Escala® environment, Symmetrix Fibre Channel can be connected with the LP7000E-N1 Fibre Channel host bus adapters.
- ◆ In the IBM RS/6000 and RS/6000 SP environments, Symmetrix Fibre Channel can be connected with the EMC Fibre Channel Interface for AIX platforms or with the IBM (6227) Gigabit Fibre Channel Adapter for PCI Bus.

Terminology

The following describes AIX terminology used in this section:

- ◆ **Logical Volume Manager (LVM)** is the part of the RS/6000 AIX system responsible for the disk subsystems. Interaction with the LVM occurs via SMIT menus or via stand-alone commands.
- ◆ In LVM terminology, each physical disk device connected to the host system is a *physical volume* (PV). Thus, each Symmetrix logical volume is considered to be a PV by AIX. By default, AIX refers to the physical volumes as `hdisk0`, `hdisk1`, `hdisk2`, etc.

- ◆ Physical volumes can be grouped into **Volume Groups**. The name of the volume group is chosen by the person who defines the group.

When a volume group is varied online, the physical devices in it are reserved.

- ◆ The control information for a volume group resides in the **volume group descriptor area (VGDA)**. At least one VGDA resides on each PV in the volume group.
- ◆ **Physical Partitions (PPs)** are the smallest unit of disk space that can be allocated in a volume group. Any disk space allocated is an integral number of PPs. By default a PP is 4 MB. When a volume group is created, you can specify a PP size of 1 MB to 1024 MB.
- ◆ **Logical Volumes (LVs)** are groups of PPs. A typical LV is simply a collection of disk space. The PPs that constitute an LV can reside on separate physical volumes and need not be contiguous when they are on the same physical volume. All the PPs in an LV must be on physical volumes in the same volume group.
- ◆ To change the configuration, you can either use SMIT or the equivalent AIX commands. These commands usually require *root* privileges. When working with SMIT, it is possible to see the AIX command that was generated by a specific SMIT menu selection by pressing <F6> (ASCII interface) or look in the command window (Motif interface).
- ◆ The RISC System/6000 has two ways of making devices available. The first way uses SMIT (System Management Interface Tool). Refer to the *RISC System/6000* documentation for more information on this tool. The second way uses individual commands and is described below.
- ◆ AIX will not vary off volume groups on a system shutdown or power off. Therefore, device reservation will remain in place. These reservations are only cleared on system reboot.

Table 9-1 AIX Operating System Features Summary

O/S	AIX version 4.2.1 and higher
Largest recognizable drive size	N/A
Supports multiple target IDs	Yes
Supports multiple LUNs per target ID	Yes
Partition representing the entire disk	N/A
To see the drive configuration, use this utility:	<code>lsdev -Cc disk</code>
To configure EMC devices use this utility:	<code>/usr/lpp/Symmetrix/bin/emc_cfgmgr</code>
Special drive definition files needed	yes - ODM definitions

Useful AIX Utilities and Functions

This section lists AIX functions and utilities you may use to define and manage Symmetrix devices. The use of these functions and utilities is optional. They are listed for reference only.

◆ `SMIT`

SMIT (System Management Interface Tool) is a menu-driven utility for managing the device configuration of the RS/6000. SMIT is available in a windows version (`xinit` command) or an ASCII interface.

SMIT can be used for a number of standard AIX functions such as:

- Installing and maintaining software
- Installing and managing devices
- Managing disk devices using the LVM
- Problem determination

◆ `cfgmgr -v`

Configuration manager with verbose parameter. This is an inline hardware resource utility. `cfgmgr`, when run, walks the bus(es) and examines all adapters and devices connected to the RS/6000. It is frequently used in conjunction with the `lsdev` and `lscfg` commands.

◆ `lsdev -Cs scsi`

Command that lists all the available SCSI device(s).

◆ `lsdev -Cs fcscsi`

Command for use with IBM driver that lists all the available Fibre Channel device(s).

◆ `lsdev -Cs fcs`

Command for use with the IBM driver that lists all the Fibre Channel device(s) in the system.

- ◆ `lsdev -Cs fcp`
Command for use with the Bull driver that lists all the host controller interfaces (`fcp`) in the system.
- ◆ `lsdev -Cc disk`
Command that lists all disk devices.
- ◆ `lsdev -C | more`
Command that lists all device(s). It directs the output to the screen one "page" at a time.
- ◆ `lscfg | grep scsi`
Command that lists the configuration (in PCI-based machines). It shows all SCSI channel host controller cards available or defined in the system.
- ◆ `lsdev -Ct efscsi`
Command for use with IBM driver to list all Fibre Channel host controller protocol devices.
- ◆ `lsdev -Ct lp7000`
Command for use with Bull driver to list all Fibre Channel host controller adapter cards available or defined in the system.
- ◆ `lsdev -Ct df1000f7`
Command for use with IBM driver to list all Fibre Channel host controller cards available or defined in the system.
- ◆ `lsattr -E -l scsi#`
Command that lists the attributes (addresses) of the `scsi#` requested on the Fast-Wide Differential host controller card.

- ◆ `lsattr -El fchan#`
Command for use with Bull driver that lists the attributes (addresses) of the `fchan#` requested on the Fibre Channel host controller card.
- ◆ `lsattr -El fcs#`
Command for use with IBM driver that lists the attributes (addresses) of the `fcs#` requested on the Fibre Channel host controller card.
- ◆ `lsattr -El fcp#`
Command for use with Bull driver that lists the attributes (addresses) of the `fcs#` requested on the Fibre Channel host controller card.
- ◆ `lsattr -El fscsi#`
Command for use with IBM driver that lists the attributes (addresses) of the `fscsi#` requested on the Fibre Channel host controller card.
- ◆ `lsvg -o`
Command that displays all volume groups that are varied on (activated).
- ◆ `lsvg -l vname`
Command that displays logical volumes, physical partitions, and the state of the logical volume within the specified volume group.
- ◆ `mkdev`
Configures a disk device making sure it is available as a physical volume.
- ◆ `mkvg`
Groups one or more physical disks into a volume group.

- ◆ `varyonvg`
Makes the specified volume group available to the host.
- ◆ `varyoffvg`
Makes the specified volumes group unavailable to the host.
- ◆ `crfs`
Command that creates a new filesystem.
- ◆ `mount`
Command that mounts a filesystem.



CAUTION

Inappropriate use of the following `rmdev` command can result in a system problem and/or loss of data. Take extra caution when using this command to ensure that the devices and/or controllers you are working with are not in use and are the correct addresses.

- ◆ `rmdev -l scsix -d`
Command that removes a SCSI controller (as specified by the "x" in scsix) from the ODM.
- ◆ `rmdev -l hdiskx`
Command that removes the specified disk device.
- ◆ `rmdev -l hdiskx -d`
Command that removes the specified disk device (`hdiskx`) and deletes it from the ODM.
- ◆ `rmdev -l fcpx -d`
Command that removes the Bull Fibre Channel Protocol driver instance from the ODM.

- ◆ `rmdev -l fscsix -d`

Command that removes the IBM Fibre Channel Protocol driver instance from the ODM.

- ◆ `rmdev -l fchanx -d`

Command for use with Bull driver that removes the Fibre Channel controller `x` from the ODM.

- ◆ `rmdev -Rdl fchan#`

Command for use with Bull driver that removes the Fibre Channel controller `x` and its associated pseudo adapter (`fcp`).

- ◆ `rmdev -l fcsx -d`

Command for use with IBM driver that removes the Fibre Channel controller `x` from the ODM.

- ◆ `errpt -aN hdiskx*`

Error report. This command generates an error report that lists all errors associated with the `hdiskx`.

- ◆ `odmget XxxXx`

where `XxxXx` is one of the following ODM (object database manager) attributes:

<code>CuAt</code>	ODM Customize Attributes
<code>PdAt</code>	Predefined Attributes
<code>CuDv</code>	Customize Devices
<code>PdDV</code>	Predefined Devices

This command allows you to see all hardware and devices configured, defined, or available in the system. Predefined attributes are supported devices in the `/dev` library. Customized attributes are all configured, defined, or available devices. All customized attributes are located in the `/etc/objrepos` library.

Useful EMC Commands and Functions

This section lists EMC commands and functions you may use to define and manage Symmetrix devices. The use of these functions and utilities is optional. They are listed for reference only.

- ◆ `emc_cfgmgr`
Configures Symmetrix devices.
- ◆ `inq`
Performs an inquiry to all available Symmetrix devices.
- ◆ `mkbcv -a`
Configures all Defined Symmetrix TimeFinder devices and makes them in an Available state.
- ◆ `mkbcv <hdisk>`
Configures a specific Defined Symmetrix TimeFinder device and makes it in an Available state.
- ◆ `rmbcv -a`
Configures all Available Symmetrix TimeFinder devices and makes them in a Defined state.
- ◆ `rmbcv <hdisk>`
Configures a specific Available Symmetrix TimeFinder device and makes it in a Defined state.

Symmetrix Configuration

In the AIX host environment, you can configure the Symmetrix disk devices into logical volumes. EMC recommends the following logical volume ratios as indicated in the table below:

Table 9-2 Disk Device Configuration

		Symmetrix Logical Volume to Physical Splits ^a							
		1:1	2:1	3:1	4:1	5:1	6:1	7:1	8:1 ^b
Physical Devices	4 GB	x	x						
	9 GB	x	x	x	x				
	18 GB		x	x	x	x	x	x	x
	23 GB		*	x	x	x	x	x	x
	36 GB			*	x	x	x	x	x
	47 GB			*	*	x	x	x	x
	50 GB ^c				x	x	x	x	x
	73 GB ^{c,d}					x	x	x	x

- a. * indicates ratios available only in microcode level 5265 or higher.
- b. With microcode revision level 5263, maximum is 8:1. With 5264, maximum is 16:1. With 5265, maximum is 32:1.
- c. The 50 GB and 73 GB drives are supported only by microcode revision level 5x66 and higher.
- d. The 73 GB drive does not support the RAID-S configuration on Symmetrix.

For specific cylinder counts and logical volume sizes, refer to Table 2-1 on page 2-5 and the following tables in that chapter.

The EMC Customer Engineer (CE) should contact the EMC Configuration Specialist for updated online information. This information is necessary to configure Symmetrix to support the customer's host environment.

Incorporating Symmetrix

With AIX 4.2 or Higher

NOTE: The procedures here are for a new installation.

Installation requires downloading the **EMC_AIX_ODM.4.3.tar.Z** fileset from **ftp.emc.com** and should be run prior to attaching SYMMETRIX to the AIX Server.

The EMC Symmetrix AIX support software fileset is installed on AIX based machines using the `installp` command.

1. Log onto the RS/6000 as **root**.
2. From the FTP server, `ftp.emc.com`, in
`pub/symm3000/aix/ODM_DEFINITIONS`
download the fileset `EMC_AIX_ODM.4.3.x.x.tar.Z`.
3. In the **/tmp** directory, uncompress the fileset. Enter
`uncompress EMC_AIX_ODM.4.3.x.x tar.Z`
4. Untar the fileset. Enter
`tar -xvf EMC_AIX_ODM.4.3.x.x tar`

This step creates the following three files:

```
EMC.AIX.4.3.x.x
README
.toc
```

5. For installing with SMIT, copy the **EMC.AIX.4.3.x.x** file into the **/usr/sys/inst.images** directory.

Using the Command Line to Install

The command line interface can be used to complete the installation.

1. From the **/tmp/EMC** directory in Step 4 above, enter the following command:

```
installp -ac -gX -d . Symmetrix.aix.rte
```

The `.` between `-d` and `symmetrix` is the current directory where the installation file resides and **symmetrix.aix.rte** is the name of the software to install.

2. Verify that the Installation Summary reports "SUCCESS."
3. Attach Symmetrix to the server and run `emc_cfgmgr` to configure all the Symmetrix devices.

Using the SMIT Utility to Install

1. If the SMIT menu interface is preferred, invoke **smit installp** following Step 5 above.
2. Select the option:
`Install and Update from LATEST Available Software`
3. Use the `List` function to select `/usr/sys/inst.images` as the Installation directory.
4. Use the `List` function to select
EMC Symmetrix AIX Support Software
5. Press Enter AFTER making all desired changes. You can use the default options.
6. Be sure to scroll all the way down to the bottom of the window to see the Installation Summary and verify that the "SUCCESS" message is displayed.
7. Attach Symmetrix to the server and run **cfgmgr** to configure all the Symmetrix devices.

High Availability Support for Symmetrix

This section applies only if you have HACMP 4.2 or above and your installation environment has concurrent resource groups configured. For more detailed updated information, see the README file.

The EMC Symmetrix HA support software files set is installed on AIX based machines using the `installp` command, as in the following example. (`installp` is described in *Using the Command Line to Install* on page 9-12.)

```
installp -ac -g -d . Symmetrix.ha.rte
```

This files set provides a utility called **symcurrent** to add and delete Symmetrix device types to and from the **clconraid.dat** configuration file. This file is a member of the HACMP concurrent files sets **cluster.base.server.diag** and **cluster.es.server.diag**.

Fibre Channel for AIX

For installation and other information, see the documentation provided with the Fibre Channel Host Bus Adapter. This manual is the *EMC Fibre Channel Interface for AIX*, EMC P/N 300-999-038.

Fibre Channel Interface support for Symmetrix is available in the following filesets:

`Symmetrix.fcscsi.rte` (for EMC)

`Symmetrix.fcp.rte` (for IBM)

Device Characteristics for Symmetrix

EMC ODM definitions set the Symmetrix devices to the default characteristics listed in Table 9-3.

Table 9-3 Device Characteristics

Characteristic	Set to
Queue depth	8
Queuing type	simple
Use QERR bit	no
Device CLEARS its queue on error	no
READ/WRITE timeout	40 sec

If you are not able to install ODM support files successfully, please contact the EMC Technical Support staff.

Configuring Disks Using Commands

This section describes how to use the AIX stand-alone commands to create volume groups, define logical volumes and filesystems, and mount the filesystems. See *Configuring Disks Using SMIT* on page 9-20 to perform these same actions using SMIT.

Make a Volume Group

In AIX you can include several physical disk devices in a “volume group”. Each AIX system can have 1 to 255 volume groups. A volume group is a collection of 1 to 128 physical volumes (depending on versions of AIX) of varying size and type. See Table 9-5, *Limitations of LVM in AIX Versions*.

More than 1016 Physical Partitions per Physical Volume

AIX Version 4.3.1 and higher supports more than the former limit of 1016 physical partitions per physical volume. To support a volume group greater than 1016 physical partitions using the same VGDA and VGSA areas, the number of disks supported in the volume group is reduced.

Add the `-t` flag to the `chvg` and `mkvg` commands to convert and create a volume group with multiples of 1016 partitions per disk. This reduces the total number of disks that can be added to the volume group by same fraction. After you change or create a volume group with more than 1016 physical partitions per disk, you cannot import it into AIX versions below 4.3.1.

The `-t` factor allows 1016 multiplied by *factor* physical partitions per physical volume. For example, a partition size of at least 16 MB is necessary to create a volume group with a 10 GB disk. With a factor of 2, a smaller partition size of 8 MB can be used. This limits the total number of disks that can be added to the volume group. If a factor value is used, a maximum of $\text{MAXPVS}/\text{factor}$ disks can be included in the volume group. See Table 9-5, *Limitations of LVM in AIX Versions*.

Relationships between factor -t, physical partition numbers per physical disk, and the number of disks allowed in one volume group are presented in the following table.

Table 9-4 Factor -t, Physical Partition Numbers, Disks per VG

Factor -t	PP numbers	Number of disks per VG
1	1016	32
2	2032	16
3	3048	10
4	4064	8
5	5080	6
6	6096	5
7	7112	4
8	8128	4
16	16256	2

When adding a new volume group in AIX 4.3.1 and later, the physical partition sizes can be 1, 2, 4, 8, 16, 32, 64, 128, 256, 512, or 1024 MB. Volume groups created with 512 MB and 1024 MB cannot be imported into previous versions of AIX.

When adding a new volume group in versions of AIX before 4.3.1, the physical partition sizes can be 1, 2, 4, 8, 16, 32, 64, 128, or 256 MB.

The volume group format in AIX 4.3.2 allows a maximum number of 128 disks (physical volumes). The maximum number of logical volumes in this format is 512.

The following table lists LVM limits for different versions of AIX.

Table 9-5 Limitations of LVM in AIX Versions

Version	Volume group (per system)	Physical volume (per volume group)	Physical partition (per physical volume)	Logical volume (per volume group)	Logical partition (per logical partition)
AIX 4.3.3	255	128	1016 multiplied by -t factor	512	32,512
AIX 4.3.2	255	128	1016 multiplied by -t factor	512	32,512
AIX 4.3.1	255	32	1016 multiplied by -t factor	256	32,512
AIX 4.3.0	255	32	1016 multiplied by -t factor	256	32,512
AIX 4.2	255	32	1016	256	32,512

The LVM format in AIX 4.3.2 and higher supports the smaller volume group configurations of the previous versions of AIX. A migration path is provided to convert old volume groups to the new volume group format. You can make this conversion only if you have enough free partitions on each of the physical volumes in the volume group to be allocated.

In the following examples, each Symmetrix disk device will be defined as a 1 volume group for simplicity. You can choose to treat each Symmetrix device as one volume group or join multiple Symmetrix devices into one volume group.

To make a volume group for a Symmetrix device, enter a statement similar to the following for each volume group to create:

```
mkvg -y testvg hdisk2
```

For example, this statement creates a volume group from the physical disk device, "hdisk2" named **testvg**.

The default physical partition size is 4 MB.

Make the Volume Group Available

Once a volume group is created it must be made available to the operating system.

To make a volume group available to the host (for example, volume group **testvg**), use SMIT or type the following statement:

```
varyonvg testvg
```

Repeat this action for each volume group. The `varyonvg` command activates a volume group making it known to the host system.

Conversely, the `varyoffvg` command deactivates a volume group, thus making it unavailable to the host.

At this point we have defined “raw” devices. If you wish to create filesystems on the raw devices follow the instructions in the next paragraphs.

Create Logical Volumes

Now that the volume groups are defined, you can either use the volume groups as is, or define logical volumes.

The following example illustrates how to make the logical volume, `testlv`, in volume group `testvg`. `testlv` is a journaled file system with 100 physical partitions, 400 MB in size.

```
mk1lv -y testlv -t jfs testvg 100
```

Creating a Mount Point

Create a mountpoint that points to a logical volume. The following command creates a mountpoint named “`emc`”.

```
mkdir /emc
```

Create and Mount Filesystems

To create and mount filesystems use SMIT or the `crfs` and `mount` commands.

For example, use the `crfs` command in a statement similar to the following to create a filesystem:

```
crfs -v jfs -d /dev/testlv -m /emc -A yes -a size=41943094
```

where:

<code>-v jfs</code>	Makes a 'journaled file system'
<code>-d</code>	Specifies the logical volume
<code>-m path></code>	Specifies the mount point for the file system
<code>-A yes</code>	Mounts filesystem automatically at system restart
<code>-a</code>	Specifies the various attributes for this file system. 'size=x' specifies the filesystem size in blocks (512 bytes per block). 41943094 is the size of a 2GB filesystem.

Check the filesystem as follows:

```
fsck /emc
```

The `fsck` command checks the integrity of the filesystem mounted at the `emc` directory.

Mount each filesystem at the desired mount point as follows:

```
mount /emc
```

Make a "lost and found" directory as follows:

```
cd /emc  
mk lost+found
```

This changes to the `emc` directory and creates a directory called `lost` and `found`.

Configuring Disks Using SMIT

The above sections added devices to the RS/6000 using stand-alone commands. These devices can also be added to the RS/6000 using the SMIT menu. To add devices using SMIT, follow these steps:

1. Log in as **root**.
2. Launch SMIT by typing `smit`. The SMIT menu appears.

Add a Volume Group

Perform the following steps to create and name a volume group and select the physical volume(s) for that group:

1. Choose **System Storage Management** from the SMIT Main Menu.
2. Choose **Logical Volume Manager** from the menu that appears.
3. Choose **Volume Groups** from the menu.
4. Choose **Add a Volume Group** from the menu.
5. Enter a volume group name (e.g., `vg01`) when prompted.
6. TAB down to the next field and enter the **Physical Partition Size** for the disk. (4 MB is the default size.)
7. TAB down, then press the F4 key. Choose the **Physical Volume Name** by pressing the F7 key when you are next to the disk(s) you want to select, then press ENTER.
8. Click on the **ENTER** button. If this action completes successfully, press the F3 key three times to return to the Logical Volume Manager menu.

Add Logical Volumes to the Volume Group

Perform the following steps to add logical volumes to the volume group:

1. Choose **Logical Volumes** from the Logical Volume Manager menu and then select **Add a Logical Volume**.
2. Press the F4 key and choose the volume group to which you are adding the logical volume(s) by pressing the F7 key when beside the volume group you wish to select. Press ENTER

If you are adding the logical volume to a new volume group, just enter a name for the volume group.

3. Enter a name for the logical volume when prompted.
4. Enter the number of **Physical Partitions**.
5. Enter the **Physical Volume name**. Pressing F4 will provide a popup list of available volume groups.
6. Change any other parameters as desired such as striping options, etc.
7. Click on the **ENTER** button. If this action completes successfully, press the F3 key four times to return to the Physical and Logical Storage menu.

Creating Filesystems

Perform the following steps to create and mount filesystems:

1. Choose **File Systems** from the Physical and Logical Storage menu.
2. Choose **Add/Change/Show/Delete File Systems** from the menu that appears.
3. Choose **Journaled File System (jfs)** to create a journaled filesystem.
4. Choose **Add a Journaled File System on a Previously Defined Logical Volume** from the menu that appears.
5. Choose **Add a Standard Journaled File System** from the menu that appears.
6. Enter the name of logical volume on which to create a filesystem. Pressing F4 will provide a popup list of available logical volumes.
7. Enter a **mount point** for that filesystem.
8. Click on the **ENTER** button. If this action completes successfully, choose **CANCEL** if you are using an X terminal (or F3 several times if using an ASCII terminal) to return to the Physical and Logical Storage menu.
9. Choose **File Systems** from the Physical and Logical Storage menu.
10. Choose **Mount a File System** from the menu that appears.
11. Select the filesystem you wish to mount. Pressing F4 will provide a popup list of available filesystems.

12. Choose **Mount Point** from the menu that appears. Pressing F4 will provide a popup list of available mount points.
13. Click on the **DO** button. If this action completes successfully, choose **CANCEL** if you are using an X terminal (or F3 if using an ASCII terminal) to exit SMIT.
14. Verify the filesystem is mounted by issuing the `mount` or `df` command.

Adding Devices Online

Whenever devices are added online to the Symmetrix unit or device channel addresses are changed, you will need to run the script described below in order to introduce the new devices to the system.

1. Change directory to `/usr/lpp/symmetrix/bin`
2. Invoke the shell script **emc_cfgmgr**

Whenever TimeFinder devices are configured by the host, they are forced to a Defined state. To make these Defined devices Available, follow one of the following steps:

1. `mkbcv -a`
Configures all Defined Symmetrix TimeFinder devices and makes them in an Available state.
2. `mkbcv <hkdisk>`
Configures a specific Defined Symmetrix TimeFinder device and makes it in an Available state.

System and Error Messages

In an AIX environment you can view system status and error messages in two ways:

- ◆ via SMIT
- ◆ via the stand-alone command `errpt`

AIX logs system status and error messages in `/var/adm/ras/errlog`. This error log is not viewable directly. Follow the steps below to view the error log according to the way you prefer:

Viewing using SMIT

Follow these steps to view the error log using SMIT:

1. Choose **Problem Determination** from the SMIT main menu.
2. Choose the **Error Log** option.
3. Choose the **Generate Error Report** option.
4. Decide where the generated report will reside.
5. Respond **Yes** or **No** at the "Concurrent Error Reporting" prompt.
6. Modify the reporting options as appropriate for your environment and press **ENTER** to generate the report.

Viewing using `errpt`

Follow these steps to view the error log using the `errpt` command. The `errpt` command has the following format:

```
errpt -options parameters
```

If you specify the `errpt` command with no options or parameters, a complete summary report displays.

The following lists the `errpt` options and parameters:

<code>-a</code>	Displays information about errors in the error log file in a detailed format.
<code>-c</code>	Formats and displays each of the error entries concurrently.
<code>-e Enddate</code>	Specifies all records posted before the Enddate variable, where the Enddate variable has the form <code>mmddhhmmyy</code> (month, day, hour, minute, and year)
<code>-s Startdate</code>	Specifies all records posted after the Startdate variable, where the Startdate variable has the form <code>mmddhhmmyy</code> (month, day, hour, minute, and year)
<code>-S ResourceClassList</code>	Generates a report of resource classes specified by the ResourceClassList variable. Some of the ResourceClassList members are disk, adapter, and tape.

The following is an example of an `errpt` output. Please note that error report entries will vary considerably and yours may not resemble this example.

```
errpt -as 1118150096
```

```
-----
LABEL:DISK_ERR4
IDENTIFIER:1581762B
```

```
Date/Time:      Mon Nov 18 15:58:34
Sequence Number: 459
Machine Id:     000009839100
Node Id:        lss5030
Class:          H
Type:           TEMP
Resource Name:   hdisk12
Resource Class:  disk
Resource Type:   osdisk
Location:        00-02-01-5,0
VPD:
```

```
Manufacturer.....EMC
Machine Type and Model.....SYMMETRIX
Part Number.....
ROS Level and ID.....35303632
Serial Number.....6800B161
EC Level.....
FRU Number.....
Device Specific.(Z0).....000002029000003A
Device Specific.(Z1).....Q
Device Specific.(Z2).....L r
Device Specific.(Z3).....
Device Specific.(Z4).....
Device Specific.(Z5).....
```

```
Description
DISK OPERATION ERROR
```

```
Probable Causes
MEDIA
DASD DEVICE
```

```
User Causes
MEDIA DEFECTIVE
```

```
Recommended Actions
FOR REMOVABLE MEDIA, CHANGE MEDIA AND RETRY
PERFORM PROBLEM DETERMINATION PROCEDURES
```

```
Failure Causes
MEDIA
DISK DRIVE
```

```
Recommended Actions
FOR REMOVABLE MEDIA, CHANGE MEDIA AND RETRY
PERFORM PROBLEM DETERMINATION PROCEDURES
```

Detail Data

SENSE DATA

```

0605 0000 0809 BDA8 0800 0000 0000 0000 0102 0000 F000 0B00 09BD A80A 0000 0000
4400 0000 0000 0000 0100 FFFF 020E 4040 000A 0000 0000 0000 0000 0000 0316 000F
0000 0000 0000 0040 0200 0000 0000 0000 4000 0000 0416 0008 750F 0000 0000 0000
0000 0000 0000 0000 1518 0000 080A 0000 0000 0000 FFFF 0000 0A06 0011 0000 0000
3206 2000 0000 0000 0002 0000 0000 0000 0000 0000 0000 303A 0004 0180

```

The preceding error appeared after Symmetrix aborted a Read command. Although the command was aborted, it was later successfully retried by the application. The following is an explanation of the important information for analyzing the error report.

- ◆ LABEL:DISK_ERR4
The error was a physical volume recovered error. (Other errors are explained in `/usr/include/sys/errids.h`)
- ◆ Type: TEMP
Temp denotes a recoverable error.
- ◆ Resource Name: hdisk12 Resource Class: disk
These two fields identify the device (and name) where the problem manifested itself. In this case hdisk12 was a virtual disk in a Symmetrix unit.
- ◆ Location: 00-02-01-5,0
This field provides the location of the failing device. In this instance the device was on controller number 2, with SCSI target 5 and LUN 0.

The next few fields called the VPD contain certain characteristics about the device in question. For example, in this case

Serial Number.....6800B161

indicates that this device is in a Symmetrix unit with a serial number that ends with 68, its device number is 00B, and it was accessed through SA16 port B.

Immediately following the VPD are several fields: Description, Probable Causes, User Causes, and Failure Causes. These fields contain preprogrammed responses to groups of errors. These fields only contain suggestions and recommendations.

In order to fully grasp the problem, analyze the Detail Data field as shown in the following example:

Detail Data

SENSE DATA

```

0605 0000 0809 BDA8 0800 0000 0000 0000 0102 0000 F000 0B00 09BD A80A 0000 0000
4400 0000 0000 0000 0100 FFFF 020E 4040 000A 0000 0000 0000 0000 0000 0316 000F
0000 0000 0000 0040 0200 0000 0000 0000 4000 0000 0416 0008 750F 0000 0000 0000
0000 0000 0000 0000 1518 0000 080A 0000 0000 0000 FFFF 0000 0A06 0011 0000 0000
3206 2000 0000 0000 0002 0000 0000 0000 0000 0000 0000 303A 0004 0180

```

Table 9-6 Detail Data

Byte ^a	Value ^b	Meaning	Explanation
0	06	Size of the cdb that resulted in the error	A 6 byte cbd caused this error.
1	05	SCSI ID where cdb was sent	SCSI ID 5
2	00	Reserved	none
3	00	Any flags associated with the cdb would be located here.	none
4	08	Bytes 4-15 contain the actual cdb that resulted in the error.	The cdb was a 6 byte read command.
5	09	Bytes 5-7 contain the starting address for the read.	The read shall start at logical block 09BDA8.
6	BD		
7	A8		
8	08	Number of blocks requested.	8 blocks were requested.
9	00	Control byte.	none
10-15	00		Reserved for longer cdb.
16	01	Byte which decies if byte 17 will contain SCSI status or if byte 18 will contain adapter status.	Look to byte 17 for SCSI status.
17	02	SCSI status.	Check condition was reported.
18	00	Adapter status.	Only valid if byte 16 had been 02.
19	00	Reserved	none
20-51		Sense Data	see
52-155		Data used by AIX to provide the user with the information in many of the above fields.	

a. The byte field is a decimal number.

b. The value field is a hexadecimal value.

In the figure below, the sense data is underlined. Table 9-7 on page 9-29 explains the meanings of the sense data.




Figure 9-1 Detail Data Format Diagrammed

Table 9-7 Sense Data Format Explained

Byte ^a	Value ^b	Meaning	Explanation
0	F0	Error Code	Bytes 3-6 of sense data are valid and this is a current error (not deferred)
1	00	Segment number	Not valid for the read cdb.
2	0B	Sense key.	The command was aborted.
3	00	Bytes 3-6 contain information bytes that are particular to each cdb.	The entire read was aborted. (9BDA8 is the same address specifies by bytes 5-7 from above.
4	09		
5	BD		
6	A8		
7	0A	Additional valid sense data bytes.	The next 10 bytes contain valid data.
8-11	00	Command specific data.	none

Table 9-7 Sense Data Format Explained (*continued*)

Byte ^a	Value ^b	Meaning	Explanation
12	44	Additional sense code (ASC)	There was an internal target failure.
13	00	Additional sense code qualifier	
14	00	Field replaceable unit code.	none
15-17	00	Sense key specific	none
18-31		Additional sense bytes (vendor specific)	Data in these bytes are not valid because of sense byte #7

a. The byte field is a decimal number.

b. The value field is a hexadecimal value.

This chapter describes how to incorporate Symmetrix in the SGI/CHALLENGE environment.

◆ Overview	10-2
◆ Useful IRIX Utilities and Functions.....	10-3
◆ Symmetrix Configuration	10-5
◆ Configuring the SGI System	10-7
◆ Creating the Filesystem	10-12
◆ Creating a Stripe Volume Set.....	10-14
◆ Channel Failover	10-20
◆ Adding Devices Online	10-24
◆ System and Error Messages.....	10-25

Overview

When using Symmetrix in the Silicon Graphics (SGI)/Challenge environment, please note the following:

- ◆ For Symmetrix models 55/3500, 52/32xx, and 51/31xx, EMC recommends that you use the Fast-Wide Differential SCSI director P/N 200-881-903.
- ◆ For Symmetrix models 57/3700, 54/34xx, and 53/33xx, EMC recommends that you use either the Fast-Wide Differential SCSI director P/N 201-207-917 (60 MHz processor) or P/N 201-207-927 (66 MHz processor) or Ultra-Wide Differential SCSI director P/N 201-277-917 (60 MHz processor) or P/N 201-277-927 (66 MHz μ P).
- ◆ If LUN addressing is required, SGI only supports multiple LUNs per target ID in IRIX 5.3 or higher with XLV Volume Manager.
- ◆ Make sure the host is powered off or the Symmetrix port is offline when connecting Symmetrix to the SGI host.



CAUTION

IMPORTANT: The SGI host uses target ID 0. Do not use target ID 0 for Symmetrix devices connected to an SGI host.

Table 10-1 SGI/Challenge Operating System Features

O/S	IRIX 5.3 or higher
Largest recognizable drive size	> 36 GB
Largest defined partition	> 36 GB
Largest filesystem	2 ⁶⁴ MB
Supports multiple target IDs	Yes
Supports multiple LUNs per target ID	Yes ^a
Partition representing the entire disk	s7
To see the drive configuration, use this utility:	fx or prtvtoc
To configure drives use this utility:	fx
Special drive definition files needed	no

a. IRIX version 5.3 XFS is required for LUN support.

Useful IRIX Utilities and Functions

This section lists IRIX functions and utilities you may use to define and manage Symmetrix devices. The use of these functions and utilities is optional. They are listed for reference only.

- ◆ `MAKEDEV`
Makes device entries for all disk configured in the system.
- ◆ `hinv -v`
Displays the current hardware configuration known to the host.
- ◆ `fx`
A menu-driven disk utility through which devices can be configured.
- ◆ `mkdir`
Makes a directory.
- ◆ `mkfs`
Makes a filesystem.
- ◆ `mount`
Mounts a filesystem.
- ◆ `df`
Lists all mounted disk filesystems (with size in 512K blocks).
- ◆ `showprods`
Lists the products installed or present in the system.

- ◆ `insight`
An SGI online documentation viewer that can be executed from an x-term window.
- ◆ `uname -a`
Obtains host information including network name, operating system type and version.
- ◆ `shutdown -g0 -y`
Gracefully shuts down the operating system prior to powering off the system.
- ◆ `scsiha`
Probes and controls I/O buses.

Symmetrix Configuration

In the SGI/Challenge host environment, you can configure the Symmetrix disk devices into logical volumes. EMC recommends the following logical volume ratios as indicated in the table below:

Table 10-2 Disk Device Configuration

		Symmetrix Logical Volume to Physical Splits ^a							
		1:1	2:1	3:1	4:1	5:1	6:1	7:1	8:1 ^b
Physical Devices	4 GB	x	x						
	9 GB	x	x	x	x				
	18 GB		x	x	x	x	x	x	x
	23 GB		*	x	x	x	x	x	x
	36 GB			*	x	x	x	x	x
	47 GB			*	*	x	x	x	x
	50 GB ^c				x	x	x	x	x
	73 GB ^{c, d}					x	x	x	x

- * indicates ratios available only in microcode level 5265 or higher.
- With microcode revision level 5263, maximum is 8:1. With 5264, maximum is 16:1. With 5265, maximum is 32:1.
- The 50 and 73 GB drives are supported only by microcode revision levels 5x66 and higher.
- The 73 GB drive does not support the RAID-S configuration on Symmetrix.

For specific cylinder counts and logical volume sizes, refer to Table 2-1 on page 2-5 and the following tables in that chapter.

The EMC Customer Engineer (CE) should contact the EMC Configuration Specialist for updated online information. This information is necessary to configure Symmetrix to support the customer's host environment.

For IRIX 5.3 or higher without XFS, EMC Symmetrix 9 GB devices should not be configured to have more than 17,476 cylinders per volume. IRIX does not support filesystems larger than 8 GB without XFS.

For IRIX 5.3, assign LUNs consecutively starting at LUN 0 per target ID.
For IRIX 6.2 and higher, there is no need to assign consecutive LUNs per target ID. LUN 0 is required, however.
Do not assign target ID 0 to Symmetrix devices.

In order to run with tag command queuing, you must enable the QERR bit in the edit directors information.

Configuring the SGI System

After all the ports are connected to the Symmetrix Fast-Wide SCSI channel interfaces and the Symmetrix devices have been configured (via AutoInstall) and the interfaces brought online, use the following procedure to verify that the host can see all the Symmetrix devices assigned to that host.

1. Power on the SGI host.
2. At the System Maintenance menu on the system console, choose **1 Start System**. The operating system will boot to IRIX.

During the boot process, you may see error messages pertaining to the newly attached disks on the controller. This is normal and you can ignore these errors. If, however, these disks are already configured in the system, you need to determine the cause of the error condition.

3. At the login prompt, log into the system as **root**.
4. At the prompt, type `hinv -c disk > disk_config`
The host checks for all controllers installed in its system and lists them in the file **disk_config**.
5. At the prompt, type `more disk_config`

The contents of the file, **disk_config**, appear on the screen similar to that shown in Figure 10-1 below, where "unit 4" is the target ID.

```
Integral SCSI Controller 41: Version WD33C95A, differential, revision 1

Disk drive: unit 4, LUN 3 on SCSI Controller 41
Disk drive: unit 4, LUN 2 on SCSI Controller 41
Disk drive: unit 4, LUN 1 on SCSI Controller 41
Disk drive: unit 4 on SCSI Controller 41
```

Figure 10-1 disk_config File

LUN0 is implied when not displayed on the screen as shown in the last entry of the listing in Figure 10-1.

6. Look for the numbers reported for the Integral SCSI controllers and the unit number of the disk drive on those controllers to ensure that all Symmetrix devices configured in the host are recognized by the system.
7. Proceed to the next section, *Partitioning and Labeling the Devices* on page 10-8 if you are running IRIX 6.4 or higher.

If you are running IRIX 6.2 or earlier, change to the `/dev` directory by typing `cd /dev`

8. At the prompt, type `./MAKEDEV dks`

`MAKEDEV dks` may run for several minutes. It makes disk device entries for any disk device the host sees.

The following is the disk naming convention (or format) the host uses:

```
dks<controller#>d<target ID>l<LUN#>s<partition#>
```

For example, you may see disk names similar to the following. You may specify a disk based on this name.

```
/dev/rdisk/dsk41d4l3s7
```

or,

```
/dev/rdisk/dsk41d4s7 (for LUN0)
```

When LUN0 is used, the device name does not include the LUN number.

Partitioning and Labeling the Devices

Tag Command Queuing Disabled

Use the `fx -x` command to partition and label the new devices.

1. At the system prompt, type `fx -x`
2. At the “fx: “device-name” = (dksc)” prompt, press `<CR>` to select SCSI disk devices.
3. At the “fx: ctlr # = (0)” prompt, enter the number of the SCSI controller for a device (from `hinv` command in previous section).
4. At the “fx: drive# = (1)” prompt, enter the target ID number for that device.

5. At the "fx: lun# = (0)" prompt, press <CR> to select LUN 0 or enter the LUN number of the Symmetrix device you wish to partition on the specified controller (step 3 above).

SCSI drive type = EMC SYMMETRIX 5060 appears.

The revision number (5060) will change over time.

If the SYMMETRIX header does not appear, the selected disk is not an EMC volume.

A bad label warning should appear along with a message stating a new label was created. If this message does not appear, then the selected disk has already been labeled.

If the message "Warning: *Failed to open * /dev/rdisk/..." appears, the host does not see the selected drive or the device entry for that devices does not exist. Check the Symmetrix disk configuration.

6. At the "fx>" menu prompt, type 1 (for label).

Step 7 and 8 below allow you to view the diskinfo and partition labels. If you do not want to view this information, skip to step 9 below.

7. At the "fx/label>" prompt, type sh (for show) to display the drive parameters.
8. At the "fx/label/show>" prompt, type all to display all drive parameters
9. At the "fx/label/show>" prompt, type . . to go back a level. Repeat this step until you reach the "fx>" prompt.
10. At the "Label info has changed for disk dksc(x.x.x) . Write out change? (Yes)" prompt, press <CR> to write the label for the disk.
11. Repeat steps 3 through 10 until all new LUNs on all new controllers are labeled.
12. Type exit

If you want to work with filesystems on Symmetrix disk devices, proceed with the steps described in the following sections.

Tag Command Queuing Enabled

Have the EMC Customer Engineer set the *QERR* bit on the SA port that is connected to the SGI. Have the EMC Customer Engineer re-configure Symmetrix in an online state for these changes to take effect.

To enable the tag command and to set the host device queue depth on all IRIX operating systems:

1. At the system prompt, type `fx -x`
2. At the "fx: "device-name" = (**dksc**)" prompt, press <CR> to select SCSI disk devices.
3. At the "fx: ctlr # = (0)" prompt, enter the number of the SCSI controller for a device (from `hinv` command in previous section).
4. At the "fx: drive# = (1)" prompt, enter the target ID number for that device.
5. At the "fx: lun# = (0)" prompt, press <CR> to select LUN 0 or enter the LUN number of the Symmetrix device you wish to partition on the specified controller (step 3 above).

The following appears:

```
SCSI drive type = EMC SYMMETRIX 5060
```

The revision number (5060) will change over time.

If the SYMMETRIX header does not appear, the selected disk is not an EMC volume.

A bad label warning should appear along with a message stating a new label was created. If this message does not appear, then the selected disk has already been labeled.

If the message "**Warning: *Failed to open * /dev/rdisk/...**" appears, the host does not see the selected drive or the device entry for that devices does not exist. Check the Symmetrix disk configuration.

6. At the "fx>" menu prompt, type 1 (for label).
7. At the "fx/label>" prompt, type `se` (for set).
8. At the "fx/label>" prompt, type `para` (for parameters).
9. Press the **ENTER** key until `Enable CTQ= (disabled)` appears.
10. Type `enable` at this prompt.
11. Enter the queue depth value of the Symmetrix.
12. Press the **ENTER** key until `About to modify drive parameter... (OK) ?` appears.

13. Type `yes` at this prompt.
14. Type `..` until you see `Label info has changed...write out changes? (yes).`
15. Press the **ENTER** key to accept the default value of "yes" for this prompt.
16. Repeat steps 1 through 10 for each Symmetrix device.

Be sure to enable the QERR bit on Symmetrix.

Creating the Filesystem

Once the physical volumes are available to the operating system, create a new filesystem for each Symmetrix disk. IRIX supports two different types of filesystems:

- ◆ `efs` - this is the legacy SGI filesystem. It has been replaced with `xfs` in IRIX 6.2 and higher.
- ◆ `xfs` - this is an optional filesystem type. It allows you to create a filesystem greater than 2 GB in size. The maximum `xfs` filesystem size is 2^{64} MB.

`XLV` is a utility that is necessary for striping volumes, concatenating volumes, and system level mirroring. See *Creating a Stripe Volume Set* on page 10-14.

The following sections describe how to create and mount `efs` and `xfs` filesystems.

Creating and Mounting an xfs Filesystem

Creating and mounting an `xfs` filesystem involves performing the following steps:

1. Create a mount directory for each of the Symmetrix volumes by typing a statement similar to the following for each physical volume:

```
mkdir /symm0
```

This example creates a mount directory called `symm0`

2. Create an `xfs` filesystem using the `mkfs` command in a statement similar to the following:

```
mkfs -d name=/dev/rdsk/dks47d3s7 -b size=1k -l internal,size=10m
```

3. Before you can mount the directory, you need to place the drives in a mount table. The best way to mount all new filesystems is to update the file, `/etc/fstab`.

Edit the file with a line similar to the one shown below for each new filesystem:

```
/dev/dsk/dks20d111s7 /syemm0 xfs rw,raw=/dev/rdisk/dks20d111s7 1 2
```

Repeat this step for all additional filesystems.

4. Mount all the filesystems described in the file **/etc/fstab** by typing a statement similar to the following:

```
mount -a
```

Repeat this step for all devices.

5. Use the **df** command to show all mounted filesystems and the available free space.

Creating a Stripe Volume Set

The `XLV` utility allows the concatenation of disk devices to create larger logical volumes, stripe data across disk devices.

Keep in mind the following when working with stripe volume sets:

- ◆ The SGI host must be running IRIX 5.3 or greater with XFS.
- ◆ All devices in a stripe volume set should be the same size and type.
- ◆ For best performance use one drive from 8 different channels rather than eight drives on one channel.
- ◆ The drive partition used to create the stripe filesystem must be of type 'xfs'. You can change the filesystem type using the `fx -x` command. However, the `mkfs` command shown in the previous example will automatically convert the specified partition to an xfs file type.

Using XLV

The `XLV` utility links multiple disks together to create larger filesystems than the filesystem limit on a single drive. For example, the `efs` filesystem limit on a single drive is 2 GB. The `xfs` filesystem limit on a single drive is 18 GB. With `XLV`, the filesystem limit is 2^{64} MB.

You must use this utility when striping or concatenating devices, or performing system level mirroring. You would use this utility to create your "disks" (with or without protection) before you create your `xfs` filesystem.

Concatenating Disks (no mirroring)

Enter the following commands to concatenate multiple disks to create a larger volume. These commands do not create a mirrored volume.

```

xlv_make
vol xlv0
data
plex
ve dks0d2s7
ve dks0d3s7
end

```

```
show
exit
yes (to confirm)
```

Repeat this sequence of commands for each disk to create.

Striping Disks (no mirroring)

Enter the following commands to create a striped volume across multiple disks. These commands do not create a mirrored volume.

```
xlsv_make
vol xlv0
data
plex
ve -stripe -stripe_unit 256 dks0d2s7 dks0d3s7
end
show
exit
yes (to confirm)
```

Repeat this sequence of commands for each disk to create.

Concatenating Disks (with mirroring)

Enter the following commands to concatenate multiple disks to create a larger volume. These commands create a mirrored volume.

```
xlsv_make
vol xlv0
data
plex
ve dks0d2s7
ve dks0d3s7
plex
ve dks0d4s7
ve dks0d5s7
end
show
exit
yes (to confirm)
```

Repeat this sequence of commands for each disk to create.

Striping Disks (with mirroring)

Enter the following commands to create a striped volume across multiple disks. These commands create a mirrored volume.

```
xlsv_make
vol xlv0
data
```

```
plex
ve -stripe -stripe_unit 256 dks0d2s7 dks0d3s7
plex
ve -stripe -stripe_unit 256 dks0d4s7 dks0d5s7
end
show
exit
yes (to confirm)
```

Repeat this sequence of commands for each disk to create.

Creating and Mounting the Filesystem

Now that you have created the disk drive, you can create either an efs or xfs filesystem.

1. Create a mount directory for the disk by typing a statement similar to the following:

```
mkdir /vol1
```

This example creates a mount directory called “vol1”.

2. Create an efs or xfs filesystem using the `mkfs` command in a statement similar to the following:

For an xfs filesystem type:

```
mkfs -d name=/dev/rdisk/ylv/ylv0 -b size=1k -l internal,size=10m
```

3. Mount the filesystem by adding the following line into the `/etc/fstab` file (for an xfs filesystem) type:

```
/dev/dsk/ylv/ylv0 /vol1 xfs rev,raw=/dev/rdisk/ylv/ylv0 1 2
```

4. Mount all of the filesystems listed in the `/etc/fstab` file by entering the following command:

```
mount -a
```

Using the XVM Volume Manager

The XVM Volume Manager provides a logical organization to disk storage that enables an administrator to combine underlying physical disk storage into a single logical unit, known as a logical volume. Logical volumes behave like standard disk partitions and can be used as arguments anywhere a partition can be specified.

A logical volume allows a filesystem or raw device to be larger than the size of a physical disk. Using logical volumes can also increase disk I/O performance because a volume can be striped across more than one disk. Logical volumes can also be used to mirror data on different disks.

The `xvm` command is a command line interface which allows the administrator to create and modify volume entities.

Creating logical volume labels

```
xvm
xvm:local> label -name mydev1 dks5d5l1vol
xvm:local> label -name mydev2 dks5d5l2vol
```

If using a device attached through a fibre channel switch, the syntax would be as follows:

```
xvm:local> label -name mydev1 50060486ce8320cd/lun72vol/c14p0
```

Defining Slices

Before a device can be concatenated, mirrored, or striped it must be sliced. In the example below, one slice is defined for the entire device.

```
xvm:local> slice -all mydev1
xvm:local> slice -all mydev2
```

After the volumes have been labeled and sliced they can then be concatenated, striped, or mirrored using the `xvm` command.

Concatenating Disks

```
xvm:local> concat -vename my_volume_ve -volname my_volume_vol \
slice/mydev1s0 slice/mydev2s0
```

The `s0` appended to the device name represents the slice number which was generated using the `slice` command.

Striping a Volume

After the devices are labeled and sliced they can then be used to make a stripe set. Below is the command sequence to create a striped volume.

```
xvm:local> stripe -volname my_volume slice/mydev1s0 slice/mydev2s0
```

Creating and mounting the filesystem

Now that you have created the XVM volume, you can create an xfs filesystem.

1. Create the mount directory for the disk by typing a statement similar to the following:

```
mkdir /vol1
```

2. Create an xfs filesystem using the `mkfs` command in a statement similar to the following:

```
mkfs /dev/lxvm/my_volume_vol
```

3. Place the following line in the `/etc/fstab` file:

```
/dev/lxvm/my_volume_vol /vol1 xfs
/dev/rlxvm/my_volume_vol /vol1 rw 0 0
```

4. Mount all of the filesystems described in the `/etc/fstab` file by typing:

```
mount -a
```

Example

The following example creates a 20 GB striped volume named “stripe_0” from five 4 GB Symmetrix volumes.

1. At the prompt, enter `xlvm_make`, and type the commands shown below to create a stripe volume set called “stripe_0”.

```
xlvm_make
vol stripe_0
data
ve -stripe -stripe_unit 256 dks47d3s7 dks46d6s7 dks43d1s7 dks42d5s7 dks41d7s7
end
show
exit
```

2. At the confirmation prompt, type `yes`.

3. Create a mount point for the new stripe volume set by typing

```
mkdir /stripe_0
```

4. Make a filesystem on each Symmetrix volume that will be used in the striped filesystem using the `mkfs` command.

For example, type a command similar to the following:

```
mkfs -d name=/dev/rdsk/ylv/stripe_0 -b size=1k -l internal,size=10m
```

5. Mount the new stripe volume set by typing

```
mount /dev/dsk/ylv/stripe_0 /stripe_0
```

6. Confirm that the new stripe volume set is the size you expect using the `df -k` command.

Deleting a Stripe Volume Set

To delete a stripe filesystem unmount the volume and then run the XLV administrator utility as shown in the example below.

1. Unmount the stripe volume set, `stripe_0`, by typing

```
umount /stripe_0
```

2. Confirm the stripe volume set is unmounted by typing `df -k`

3. At the prompt, enter `ylv_admin` to start the XLV administrator utility and type in the commands shown below to delete the stripe volume set, `stripe_0`.

```
ylv_admin  
delete object stripe_0
```

4. Remove the “block” xlv device entries by typing

```
rm /dev/ylv/stripe_0
```

5. Remove the “raw” xlv device entries by typing

```
rm /dv/ylv/stripe_0
```

Channel Failover

IRIX has the ability to support up to four paths to the same set of logical volumes. If the primary path to the volume fails with this feature, the host automatically switches the path of the volume to an alternate connection without affecting the application.

This feature provides channel failover capabilities only and does not provide load balancing between paths.

Setting up Alternate Paths

To set up alternate paths, you must first create XLV or XVM logical volumes. Refer to the section titled *Creating the Filesystem* on page 10-12 for instructions.

The volumes must be of type XLV or XVM and not XFS.

Edit the `/etc/failover.conf` file to define the alternate path configuration.

This file defines the primary path and up to three alternate paths for each device. An example `/etc/failover.conf` file is below.

# Name	Pri path	alt path	alt path	alt path
GroupA	sc6d5010	sc7d5010	sc8d5010	sc9d5010
GroupB	sc6d5110	sc7d5110	sc8d5110	sc9d5110
GroupC	sc6d5211	sc7d5211	sc8d5211	sc9d5211

Below is a sample entry for devices on a Fibre Channel fabric:

fabric1	50060486ce8320cd/lun66/c14p0 \
	50060486ce8320dd/lun66/c59p0
fabric2	50060486ce8320cd/lun67/c14p0 \
	50060486ce8320dd/lun67/c59p0
fabric3	50060486ce8320cd/lun68/c14p0 \
	50060486ce8320dd/lun68/c59p0

The first field is the group name for the path. The remaining fields list the primary and alternate paths, respectively. Note that the device names are of type `/dev/scsi` and not `/dev/rdisk`. In order for channel failover to work, the devices must exist under `/dev/scsi`.

Another important rule for configuration is that both target and LUN must be the same for all of the paths. Below is an example path that will not work.

```
GroupD  sc6d5112  sc7d5113  sc8d5114  sc9d5115
```

In this example the host will report an error stating that the path contains different devices even though in reality they may be the same Symmetrix logical device.



CAUTION

The software does not check the device serial number to ensure that the devices are the same. The only form of a consistency check is that the target and LUN are the same. Therefore, it is possible that an alternate path could point to a different logical device. In this case a channel failure could result in data corruption as the application begins writing to a different volume.

Starting Channel Failover

After editing the `/etc/failover.conf` file the channel failover configuration must be activated. Type the following command:

```
/etc/init.d/failover start
```

This command is run during boot up. Therefore, channel failover will automatically be started on future reboots.

Viewing the state of paths

There are two commands to view the state of the primary versus alternate paths. The first is the `hinv` (hardware inventory) command. This command will show primary and alternate paths. There is no correlation between the paths, however. That is, there is no way to tell the primary to alternate relationship. The key information displayed by the `hinv` command is the state of the path.

Below is example output from the `hinv` command.

```
Integral SCSI controller 3: Version QL1040B (rev. 2), differential
  Disk drive: unit 0 on SCSI controller 3 (primary path)
  Disk drive: unit 0, lun 1 on SCSI controller 3 (primary path)
  Disk drive: unit 0, lun 2 on SCSI controller 3 (primary path)

Integral SCSI controller 4: Version QL1040B (rev. 2), differential
  Disk drive: unit 0 on SCSI controller 4 (alternate path)
  Disk drive: unit 0, lun 1 on SCSI controller 4 (alternate path)
  Disk drive: unit 0, lun 2 on SCSI controller 4 (alternate path)
```

Below is example `hinv` output from a failed path.

```
Integral SCSI controller 3: Version QL1040B (rev. 2), differential
  Disk drive: unit 0 on SCSI controller 3 (alternate path) DOWN
  Disk drive: unit 0, lun 1 on SCSI controller 3 (alternate path) DOWN
  Disk drive: unit 0, lun 2 on SCSI controller 3 (alternate path) DOWN

Integral SCSI controller 4: Version QL1040B (rev. 2), differential
  Disk drive: unit 0 on SCSI controller 4 (primary path)
  Disk drive: unit 0, lun 1 on SCSI controller 4 (primary path)
  Disk drive: unit 0, lun 2 on SCSI controller 4 (primary path)
```

As you can see in the above example the once alternate path now shows up as primary and the once primary path now shows up as alternate and DOWN.

The other useful command for displaying the alternate path configuration is `scsifo`. The output from this command is more specific, as far as the relationship between primary and alternate paths, than the `hinv` command. Example output of the command follows.

```
# scsifo -d

Group C:
  [P] sc3d012 (184)
  [ ] sc4d012 (430)
Group B:
  [P] sc3d011 (175)
  [ ] sc4d011 (421)
Group A:
  [P] sc3d010 (166)
  [ ] sc4d010 (412)
```

After a path failure, the output of the `scsifo` command will look like the following.

```
# scsifo -d

Group C:
  [ ] sc3d012 (184)
  [P] sc4d012 (430)
Group B:
  [ ] sc3d011 (175)
  [P] sc4d011 (421)
Group A:
  [ ] sc3d010 (166)
  [P] sc4d010 (412)
```

Adding Devices Online

Whenever devices are added online to the Symmetrix unit or device channel addresses are changed, you will need to perform the actions described below in order to introduce the new devices to the system.

The SGI operating system (6.2 and earlier) does not support adding devices online.

The EMC PSE Configuration Group can add new devices to the Symmetrix configuration while Symmetrix remains online to the host. Once this is complete, stop all host activity and reboot the host using `shutdown -y -g0 -i6`. See *Configuring the SGI System* on page 10-7 to introduce the new devices to the host environment.

IRIX 6.4 and higher enables users to add devices online using the following commands:

```
# hinv (to see the I/O buses)
# scsiha -rp bus#
# ioconfig -f /hw
# hinv (to see the devices)
```

System and Error Messages

SGI logs system and error messages including SCSI error messages to a file called **/var/adm/SYSLOG** and also displays these messages at the system console.

For old error messages, SGI archives these messages into **/var/adm/oSYSLOG**.

`syslog` writes the message onto the system log and tags it with a priority. This message can be written to the host console, log files, or forward to another host as specified.

Priorities are encoded as a facility and a level. The facility describes the part of the system generating the message. The level is selected from an ordered list:

- ◆ **LOG_EMERG** - panic condition. The message is broadcast to all users.
- ◆ **LOG_ALERT** - condition has occurred that should be corrected immediately - such as a corrupted system database.
- ◆ **LOG_CRIT** - a critical condition has occurred such as a hard device error.
- ◆ **LOG_ERR** - error
- ◆ **LOG_WARNING** - warning message.
- ◆ **LOG_NOTICE** - condition is not an error condition, but is one that should be specially handled.
- ◆ **LOG_INFO** - information message.
- ◆ **LOG_DEBUG** - debug message.

If `syslog` cannot write the message to its **SYSLOG**, it will attempt to write the message to the host console if the **LOG_CONS** option is set. `logopt` specifies the logging options. Some logging options include:

- ◆ **LOG_PID** - log the process ID with each message.
- ◆ **LOG_CONS** - force writing messages to the console if unable to send it to **SYSLOG**.
- ◆ **LOG_ODELAY** - delay opening the connection to **SYSLOG** until the first `syslog` call.
- ◆ **LOG_NDELAY** - open the connection to **SYSLOG** immediately.

The facility parameter encodes a default facility to be assigned to all messages that do not have an explicit facility encoded:

- ◆ LOG_KERN - messages generated by the kernel.
- ◆ LOG_USER - messages generated by random user processes.
- ◆ LOG_MAIL - the mail system.
- ◆ LOG_DAEMON - system daemons.
- ◆ LOG_AUTH - the authorization system.
- ◆ LOG_LPR - the line printer spooling system.
- ◆ LOG_NEWS - the network news system.
- ◆ LOG_UUCP - the UNIX-to-UNIX copy system.

This chapter describes how to incorporate Symmetrix into the Silicon Graphics (SGI) environment using Origin systems running IRIX 6.4 or greater.

◆ Overview	11-2
◆ Useful IRIX Utilities and Functions.....	11-3
◆ Symmetrix Configuration	11-5
◆ Configuring the SGI System	11-6
◆ Creating the Filesystem	11-10
◆ Creating a Stripe Volume Set.....	11-12
◆ Channel Failover	11-18
◆ Adding Devices Online	11-22
◆ System and Error Messages	11-23

Overview

When using Symmetrix in the Silicon Graphics (SGI)/Origin environment, please note the following:

- ◆ Symmetrix attaches to Silicon Graphics Corporation's SGI Fast-Wide Differential SCSI (IO4) card, Ultra-wide Differential SCSI (XIO) card, and the FWD SCSI card in the SGI S-Series.
- ◆ For Symmetrix models 55/3500, 52/32xx, and 51/31xx, EMC recommends that you use the FWD SCSI director P/N 200-881-903.
- ◆ For Symmetrix models 57/3700, 54/34xx, and 53/33xx, EMC recommends that you use either the FWD SCSI director P/N 201-207-917 (60 MHz processor) or P/N 201-207-927 (66 MHz processor) or UWD SCSI director P/N 201-277-917 (60 MHz processor) or P/N 201-277-927 (66 MHz processor).
- ◆ IRIX supports target IDs 1-15. Each target can support LUNs 0-7. LUN 0 must be assigned to each target.

IRIX 6.4 uses target 0 as the initiator. If you want to assign devices to target 0, use the `nvr` command to change the initiator ID. For example, the command `nvr scsihostid 7` changes the initiator ID to 7. You must reboot to have this change take effect.

Table 11-1 SGI/Origin Operating System Features

O/S	IRIX 6.4 or higher
Largest recognizable drive size	>36 GB
Largest defined partition	>36 GB
Largest filesystem	2 ⁶⁴ MB
Supports multiple target IDs	Yes
Supports multiple LUNs per target ID	Yes. (LUN 0 must be assigned to each target.)
Partition representing the entire disk	s7
To see drive configuration, use utility:	<code>fx</code> or <code>prtvtoc</code>
To configure drives use this utility:	<code>fx</code>
Special drive definition files needed	No

Useful IRIX Utilities and Functions

This section lists IRIX functions and utilities you may use to define and manage Symmetrix devices. The use of these functions and utilities is optional. They are listed for reference only.

- ◆ `MAKEDEV`
Makes devices entries for all disk configured in the system.
- ◆ `hinv -v`
Displays the current hardware configuration known to the host.
- ◆ `fx`
A menu-driven disk utility through which devices can be configured.
- ◆ `mkdir`
Makes a directory.
- ◆ `mkfs`
Makes a filesystem.
- ◆ `mount`
Mounts a filesystem.
- ◆ `df`
Lists all mounted disk filesystems (with size in 512k blocks).
- ◆ `showprods`
Lists the products installed or present in the system.

- ◆ `insight`
An SGI online documentation viewer that can be executed from an x-term window.
- ◆ `uname -a`
Obtains host information including network name, operating system type and version.
- ◆ `shutdown -g0 -y`
Gracefully shuts down the operating system prior to powering off the system.
- ◆ `xlvmake`
Used to create logical volumes.
- ◆ `xlvmgr`
Used to manage already defined logical volumes.
- ◆ `scsiha`
Used to probe and control SCSI buses.
- ◆ `ioconfig`
Used to configure the kernel.
- ◆ `scsiha`
Probes and controls I/O buses.

Symmetrix Configuration

In the SGI/Origin host environment, you can configure the Symmetrix disk devices into logical volumes. EMC recommends the following logical volume ratios as indicated in the table below:

Table 11-2 Disk Device Configuration

		Symmetrix Logical Volume to Physical Splits ^a							
		1:1	2:1	3:1	4:1	5:1	6:1	7:1	8:1 ^b
Physical Devices	4 GB	x	x						
	9 GB	x	x	x	x				
	18 GB		x	x	x	x	x	x	x
	23 GB		*	x	x	x	x	x	x
	36 GB			*	x	x	x	x	x
	47 GB			*	*	x	x	x	x
	50 GB ^c				x	x	x	x	x
	73 GB ^{c, d}					x	x	x	x

- a. * indicates ratios available only in microcode level 5265 or higher.
- b. With microcode revision level 5263, maximum is 8:1. With 5264, maximum is 16:1. With 5265, maximum is 32:1.
- c. The 50 and 73 GB drives are supported only by microcode revision levels 5x66 and higher.
- d. The 73 GB drive does not support the RAID-S configuration on Symmetrix.

For specific cylinder counts and logical volume sizes, refer to Table 2-1 on page 2-5 and the following tables in that chapter.

The EMC Customer Engineer (CE) should contact the EMC Configuration Specialist for updated online information. This information is necessary to configure Symmetrix to support the customer's host environment.

Configuring the SGI System

After all the ports are connected to the Symmetrix Fast-Wide SCSI channel interfaces and the Symmetrix devices have been configured (via AutoInstall) and the interfaces brought online, use the following procedure to verify that the host can see all the Symmetrix devices assigned to that host.

1. Power on the SGI host. The host will start to boot.

During the boot process, you may see error messages pertaining to the newly attached disks on the controller. This is normal and you can ignore these errors. If, however, these disks are already configured in the system, you need to determine the cause of the error condition.

2. At the login prompt, log into the system as **root**.
3. At the prompt, type

```
hinv -c disk > disk_config
```

The host checks for all controllers installed in its system and lists them in the file **disk_config**.

4. At the prompt, type

```
more disk_config
```

The contents of the file, **disk_config**, appears on the screen like the listing shown below. The target ID is unit 4.

```
Integral SCSI Controller 41: Version WD33C95A, differential, revision 1
```

```
Disk drive: unit 4, LUN 3 on SCSI Controller 41
Disk drive: unit 4, LUN 2 on SCSI Controller 41
Disk drive: unit 4, LUN 1 on SCSI Controller 41
Disk drive: unit 4 on SCSI Controller 41
```

LUN0 is implied when not displayed on the screen as shown in the last entry of the listing.

If devices are on a Fibre Channel fabric, the device listing will look like the following:

```
Integral SCSI controller 58: Version Fibre Channel QL2200
Fabric Disk: node 50060486ce8320dd port 0 lun 0 on SCSI controller 58
Fabric Disk: node 50060486ce8320c3 port 0 lun 0 on SCSI controller 58
Fabric Disk: node 50060486ce8320c3 port 0 lun 1 on SCSI controller 58
```

Refer to the *Symmetrix Fibre Channel Product Guide*, P/N 200-999-642 for additional information on Fabric path names.

5. Look for the numbers reported for the Integral SCSI controllers and the unit number of the disk drive on those controllers to ensure that all Symmetrix devices configured in the host are recognized by the system.

Partitioning and Labeling the Devices

Use the `fx -x` command to partition and label the new devices.

1. At the system prompt, type `fx -x`

If you are partitioning devices on a Fibre Channel fabric, then the device path must be specified on the command line. For example,

```
fx -x -d /dev/rdisk/50060486ce8320c3/lun12vol/c58p0
```

2. At the "fx: "device-name" = (dksc)" prompt, press <CR> to select SCSI disk devices.
3. At the "fx: ctlr # = (0)" prompt, enter the number of the SCSI controller for a device (from `hinv` command in previous section).
4. At the "fx: drive# = (1)" prompt, enter the target ID number for that device.
5. At the "fx: lun# = (0)" prompt, press <CR> to select LUN 0 or enter the LUN number of the Symmetrix device you wish to partition on the specified controller (step 3 above).

The following appears

```
SCSI drive type = EMC SYMMETRIX 5060
```

The revision number (5060) will change over time.

If the SYMMETRIX header does not appear, the selected disk is not an EMC volume.

A bad label warning should appear along with a message stating a new label was created. If this message does not appear, then the selected disk has already been labeled.

If the message “Warning: *Failed to open * /dev/rdisk/...” appears, the host does not see the selected drive or the device entry for that devices does not exist. Check the Symmetrix disk configuration.



CAUTION

The following steps will destroy any existing data on the affected drives. Use only when configuring new disks that do not yet have any data on them.

6. At the `fx>` menu prompt, type `r` (repartition).
7. At the `fx repartition>` prompt, type `o` (option drive).
8. FX then prompts you to specify a filesystem. Type `xfss`
9. The system then prompts you to specify whether you want to create a log partition. Type `yes`
10. Type `..` (two dots) to create another disk type. Repeat this step to create as many disk types as you need to.
11. Repeat steps 6 through 10 to label all the drives.
12. When you are finished labeling all the drives, type `exit`

Tag Command Queuing Enabled

Perform the following actions prior to enabling the tag command queuing.

1. Perform steps 1 through 6 in the previous section, *Partitioning and Labeling the Devices*.
2. At the `fx>` prompt, type `L` (label).
3. At the `fx/label>` prompt, type `se` (set).
4. At the `fx/label` prompt, type `para` (parameters).

5. Press **<Enter>** repeatedly until the message `Enable CTQ= (disabled)` appears.
6. Type `enable` at this prompt.
7. Enter the queue depth value of the Symmetrix; for example, `CTQ=24`.
8. Press `Enter` repeatedly until the prompt `About to modify drive parameter... (OK)?` appears.
9. Type `yes` at this prompt.

If you want to work with filesystems on Symmetrix disk devices, proceed with the steps described in the following sections.

Creating the Filesystem

Once the physical volumes are available to the operating system, you can create a new filesystem for each Symmetrix disk.

XLV is a utility that is necessary for striping volumes, concatenating volumes, and system level mirroring. See *Creating a Stripe Volume Set* on page 11-12.

The following describe how to create and mount xfs filesystems.

Creating and Mounting an xfs Filesystem

Creating and mounting an xfs filesystem involves performing the following steps:

1. Create a mount directory for each of the Symmetrix volumes by typing a statement similar to the following for each physical volume:

```
mkdir /symm0
```

This example creates a mount directory called “symm0”.

2. Create an xfs filesystem using the `mkfs` command in a statement similar to the following:

```
mkfs -d name=/dev/rdsk/dks47d3s7 -b size=1k -l internal,size=10m
```

This statement creates a filesystem for drive LUN 0 attached to controller 3 in section 7. Repeat this statement for all Symmetrix devices.

Your parameters will vary depending on your particular application.

3. Before you can mount the directory, you need to place the drives in a mount table. The best way to mount all new filesystems is to update the file, **/etc/fstab**. Edit the file with a line similar to the one shown below for each new filesystem:

```
/dev/dsk/dks20d111s7 /symm0 xfs rw,raw=/dev/rdsk/dks20d111s7 1 2
```

Repeat this step for all additional filesystems.

4. Mount all the filesystems described in the file **/etc/fstab** by typing a statement similar to the following:

```
mount -a
```

5. Use the **df** command to show all mounted filesystems and the available free space.

Creating a Stripe Volume Set

The XLV utility allows you to concatenate disk devices to create larger logical volumes, stripe data across disk devices, or change the configuration of volumes while that volume is actively in use as a filesystem.

Keep in mind the following when working with stripe volume sets:

- ◆ All devices in a stripe volume set should be the same size and type.
- ◆ For best performance use one drive from 8 different channels rather than eight drives on one channel.

Using XLV

The XLV utility links multiple disks together to create larger filesystems than possible on a single drive. While filesystems can be made much larger using XLV, the XFS filesystem limit of 9 million terabytes cannot be exceeded.

You must use this utility when striping or concatenating devices, or performing system level mirroring. You would use this utility to create your volumes (with or without protection) before you create your xfs filesystem.

Concatenating Disks (no mirroring)

Enter the following commands to concatenate multiple disks to create a larger volume. These commands do not create a mirrored volume.

```

xlv_make
vol xlv0
data
plex
ve dks0d2s7
ve dks0d3s7
end
show
exit
yes (to confirm)

```

The `vol xlv0` command defines the logical volume name.

Repeat this sequence of commands for each disk to create.

Striping Disks (no mirroring)

Enter the following commands to create a striped volume across multiple disks. These commands do not create a mirrored volume.

```

xlv_make
vol xlv0
data
plex
ve -stripe -stripe_unit 256 dks0d2s7 dks0d3s7
end
show
exit
yes (to confirm)

```

Repeat this sequence of commands for each disk to create.

Concatenating Disks (with mirroring)

Enter the following commands to concatenate multiple disks to create a larger volume. These commands create a mirrored volume.

```

xlv_make
vol xlv0
data
plex
ve dks0d2s7
ve dks0d3s7
plex
ve dks0d4s7
ve dks0d5s7
end
show
exit
yes (to confirm)

```

Repeat this sequence of commands for each disk to create.

Striping Disks (with mirroring)

Enter the following commands to create a striped volume across multiple disks. These commands create a mirrored striped volume.

```

xlv_make
vol xlv0
data
plex
ve -stripe -stripe_unit 256 dks0d2s7 dks0d3s7
plex
ve -stripe -stripe_unit 256 dks0d4s7 dks0d5s7
end
show
exit
yes (to confirm)

```

Repeat this sequence of commands for each disk to create.

Creating and Mounting the Filesystem

Now that you have created the volume, you can create an xfs filesystem.

1. Create a mount directory for the disk by typing a statement similar to the following:

```
mkdir /vol1
```

2. Create an xfs filesystem using the `mkfs` command in a statement similar to the following:

```
mkfs -d name=/dev/xlv/xlv0 -b size=1k -l internal,size=10m
```

3. Place the following line into the `/etc/fstab` file:

```
/dev/xlv/xlv0 /vol1 xfs rev,raw=/dev/rxlv/xlv0 1 2
```

4. Mount all the filesystems described in the `/etc/fstab` file by typing `mount -a`

Deleting a Stripe Volume Set

Before you can build a new stripe filesystem with drives used for stripe filesystems, or use the same stripe filesystem name, you must delete the old stripe filesystem.

To delete a stripe filesystem you must unmount the volume and then run the XLV administrator utility as shown in the example below.

1. Unmount the stripe volume set, `stripe_0`, by typing

```
umount /stripe_0
```

2. Confirm the stripe volume set is unmounted by typing

```
df -k
```

3. At the prompt, enter `xlv_mgr` to start the XLV administrator utility and type in the commands shown below to delete the stripe volume set, `stripe_0`.

```
xlv_mgr> delete object stripe_0
```

4. Remove the “block” XLV device entries by typing

```
#rm /dev/xlv/stripe_0
```

5. Remove the “raw” XLV device entries by typing

```
#rm /dev/xlv/stripe_0
```

Using the XVM Volume Manager

The XVM Volume Manager provides a logical organization to disk storage that enables an administrator to combine underlying physical disk storage into a single logical unit, known as a logical volume. Logical volumes behave like standard disk partitions and can be used as arguments anywhere a partition can be specified.

A logical volume allows a filesystem or raw device to be larger than the size of a physical disk. Using logical volumes can also increase disk I/O performance because a volume can be striped across more than one disk. Logical volumes can also be used to mirror data on different disks.

The `xvm` command is a command line interface which allows the administrator to create and modify volume entities.

Creating logical volume labels

```
xvm
xvm:local> label -name mydev1 dks5d5l1vol
xvm:local> label -name mydev2 dks5d5l2vol
```

If using a device attached through a fibre channel switch, the syntax would be as follows:

```
xvm:local> label -name mydev1 50060486ce8320cd/lun72vol/c14p0
```

Defining Slices

Before a device can be concatenated, mirrored, or striped it must be sliced. In the example below, one slice is defined for the entire device.

```
xvm:local> slice -all mydev1
xvm:local> slice -all mydev2
```

After the volumes have been labeled and sliced they can then be concatenated, striped, or mirrored using the `xvm` command.

Concatenating Disks

```
xvm:local> concat -vename my_volume_ve -volname my_volume_vol \
slice/mydev1s0 slice/mydev2s0
```

The `s0` appended to the device name represents the slice number which was generated using the `slice` command.

Striping a Volume

After the devices are labeled and sliced they can then be used to make a stripe set. Below is the command sequence to create a striped volume.

```
xvm:local> stripe -volname my_volume slice/mydev1s0 slice/mydev2s0
```

Creating and mounting the filesystem

Now that you have created the XVM volume, you can create an xfs filesystem.

1. Create the mount directory for the disk by typing a statement similar to the following:

```
mkdir /vol1
```

2. Create an xfs filesystem using the `mkfs` command in a statement similar to the following:

```
mkfs /dev/lxvm/my_volume_vol
```

3. Place the following line in the `/etc/fstab` file:

```
/dev/lxvm/my_volume_vol /vol1 xfs
/dev/rlxvm/my_volume_vol /vol1 rw 0 0
```

4. Mount all of the filesystems described in the `/etc/fstab` file by typing:

```
mount -a
```

Example

The following example creates a 20 GB striped volume named “stripe_0” from five 4 GB Symmetrix volumes.

1. At the prompt, enter `xlvm_make`, and type the commands shown below to create a stripe volume set called “stripe_0”.

```
xlvm_make
vol stripe_0
data
ve -stripe -stripe_unit 256 dks47d3s7 dks46d6s7 dks43d1s7 dks42d5s7 dks41d7s7
end
show
exit
```

2. At the confirmation prompt, type `yes`.
3. Create a mount point for the new stripe volume set by typing

```
mkdir /stripe_0
```


4. Make a filesystem on each Symmetrix volume that will be used in the striped filesystem using the `mkfs` command. For example, type a command similar to the following:

```
mkfs -d name=/dev/xlv/stripe_0 -b size=1k -l internal,size=10m
```

5. Mount the new stripe volume set by typing

```
mount /dev/xlv/stripe_0 /stripe_0
```

6. Confirm that the new stripe volume set is the size you expect using the `df -k` command.

Channel Failover

IRIX has the ability to support up to four paths to the same set of logical volumes. If the primary path to the volume fails with this feature, the host automatically switches the path of the volume to an alternate connection without affecting the application.

This feature provides channel failover capabilities only and does not provide load balancing between paths.

Setting up Alternate Paths

To set up alternate paths, you must first create XLV or XVM logical volumes. Refer to the section titled *Creating the Filesystem* on page 11-10 for instructions.

The volumes must be of type XLV or XVM and not XFS.

Edit the **/etc/failover.conf** file to define the alternate path configuration.

This file defines the primary path and up to three alternate paths for each device. An example **/etc/failover.conf** file is below.

#	Name	Pri path	alt path	alt path	alt path
GroupA	sc6d5010	sc7d5010	sc8d5010	sc9d5010	
GroupB	sc6d5110	sc7d5110	sc8d5110	sc9d5110	
GroupC	sc6d5211	sc7d5211	sc8d5211	sc9d5211	

Below is a sample entry for devices on a Fibre Channel fabric:

fabric1	50060486ce8320cd/lun66/c14p0 \	50060486ce8320dd/lun66/c59p0
fabric2	50060486ce8320cd/lun67/c14p0 \	50060486ce8320dd/lun67/c59p0
fabric3	50060486ce8320cd/lun68/c14p0 \	50060486ce8320dd/lun68/c59p0

The first field is the group name for the path. The remaining fields list the primary and alternate paths, respectively. Note that the device names are of type **/dev/scsi** and not **/dev/rdisk**. In order for channel failover to work, the devices must exist under **/dev/scsi**.

Another important rule for configuration is that both target and LUN must be the same for all of the paths. Below is an example path that will not work.

```
GroupD  sc6d5112  sc7d5113  sc8d5114  sc9d5115
```

In this example the host will report an error stating that the path contains different devices even though in reality they may be the same Symmetrix logical device.



CAUTION

The software does not check the device serial number to ensure that the devices are the same. The only form of a consistency check is that the target and LUN are the same. Therefore, it is possible that an alternate path could point to a different logical device. In this case a channel failure could result in data corruption as the application begins writing to a different volume.

Starting Channel Failover

After editing the `/etc/failover.conf` file the channel failover configuration must be activated. Type the following command:

```
/etc/init.d/failover start
```

This command is run during boot up. Therefore, channel failover will automatically be started on future reboots.

Viewing the State of Paths

There are two commands to view the state of the primary versus alternate paths. The first is the `hinv` (hardware inventory) command. This command will show primary and alternate paths. There is no correlation between the paths, however. That is, there is no way to tell the primary to alternate relationship. The key information displayed by the `hinv` command is the state of the path.

Below is example output from the `hinv` command.

```
Integral SCSI controller 3: Version QL1040B (rev. 2), differential
  Disk drive: unit 0 on SCSI controller 3 (primary path)
  Disk drive: unit 0, lun 1 on SCSI controller 3 (primary path)
  Disk drive: unit 0, lun 2 on SCSI controller 3 (primary path)

Integral SCSI controller 4: Version QL1040B (rev. 2), differential
  Disk drive: unit 0 on SCSI controller 4 (alternate path)
  Disk drive: unit 0, lun 1 on SCSI controller 4 (alternate path)
  Disk drive: unit 0, lun 2 on SCSI controller 4 (alternate path)
```

Below is example `hinv` output from a failed path.

```
Integral SCSI controller 3: Version QL1040B (rev. 2), differential
  Disk drive: unit 0 on SCSI controller 3 (alternate path) DOWN
  Disk drive: unit 0, lun 1 on SCSI controller 3 (alternate path) DOWN
  Disk drive: unit 0, lun 2 on SCSI controller 3 (alternate path) DOWN

Integral SCSI controller 4: Version QL1040B (rev. 2), differential
  Disk drive: unit 0 on SCSI controller 4 (primary path)
  Disk drive: unit 0, lun 1 on SCSI controller 4 (primary path)
  Disk drive: unit 0, lun 2 on SCSI controller 4 (primary path)
```

As you can see in the above example the once alternate path now shows up as primary and the once primary path now shows up as alternate and DOWN.

The other useful command for displaying the alternate path configuration is `scsifo`. The output from this command is more specific, as far as the relationship between primary and alternate paths, than the `hinv` command. Example output of the command follows.

```
# scsifo -d

Group C:
  [P] sc3d012 (184)
  [ ] sc4d012 (430)
Group B:
  [P] sc3d011 (175)
  [ ] sc4d011 (421)
Group A:
  [P] sc3d010 (166)
  [ ] sc4d010 (412)
```

After a path failure, the output of the `scsifo` command will look like the following.

```
# scsifo -d

Group C:
  [ ] sc3d012 (184)
  [P] sc4d012 (430)
Group B:
  [ ] sc3d011 (175)
  [P] sc4d011 (421)
Group A:
  [ ] sc3d010 (166)
  [P] sc4d010 (412)
```

Adding Devices Online

Whenever devices are added online to the Symmetrix unit or device channel addresses are changed, you will need to perform the actions described below in order to introduce the new devices to the system.

The EMC PSE Configuration Group can add new devices to the Symmetrix configuration while Symmetrix remains online to the host.

IRIX version 6.4 and greater allow online configuration. Once the Symmetrix devices have been added, issue the following commands:

```
# hinv (to see the I/O buses)
# scsiha -rp bus#
# ioconfig -f /hw
# hinv (to see the devices)
```

System and Error Messages

SGI logs system and error messages including SCSI error messages to a file called **/var/adm/SYSLOG** and also displays these messages at the system console.

For old error messages, SGI archives these messages into **/var/adm/oSYSLOG**.

`syslog` writes the message onto the system log and tags it with a priority. This message can be written to the host console, log files, or forward to another host as specified.

Priorities are encoded as a facility and a level. The facility describes the part of the system generating the message. The level is selected from an ordered list:

- ◆ **LOG_EMERG** - panic condition. The message is broadcast to all users.
- ◆ **LOG_ALERT** - condition has occurred that should be corrected immediately - such as a corrupted system database.
- ◆ **LOG_CRIT** - a critical condition has occurred such as a hard device error.
- ◆ **LOG_ERR** - error
- ◆ **LOG_WARNING** - warning message.
- ◆ **LOG_NOTICE** - condition is not an error condition, but is one that should be specially handled.
- ◆ **LOG_INFO** - information message.
- ◆ **LOG_DEBUG** - debug message.

If `syslog` cannot write the message to its **SYSLOG**, it will attempt to write the message to the host console if the **LOG_CONS** option is set. `logopt` specifies the logging options. Some logging options include:

- ◆ **LOG_PID** - log the process ID with each message.
- ◆ **LOG_CONS** - force writing messages to the console if unable to send it to **SYSLOG**.
- ◆ **LOG_ODELAY** - delay opening the connection to **SYSLOG** until the first `syslog` call.
- ◆ **LOG_NDELAY** - open the connection to **SYSLOG** immediately.

The facility parameter encodes a default facility to be assigned to all messages that do not have an explicit facility encoded:

- ◆ LOG_KERN - messages generated by the kernel.
- ◆ LOG_USER - messages generated by random user processes.
- ◆ LOG_MAIL - the mail system.
- ◆ LOG_DAEMON - system daemons.
- ◆ LOG_AUTH - the authorization system.
- ◆ LOG_LPR - the line printer spooling system.
- ◆ LOG_NEWS - the network news system.
- ◆ LOG_UUCP - the UNIX-to-UNIX copy system.

This chapter describes how to incorporate Symmetrix in the Sun/Solaris environment.

◆ Overview	12-2
◆ Useful Solaris Utilities and Functions	12-4
◆ Hardware Mapping	12-5
◆ Symmetrix Configuration	12-6
◆ Incorporating Symmetrix	12-9
◆ Format, Partition, Label	12-12
◆ Adding Devices Online	12-14
◆ Solstice DiskSuite	12-15
◆ Veritas Volume Manager	12-25
◆ Create and Mount a File System	12-40
◆ System and Error Messages	12-48

Overview

This chapter covers the Sun/Solaris environment. See other chapters for other Sun operating systems. When using Symmetrix in the Sun/Solaris environment, please note the following:

- ◆ For Symmetrix models 55/3500, 52/32xx, and 51/31xx, EMC recommends that you use the Fast-Wide Differential SCSI director P/N 200-881-903.
- ◆ For Symmetrix models 57/3700, 54/34xx, 53/33xx and Symmetrix 5 (8130, 8430, and 8730), EMC recommends that you use the Fast-Wide Differential SCSI director P/N 201-207-917 (60 MHz processor), P/N 201-207-927 (66 MHz processor), Ultra-Wide Differential SCSI director P/N 201-277-917 (60 MHz processor), P/N 201-277-927 (66 MHz processor), or Fibre Channel director P/N 200-520-910 or P/N 200-520-920.
- ◆ Symmetrix attaches to a Sun Sbus Fast-Wide Differential SCSI Controller (Model X1062A), Sun Sbus Ultra-Wide Differential SCSI Controller (Model X1065A), Sun PCI Ultra-Wide Differential SCSI Controller (Model X6541A), or Jaycor Sbus Fibre Channel (Models FC-1063-EMC and FC64-1063-EMC), Jaycor PCI Fibre Channel (Model FCIN-1063-EMC), Emulex PCI Fibre Channel (Models LP7000E-EMC and LP8000-EMC) or QLogic PCI Fibre Channel (Models QLA2100F-EMC and QLA2200F-EMC).
- ◆ The Sun host must have Wide, Ultra Wide SCSI Differential, or Fibre Channel cards installed in order to connect to the Symmetrix unit.
- ◆ When using Symmetrix with the Sun host in the Solaris environment, the Sun host must be running Solaris 2.3 or greater.
If Jaycor Fibre Channel cards are installed in the host, the host must be running Solaris 2.5.1 or greater.
If Emulex or QLogic Fibre Channel cards are installed in the host, the host must be running Solaris 2.6 or greater.
- ◆ With Solaris 2.3, do not use Command Tag Queueing with Symmetrix. With Solaris 2.4 or above, use Command Tag Queueing with Symmetrix.

Table 12-1 below summarizes the operating system features in the Sun/Solaris environment.

Table 12-1 Sun/Solaris Operating System Features

O/S	Sun Solaris 2.3 or greater
Largest recognizable drive size	N/A (unlimited)
Largest defined partition	N/A (unlimited)
Largest data file	2 GB (4,194,304 blks)
Largest filesystem	N/A (unlimited)
Supports multiple Target IDs	Yes
Supports multiple LUNs per target ID	Yes
Partition representing the entire disk	s2
To see the drive configuration, use this utility:	format
To configure drives use this utility:	format
Special drive definition files needed	yes

Useful Solaris Utilities and Functions

This section lists Solaris functions and utilities you may use to define and manage Symmetrix devices. The use of these functions and utilities is optional. They are listed for reference only.

◆ `format`

The Solaris disk format utility. Allows you to format, partition, and label disk drives.

◆ `newfs`

Creates a filesystem.

◆ `shutdown`

Gracefully shutdown the system.

shutdown is the preferred command for system shutdown. The **halt** command is not recommend.

Hardware Mapping

Use the following procedure to determine the target IDs available to the Symmetrix unit.

If the host is not powered:

1. Power on the Sun host system and wait for it to perform its self-tests.
2. At the **Open Firmware** prompt (**OK>**), type `printenv`
3. Review the listing that appears and determine the value of **scsi_initiator_id**. This target ID is reserved for the system and may not be used by any other device.

If the host is powered up:

1. Type `eeeprom` at the system prompt.
2. Review the listing that appears and determine the value of **scsi_initiator_id**. This target ID is reserved for the system and may not be used by any other device.

Do NOT use the **scsi_initiator_id** value as a target ID value for Symmetrix devices.

Symmetrix Configuration

In the Solaris host environment, you can configure the Symmetrix disk devices into logical volumes. EMC recommends the following logical volume ratios as indicated in the table below:

Table 12-2 Disk Device Configuration

		Symmetrix Logical Volume to Physical Splits ^a							
		1:1	2:1	3:1	4:1	5:1	6:1	7:1	8:1 ^b
Physical Devices	4 GB	x	x						
	9 GB	x	x	x	x				
	18 GB		x	x	x	x	x	x	x
	23 GB		*	x	x	x	x	x	x
	36 GB			*	x	x	x	x	x
	47 GB			*	*	x	x	x	x
	50 GB ^c				x	x	x	x	x
	73 GB ^{c, d}					x	x	x	x

- a. * indicates ratios available only in microcode level 5265 or higher.
- b. With microcode revision level 5263, maximum is 8:1. With 5264, maximum is 16:1. With 5265, maximum is 32:1.
- c. The 50 and 73 GB drives are supported only by microcode revision levels 5x66 and higher.
- d. The 73 GB drive does not support the RAID-S configuration on Symmetrix.

For specific cylinder counts and logical volume sizes, refer to Table 2-1 on page 2-5 and the following tables in that chapter.

The EMC Customer Engineer (CE) should contact the EMC Configuration Specialist for updated online information. This information is necessary to configure Symmetrix to support the customer's host environment.

SPARCstations and SPARCservers

Perform this procedure prior to booting the operating system when running Solaris on SPARCstations and SPARCservers.

1. Power on the Sun server and wait for it to perform its self-tests.
2. When the system banner appears on the screen display, press **<BREAK>** if you are using an ASCII terminal or **<STOP> <A>** if you are using a workstation terminal.
3. At the **OK>** prompt, type

```
probe-scsi-all
```

The host looks for all attached SCSI devices and displays information similar to the following (Units 0-3 are LUNs):

```
Target 0
Unit 0 Disk EMC SYMMETRIX 50607801E153
      Copyright (c) 1994
      EMC Corp. All rights reserved.
Unit 1 Disk EMC SYMMETRIX 50607801D153
      Copyright (c) 1994
      EMC Corp. All rights reserved.
Unit 2 Disk EMC SYMMETRIX 50607801F160
      Copyright (c) 1994
      EMC Corp. All rights reserved.
Unit 3 Disk EMC SYMMETRIX 506078020161
      Copyright (c) 1994
      EMC Corp. All rights reserved.
```

This disk value may change based on microcode rev in use

The disk value returned (e.g., 50607810E153) provides additional detail. The syntax of this number is MMMMWXXYYZ, where:

MMMM = Symmetrix microcode revision level

WW = Last two digits of the Symmetrix serial number

XXX = Symmetrix device number

YY = Symmetrix SCSI Director number

Z = Symmetrix SCSI Direct port number

4. Make sure the target ID and LUNs listed agree with the device information configured in the Symmetrix unit.

If this information does not agree, check the Symmetrix-host port connection and the AutoInstall configuration.

5. At the **OK>** prompt, type

```
boot -r
```

The host will boot to a login prompt. The host will also issue a warning message for each Symmetrix logical unit attached but not yet labeled.

UltraSPARC Series

On certain SUN platforms such as the SUN UltraSPARC Series, the following message may be displayed when you run `probe-scsi-all`:

```
OK> probe-scsi-all
```

```
This command may hang the system if a STOP-A or halt
command has been executed. Please type reset-all to
reset the system before executing this command.
```

```
Do you wish to continue? (Y/N) N
```

Before running the `reset-all` command, the `auto-boot?` parameter must be set to "false" to prevent automatic reboot of the system after reset as follows:

```
OK> setenv auto-boot? false
{use "printenv" command to verify}
```

Next, run the `reset-all` command:

```
OK> reset-all
```

Then, run the `probe-scsi-all` command:

```
OK> probe-scsi-all
```

Set the `auto-boot?` parameter back to TRUE if desired:

```
OK> setenv auto-boot? true
```

Resume system operation:

```
OK> go
```


Incorporating Symmetrix

The following sections specify the actions necessary to incorporate Symmetrix volumes into a Solaris system.

To Recognize LUNs

Symmetrix devices are addressed on the SCSI bus using SCSI target IDs and SCSI LUNs. Each target ID can have up to 8 LUNs associated with it. By default, Solaris only searches for target IDs on the SCSI bus. If LUN addressing is required in addition to target ID addressing, you will need to modify the file `/kernel/drv/sd.conf` to have the following definitions for each target ID which needs to support multiple LUNs.

```
name="sd" class="scsi"
  target=0 lun=0;
name="sd" class="scsi"
  target=0 lun=1;
name="sd" class="scsi"
  target=0 lun=2;
name="sd" class="scsi"
  target=0 lun=3;
name="sd" class="scsi"
  target=0 lun=4;
name="sd" class="scsi"
  target=0 lun=5;
name="sd" class="scsi"
  target=0 lun=6;
name="sd" class="scsi"
  target=0 lun=7;
```

Each target ID can have up to 8 LUNs.

This addition to `sd.conf` is necessary when using LUN addressing. Make sure that for target IDs that do not need LUN support, only the `target=x lun=0` line is specified. This decreases the time needed to reboot the system.

To enable Solaris to support the Symmetrix system, EMC provides Symmetrix device definition files. The device definition files are available on EMC's FTP server, <ftp.emc.com>, in `/pub/symm3000/solaris`.

Appendix A, *Symmetrix Device Definition Files* has instructions on how to transfer these files to your host.

System Specification File Modification

There are three parameters in the file, */etc/system*, that require modification when operating in a Solaris environment. These parameters are

1. `scsi_options`
2. `sd_io_time`
3. `sd_max_throttle`

This section describes the necessary modifications to the three parameters.

1. Make sure the mask `scsi_options` has the following values:

- ◆ For Solaris 2.3, set the following in the file:

```
set scsi_options = 0x378
```

- ◆ For hosts configured with the Sbus Fast-Wide Differential SCSI Controller X1062A running Solaris 2.4 or higher, set the following in the file:

```
set scsi_options = 0x3F8
```

- ◆ For hosts configured with the Sbus Ultra-Wide Differential SCSI Controller (X1065A or X6541A) running Solaris 2.5 or higher, set the following in the file:

```
set scsi_options = 0x7F8
```

The bit meaning for this parameter is as follows:

Table 12-3 Bit Settings

Bit	Mask	Meaning
3	0x8	Disconnect enable
4	0x10	Link enable
5	0x20	Sync xfer enable
6	0x40	Parity support enable
7	0x80	Tag commands enable
8	0x100	Fast SCSI enable
9	0x200	Wide SCSI enable
10	0x400	Ultra SCSI enable

Since the **/etc/system** file can also affect other disks in the system, consider the effects of implementing the `scsi_options=0x7f8` mask on the other disks. If this mask is not set, Symmetrix operations will not benefit from these options (e.g., Symmetrix will not use Tag Command Queuing).

With Solaris 2.4 and later, `scsi_options=0x7f8` should be the system default. Solaris 2.3 does not support Tag Command Queuing.

2. Set `sd_io_time` to 120 seconds. This setting will prevent the host from issuing warning messages while non-disruptive operations are performed on Symmetrix. The entry should read:

```
set sd:sd_io_time = 0x78
```

3. Set `sd_max_throttle` to 20. This setting prevents the host from over-sending tag queuing commands which may cause `scsi cmd timeout` and `scsi bus reset`. The edited entry should read:

```
set sd:sd_max_throttle = 20
```

Format, Partition, Label

Use the `format` command to partition and label the new devices. The devices will appear under `/dev/dsk`.

Internal Sun drives on SPARC 5, 10, 20, Ultra1, Ultra2, Ultra30, Ultra60 and Ultra450 are usually found on SCSI controller 0. Symmetrix drives typically start at SCSI controller 1 or SCSI controller 2. This may vary on SPARC 1000 and 2000, Ultra 3X00, 4X00, 5X00, 6X00, and 10000.

If you are running Solaris 2.3, follow the instructions below. If you are running Solaris 2.4 or higher, go to the the next section.

Solaris 2.3

To partition and label new devices:

1. At the system prompt, type `format`
The host searches for all disks. It generates a display placing all unlabeled disks at the beginning of the listing.
2. At the `Specify Disk` prompt, enter the number of the first EMC drive. The Format menu appears.
3. At the `format>` prompt, enter `type`
4. At the `Specify disk type (enter its number):` prompt, enter the disk type number from the list that appears on the screen or choose **0** for **Auto-configuration**.
5. At the `format>` prompt, type `1` (for label).
6. At the `Ready to label disk, continue?` prompt, enter `Y`
7. At the `format>` prompt, enter `disk`.
8. Repeat steps 3 through 7 above for all remaining *unknown* Symmetrix devices.

If you wish to use raw devices, there is no need to create filesystems for each device. Otherwise, create new filesystems for each device as described in *Create and Mount a File System* on page 12-40.

Solaris 2.4 and Later

To partition and label new devices:

1. At the `root` prompt, type `format`
The host searches for all disks. It generates a display placing all unlabeled disks at the beginning of the listing.
2. At the `Specify Disk` prompt, enter the number of the first EMC drive. The Format menu appears.
3. At the `format>` prompt, type `label`
4. At the `Disk not labeled. Label it now?` prompt, type `Y`
5. Type `disk` at the prompt to display a listing of the disks.
6. Repeat steps 2 through 5 for all Symmetrix disks.

If you wish to use raw devices, there is no need to create filesystems for each device. Otherwise, create new filesystems for each device as described in *Create and Mount a File System* on page 12-40.

Adding Devices Online

Whenever devices are added online to the Symmetrix unit or device channel addresses are changed, you will need to perform the actions described below in order to introduce the new devices to the system.

To add new Symmetrix devices while online in the Solaris environment, perform the following steps:

1. Use the **online upgrade** feature at the Symmetrix service processor to add/map new drives to the SCSI host channel.
2. Execute the following two utilities:

```
drvconfig  
disks
```
3. Follow the instructions in *Format, Partition, Label* on page 12-12 to introduce new devices to the host environment.

Solstice DiskSuite

Solstice DiskSuite is a tool for disk and file management. This tool can be used to create and manage logical disks, mirrored or striped volumes, and file systems. DiskSuite supports large file systems, file system expansion and volume manager level intent logging for fast file system recovery.

The following paragraphs provide instructions for using Solstice DiskSuite (SDS) to create concatenated, striped, and mirrored logical volumes for use with either standard UNIX file systems (UFS) or UFS with intent logging.

See *Intent Logging* on page 12-40 for information on the SDS UFS logging feature

For additional information on using DiskSuite in the Solaris environment refer to *Solstice Disk Suite Administrator's Guide*.

For readability purposes, many of the example lines in this section wrap as indicated by "\" at the end of a line. Your entries should be typed on a single line.

Creating a Concatenated Volume

Use this procedure to add a concatenated metadvice (volume). The example below uses four disks: c1t0d0, c1t0d1, c1t0d2, and c1t0d3. The metadvice is configured to be contained in slice s0 on all devices.

To create a concatenated metadvice comprising slice s0 on four disks, perform the following steps:

1. Edit the `/etc/opt/SUNWmd/md.tab` and add the entry below to define the metadvice:

```
/dev/md/dsk/d0 4 1 /dev/dsk/c1t0d0s0 1 /dev/dsk/c1t0d1s0
                  1 /dev/dsk/c1t0d2s0 1 /dev/dsk/c1t0d3s0
```

where:

d0 = name of concatenated metadvice
4 = number of device rows to be concatenated
1 = number of components in each row

2. To create the concatenated metadvice, type

```
metainit d0
```

3. To verify configuration, type

```
metastat
or
metastat d0
```

Concatenated Volumes for UFS Logging

Use this procedure to add a concatenated metadevice (volume) for use with UFS logging. See *Intent Logging* on page 12-40 for a description of the UFS logging feature.

To use UFS logging, you must create a metatrans device. A metatrans device consist of two components (metadevices):

- ◆ Master Metadevice - user data space
- ◆ Logging Metadevice - logging space

You must create one or more slice on one or more disk to be the logging slice. The general guideline of space to allocate for the UFS logging is 10 MB per 1 GB of user space up to a maximum of 64 MB.

The example below uses four disks: c1t0d0, c1t0d1, c1t0d2, and c1t0d3.

- ◆ All disks are on same controller c1
- ◆ Slice s0 is configured to contain the master user data
- ◆ Slice s3 is configured to contain the UFS logging space (16MB/disk)

This example spreads the logging across four disks. You may want to put logging on a single disk to minimize risk of failure when any one of the four volumes fails. Trade-off is a penalty in wasted space on other disks: if disk1 logging is 64 MB, then disk 2, 3, and 4 will each waste 64 MB.

To create a concatenated metadevice for a UFS logging file system comprising slice s0 and s3 on four disks, perform the following steps:

1. Edit the **/etc/opt/SUNWmd/md.tab** and add the entries below to define the metadevices:

```
/dev/md/dsk/d1  4  1 /dev/dsk/c1t0d0s0 1 /dev/dsk/c1t0d1s0 \
                  1 /dev/dsk/c1t0d2s0 1 /dev/dsk/c1t0d3s0
/dev/md/dsk/d2  4  1 /dev/dsk/c1t0d0s3 1 /dev/dsk/c1t0d1s3 \
                  1 /dev/dsk/c1t0d2s3 1 /dev/dsk/c1t0d3s3
/dev/md/dsk/d0 -t /dev/md/dsk/d1 /dev/md/dsk/d2
```

where:

d0 = name of metatrans device consisting of:
d1 = name of master concatenated metadvice
d2 = name of logging concatenated metadvice
4 = number of device rows to be concatenated
1 = number of components in each row
-t = option to define a metatrans device.

2. To create the metatrans device, type

```

metainit d2
(create the logging device)

metainit d1
(create the master device)

metainit d0
(create the metatrans device)
  
```

3. To verify configuration, type

```

# metastat
(or metastat d0)
  
```

Creating a Striped Volume

Use this procedure to add a striped metadvice (volume). The example below uses four disks: c1t0d0 (disk1), c2t0d0 (disk2), c1t0d1 (disk3), and c2t0d1 (disk4).

- ◆ The metadvice is configured to be contained in slice s0 on all devices.
- ◆ Stripe width is 64 KB.
- ◆ Striping is across controllers for performance. Trade-off consideration is that a single controller failure will disable the whole volume.
- ◆ Stripe sequence is disk1, disk2, disk3, disk4, then back to disk1 and so on.

To create a striped metadvice comprising slice s0 on four disks, perform the following steps:

1. Edit the `/etc/opt/SUNWmd/md.tab` and add the entry below to define the striped metadvice:

```
/dev/md/dsk/d0 1 4 /dev/dsk/c1t0d0s0 /dev/dsk/c2t0d0s0 \
/dev/dsk/c1t0d1s0 /dev/dsk/c2t0d1s0 -i 64k
```

Where:

d0 = name of striped metadvice

1 = number of device rows to be striped

4 = number of components in each row

-i 64k = Specify stripe width size (8k, 16k [default], 32k, or 64k)

2. To create the striped metadvice, type

```
metainit d0
```

3. To verify configuration

```
metastat (or metastat d0)
```

Creating a Striped Volume for UFS Logging

Use this procedure to add a striped metadvice (volume) for use with UFS logging. See *Intent Logging* on page 12-40 for a description of the UFS logging feature.

To use UFS logging, you must create a metatrans device. A metatrans device consist of two components (metadevices):

- ◆ Master Metadvice - user data space
- ◆ Logging Metadvice - logging space

You must create one or more slice on one or more disk to be the logging slice. The general guideline of space to allocate for the UFS logging is 10 MB per 1 GB of user space up to a maximum of 64 MB.

Logging metadevices should be placed on disks separate from the data disks to avoid hot spots resulting from constant access to the logging area.

The example below uses four disks: c1t0d0 (disk1), c2t0d0 (disk2), c1t0d1 (disk3), and c2t0d1 (disk4).

- ◆ Disks are on two controllers: c1 and c2
- ◆ Slice s0 is configured to contain the master user data on all devices.
- ◆ Slice s3 is configured to contain UFS logging space (16MB/disk).

- ◆ Stripe width is 64 KB.
- ◆ Striping is across controller for performance. Trade-off consideration is that a single controller failure will disable the whole volume. Stripe sequence is disk1, disk2, disk3, disk4, then back to disk1 and so on.

To create a striped metadvice for UFS logging comprising slices s0 and s3 on four disks, perform the following steps:

1. Edit the **/etc/opt/SUNWmd/md.tab** and add the entries below to define the striped metatrans device:

```
/dev/md/dsk/d1  1  4  /dev/dsk/c1t0d0s0  /dev/dsk/c2t0d0s0 \
                  /dev/dsk/c1t0d1s0  /dev/dsk/c2t0d1s0 -i 64k
/dev/md/dsk/d2  1  4  /dev/dsk/c1t0d0s3  /dev/dsk/c2t0d0s3 \
                  /dev/dsk/c1t0d1s3  /dev/dsk/c2t0d1s3 -i 64k
/dev/md/dsk/d0 -t  /dev/md/dsk/d1  /dev/md/dsk/d2
```

Where:

d0 = name of metatrans device consisting of:

d1 = name of master striped metadvice

d2 = name of logging striped metadvice

1 = number of device rows to be striped

4 = number of components in each row

-i 64k = Specify the stripe width (8k, 16k [default], 32k, 64k)

-t = option to define a metatrans device

2. To create the striped metatrans device, type

```
metainit d2
(create the striped logging device)
```

```
metainit d1
(create the striped master device)
```

```
metainit d0
(create the striped metatrans device)
```

3. To verify configuration, type

```
metastat
(or metastat d0)
```

Creating a Mirrored Striped Volume

Use this procedure to add a mirrored and striped metadvice (volume). The example below uses eight disks:

Submirror 1 - c1t0d0, c1t0d1, c1t0d2, and c1t0d3

Submirror 2 - c2t0d0, c2t0d1, c2t0d2, and c2t0d3

- ◆ Slice s0 is configured to contain data on all devices.
- ◆ Stripe width is 64 KB.
- ◆ Metamirror consists of two submirrors.
Each submirror comprises four disks on the same controller: c1 for submirror 1 and c2 for submirror 2
- ◆ Metamirror is created as a 1-way mirror, then a second mirror is added using the **metattach** command
- ◆ Stripe sequence is disk1, disk2, disk3, disk4, then back to disk1 and so on.

To create a mirrored and striped metamirror device comprising slice s0 on eight disks, perform the following steps:

1. Edit the **/etc/opt/SUNWmd/md.tab** and add the entries below to define the metamirror:

```
/dev/md/dsk/d1  1  4 /dev/dsk/c1t0d0s0 /dev/dsk/c1t0d1s0 \
                  /dev/dsk/c1t0d2s0 /dev/dsk/c1t0d3s0 -i 64k
/dev/md/dsk/d2  1  4 /dev/dsk/c2t0d0s0 /dev/dsk/c2t0d1s0 \
                  /dev/dsk/c2t0d2s0 /dev/dsk/c2t0d3s0 -i 64k
/dev/md/dsk/d0 -m /dev/md/dsk/d1 1
```

Where:

- d0** = name of metamirror consisting of:
- d1** = name of striped metadvice (submirror 1)
- d2** = name of striped metadvice (submirror 2)
- 1** = number of device rows to be striped
- 4** = number of components in each row.
- i 64k** = Specify the stripe width (8k, 16k [default], 32k, 64k)
- m** = option to define a metamirror device.
- 1** = Define as a 1-way mirror.

2. To create the striped metatrans device, type

```
metainit d1
(create the striped (submirror 1) metadvice)
```

```
metainit d2
(create the striped (submirror 2) metadvice)
```

```
metainit d0
(create the 1-way Metamirror)
```

3. To add the second submirror to the metamirror, type

```
metattach /dev/md/dsk/d0 /dev/md/dsk/d2
```

4. To verify configuration, type

```
metastat
```

```
(or metastat d0)
```

Mirrored Striped Volume for UFS Logging

Use this procedure to add a mirrored and striped metadvice (volume) for use with UFS logging. Refer to *Intent Logging* on page 40 for a description of the UFS logging feature.

To use UFS logging, you must create a metatrans device. A metatrans device consist of two components (metadevices):

- ◆ Master Metadvice - user data space
- ◆ Logging Metadvice - logging space

You must create one or more slice on one or more disk to be the logging slice. The general guideline of space to allocate for the UFS logging is 10 MB per 1 GB of user space up to a maximum of 64 MB.

Logging metadevices should be placed on disks separate from the data disks to avoid hot spots resulting from constant access to the logging area.

The example below uses eight disks:

Submirror 1 - c1t0d0 (disk1), c1t0d1 (disk2), c1t0d2 (disk3), and c1t0d3 (disk4)

Submirror 2 - c2t0d0 (disk1), c2t0d1 (disk2), c2t0d2 (disk3), and c2t0d3 (disk4)

- The metadvice is configured to contain master user data on slice s0 on all devices.
- Slice s3 is configured to contain the UFS logging Space (16MB/disk)
- Stripe width is 64 KB

- Metamirror consists of two submirrors
- Each submirror comprises four disks on two controllers: c1 for submirror 1 and c2 for submirror 2
- Metamirror is created as a 1-way mirror, then a second mirror is added using the `metattach` command
- Stripe sequence is disk1, disk2, disk3, disk4, then back to disk1 and so on.

To create a mirrored and striped metamirror device for use with UFS logging, and comprising slice s0 and s3 on eight disks, perform the following steps:

1. Edit the `/etc/opt/SUNWmd/md.tab` and add the entries below to define the metamirror:

```
/dev/md/dsk/d1  1  4 /dev/dsk/c1t0d0s0 /dev/dsk/c1t0d1s0 \
                  /dev/dsk/c1t0d2s0 /dev/dsk/c1t0d3s0 -i 64k
/dev/md/dsk/d2  1  4 /dev/dsk/c2t0d0s0 /dev/dsk/c2t0d1s0 \
                  /dev/dsk/c2t0d2s0 /dev/dsk/c2t0d3s0 -i 64k
/dev/md/dsk/d3  1  4 /dev/dsk/c1t0d0s3 /dev/dsk/c1t0d1s3 \
                  /dev/dsk/c1t0d2s3 /dev/dsk/c1t0d3s3 -i 64k
/dev/md/dsk/d4  1  4 /dev/dsk/c2t0d0s3 /dev/dsk/c2t0d1s3
                  /dev/dsk/c2t0d2s3 /dev/dsk/c2t0d3s3 -i 64k
/dev/md/dsk/d0 -m /dev/md/dsk/d1 1
/dev/md/dsk/d5 -m /dev/md/dsk/d3 1
/dev/md/dsk/d6 -t /dev/md/dsk/d0 /dev/md/dsk/d5
```

where:

d6 = name of mirrored metatrans device consisting of:

d1 = name of striped master metadvice (master submirror 1)

d2 = name of striped master metadvice (master submirror 2)

d3 = name of striped UFS log metadvice (log submirror 1)

d4 = name of striped UFS log metadvice (log submirror 2)

d0 = name of master metamirror

d5 = name of log metamirror

1 = number of device rows to be striped

4 = number of components in each row

-i 64k = Specify the stripe width (8k, 16k [default], 32k, 64k)

-m = option to define a metamirror device

-t = option to define a metatrans device

2. To create the mirrored metatrans device, type

```
metainit d1
```

(Create the striped (submirror 1) master)

```
metainit d2
```

(Create the striped (submirror 2) master)

```
metainit d3
```

(Create the striped (submirror 1) logging)

```
metainit d4
```

(Create the striped (submirror 2) logging)

```
metainit d0
```

(Create the 1-way master metamirror d0)

```
metainit d5
```

(Create the 1-way logging metamirror d5)

```
metattach d0 d2
```

(Attach second submirror d2 to master metamirror d0)

```
metattach d5 d4
```

(Attach second submirror d4 to log metamirror d5)

```
metainit d6
```

(Create the metatrans device d6 consisting of master device d0 and logging device d5)

3. To verify configuration, type

```
metastat
```

(or `metastat d0`)

DiskSuite State Database Replicas

The *state database* stores all configuration and status information for DiskSuite objects. Without the state database DiskSuite is unable to access any devices and all data could be lost. Replicas of the database are created so that DiskSuite can compare copies to verify the current configuration and running state of all objects.

It is recommended that at least three replicas be created. If one replica is not available, the remaining two can still be compared to verify configuration and state information. Three replicas can be stored on a system boot disk, however, this creates a single point of failure at the boot disk. Additional replicas should be created on other system disks including Symmetrix devices to protect against the loss of the boot disk. The following considerations apply when planning locations for state database replicas.

- ◆ Always create at least three replicas. DiskSuite will not function if it has only one state database.
- ◆ Replicas can be stored on any unused partition or on partitions that are also part of a metadevice or logging device with the exceptions of **root**, **swap**, **/usr** or file system.
- ◆ Replicas should be spread evenly across host controllers. In Symmetrix, replica devices should be behind separate even and odd SCSI adapters.
- ◆ Store replicas on least two disks on each controller.
- ◆ Each disk should contain only one replica.

Useful DiskSuite Commands

This section lists DiskSuite functions and utilities you may use to define and manage Symmetrix devices. The use of these functions and utilities is optional. They are listed for reference only.

- ◆ `metaclear devicename`
Delete a metadevice. When removing a metadevice hierarchy, use **metaclear** in top down sequence, for example


```
metaclear meta trans
metaclear master_mirror
metaclear log_mirror
metaclear submirrora1
metaclear submirrora2
metaclear submirrorb1
metaclear submirrorb2
```
- ◆ `metadb -i`
Check status of metastate database replicas.
- ◆ `metareplace`
Fix a defective mirrored or RAID device.
- ◆ `metastat`
Display status of mirrored or RAID device.
- ◆ `metahs`
Command to add a hot spare.

Veritas Volume Manager

Veritas VxVM and VxFS are tools for disk and file management. VxVM can be used to create logical disks, mirrored and striped volumes. VxFS supports large file systems, file system expansion and a journaling file system.

The following paragraphs provide instructions for using VxVM to create concatenated, striped, and mirrored logical volumes for use with either standard UNIX file systems (UFS) or VxFS journaled file systems.

See *Intent Logging* on page 12-40 for a description of journaling file systems.

For additional information on using Veritas in the Solaris environment refer to *Veritas Volume Manager (VxVM) System Administrator's Guide*.

Initializing Disks

All disks must be initialized before they can be used by VxVM. The initialization process stores configuration and identification data on the disk for use by VxVM.

Disks must be formatted before being initialized by VxVM. Formatting is done outside the volume manager using the **format** command. See *Format, Partition, Label* on page 12-12.

Use the **vxdiskadm** menu-driven interface to initialize disks before attempting to create VxVM volumes. To start **vxdiskadm**, log in as **root** and type

```
vxdiskadm
```

Select **1 Add or Initialize one or more disk**, and follow the system prompts to select devices to be initialized.

Creating a Concatenated Volume

Use this procedure to add a 2 GB concatenated volume of four disks. The configuration uses 500 MB on each 2 GB device.

There are two methods for creating a concatenated volume presented below. The manual **vxmake** method with more control over volume subdisk and plex creation, and the more convenient **vxassist** method with less control (automatic creation and allocation of volume components).

The disk must be initialized before attempting to create volumes (refer to *Initializing Disks* on page 12-25). The initialization process creates slice 02 on the disk. When VxVM puts the disk under its control it must first detect slice 02. VxVM can then create slice 03 as a private region and slice 04 as a public region.

Creating a Concatenated Volume using vxmake

Use the manual **vxmake** method to create a concatenated volume as follows:

1. To verify that the disk are online and not being used, type

```
vxdisk list
```

Check the displayed list for available disks. For example:

DEVICE	TYPE	DISK	GROUP	STATUS
c0t1d0s2	sliced	-	-	online
c0t3d0s2	sliced	-	-	error
c1t0d0s2	sliced	-	-	error
c1t0d1s2	sliced	-	-	error
c1t0d2s2	sliced	-	-	online
c1t0d3s2	sliced	-	-	online
c1t0d4s2	sliced	-	-	online
c1t0d5s2	sliced	-	-	online
c1t0d6s2	sliced	-	-	online
c1t0d7s2	sliced	-	-	online
* c1t1d0s2	sliced	-	-	online
* c1t1d1s2	sliced	-	-	online
* c1t1d2s2	sliced	-	-	online
* c1t1d3s2	sliced	-	-	online
c1t1d4s2	sliced	-	-	online
c1t1d5s2	sliced	-	-	online
c1t1d6s2	sliced	-	-	online
c1t1d7s2	sliced	-	-	online

* Disks used in this example

2. To create disk group named **cust** and add four disks to disk group **cust**, type

```
vxvg init cust cust-00=c1t1d0
vxvg -g cust adddisk cust-01=c1t1d1
vxvg -g cust adddisk cust-02=c1t1d2
vxvg -g cust adddisk cust-03=c1t1d3
```

3. To verify subdisk creation, type

```
vxdisk list
```

Check the displayed list for subdisk and group. For example:

DEVICE	TYPE	DISK	GROUP	STATUS
		.		
		.		
		.		
c1t1d0s2	sliced	cust-00	cust	online
c1t1d1s2	sliced	cust-01	cust	online
c1t1d2s2	sliced	cust-02	cust	online
c1t1d3s2	sliced	cust-03	cust	online

4. To create 500 MB subdisks on each disks, type

```
vxmake -g cust sd cust-sd0 cust-00,0,500m
vxmake -g cust sd cust-sd1 cust-01,0,500m
vxmake -g cust sd cust-sd2 cust-02,0,500m
vxmake -g cust sd cust-sd3 cust-03,0,500m
```

5. To verify configuration database information, type

```
vxprint -st
```

Check display for correct configuration data.

For example:

Disk group: cust

SD NAME	PLEX	DISK	DISKOFFS	LENGTH	[COL/]OFF	DEVICE	MODE
sd cust-sd0	-	cust-00	0	1024000	-	c1t1d0	ENA
sd cust-sd1	-	cust-01	0	1024000	-	c1t1d1	ENA
sd cust-sd2	-	cust-02	0	1024000	-	c1t1d2	ENA
sd cust-sd3	-	cust-03	0	1024000	-	c1t1d3	ENA

6. To create striped plex named **cust-plex** out of the subdisks, type

```
vxmake -g cust plex cust-plex sd=cust-sd0,cust-sd1, \ cust-sd2,cust-sd3
layout=concat stwidth=64k
```

7. To verify configuration database information, type

```
vxprint -pt
```

Check display for correct configuration data.

For example:

Disk group: cust

PL NAME	VOLUME	KSTATE	STATE	LENGTH	LAYOUT	NCOL/WID	MODE
pl cust-plex	cust-vol	ENABLED	ACTIVE	4096000	STRIPE	4/128	

8. To create volume out of the concatenated plex **cust-plex**, type

```
vxmake -g cust -U fsgen vol cust-vol \ plex=cust-plex
```

9. To verify configuration database information, type

```
vxprint -vt
```

Check display for correct configuration data.

For example:

Disk group: cust

V NAME	USETYPE	KSTATE	STATE	LENGTH	READPOL	PREFPLEX
v cust-vol	fsgen	DISABLED	EMPTY	4096000	ROUND	-

10. To initialize the volume to the CLEAN state, type

```
vxvol -g cust init clean cust-vol
```

11. To start the concatenated volume, type

```
vxvol -g cust start cust-vol
```

12. To verify configuration database information, type

```
vxprint -vt
```

Check display for correct configuration data.

For example:

Disk group: cust

V NAME	USETYPE	KSTATE	STATE	LENGTH	READPOL	PREFPLEX
v cust-vol	fsgen	ENABLED	ACTIVE	4096000	ROUND	-

Creating a Concatenated Volume using vxassist

Creating volumes with **vxassist** is the recommended way of creating volumes since the user does not deal with the complicated commands as seen in the manual **vxmake** method. This example creates 2 GB concatenated volumes **cust-vol1** and **cust-vol2** on disk group **cust**. **vxassist** creates the subdisks and plexes, and initializes and starts the volume automatically.

1. To create disk group named **cust** and add four disks to disk group **cust**, type

```
vxdbg init cust cust-00=c1t1d0
vxdbg -g cust adddisk cust-01=c1t1d1
vxdbg -g cust adddisk cust-02=c1t1d2
vxdbg -g cust adddisk cust-03=c1t1d3
```

2. To create concatenated volume named **cust-vol1** with 2.0 GB size using disk in disk group **cust**, type

```
vxassist -g cust -fsgen make cust-vol1 2g \ layout=span
```

3. To verify configuration database information, type

```
vxprint -vt cust-vol1
```

Check display for correct configuration data.

For example:

Disk group: cust

V NAME	USETYPE	KSTATE	STATE	LENGTH	READPOL	PREFPLEX
v cust-vol	fsgen	ENABLED	ACTIVE	4194394	SELECT	-

Creating a Striped Volume

Use this procedure to add a 2 GB striped volume of four disks. The configuration uses 500 MB on each 2 GB device.

There are two methods for creating a striped volume presented below. The manual **vxmake** method with more control over volume subdisk and plex creation, and the more convenient **vxassist** method with less control (automatic creation and allocation of volume components).

Creating a Striped Volume with vxmake

Use the manual **vxmake** method to create a striped volume as follows:

1. To verify that the disk are online and not being used, type

```
vxdisk list
```

Check the displayed list for available disks. For example:

	DEVICE	TYPE	DISK	GROUP	
	c0t1d0s2	sliced	-	-	online
	c0t3d0s2	sliced	-	-	error
*	c1t0d0s2	sliced	-	-	online
*	c1t0d1s2	sliced	-	-	online
	c1t0d2s2	sliced	-	-	online
	c1t0d3s2	sliced	-	-	online
	c1t0d4s2	sliced	-	-	online
	c1t0d5s2	sliced	-	-	online
	c1t0d6s2	sliced	-	-	online
	c1t0d7s2	sliced	-	-	online
*	c2t0d0s2	sliced	-	-	online
*	c2t0d1s2	sliced	-	-	online
	c2t0d2s2	sliced	-	-	online
	c2t0d3s2	sliced	-	-	online
	c2t0d4s2	sliced	-	-	online
	c2t0d5s2	sliced	-	-	online
	c2t0d6s2	sliced	-	-	online
	c2t0d7s2	sliced	-	-	online

* Disks used in this example

2. To create disk group named **cust** and add four disks to disk group **cust**, type

```
vxvg init cust cust-00=c1t0d0
vxvg -g cust adddisk cust-01=c2t0d0
vxvg -g cust adddisk cust-02=c1t0d1
vxvg -g cust adddisk cust-03=c2t0d1
```

3. To verify subdisk creation, type

```
vxdisk list
```

Check the displayed list for subdisk and group. For example:

DEVICE	TYPE	DISK	GROUP	STATUS
			.	
			.	
			.	
c1t0d0s2	sliced	cust	-00	customonline
c2t0d0s2	sliced	cust	-01	customonline
c1t0d1s2	sliced	cust	-02	customonline
c2t0d1s2	sliced	cust	-03	customonline

4. To create 500 MB subdisks on each disks, type

```
vxmake -g cust sd cust-sd0 cust-00,0,500m
vxmake -g cust sd cust-sd1 cust-01,0,500m
vxmake -g cust sd cust-sd2 cust-02,0,500m
vxmake -g cust sd cust-sd3 cust-03,0,500m
```

5. To verify configuration database information, type

```
vxprint -st
```

Check display for correct configuration data.

For example:

Disk group: cust

SD	NAMEPLEX	DISK	DISKOFFS	LENGTH	[COL/] OFF	DEVICE	MODE
sd	cust-sd0-cust-000			1024000	-	c1t0d0	ENA
sd	cust-sd1-cust-010			1024000	-	c2t0d0	ENA
sd	cust-sd2-cust-020			1024000	-	c1t0d1	ENA
sd	cust-sd3-cust-030			1024000	-	c2t0d1	ENA

6. To create striped plex named **cust-plex** out of the subdisks, type

```
vmake -g cust plex cust-plex sd=cust-sd0,cust-sd1,cust-sd2,cust-sd3 \
layout=stripe stwidth=64k
```

Stripe width values are 1Kbytes to 256 Kbytes. The default value is 64 Kbytes.

7. To verify configuration database information, type

```
vxprint -pt
```

Check display for correct configuration data.

For example:

Disk group: cust

PL NAME	VOLUME	KSTATE	STATE	LENGTH	LAYOUT	NCOL/WID	MODE
pl cust-plex	cust-vol	ENABLED	ACTIVE	4096000	STRIPE	4/128	

8. To create volume **cust-vol** out of the plex **cust-plex**, type

```
vxmake -g cust -Ufsgen vol cust-vol plex=cust-plex
```

9. To verify configuration database information, type

```
vxprint -vt
```

Check display for correct configuration data.

For example:

Disk group: cust

V NAME	USETYPE	KSTATE	STATE	LENGTH	READPOL	PREFPLEX
v cust-vol	fsgen	DISABLED	EMPTY	4096000	ROUND	-

10. To initialize the volume to the CLEAN state, type

```
vxvol -g cust init clean cust-vol
```

11. To start the concatenated volume, type

```
vxvol -g cust start cust-vol
```

12. To verify configuration database information, type

```
vxprint -vt
```

Check display for correct configuration data.

For example:

Disk group: cust

V NAME	USETYPE	KSTATE	STATE	LENGTH	READPOL	PREFPLEX
v cust-vol	fsgen	ENABLED	ACTIVE	4096000	ROUND	-

Creating a Striped Volume with vxassist

Creating volumes with **vxassist** is the recommended way of creating volumes since the user does not deal with the complicated commands as seen in the manual **vxmake** method. This example, creates a 2 GB striped volume named **cust-vol1**. **vxassist** creates the subdisk and plex, and initializes and starts the volume automatically.

1. To create disk group named **cust** and add four disks to disk group **cust**, type


```
vxdbg init cust cust-00=c1t0d0
vxdbg -g cust adddisk cust-01=c2t0d0
vxdbg -g cust adddisk cust-02=c1t0d1
vxdbg -g cust adddisk cust-03=c2t0d1
```
2. To create striped volume named **cust-vol1** with 2.0 GB size using disks in disk group **cust**, type

```
vxassist -g cust -Ufsgen make cust-vol1 2g layout=stripe stwidth=64k \ cust-00
cust-01 cust-02 cust-03
```

Stripe width values are 1Kbytes to 256 Kbytes. The default value is 64 Kbytes.

3. To verify configuration database information, type

```
vxprint -vt cust-vol1
```

Check display for correct configuration data. For example:

Disk group: cust

V NAME	USETYPE	KSTATE	STATE	LENGTH	READPOL	PREFPLEX
v cust-vol1	fsgen	ENABLED	ACTIVE	4194394	ROUND	-

Creating a Striped Mirrored Volume

Use this procedure to add a 2 GB striped volume of four disks with a mirror on another four disks. The configuration uses 500 MB on each 2 GB device.

There are two methods for creating a striped volume presented below. The manual **vxmake** method with more control over volume subdisk and plex creation, and the more convenient **vxassist** method with less control (automatic creation and allocation of volume components).

Manual **vxmake** Method

Use the manual **vxmake** method to create a concatenated volume as follows:

1. To verify that the disk are online and not being used, type

```
vxdisk list
```

Check the displayed list for available disks.

For example:

DEVICE	TYPE	DISK	GROUP	STATUS
c0t1d0s2	sliced	-	-	online
c0t3d0s2	sliced	-	-	error
c1t0d0s2	sliced	-	-	error
c1t0d1s2	sliced	-	-	error
c1t0d2s2	sliced	-	-	online
c1t0d3s2	sliced	-	-	online
c1t0d4s2	sliced	-	-	online
c1t0d5s2	sliced	-	-	online
c1t0d6s2	sliced	-	-	online
c1t0d7s2	sliced	-	-	online
* c1t1d0s2	sliced	-	-	online
* c1t1d1s2	sliced	-	-	online
* c1t1d2s2	sliced	-	-	online
* c1t1d3s2	sliced	-	-	online
c1t1d4s2	sliced	-	-	online
c1t1d5s2	sliced	-	-	online
c1t1d6s2	sliced	-	-	online
c1t1d7s2	sliced	-	-	online
* c2t1d0s2	sliced	-	-	online
* c2t1d1s2	sliced	-	-	online
* c2t1d2s2	sliced	-	-	online
* c2t1d3s2	sliced	-	-	online

* Disks used in this example

2. To create Disk group named **cust** and add eight disks to disk group **cust**, type

```
vx dg init cust cust-00=c1t1d0
vx dg -g cust adddisk cust-01=c1t1d1
vx dg -g cust adddisk cust-02=c1t1d2
vx dg -g cust adddisk cust-03=c1t1d3
vx dg -g cust adddisk cust-04=c2t1d0
vx dg -g cust adddisk cust-05=c2t1d1
vx dg -g cust adddisk cust-06=c2t1d2
vx dg -g cust adddisk cust-07=c2t1d3
```

3. To verify subdisk creation, type

```
vxdisk list
```

Check the displayed list for subdisk and group. For example:

DEVICE	TYPE	DISK	GROUP	STATUS
		.		
		.		
c1t1d0s2	sliced	cust-00	cust	online
c1t1d1s2	sliced	cust-01	cust	online
c1t1d2s2	sliced	cust-02	cust	online
c1t1d3s2	sliced	cust-03	cust	online
c1t1d4s2	sliced	-	-	online
c2t1d0s2	sliced	cust-04	cust	online
c2t1d1s2	sliced	cust-05	cust	online
c2t1d2s2	sliced	cust-06	cust	online
c2t1d3s2	sliced	cust-07	cust	online

4. To create 500 MB subdisks on each disks, type

```
vxmake -g cust sd cust-sd0 cust-00,0,500m
vxmake -g cust sd cust-sd1 cust-01,0,500m
vxmake -g cust sd cust-sd2 cust-02,0,500m
vxmake -g cust sd cust-sd3 cust-03,0,500m
vxmake -g cust sd cust-sd4 cust-04,0,500m
vxmake -g cust sd cust-sd5 cust-05,0,500m
vxmake -g cust sd cust-sd6 cust-06,0,500m
vxmake -g cust sd cust-sd7 cust-07,0,500m
```

5. To verify configuration database information, type

```
vxprint -st
```

Check display for correct configuration data.

For example:

Disk group: cust

SD NAME	PLEX	DISK	DISKOFFS	LENGTH	[COL/]OFF	DEVICE	MODE
sd cust-sd0	-	cust-00	0	1024000	-	c1t1d0	ENA
sd cust-sd1	-	cust-01	0	1024000	-	c1t1d1	ENA
sd cust-sd2	-	cust-02	0	1024000	-	c1t1d2	ENA
sd cust-sd3	-	cust-03	0	1024000	-	c1t1d3	ENA
sd cust-sd4	-	cust-00	0	1024000	-	c2t1d0	ENA
sd cust-sd5	-	cust-01	0	1024000	-	c2t1d1	ENA
sd cust-sd6	-	cust-02	0	1024000	-	c2t1d2	ENA
sd cust-sd7	-	cust-03	0	1024000	-	c2t1d3	ENA

- To create striped plexes **cust-plex1** and **cust-plex2** out of the subdisks, type

```

cust-plex1 make up submirror 1, cust-plex2 make up submirror 2

vxmake -g cust plex cust-plex1 sd=cust-sd0,cust-sd1,cust-sd2,cust-sd3\
  layout=stripe stwidth=64k
vxmake -g cust plex cust-plex2 sd=cust-sd4,cust-sd5,cust-sd6,cust-sd7\
  layout=stripe stwidth=64k
  
```

Stripe width values are 1Kbytes to 256 Kbytes. The default value is 64 Kbytes.

- To verify configuration database information, type

```
vxprint -pt
```

Check display for correct configuration data.

For example:

Disk group: cust

PL NAME	VOLUME	KSTATE	STATE	LENGTH	LAYOUT	NCOL/WID	MODE
pl	cust-plex1-vol	DISABLED	-	4096000	STRIPE	4/128RW	

- To create mirror volume out of the **cust-plex1** and then attach second **cust-plex2**

```

vxmake -g cust -U fsgen vol cust-mirvol plex=cust-plex1
vxplex -g cust att cust-mirvol cust-plex2
  
```

- To verify configuration database information, type

```
vxprint -vt
```

Check display for correct configuration data.

For example:

Disk group: cust

V NAME	USETYPE	KSTATE	STATE	LENGTH	READPOL	PREFPLEX
v	cust-mirvol	fsgen	DISABLED	EMPTY	4096000	ROUND -

10. To initialize the volume to the CLEAN state, type

```
vxvol -g cust init clean cust-mirvol
```

11. To start the concatenated volume, type

```
vxvol -g cust start cust-mirvol
```

12. To verify configuration database information, type

```
vxprint -vt
```

Check display for correct configuration data.

For example:

Disk group: cust

V NAME	USETYPE	KSTATE	STATE	LENGTH	READPOL	PREFFLEX
v cust-mirvol	fsgen	ENABLED	ACTIVE	4096000	ROUND	-

13. To add two copies of dirty region logging (DRL) to the mirror volume, type

```
vxassist -g cust addlog cust-mirvol
(Add first log)
```

```
vxassist -g cust addlog cust-mirvol
(Add second log)
```

Dirty region logging (DRL is a bit map mechanism used to facilitate mirror resynchronization after a system failure. Vxassist is the easy way to add a DRL log to the volume. Vxassist creates a subdisk in the diskgroup; its size based on the volume size (1 KB for every 2 GB). The log for this volume may be 4 sectors or 2 Kbytes (1 KB for each 2 GB mirror). When possible, the logging subdisk is created on a disk separate from the data volume. Vxassist creates a plex out of the subdisk and attaches the log plex to the mirrored volume.

Creating Mirrored Volume with vxassist

Creating volumes with **vxassist** is the recommended way of creating volumes since the user does not deal with the complicated commands as seen in the manual **vxmake** method. This example creates a mirrored volume named **cust-mirvol1** with four disks on controller c1 and four disks on controller c2. **vxassist** creates the subdisk and plexes, and initializes and starts the volume automatically.

1. To create disk group named **cust** and add eight disks to disk group **cust**, type

```
vxdbg init cust cust-00=c1t1d0
vxdbg -g cust adddisk cust-01=c1t1d1
vxdbg -g cust adddisk cust-02=c1t1d2
vxdbg -g cust adddisk cust-03=c1t1d3
vxdbg -g cust adddisk cust-04=c2t1d0
vxdbg -g cust adddisk cust-05=c2t1d1
vxdbg -g cust adddisk cust-06=c2t1d2
vxdbg -g cust adddisk cust-07=c2t1d3
```

2. To create mirrored volume named **cust-mirvol** with 2.0 GB size using disks **cust-00**, **cust-01**, **cust-02** and **cust-03** (controller c1) in disk group **cust**, type

```
vxassist -g cust -Ufsgen make cust-mirvol1
          2g layout=stripe stwidth=64k \
          cust-00 cust-01 cust-02 cust-03
```

Stripe width values are 1Kbytes to 256 Kbytes. The default value is 64 Kbytes.

3. To verify configuration database information, type

```
vxprint -vt cust-mirvol
```

Check display for correct configuration data.
For example:

Disk group: cust

V NAME	USETYPE	KSTATE	STATE	LENGTH	READPOL	PREFPLEX
v cust-mirvol	fsgen	ENABLED	ACTIVE	4194394	ROUND	-

4. To create submirror for mirrored volume **cust-mirvol** of 2.0 GB size with disks **cust-04**, **cust-05**, **cust-06** and **cust-07** on second controller c2, type

```
vxassist -g cust mirror cust-mirvol layout=stripe cust-04 cust-05\
          cust-06 cust-07 &
```

The mirror volume automatically synchronizes the second submirror. This can take a while.

5. To add two copies of dirty region logging (DRL) to the mirror volume, type

```
vxassist -g cust addlog cust-mirvol1  
(Add first log)
```

```
vxassist -g cust addlog cust-mirvol1  
(Add second log)
```

Dirty region logging (DRL) is a bit map mechanism used to facilitate mirror resynchronization after a system failure. **Vxassist** is the easy way to add a DRL log to the volume. **Vxassist** creates a subdisk in the diskgroup; its size based on the volume size (1 KB for every 2 GB). The log for this volume may be 4 sectors or 2 Kbytes (1 KB for each 2 GB mirror). When possible, the logging subdisk is created on a disk separate from the data volume. **Vxassist** creates a plex out of the subdisk and attaches the log plex to the mirrored volume.

Create and Mount a File System

Volumes created and managed by Solstice DiskSuite or VxVM volume manager may be used as raw devices, with the standard UNIX file system, or with the VxFS journaled file system.

Intent Logging

Intent logging records pending changes to the file system structure in an intent log. The intent log is replayed during system failure recovery to complete or abandon changes to the structure that were pending at the time of system failure. The file system can then be mounted without completing a full structural check (fsck). An intent logging system can significantly reduce recovery time following a system failure.

Intent logging may be preformed at the file system level or at the volume manager level. Intent logging at the file system level, generally known as a journaling file system (JFS), is usually more effective than volume manager intent logging.

VxFS and VxVM Intent Logging

The Veritas File System (VxFS) includes a journaling file system that provides intent logging at the file system level. Pending changes to the file system, written to an intent log, are scanned during recovery from a system failure. Changes that were active at the time of failure are completed, and the file system is mounted without requiring a fsck of the entire file system. File system recovery is done in a few seconds; much faster than a standard UFS recovery that requires a complete fsck.

VxVM Dirty Region Logging (DRL) provides intent logging at the volume manager level to reduce the time required to resynchronize mirrored volumes after a system failure. DRL is applied to VxVM mirrored volumes only. Striped or concatenated volumes do not use intent logging, but may rely on VxFS for fast recovery.

SDS Intent Logging

Solstice DiskSuite (SDS) uses a standard UNIX file system (UFS) that does not provide intent logging at the file system level. However, the VxFS Veritas file system may be used in the SDS environment to provide a journaling file system for SDS volumes.

At the volume manger level, SDS uses a method called the *UFS logging feature* to provide intent logging for all volume types (striped, concatenated, mirrored, and RAID). UFS logging is not a journaling file system. It uses the standard UFS and does intent logging at the volume manager level.

Procedures for creating volumes with UFS logging are provided in *Solstice DiskSuite* on page 12-15.

Procedures for creating mirrored volumes with DRL are provided in *Veritas Volume Manager* on page 12-25.

The following paragraphs outline the steps required to add, create and mount a standard UNIX file system or a VxFS journaled file system for volumes and raw devices.



CAUTION

You may wish to place the Symmetrix devices in the mount table. This requires `/etc/vfstab` to be edited. `/etc/vfstab` is syntax-sensitive, and if not edited properly may prevent the system from booting.

UFS File System on Raw Device

Create New File System

To create standard UNIX file systems under Solaris OS, log in as **root** and proceed as follows for each new device.

Once you have formatted, partitioned, and labeled each Symmetrix disk device, create a new file system for each Symmetrix disk. To do this, use the **newfs** command in a statement similar to the following:

```
newfs -v /dev/rdisk/c1t0d0s0
```

The **Construct new filesystem?** prompt appears. Type **y** at the prompt.

The actions above created a new file system for the Symmetrix disk connected to SCSI controller 1, target ID 0, lun 0, partition 0.

Create Mount Directory

Once the file systems for each Symmetrix device are created, create a mount directory for each device. To do this, type a statement similar to the following for each Symmetrix device:

```
mkdir /fs/c1t0d0s0
```

where **/fs/c1t0d0s0** is the complete path for the new file system directory.

The **/fs** directory must exist prior to creating the mount directories.

Mount the Filesystem

To mount each filesystem, type a statement similar to the following:

```
mount /dev/dsk/c1t0d0s0 /fs/c1t0d0s0
```

VxFS File System on Raw Device

A VxFS journaling file system is created using the **mkfs** command with arguments provided for block size, log size, device name and size. To create a VxFS file system first determine the size in sectors of the volume. The size of the volume is displayed under the **Sector Count** field of the **prtvtoc** output.

Once you have formatted, partitioned, and labeled each Symmetrix disk device, create a new file system for each Symmetrix disk. To create VxVM journaling file systems under VxFS and Solaris OS, log in as **root** and proceed as follows for each new device.

Create New File System

1. To display sector count information, type

```
prtvtoc /dev/rdisk/clt0d0s0
```

The size of the disk is displayed under the **Sector Count** field.
(Assume 4099000 for this example.)

2. To create the VxFS file system for the volume, type

```
mkfs -F vxfs -o bsize=4096,logsize=256 \ /dev/rdisk/clt0d0s0 4099000
```

Where:

bsize = block size in bytes (1k, 2k, 4k, or 8k - 1k default for filesystems < 4 GB, 4k default for filesystems > 4 GB)

logsize = size of VxFS file system logging in blocks (256 blocks default, 32 to 1024 blocks)

4099000 = file system size in sectors (from **prtvtoc** command)

Create Mount Directory

Once the file systems for each Symmetrix device are created, create a mount directory for each device. To do this, type a statement similar to the following for each Symmetrix device:

```
mkdir /fs/clt0d0s0
```

Where **/fs/clt0d0s0** is the complete path for the new file system directory.

Mount the Filesystem

To mount each filesystem, type a statement similar to the following:

```
mount -F vxfs /dev/dsk/clt0d0s0 /fs/clt0d0s0
```

UFS File System on SDS Device

Create New File System

To create standard UNIX file systems under Solstice DiskSuite, log in as **root** and proceed as follows for each new device.

Once you have formatted, partitioned, and labeled each Symmetrix disk device, and created SDS volumes, create a new file system on each volume. To do this, use the **newfs** command in a statement similar to the following:

```
newfs -v /dev/md/rdisk/d0
```

The **Construct new filesystem?** prompt appears. Type **y** at the prompt.

The actions above created a new file system for the Symmetrix disk defined as SDS metadvice **d0**.

Create Mount Directory

Once the file systems for each device are created, create a mount directory for each device. To do this, type a statement similar to the following for each device:

```
mkdir /fs/d0
```

Where **/fs/d0** is the complete path for the new file system directory.

You can assume that the **/fs** directory existed prior to creating the **d0** directory.

Mount the Filesystem

To mount each filesystem, type a statement similar to the following:

```
mount /dev/md/dsk/d0 /fs/d0
```

VxFS File System on SDS Device

A VxFS journaling file system is created using the **mkfs** command with arguments provided for block size, log size, device name and size. To create a VxFS file system first determine the size in sectors of the volume. The size of the volume is displayed under the **Sector Count** field of the **prtvtoc** output.

To create VxFS journaling file systems under Solstice DiskSuite, log in as **root** and proceed as follows for each new device.

Create New File System

1. To display sector count information, type

```
prtvtoc /dev/md/rdisk/d0
```

The size of the disk is displayed under the Sector Count field.
(Assume 4099000 for this example.)

2. To create the VxFS file system for the volume, type

```
mkfs -F vxfs -o bsize=1024 logsize=512\ /dev/md/rdisk/d0 4099000
```

Where:

bsize = block size in bytes (1k, 2k, 4k, or 8k - 1k default for filesystems < 4 GB, 4k default for filesystems > 4 GB)

logsize = size of VxFS file system logging in blocks (256 blocks default, 32 to 1024 blocks)

4099000 = file system size in sectors (from **prtvtoc** command)

Create Mount Directory

Once the file systems for each device are created, create a mount directory for each device. To do this, type a statement similar to the following for each device:

```
mkdir /fs/d0
```

Where **/fs/d0** is the complete path for the new file system directory.

Mount the Filesystem

To mount each filesystem, type a statement similar to the following:

```
mount -F vxfs /dev/md/dsk/d0 /fs/d0
```

UFS File System on VxVM Device

To create standard UNIX file systems under Veritas Volume Manager, log in as **root** and proceed as follows for each new device.

Create New File System

Once you have formatted, partitioned, and labeled each Symmetrix disk device, and created VxVM volumes, create a new file system on each volume. To do this, use the **newfs** command in a statement similar to the following:

```
newfs -v /dev/vx/rdisk/dskgrp/vol1
```

The **Construct new filesystem?** prompt appears. Type **y** at the prompt.

The actions above created a new file system for the Symmetrix disk defined as **vol1** in the diskgroup named **dskgrp**.

Create Mount Directory

Once the file systems for each device are created, create a mount directory for each device. To do this, type a statement similar to the following for each device:

```
mkdir /fs/vol1
```

Where **/fs/vol1** is the complete path for the new file system directory.

Mount the Filesystem

To mount each filesystem, type a statement similar to the following:

```
mount -F vsfs /dev/vx/dsk/dskgrp/vol1 /fs/vol1
```

VxFS File System on VxVM Device

A VxFS journaling file system is created using the **mkfs** command with arguments provided for block size, log size, device name and size. To create a VxFS file system first determine the size in sectors of the volume. The size of the volume is displayed under the **Sector Count** field of the **prvtoc** output.

To create VxFS journaling file systems under Veritas Volume Manager, log in as **root** and proceed as follows for each new device.

Create New File System

1. To display sector count information, type

```
prvtoc /dev/vx/rdisk/dskgrp/vol1
```

The size of the disk is displayed under the **Sector Count** field. (Assume 4099000 for this example.)

2. To create the VxFS file system for the volume, type

```
mkfs -F vxfs -o bsize=4096 logsize=512\ /dev/vx/rdisk/dskgrp/vol1 4099000
```

Where:

bsize = block size in bytes (1k, 2k, 4k, or 8k - 1k default for filesystems < 4 GB, 4k default for filesystems > 4 GB)

logsize = size of VxFS file system logging in blocks (256 blocks default, 32 to 1024 blocks)

4099000 = file system size in sectors (from **prvtoc** command)

Create Mount Directory

Once the file systems for each device are created, create a mount directory for each device. To do this, type a statement similar to the following for each device:

```
mkdir /fs/vol1
```

Where **/fs/vol1** is the complete path for the new file system directory.

Mount the Filesystem

To mount each filesystem, type a statement similar to the following:

```
mount -F vxfs /dev/vx/dsk/dskgrp/vol1 /fs/vol1
```

File System Expansion

DiskSuite volume manager allows you to increase the available storage space of an existing volume by concatenating additional volumes to the metadvice using **metattach**. A file system can then be expanded to fill all or part of the additional space using the **growfs** command. The file system can remain mounted, but will be locked (**lockfs**) during the expansion. For detailed information on expanding a file system under DiskSuite refer to *Solstice DiskSuite 4.0 /4.1 Administrator's Guide*.

VxVM allows *growing* of a mounted file system using **vxassist** commands. The following steps outline the procedure.

1. Log in as **root**.
2. To determine how large the volume can grow, type

```
vxassist maxgrow vol
```

Where **vol** is the volume name

Volume vol can be extended by 12533760 to 16629760 (8120Mb)

3. Determine the size of the current volume, type

```
vxprint -vt
```

```
Disk group: cust
```

V	NAME	USETYPE	KSTATE	STATE	LENGTH	READPOL	PREFPLEX
v	vol	fsген	ENABLED	ACTIVE	4096000	ROUND	-

The length unit is in sectors (1 sectors = 512 bytes). Therefore **vol** is approximately 2 GB.

4. To grow the volume **vol** to 4 GB, type

```
vxassist growto cust-mirvol 4g
vxprint -vt
```

Disk group: cust

V NAME	USETYPE	KSTATE	STATE	LENGTH	READPOL	PREFPLEX
v cust-mirvol	fsgen	ENABLED	ACTIVE	8388608	ROUND	-

System and Error Messages

Solaris logs system and error messages including SCSI error messages to a file called **/var/adm/messages** and also displays these messages at the system console.

This chapter describes how to incorporate Symmetrix in the Solaris x86 environment.

◆ Overview	13-2
◆ Useful Solaris x86 Utilities and Functions	13-4
◆ Host Bus Adapters	13-6
◆ Symmetrix Configuration	13-9
◆ Incorporating Symmetrix	13-11
◆ Partitioning and Labeling a New Device	13-13
◆ Creating and Mounting a UFS File System	13-18
◆ Adding Devices Online	13-19
◆ System and Error Messages	13-20

Overview

This chapter covers the Sun/Solaris x86 OS environment for the Intel processor. See other chapters for other Sun operating systems.

When using Symmetrix in the Solaris x86 environment, please note the following:

- ◆ The Solaris x86 host must have either Wide, Ultra Wide SCSI Differential, or Fibre Channel cards installed in order to connect to the Symmetrix unit.
- ◆ Targets 0 through F are valid for Symmetrix SCSI devices, with the conventional exception of target 7 which is reserved for Initiator ID use.

The table below summarizes the operating system features in the Solaris x86 environment.

Table 13-1 Solaris x86 Operating System Features

O/S	Sun Solaris 2.7 x86
Largest recognizable drive size	N/A
Largest defined partition	N/A
Largest data file	Limited to device size
Largest filesystem	N/A
Number of SCSI target IDs supported	Hex F
Number of SCSI LUNs per target ID supported	Hex F
Number of Fibre Channel targets supported	1
Number of Fibre Channel LUNs supported	0-127
Partition representing the entire disk	s2
To see the drive configuration, use this utility:	format or prtvtoc
To configure drives use this utility:	format
Special drive definition files needed	yes

Host Setup Overview for New Symmetrix

The following is a general overview of the Symmetrix integration to the Solaris x86 environment. Please see references within each step below for specific procedures.

1. Verify that the host bus adapter is installed in the host system with the proper driver version, firmware and system BIOS. Verify that the Symmetrix is connected to the host. See *Host Bus Adapters* on page 13-6 for EMC-qualified HBA driver and firmware version. See also Host Bus Adapter vendor documentation for HBA and driver installation procedures.
2. During boot (if available) use the Host Bus Adapter utility to verify configuration settings. If available use the scan function of the HBA utility to verify the SCSI or Fibre Channel target LUN configuration. See *Host Bus Adapters* on page 13-6, Appendix C, *Host Bus Adapter Configuration*, and HBA vendors for more information.
3. For SCSI host board adapters, verify that the **/etc/system** and **/kernel/drv/sd.conf** files are configured in accordance with *Incorporating Symmetrix* on page 13-11.

For Fibre Channel host bus adapters, verify **/etc/system** and **/kernel/drv/sd.conf** and **/kernel/drv/qla2100.conf** files are configured in accordance with *Partitioning and Labeling a New Device* on page 13-13.

If any changes are made to any of these files the system must be rebooted with the command `reboot -- -rv`

After steps 1 through 3 have been completed and the system has been rebooted you will be able to see all Symmetrix raw devices in the **/dev/rdisk/** directory. To see all targets and LUNs in the Symmetrix configuration use the command `ls /dev/rdisk/*s0`

You may also view Symmetrix devices by issuing the `format` command.

4. To partition and label raw devices follow the procedures in *Partitioning and Labeling a New Device* on page 13-13.
5. To create and mount a file system follow the procedures in *Creating and Mounting a UFS File System* on page 13-18.
6. To add a Symmetrix device online, follow the procedures in *Adding Devices Online* on page 13-19.

Useful Solaris x86 Utilities and Functions

This section lists Solaris functions and utilities you may use to define and manage Symmetrix devices. The use of these functions and utilities is optional. They are listed for reference only.

- ◆ `format`

The Solaris x86 disk format utility allows you to format, partition, and label disk drives via an interactive menu.

- ◆ `newfs`

This command creates a new ufs file system.

- ◆ `drvconfig`

This command detects and configures any new devices.

- ◆ `disks`

This command creates symbolic links in the `/dev/dsk/` and `/dev/rdsk/` directories pointing to the actual device special file. This command should be run after executing `drvconfig`.

- ◆ `disks -C`

This command removes symbolic links in the `/dev/dsk/` and `/dev/rdsk/` directories. If a disk device is removed the symbolic links are not automatically removed from these directories until this command is issued.

- ◆ `fdisk`

This command is used to create and modify the partition table, and to install the master boot record that is put in the first sector of the fixed disk. The `fdisk` command can be used in command mode or through the interactive `format` command.

- ◆ `prtvtoc`

This command displays disk information from the VTOC table (volume table of contents).

- ◆ `fsck`

This command audits and interactively compares inconsistent conditions for file systems.

- ◆ `mkdir`

This command creates a specified directory.

- ◆ `mount`

This command attaches a file system to a specified directory called the mount point.

- ◆ `umount`

This command un-mounts a mounted file system.

- ◆ `shutdown`

This command shuts down the system gracefully.

- ◆ `reboot -- -rv`

This command stops and restarts the operating system. Also verbosely checks and configures new devices.

Host Bus Adapters

Consult your EMC Sales representative for the most current listing of supported hosts and supported HBAs.

Table 13-2 Adaptec AHA2944UW^a

Hosts Qualified by EMC:	For qualified hosts and BIOS, see your EMC representative for information in the latest revision of the Open Systems Support Matrix.
Targets supported	0-F
LUNs supported	0-F
OS Patches Needed	None
Firmware	2.11
Driver	Solaris x86 Release driver
HBA Driver configuration	Max Sync Transfer Rate (MB/Sec): 40.0
	Initiate Wide Negotiation: Yes
	Disable BIOS
Notes	To access HBA configuration utility, press <Ctrl> A at boot time.
	No LUN skipping
	No LUN out of sequence

a. For more information regarding SCSI parameter settings for this card, refer to Appendix C, *Host Bus Adapter Configuration*.

Table 13-3 NCRpqs 2.0 Quad Card

Hosts Qualified by EMC	NCR 4300 BIOS 1.00.10.CD0
Targets supported	0-F
LUNs supported	0-F
Special Patches Needed	None
Firmware	Firmware BIOS version 2.13 FPGA memory Memory Controller Version September 22, 1996
Driver	Version 1.1, Rev=1998.09.09.11.23
HBA Driver configuration	None

Table 13-4 NCRpqs 2.1 Quad Card

Hosts Qualified by EMC	NCR 4400 BIOS Phoenix BIOS 4.0 Release 6.0
Targets supported	0-F
LUNs supported	0-F
Special Patches Needed	None
Firmware	Firmware BIOS version 3.02 FPGA memory Memory Controller Version September 22, 1996
Driver	Version 1.1, Rev=1998.09.09.11.23
HBA Driver configuration	Default
Notes	NCR requires using only bus 1, Slot 6-10 in NCR 4400 with NCRpqs 2.1 with BIOS 3.02. To access the config utility press <CTRL> P at boot time.

Table 13-5 Qlogic QLA2100F (Fibre Channel) Card^a

Hosts Qualified by EMC	For qualified hosts and BIOS, see your EMC representative for information in the latest revision of the Open Systems Support Matrix.
Targets supported	
LUNs supported	0-127
Special OS Patches Needed	None
Firmware	Version 1.37
Driver	Contact your EMC service representative for updated online information. See also the latest revision of the Symmetrix Fibre Channel Product Guide, P/N 200-999-642.
HBA Utility configuration	At Boot time Press <alt> Q when the qla2100F banner is displayed to enter Fast!Util Options menu. Under Configuration Settings: 1) Change all settings to default Under Advanced Adapter Settings: 2) Change Max throttle to 20 3) Change Max Luns to 128
Notes	Fibre utility will not display all Symmetrix devices configured to the host. Use Scan Utility only for connectivity verification. For parameter settings for this card, contact EMC for updated online information. ^a

a. Refer also to the latest revision of the *Symmetrix Fibre Channel Product Guide*, Part Number 200-999-642.

Table 13-6 QLogic QLA2200-EMC

Hosts Qualified by EMC	For qualified hosts and BIOS, see your EMC representative for information in the latest revision of the Open Systems Support Matrix.
Targets supported	
LUNs supported	0-127
Special OS Patches Needed	None
Firmware	Version 1.59
Driver	3.11
HBA Driver configuration	Press <Alt> Q at boot time
Notes	For complete setup procedures, refer to the latest revision of the Symmetrix Fibre Channel Product Guide, Part Number 200-999-642.

Symmetrix Configuration

In the Solaris x86 host environment, you can configure the Symmetrix disk devices into logical volumes. EMC recommends the following logical volume ratios as indicated in the table below:

Table 13-7 Disk Device Configuration

		Symmetrix Logical Volume to Physical Splits ^a							
		1:1	2:1	3:1	4:1	5:1	6:1	7:1	8:1 ^b
Physical Devices	4 GB	x	x						
	9 GB	x	x	x	x				
	18 GB		x	x	x	x	x	x	x
	23 GB		*	x	x	x	x	x	x
	36 GB			*	x	x	x	x	x
	47 GB			*	*	x	x	x	x
	50 GB ^c				x	x	x	x	x
	73 GB ^{c, d}					x	x	x	x

- * indicates ratios available only in microcode level 5265 or higher.
- With microcode revision level 5263, maximum is 8:1. With 5264, maximum is 16:1. With 5265, maximum is 32:1.
- The 50 and 73 GB drives are supported only by microcode revision levels 5x66 and higher.
- The 73 GB drive does not support the RAID-S configuration on Symmetrix.

For specific cylinder counts and logical volume sizes, refer to Table 2-1 on page 2-5 and the following tables in that chapter.

The EMC Customer Engineer (CE) should contact the EMC Configuration Specialist for updated online information. This information is necessary to configure Symmetrix to support the customer's host environment. The EMC Customer Engineer will assign target IDs and LUNs to the Symmetrix devices. The CE will also bring the Symmetrix channel and disk directors online.

Use the following procedure to verify that the host can see all the target Symmetrix devices. After configuring the host according to the following section, *Incorporating Symmetrix*, you will be able to see all

Symmetrix raw devices in the **/dev/rdsk/** directory. To see all targets and LUNs in the Symmetrix configuration, use the command `ls /dev/rdsk/*s0`. You may also view Symmetrix devices by issuing the `format` command.

Incorporating Symmetrix

The following tells how to incorporate Symmetrix volumes into Solaris x86 systems with SCSI HBAs. See page 13-13 for incorporating Fibre Channel HBAs. For information on x86 with Fibre Channel, see EMC's *Symmetrix Fibre Channel Product Guide*, P/N 200-999-642

If any of the following changes are made to the `/etc/system` or `/kernel/drv/sd.conf` files, the system must be rebooted with the `reboot -- -rv` command to incorporate the changes.

To Recognize LUNs

Symmetrix devices are addressed on the SCSI bus using SCSI target IDs and SCSI LUNs. Each target ID can have 0-15 LUNs. By default, Solaris x86 only searches for target IDs on the SCSI bus. If LUN addressing is required, you will need to modify the file `/kernel/drv/sd.conf`. Add the following definitions for each target ID which needs to support multiple LUNs.

```
name="sd" class="scsi"
target=0 lun=0;
name="sd" class="scsi"
target=0 lun=1;
name="sd" class="scsi"
target=0 lun=2;
name="sd" class="scsi"
target=0 lun=3;
.
.
.
name="sd" class="scsi"
target=0 lun=12;
name="sd" class="scsi"
target=0 lun=13;
name="sd" class="scsi"
target=0 lun=14;
name="sd" class="scsi"
target=0 lun=15;
```

This addition to `sd.conf` is necessary when using LUN addressing. Make sure for target IDs that do **not** need LUN support, specify **only** the `target=x lun=0` line. This decreases the reboot time.

To enable Solaris x86 to support the Symmetrix system, EMC provides Symmetrix device definition files. These files are available on EMC's FTP server, ftp.emc.com, in `/pub/symm3000/solaris`. See Appendix A for instructions to transfer these files to your host.

System Specification File Modification

There are two parameters in the file **/etc/system** that require modification when operating in a Solaris x86 environment. These parameters are `scsi_options` and `sd_io_time`. This section describes the necessary modifications.

- ◆ For hosts configured with the Ultra-Wide Differential SCSI Controller, include the following entry in the **/etc/system** file:

```
set scsi_options = 0x7F8
```

The Solaris x86 default setting for SCSI options is 0x7F8. Since the **/etc/system** file can also affect other disks in the system, consider the effects of implementing the `scsi_options=0x7F8` mask on the other disks.

Bit meanings for the settings are listed in Table 13-8.

Table 13-8 Bit settings for SCSI Options

Bit	Mask	Meaning
3	0x8	Disconnect enable
4	0x10	Link enable
5	0x20	Sync xfer enable
6	0x40	Parity support enable
7	0x80	Tag commands enable
8	0x100	Fast SCSI enable
9	0x200	Wide SCSI enable
10	0x400	Ultra SCSI enable

- ◆ In the file **/etc/system**, set `sd_io_time` to 120 seconds. This setting will prevent the host from issuing warning messages while non-disruptive operations are performed on Symmetrix:

```
set sd:sd_io_time = 0x78
```

If either the **/etc/system** or **/kernel/drv/sd.conf** files is changed, the system must be rebooted with the `reboot -- -rv` command to incorporate the changes.

Partitioning and Labeling a New Device

Disk Device File Naming Conventions

Solaris x86 platform disk partitioning has one more level than that of Solaris for SPARC. On the Solaris SPARC platform, slices are the primary method of partitioning a disk. For the x86 platform, slices are sub-partitions of PC partitions. Using the `fdisk` command allows up to 4 PC partitions on a disk device. This is done so that more than one operating system can reside on or be booted from the same disk.

Raw disk device filenames take two forms:

1. `/dev/rdisk/cCtTdDsS`

Example: `/dev/rdisk/c0t0d0s0`

2. `/dev/rdisk/cCtTdDpP`

Example: `/dev/rdisk/c0t0d0p0`

where:

C is the occurrence of the host adapter board.

T is the SCSI target ID.

D is the logical unit number (LUN) of the disk device.

S is the disk slice number.

P is the `fdisk` partition number. **P** can range from 0 to 4. The `p0` node corresponds to the whole disk, and the `p1` to `p4` nodes correspond to partitions displayed by the `fdisk` command.

The `fdisk` command will only work with device `p0` representing the entire device.

Partitioning and Slicing Symmetrix Devices

To partition and slice a Symmetrix device for the use of Solaris x86, log in as **root** and follow the procedure below for each new device. Once the Symmetrix and the host have been integrated there will be raw device files in the `/dev/rdisk/` directory representing the target and LUNs specified by the EMC customer engineer.

To partition and label new devices:

1. At the root prompt, type **format <Enter>**.

The host searches for all disks and displays a list, placing all unlabeled disks at the beginning of the listing. Here is an example:

AVAILABLE DISK SELECTIONS:

```
0. c0t0d0 <DEFAULT cyl 5512 alt 2 hd 8 sec 193> /pci@1,0/pci9004,8178@a/sd@0,0
1. c1t0d0 <DEFAULT cyl 4093 alt 2 hd 15 sec 64> /pci@0,0/pci9004,8478@8/sd@0,0
2. c1t0d1 <DEFAULT cyl 4093 alt 2 hd 15 sec 64> /pci@0,0/pci9004,8478@8/sd@0,1
Specify disk (enter its number):
```

2. At the Specify disk prompt, enter the number of the first EMC drive. The following menu appears:

FORMAT MENU:

```
disk          - select a disk
type          - select (define) a disk type
partition     - select (define) a partition table
current       - describe the current disk
format        - format and analyze the disk
fdisk         - run the fdisk program
repair        - repair a defective sector
label         - write label to the disk
analyze       - surface analysis
defect        - defect list management
backup        - search for backup labels
verify        - read and display labels
save          - save new disk/partition definitions
inquiry       - show vendor, product and revision
volname       - set 8-character volume name
!<cmd>        - execute <cmd>, then return
quit
format>
```

3. At the `format>` prompt, type the word **type** and press **<Enter>**. The following menu appears:

```
AVAILABLE DRIVE TYPES:
```

- 0. Auto configure
- 1. DEFAULT
- 2. DEFAULT
- 3. DEFAULT
- 4. Other

```
Specify disk type (enter its number) [3]:
```

- Type **0 <Enter>**. A display similar to the following appears:

```
clt3d6: configured with capacity of 1.05GB  
<EMC-SYMMETRIX-5265 cyl 2296 alt 2 hd 15 sec 64>
```

```
Format>
```

4. At the `format>` prompt, type **fdisk <Enter>**.

- If no PC partition table exists on the Symmetrix, a message similar to the following appears:

```
No fdisk table exists. The default partition for the disk is:  
a 100% "SOLARIS System" partition  
Type y to accept the default partition, otherwise type n to edit the  
partition table.
```

You must create a default partition or specify the initial table values. Type **Y <Enter>** to accept the default option. This allocates 100% of the disk for the Solaris X86 system and makes the Solaris X86 system partition active.

- If a partition table already exists, the command program displays a partition table and a menu as shown below, with options allowing you to modify the table:

```

                                Cylinders
Partition Status      Type   Start  End    Length %
=====
      1      Active   Solaris  1    4094   4094  100

SELECT ONE OF THE FOLLOWING:

1. Create a partition
2. Specify the active partition
3. Delete a partition
4. Exit (update disk configuration and exit)
5. Cancel (exit without updating disk configuration)
Enter Selection:

```

Figure 13-1 Solaris X86 fdisk Menu Example



CAUTION

Changing the partition table results in loss of any previous data.

- If Type is Solaris and Status is Active, type **4** <Enter> (Exit), then proceed to Step 5.
- If Type is Solaris and Status is *not* Active, type **2** <Enter>. When prompted for the partition to activate or boot from, enter the number of the Solaris X86 system partition. If the **fdisk** main menu (see Figure 13-1) shows a Solaris X86 partition with a status of **Active**, type **4** <Enter>, then proceed to Step 5.
- If there is an unwanted partition, type **3** <Enter>. When prompted for the partition to delete, enter the number and press <Enter>. When prompted to verify the deletion, press <Y>, then proceed to Step 5.
- If you want to create a Solaris system partition, type **1** <Enter>. When prompted with a list of partition types, type **1** <Enter> to select Solaris system partition.

When prompted for the percentage of the disk you want to use for the partition, type **100** <Enter>. If 100% of the disk is not available, use all of the disk that is available.

When prompted to activate the partition, type **Y** <Enter>, then proceed to Step 5. (Solaris X86 partitions are accessible only if they are in the Active state.)

5. At the `format>` prompt, type **label** **<Enter>**.
6. At the `Disk not labeled. Label it now?` prompt, type **Y** **<Enter>**.
7. When prompted to display a listing of the disks, type **di** **<Enter>**.
8. Repeat steps 2 through 7 for each remaining Symmetrix disk.

Creating and Mounting a UFS File System

To create standard UNIX file systems under Solaris x86 OS, log in as **root** and proceed as follows for each previously labeled device.

Create a New File System

Once you have formatted, partitioned, and labeled each Symmetrix disk device (see Section 12.6), create a new file system for each Symmetrix device. To do this, use the `newfs` command in a statement similar to the following:

```
newfs /dev/rdisk/c1t0d0s2
```

You will be prompted to confirm the construction of a new file system. Type `Y` and press `<enter>`.

For default labels as specified in *Partitioning and Labeling a New Device* on page 13-13, always use slice 2 for mounted file systems. The example above creates a new file system for the Symmetrix disk connected to SCSI controller 1, target ID 0, lun 0, slice 2.

Create a Mount Directory

Once the file systems for each Symmetrix device are created, create a mount directory for each device. To do this, type a statement similar to the following for each Symmetrix device:

```
mkdir /fs/c1t0d0s0
```

where `/fs/c1t0d0s0` is the complete path for the new file system directory.

The `/fs` directory must exist prior to creating the mount directories.

Mount the Filesystem

To mount each filesystem, type a statement similar to the following:

```
mount -F ufs /dev/dsk/c1t0d0s2 /fs/c1t0d0s2
```

This action will mount the raw device `c1t0d0s2` onto the mount point `/fs/c1t0d0s2`

Adding Devices Online

Whenever devices are added online to the Symmetrix unit or device channel addresses are changed, you will need to perform the actions described below in order to introduce the new devices to the system.

To add newly configured Symmetrix devices while online in the Solaris x86 environment, perform the following steps:

Before trying to do an online upgrade verify that the targets and LUNs being added are defined in `/kernel/drv/sd.conf` in accordance with *Incorporating Symmetrix* on page 13-11. If the LUNs being added are not specified in `/kernel/drv/sd.conf` the definitions must be added and the system rebooted to incorporate changes.

1. Use the **online upgrade** feature at the Symmetrix service processor to add/map new drives to the SCSI host channel. Consult your CE for this procedure.
2. Execute the following three utilities in order to detect and create symbolic links for new devices:

```
drvconfig
```

```
disks -C
```

```
disks
```

3. Follow the instructions in *Partitioning and Labeling a New Device* on page 13-13 to introduce new devices to the host environment.

System and Error Messages

Solaris x86 logs system and error messages including SCSI error messages to a file called **/var/adm/messages** and also displays these messages at the system console.

Compaq Tru64 UNIX Environment

This chapter describes how to incorporate Symmetrix in the Compaq (formerly Digital Equipment Corporation) Tru64 UNIX environment.

◆ Overview	14-2
◆ Compaq Tru64 UNIX Utilities & Functions	14-4
◆ Symmetrix Configuration	14-6
◆ Host Configuration	14-8
◆ Device Management	14-11
◆ Using File Systems	14-15
◆ LSM	14-18
◆ System and Error Messages	14-22

Overview

When using Symmetrix in the Compaq Tru64 UNIX environment, please note the following:

- ◆ The minimum version of Tru64 UNIX (Digital UNIX) supported is V3.2A. V4.0D or higher is strongly recommended, however.
- ◆ The most current information regarding any restrictions, exceptions, firmware/driver versions, and other requirements is listed in the updated *EMC Symmetrix Open Systems Support Matrix*, available through your EMC representative.
- ◆ The following host bus adapters (HBAs) are supported:
 - KZPSA-BB (SCSI Fast-Wide Differential)
 - KZPBA-CB (Ultra-SCSI Differential)
 - KGPSA-BC (Fibre Channel)
 - KGPSA-CA (Fibre Channel)
- ◆ Symmetrix supports TURBOchannel-based DECstations and DECservers DEC 3000 Models 300-900, PCI-based AlphaStations, and the following PCI-based AlphaServers:
 - AlphaServer 1200
 - AlphaServer 1000/1000A
 - AlphaServer 2000/2000A
 - AlphaServer 2100/2100A
 - AlphaServer 4000/4100
 - AlphaServer 8200/8400
 - AlphaServer GS60/GS140
 - AlphaServer DS20/ES40
 - AlphaServer GS80/GS160/GS320
- ◆ For SCSI directors for your model of the Symmetrix, see Appendix B, *SCSI Directors and Cabling*.

EMC Files, Compaq Patches & Online Docs

- ◆ For information about Fibre Channel connectivity, refer to the latest revision of the *Symmetrix Fibre Channel Product Guide*, EMC part number 200-999-642. Other information regarding high availability configurations is listed in the online *EMC Symmetrix High Availability Product Guide*, which you can obtain through your EMC representative.
- ◆ Compaq patches for Tru64 UNIX are available from
`ftp1.support.Compaq.com/public/unix/`
- ◆ Compaq Tru64 UNIX online documentation can be found on the Compaq website:

`www.tru64unix.compaq.com/faqs/publications/pub_page/doc_list.html`

Compaq Tru64 UNIX Utilities & Functions

This section lists Compaq Tru64 UNIX functions and utilities you may use to define and manage Symmetrix devices. The use of these functions and utilities is optional. They are listed for reference only.

◆ `disklabel`

Displays and partitions a disk device. Some useful parameters for this command are:

- `-r` read label from disk
- `-rw` write label to disk
- `-re` puts you into your default editor so you can change the default partition sizes

◆ `scu`

An online SCSI utility program that provides a listing and scanning of devices connected to the host.

◆ `iostat [device name] <interval> <count>`

Displays I/O statistics.

`sar` may be used if the System V habitat has been installed. (The System V habitat requires a separate license.)

By default, `iostat` shows the first four disks in the system (internal) or the LUNs associated with a specified device. For example:

```
iostat rz16 rz32 5 0
```

◆ **LSM (Logical Storage Manager)**

Based on Veritas Volume Manager. LSM supports mirroring and striping.

You can use three interfaces with LSM commands:

- a. Command line: `vol*` commands
- b. Character_cell: `voldiskadm`
- c. Motif-based GUI: `dxlsm` (requires separate license)

Use of LSM system-level mirroring and striping requires the same license as the GUI interface.

◆ `newfs`

Creates a new UNIX File System (UFS).

◆ `mkfsmn, mkfset`

Creates a new Advanced File System (AdvFS). AdvFS provides rapid crash recovery, high performance, and the ability to manage the file system while it is online.

◆ `scsimgr`

This command is used to create device special files for newly attached disk and tape devices. This utility is automatically invoked at system boot time of the V4.0x operating system.

◆ `hwmgr`

V5.x operating system command used to display and manage hardware components.

◆ `dsfmgr`

V5.x operating system command used to create and manage device special files.

Symmetrix Configuration

In the Compaq Tru64 UNIX host environment, you can configure the Symmetrix disk devices into logical volumes. EMC recommends the following logical volume ratios as indicated in the table below:

Table 14-1 Disk Device Configuration

		Symmetrix Logical Volume to Physical Splits ^a							
		1:1	2:1	3:1	4:1	5:1	6:1	7:1	8:1 ^b
Physical Devices	4 GB	x	x						
	9 GB	x	x	x	x				
	18 GB		x	x	x	x	x	x	x
	23 GB		*	x	x	x	x	x	x
	36 GB			*	x	x	x	x	x
	47 GB			*	*	x	x	x	x
	50 GB ^c				x	x	x	x	x
	73 GB ^{c, d}					x	x	x	x

- a. * indicates ratios available only in microcode level 5265 or higher.
- b. With microcode revision level 5263, maximum is 8:1. With 5264, maximum is 16:1. With 5265, maximum is 32:1.
- c. The 50 and 73 GB drives are supported only by microcode revision levels 5x66 and higher.
- d. The 73 GB drive does not support the RAID-S configuration on Symmetrix.

For specific cylinder counts and logical volume sizes, refer to Table 2-1 on page 2-5 and the following tables in that chapter.

The EMC Customer Engineer (CE) should contact the EMC Configuration Specialist for updated online information. This information is necessary to configure Symmetrix to support the customer's host environment.

Table 14-2 Tru64 UNIX SCSI Device Support

	Tru64 UNIX V4.0	Tru64 UNIX V5.x
Valid Target IDs ^a	8 (0-7)	16 (0-F)
LUNs per Target ID	8 (0-7)	8 (0-7)

- a. Most SCSI host bus adapters have a default target ID of 7. Unless the default target ID of the adapter is different or changed, do not use target ID 7.

Have the EMC Customer Engineer configure the Symmetrix interfaces with the target IDs and LUNs required.

The 5.x version of the Tru64 operating system can support multipath configurations to Symmetrix devices. The C (Common Serial Number) Director flag should be set in V5.x multipath configurations.

The most current information regarding Director flag settings and microcode levels is listed in the updated *EMC Symmetrix Open Systems Support Matrix*, available through your EMC representative.

Host Configuration

This section describes the tasks required to install one or more Compaq host bus adapters into the AlphaServer host and configure the Tru64 UNIX environment.

Installing the HBA

Follow the instructions included with the adapter. The adapter installs into a single PCI bus slot.

If necessary, load the required HBA firmware level. This information is specified in the *EMC Symmetrix Open Systems Support Matrix*, available through your EMC representative.

Also, verify that the AlphaServer SRM console firmware is current. (You can check using `show version` at the AlphaServer console prompt).

Firmware updates can be downloaded from `ftp.digital.com/pub/DEC/Alpha/firmware/`. In addition, *Firmware Update* CDs are shipped in Compaq's Tru64 UNIX Release Kits.

To display the host IDs of the SCSI HBAs on the system, use `show pk*` or `show isp*` at the AlphaServer console prompt.

The HBA host ID can be changed if an ID conflict needs to be resolved. For example, to change HBA instance `c` to ID 6, use `set pkc0_host_id 6`.

Console-Level Scan and Configuring a Boot Device

To perform an initial scan of all devices visible to the host system before booting, use `show devices | more` at the AlphaServer console prompt. Verify that all configured Symmetrix SCSI devices are reported correctly in the display output.

Bootting from Symmetrix devices is supported for V5.0A and V5.1 of Tru64 UNIX. Information regarding requirements and restrictions is specified in the *EMC Symmetrix Open Systems Support Matrix*, available through your EMC representative.

To specify a boot device, set the `bootdef_dev` console variable after identifying the intended boot device in the `show device` output. For example:

```
set bootdef_dev dkc301
```

Rebuilding the Tru64 UNIX Kernel

When a new HBA is installed in the host system, the Tru64 UNIX kernel must be rebuilt to identify and support the new adapter.

To rebuild the kernel:

1. Type `boot -file genvmunix` at the AlphaServer console prompt.
2. After the host boots, run `bcheckrc`, then `doconfig`.
3. At the prompts, use the default configuration file and replace it, then skip the config file edit.
4. When the kernel rebuild completes, copy the new file to root:
`cp /sys/<hostname>/vmunix /vmunix.`
5. Reboot the host system.

Adding the Symmetrix Device Entry

1. Include the following entries for Symmetrix devices in the host `/etc/ddb.dbase` file.

```
SCSIDEVICE
  Type = disk
  Name = "EMC" "SYMMETRIX"
  PARAMETERS:
    TypeSubClass      = hard_disk, raid
    BlockSize         = 512
    BadBlockRecovery   = disabled
    DynamicGeometry    = true
    LongTimeoutRetry   = enabled
    DisperseQueue      = false
    TagQueueDepth      = 20
    ReadyTimeSeconds   = 45
    InquiryLength      = 160
    RequestSenseLength = 160
    PwrMgmt_Capable    = false
# (The following should be uncommented for TruCluster V5.x
# configurations only)
#   ATTRIBUTE:
#     AttributeName    = "DSBLflags"
#     Length           = 4
#     ubyte[0]         =

SCSIDEVICE
  Type = disk
  Stype = 2
  Name = "EMC" "SYMMETRIX"
  PARAMETERS:
    TypeSubClass      = hard_disk, raid
```

```

        BlockSize                = 512
        BadBlockRecovery          = disabled
        DynamicGeometry           = true
        LongTimeoutRetry          = enabled
        DisperseQueue             = false
        TagQueueDepth             = 20
        ReadyTimeSeconds          = 45
        InquiryLength             = 160
        RequestSenseLength        = 160
        PwrMgmt_Capable           = false
# (The following should be uncommented for TruCluster V5.x
# configurations only)
#   ATTRIBUTE:
#       AttributeName    = "DSBLflags"
#       Length           = 4
#       ubyte[0]         =

```

2. Recompile the the edited `/etc/ldr.dbase` file, using `ldr_config -c`.
3. Reboot the host system.
4. Verify the changes, using `ldr_config -s disk EMC SYMMETRIX`.

Device Management

Device Naming Conventions

V4.0x

In the 4.0x versions of the Tru64 UNIX operating system, device names (device special files) are determined by bus-target-lun location, as in

`rz[lun][unit][partition]`

where

[lun] is a letter ranging from "b" to "h", corresponding to LUNs 1 through 7. For LUN 0, the LUN letter is omitted.

[unit] is bus number * 8 + target ID number.

[partition] is the letter of the disk partition, from "a" to "h".

For example:

`/dev/rrzb16c` is the raw device name for partition **c** of the disk at bus 2 target 0 LUN 1

`/dev/rz28h` is the block device name for partition **h** of the disk at bus 3 target 4 LUN 0

V5.x

In the 5.x versions of the Tru64 UNIX operating system, device names (device special files) are based on the WWID identifier of a device, and not the device's location, as in

`dsk[unit][partition]`

where

[unit] is a number assigned sequentially to new devices (with unique WWIDs) when they are discovered and configured by the operating system.

[partition] is the letter of the disk partition, from "a" to "h"

For example:

`/dev/rdisk/dsk2c` is the raw device name for partition **c** of the 2nd unique device configured by the host.

`/dev/disk/dsk853g` is the block device name for partition **g** of the 853rd unique device configured.

With the ID based device naming in V5.x, a device will retain its original device name (device special file) if moved around in the configuration through physical connection, adapter or bus changes. However, if an existing device is replaced by configuring a new Symmetrix logical device or copy of the original at the same bus-target-lun location, the new device or copy will likely have a different WWID and a different device name than the original. For data migration or BCV/SRDF recovery procedures where another Symmetrix logical device will be used in place of the original, it will be necessary to rename or relink the new device using host utilities and commands.

Disklabel and Device Partitions

Before using specific device partitions, file systems, or Logical Storage Manager, devices should be labeled and partitioned with the `disklabel` utility.

For example:

- ◆ Clear or zero-out any existing label: `disklabel -z dsk853`
- ◆ Label the disk using default or unknown, as in
`disklabel -rw dsk853 default`

The following figure shows a default disk partition scheme.




Figure 14-1 Partitioning Scheme Using "unknown" Disk Type

- ◆ To edit the label and disk partitions with customized parameters, use `disklabel -re dsk853`.
- ◆ To read and check the existing label and disk partitions, use `disklabel -r dsk853`.

Adding and Managing Devices Online

V4.0x

- ◆ To rescan the devices on the host system:
`scu scan edt`
- ◆ To show the devices visible to the host:
`scu show edt`
- ◆ To create device files for new devices, first calculate the new device name based on bus-target-lun location. For example, if a device on bus 3 target 3 lun 3 was newly added, create rzc27 device files:
`cd /dev`
`./MAKEDEV rzc27`
- ◆ The `scsimgr` utility can also scan and create device files automatically:
`scsimgr -scan_all`

or
`scsimgr -scan_bus bus=N`

V5.x

- ◆ To rescan the devices on the host system:
`hwmgr -scan component`
`hwmgr -scan scsi`
- ◆ To view all devices visible to the host:
`hwmgr -view devices`

Or to show more detailed output:
`hwmgr -show scsi -full`

For example:

```

      SCSI          DEVICE    DEVICE  DRIVER NUM  DEVICE  FIRST
HWID:  DEVICEID  HOSTNAME   TYPE    SUBTYPE  OWNER   PATH  FILE   VALID  PATH
-----
1664:  138        losaz215   disk    none     2       2     dsk1366 [3/1/0]

      WWID:04100024:"EMC      SYMMETRIX      600025068000"

      BUS   TARGET  LUN    PATH STATE
      -----
      3     1       0      valid
      4     1       0      valid

```

- ◆ To show any failed (or removed) devices/paths:

```
hwmgr -show scsi -stale
```

- ◆ To permanently remove old devices from the hardware device databases, use the following commands:

```
hwmgr -refresh component
```

```
hwmgr -refresh scsi
```

```
hwmgr -delete scsi -did <N>
```

- ◆ To create device files for newly added devices:

```
dsfmgr -k
```

The `dsfmgr` command can also be used to assign the device name (device special file) of a failed or removed device to a new replacement device or copy:

```
dsfmgr -m dsk1355 dsk853
```

To check device special files for inconsistencies or errors:

```
dsfmgr -v
```

Using File Systems

Creating & Mounting UNIX File System

Now that the physical volumes and their respective mount directories are available to the operating system, a new file system can be created for each Symmetrix disk.

1. Use the `newfs` command in a statement similar to the following to create a new file system:

```
newfs /dev/rz3c
```

2. Create a directory by typing:

```
mkdir /symm
```

3. Mount the file system by typing a statement similar to the following:

```
mount /dev/rz3c /symm
```

4. Assign the ownership by typing:

```
chown oracle:dba /symm
```

Use the `df` command to show all mounted file systems and the available free space.

Advanced File System (AdvFS)

Understanding the following concepts prepares you for planning, creating, and maintaining an AdvFS file system.

- ◆ Volumes

A volume is any mechanism that behaves like a UNIX block device, such as a logical volume that is configured with the Logical Storage Manager (LSM), or a disk partition.

◆ File Domain

A file domain is a named set of one or more volumes that provides a shared storage pool for one or more file sets.

When you create a file domain using the `mkfdmn` command, you must specify a domain name and one initial volume. The `mkfdmn` command creates a subdirectory in the `/etc/fdmns` directory for each new file domain. The file domain subdirectory contains a symbolic link to the initial volume.

You can add additional volumes to an existing file domain by using the `addvol` utility. With each added volume, the `addvol` utility creates a new symbolic link in the appropriate file domain subdirectory of `/etc/fdmns`.

◆ Filesets

A fileset is both the logical file structure that the user recognizes and a unit that you can mount. Whereas you typically mount a whole UNIX file system, with the Advanced File System you mount the individual filesets of a file domain.

An Advanced File System consists of a file domain with at least one fileset that you create using the `mkfset` command.

Creating & Mounting Advanced File System

To create an AdvFS domain and fileset for each Symmetrix disk, follow these steps:

1. Create a new file domain by typing, for example:

```
mkfdmn /dev/rz112c domain1
```

2. Create a fileset in the domain just created in step 1 by typing:

```
mkfset domain1 fileset1
```

3. Create a directory by typing:

```
mkdir /symm1
```

4. Mount the AdvFS fileset by typing a command similar to the following:

```
mount -t advfs domain1#fileset1 /syml1
```

5. Assign the ownership by typing:

```
chown oracle:dba /syml1
```

To mount a directory for an advanced file system every time the system boots, follow these steps:

1. Edit `/etc/fstab` by typing:

```
vi /etc/fstab
```

2. Specify each file system to mount using a statement similar to the following:

```
domain1#fileset1 /syml1 advfs rw
```

Use the `df` command to show all mounted file systems and the available free space.

Reconstructing an AdvFS Domain

If the device special file of an AdvFS domain device is no longer valid, or a copy of an existing AdvFS domain device is newly added to the host system, perform the following procedure to create or relink the AdvFS domain.

1. Verify that the new device is a valid AdvFS volume by checking the partition **fstype**, for example:

```
disklabel -r dsk1355
```

2. Recreate (if necessary) and change to the domain directory

```
mkdir /etc/fdmns/<domain_name>
cd /etc/fdmns/<domain_name>
```

3. Reconstruct the device link(s) of the AdvFS domain using the new device special file(s):

```
ln -s /dev/disk/dsk1355c
```

LSM

The following sections describe three different ways of using LSM (a Veritas Volume Manager):

- ◆ Setting up LSM for the first time.
- ◆ Creating a mirrored (2-way) volume.
- ◆ Creating a striped volume.

Example 1: Setting Up LSM

This example takes you through the steps for setting up LSM the first time.

1. Use the `volsetup` command to add disks to your setup. For example, type:

```
volsetup rz16 rzb32 rzb40
```

These three disks are added to your setup as disk01, disk02, and disk03.

2. To add more disks:

For an entire disk, type a command similar to the following:

```
voldiskadd rzb16
```

This command adds the entire disk as disk04.

For certain partitions, type a command similar to the following:

```
voldiskadd rzc16g
```

This command adds the disk partition as disk05.

3. To create a volume on a disk:

To create a 100 MB volume on disk01 type a command similar to the following:

```
volassist make myvol1 100m disk01
```

To create a 100 MB volume anywhere but on disk02, type a command similar to the following:

```
volassist make myvol2 100m !disk02
```

To create a 10 GB volume anywhere (in `rootdg`), type a command similar to the following:

```
volassist make mybigvol 10gb
```

UNIX File System

4. To put a UFS on a volume:

To put a UFS on the 10 GB volume type a command similar to the following:

```
newfs /dev/rvol/rootdg/mybigvol
```

5. To mount a volume type a command similar to the following:

```
mount /dev/vol/rootdg/mybigvol /syml
```

Advanced File System

6. To create a new AdvFS domain type a command similar to the following:

```
mkfdmn /dev/vol/rootdg/mybigvol domain1
```

7. To create a new fileset type a command similar to the following:

```
mkfset domain1 fset1
```

8. To mount the fileset type a command similar to the following:

```
mount -t advfs domain1#fset1 /syml
```

Example 2: Creating mirrored Volume

This example demonstrates a bottom-up approach to create a mirrored (2-way) volume.

1. Create two subdisks (for example, `sd1` and `sd2`) by typing the following commands. The subdisks in this example will be 100 MB in size and start at offset 0 of each disk.

```
volmake sd sd1 rzb16,0,100m
```

```
volmake sd sd2 rzb32,0,100m
```

2. Create a plex on each subdisk by typing commands similar to the following:

```
volmake plex plx1 sd-sd1
```

```
volmake plex plx2 sd-sd2
```

3. Create a mirrored volume by typing a command similar to the following:

An unmirrored volume consists of one plex. A two-way mirrored volume consists of two plexes. A three-way mirrored volume consists of three plexes. LSM supports up to eight-way mirrors.

```
volmake -U gen vol vol01 plex=plx1,plx2
```

4. Start the volume by typing a command similar to the following:

```
volume start vol01
```

This command will take some time as it synchronizes both plexes.

Use the `volprint` command to view the result of these commands:

```
volprint -ht s-vol
```

To set volume attributes and make them permanent for an LSM volume:

```
voledit set user=oracle group=dba mode=0640 vol01
```

Example 3: Creating 4-way striped volume

This example demonstrates a bottom-up approach to create a 4-way striped volume with a striped width of 64K.

1. Create four subdisks (for example, of 500 MB each starting at offset 0 of each disk) by typing the following commands:

```
volmake sd s1-sd rz16,0,500m
volmake sd s2-sd rz32,0,500m
volmake sd s3-sd rz40,0,500m
volmake sd s4-sd rz48,0,500m
```

2. Create a striped plex by typing a command similar to the following:

```
volmake plex s-pl sd=s1-sd,s2-sd,s3-sd,s4-sd \ layout=stripe stwidth=64k
```


3. Create a volume on the striped plex by typing a command similar to the following:

```
volmake -U gen vol s-vol plex=s-pl
```

4. Start the volume by typing a command similar to the following:

```
volume start s-vol
```

Use the `volprint` command to view the result of these commands:

```
volprint -ht s-vol
```

To set volume attributes and make them permanent for an LSM volume:

```
voledit set user=oracle group=dba mode=0640 s-vol
```

System and Error Messages

Compaq Tru64 UNIX logs system and error messages to the **/var/adm/messages** file. Tru64 UNIX V5.x also logs errors to the Event Manager (EVM). EVM messages can be viewed using commands such as `evmget` and `evmshow`. For example:

```
evmget -f "[since 2000:10:21:05:00:00]" | evmsort | evmshow -t "@timestamp @@"
```

To check SCSI errors, use either the `dia` command or the `uerf` command. Before you can use the `dia` command, the DECEvent software subset should be loaded. The subset can be found on the *Associated Products Volume 2 CD*.

Here are some `dia` usage examples:

- ◆ `dia -o full -i ios`
Provides a full report of all I/O-related error events.
- ◆ `dia -o terse -i -ios -R`
Provides a detailed report all I/O related error events in reverse order.
- ◆ `dia -o brief -i disk -c`
Provides a short summary of all disk-related events as they occur.

Here are some `uerf` usage examples:

- ◆ `uerf -c err -o full`
Provides a full report of all error events.
- ◆ `uerf -c err -o terse -R`
Provides a detailed report of all error events in reverse order.
- ◆ `uerf -c err -o brief -n`
Provides a short summary of all error events as they occur.

This chapter describes how to incorporate Symmetrix in the Fujitsu environment.

◆ Overview	15-2
◆ Useful Fujitsu Utilities and Functions.....	15-5
◆ Symmetrix Configuration	15-7
◆ Incorporating Symmetrix.....	15-8
◆ Format, Partition, Label	15-12
◆ Format, Partition, Label using sysadm	15-15
◆ Adding Devices Online	15-16
◆ Solstice DiskSuite	15-17
◆ Veritas Volume Manager	15-27
◆ Create and Mount a File System	15-42
◆ System and Error Messages	15-50

Overview

When using Symmetrix in a Fujitsu environment, please note the following:

For the most current information about supported hosts, HBAs and firmware listed in the updated *EMC Symmetrix Open Systems Support Matrix*, contact your EMC representative. For information about Fibre Channel connectivity, refer to the latest revision of the *Symmetrix Fibre Channel Product Guide*, Part Number 200-999-642.

- ◆ Symmetrix supports connection to Fujitsu DS-90, Fujitsu GP7000D and Fujitsu PRIMEPOWER (GP7000F) hosts.
- ◆ Symmetrix attaches to a Fujitsu DS-90/GP7000D host via the Fast WSA Wide Differential SCSI-2 Controller (Model F7958HS1).
- ◆ When using Symmetrix with a Fujitsu DS-90 host in the UXP/DS operating system environment, the Fujitsu host must be running UXP/DS, version V20L20 (patch Y96111) or greater.
- ◆ When using Symmetrix in a Fujitsu GP7000D host environment, the host must be running the UXP/DS operating system, version V20L20 (patch Y97051) or greater.
- ◆ Fujitsu PRIMEPOWER (GP7000F) attaches to Symmetrix via the Ultra Fast Wide SCSI controller available as either Sun Model X6541A-X or Fujitsu Model GP70F-CS02.
- ◆ Fujitsu WSA Fast Wide Differential SCSI controllers support target IDs 0 to 7. Target ID 0 is reserved for use by the WSA controller. Target ID 1 to 7 are available to the user.
- ◆ Fujitsu WSA FWD SCSI controller supports 4 LUNs (0 - 3).
- ◆ For Symmetrix models 55/3500, 52/32xx, and 51/31xx, EMC recommends that you use the Fast-Wide Differential SCSI director P/N 200-881-903.
- ◆ For Symmetrix models 57/3700, 54/34xx, and 53/33xx, EMC recommends that you use either the Fast-Wide Differential SCSI director P/N 201-207-917 (60 MHz processor) or P/N 201-207-927 (66 MHz processor) or Ultra-Wide Differential SCSI director P/N 201-277-917 (60 MHz processor) or P/N 201-277-927 (66 MHz processor).

The tables below summarize the operating system features in the UXP/DS environment.

Table 15-1 Fujitsu DS-90 Operating System

O/S	UXP/DS V20L20 patch Y96111 or greater
Largest recognizable drive size	> 8.7 GB
Largest defined partition	2.15 GB (4,194,240 sectors)
Largest data file	2.07 GB (2,070,077,440 bytes)
Largest filesystem	2.15 GB (4,194,204 blks)
Supports multiple Target IDs	Yes
Supports multiple LUNs per target ID	Yes
Partition representing the entire disk	tbd
To see the drive configuration, use this utility:	prvtoc
To configure drives use this utility:	fmthard
Special drive definition files needed	yes

Table 15-2 Fujitsu GP7000D Operating System

O/S	UXP/DS V20L10 patch Y97051 or greater
Largest recognizable physical drive size	> 8.7 GB
Largest defined partition	1.0 TB (with LVCF software)
Largest data file	1.0 TB (with LVCF software)
Largest filesystem	1.0 TB (with LVCF software)
Supports multiple target IDs	Fast/Wide SCSI-2(1 to 15),others (1 to 7)
Supports multiple LUNs per target ID	Yes (0 to 3)
Partition representing the entire disk	prvtoc
To see the drive configuration, use this utility:	prvtoc
To configure drives use this utility:	fmthard
Special drive definition files needed	yes

Table 15-3 below summarizes the operating system features in the Fujitsu Siemens PRIMEPOWER Solaris environment.

Table 15-3 Fujitsu Siemens PRIMEPOWER Solaris Operating System Features

O/S	Solaris 7 or greater
Largest recognizable drive size	N/A (unlimited)
Largest defined partition	N/A (unlimited)
Largest data file	2 GB (4,194,304 blks)
Largest filesystem	N/A (unlimited)
Supports multiple Target IDs	Yes
Supports multiple LUNs per target ID	Yes
Partition representing the entire disk	s2
To see the drive configuration, use this utility:	format
To configure drives use this utility:	format
Special drive definition files needed	No

Useful Fujitsu Utilities and Functions

This section lists Fujitsu functions and utilities you may use to define and manage Symmetrix devices. The use of these functions and utilities is optional. They are listed for reference only.

- ◆ `fmthard`

The Fujitsu operating system disk format utility. Allows you to format, partition, and label disk drives.

- ◆ `mkfs`

Creates a filesystem.

- ◆ `shutdown -i0 -g0 -y`

Gracefully shutdown the system.

- ◆ `shutdown -i5 -g0 -y`

Shutdown to boot prompt (ROM) mode.

- ◆ `prtconf`

Displays system configuration such as SCSI controller name, slot number, and unit number.

- ◆ `pntdev`

Command to add device attributes such as defining partition number, disk label, etc.

- ◆ `devattr`

Lists device attributes such as number of partitions assigned for a disk.

- ◆ `disks`
Command to create device path for all partitions as as defined in **/etc/device.tab**.
- ◆ `sysadm`
System administration utility to manage/configure system resources.
- ◆ `df`
Command to report number of free disk blocks and files.
- ◆ `prtvto`
Command to print VTOC of a block device.

Symmetrix Configuration

In the UXP/DS host environment, you can configure the Symmetrix disk devices into logical volumes. EMC recommends the following logical volume ratios as indicated in the table below:

Table 15-4 Disk Device Configuration

		Symmetrix Logical Volume to Physical Splits ^a							
		1:1	2:1	3:1	4:1	5:1	6:1	7:1	8:1 ^b
Physical Devices	4 GB	x	x						
	9 GB	x	x	x	x				
	18 GB		x	x	x	x	x	x	x
	23 GB		*	x	x	x	x	x	x
	36 GB			*	x	x	x	x	x
	47 GB			*	*	x	x	x	x
	50 GB ^c				x	x	x	x	x
	73 GB ^{c, d}					x	x	x	x

- a. * indicates ratios available only in microcode level 5265 or higher.
- b. With microcode revision level 5263, maximum is 8:1. With 5264, maximum is 16:1. With 5265, maximum is 32:1.
- c. The 50 and 73 GB drives are supported only by microcode revision levels 5x66 and higher.
- d. The 73 GB drive does not support the RAID-S configuration on Symmetrix.

For specific cylinder counts and logical volume sizes, refer to Table 2-1 on page 2-5 and the following tables in that chapter.

The EMC Customer Engineer (CE) should contact the EMC Configuration Specialist for updated online information. This information is necessary to configure Symmetrix to support the customer's host environment.

Incorporating Symmetrix

The following sections specify the actions necessary to incorporate Symmetrix volumes into the

1. Fujitsu DS-90 or GP7000D host system running the UXP/DS operating system.
2. Fujitsu PRIMEPOWER (GP7000F) running the Solaris operating system.

Fujitsu DS-90 or GP7000D Host

To Recognize LUNs

Symmetrix devices are addressed on the SCSI bus using SCSI target IDs and SCSI LUNs as defined in the **/config/ioconf** file. Both the DS-90 and GP7000D support target IDs 1 to 7. Target ID 0 is reserved for use by the WSA SCSI controller. Each target ID can have up to 4 LUNs (LUN 0-3 only) associated with it. By default, the DS-90 only searches for target IDs and LUN addressing specified in the system kernel on the SCSI bus.

Configuring the System Kernel

Fujitsu requires modifications to the **/config/ioconf** file and a kernel rebuild in order to fully support Symmetrix devices. Perform the following steps to configure the kernel.

1. Edit the **/config/ioconf** file to define the disk name, the controller, the target and LUN in statements similar to the following:

```
cp/config/ioconf /config/ioconf.old
vi /config/ioconf
```

To configure 8 disks on SCSI bus (WSA-0)

```
hd-10 on WSA-0 at id=1 LUN=0
hd-11 on WSA-0 at id=2 LUN=0
hd-12 on WSA-0 at id=3 LUN=0
hd-13 on WSA-0 at id=4 LUN=0
hd-14 on WSA-0 at id=5 LUN=0
hd-15 on WSA-0 at id=5 LUN=1
hd-16 on WSA-0 at id=5 LUN=2
hd-17 on WSA-0 at id=5 LUN=3
```

where WSA-0 is the FWD controller name and unit number.



CAUTION

Do not use TID = 0. TID = 0 is reserved for the host controller. LUNs supported by UXP/DS are LUN 0-3.

2. Rebuild the Fujitsu host kernel by typing statements similar to the following:

```
cd /stand
ls -al system {Note current date/time}
touch system
ls -al system {Note new date/time}
shutdown -y -i6 -g0
```

Upon reboot, the Fujitsu host detects the new date/time in **/stand/system** and will start the kernel rebuild. The Fujitsu host first saves the current **/stand/unix** file to **/stand/OLDunix** file. Then it creates the new **/stand/unix** file. Then the Fujitsu host reboots with the new **/stand/unix** file.

Fujitsu PRIMEPOWER (GP7000F) Host

To Recognize LUNs

Symmetrix devices are addressed on the SCSI bus using SCSI target IDs and SCSI LUNs. Each target ID can have up to 8 LUNs associated with it. By default, Solaris only searches for target IDs on the SCSI bus. If LUN addressing is required in addition to target ID addressing, you will need to modify the file **/kernel/drv/sd.conf** to have the following definitions for each target ID which needs to support multiple LUNs.

```
name="sd" class="scsi"
target=0 lun=0;
name="sd" class="scsi"
target=0 lun=1;
name="sd" class="scsi"
target=0 lun=2;
name="sd" class="scsi"
target=0 lun=3;
name="sd" class="scsi"
target=0 lun=4;
name="sd" class="scsi"
target=0 lun=5;
```

```
name="sd" class="scsi"
target=0 lun=6;
name="sd" class="scsi"
target=0 lun=7;
```

Each target ID can have up to 8 LUNs.

This addition to **sd.conf** is necessary when using LUN addressing. Make sure that for target IDs that do not need LUN support, only the `target=x lun=0` line is specified. This decreases the time needed to reboot the system.

To enable Solaris to support the Symmetrix system, EMC provides Symmetrix device definition files. The device definition files are available on EMC's FTP server, **ftp.emc.com**, in **/pub/symm3000/solaris**.

Appendix A, *Symmetrix Device Definition Files* has instructions on how to transfer these files to your host.

System Specification File Modification

There are three parameters in the file, **/etc/system**, that require modification when operating in a Solaris environment. These parameters are

1. `scsi_options`
2. `sd_io_time`
3. `sd_max_throttle`

This section describes the necessary modifications to the three parameters.

1. For hosts configured with the PCI Ultra-Wide Differential SCSI Controller (X6541A) running Solaris 7 or higher, set the following in the file:

```
set scsi_options = 0x7F8
```

The bit meaning for this parameter is as follows:

Table 15-5 Bit Settings

Bit	Mask	Meaning
3	0x8	Disconnect enable
4	0x10	Link enable
5	0x20	Sync xfer enable
6	0x40	Parity support enable
7	0x80	Tag commands enable
8	0x100	Fast SCSI enable
9	0x200	Wide SCSI enable
10	0x400	Ultra SCSI enable

Since the **/etc/system** file can also affect other disks in the system, consider the effects of implementing the `scsi_options=0x7f8` mask on the other disks. If this mask is not set, Symmetrix operations will not benefit from these options (e.g., Symmetrix will not use Tag Command Queuing).

2. Set `sd_io_time` to 120 seconds. This setting will prevent the host from issuing warning messages while non-disruptive operations are performed on Symmetrix. The entry should read:

```
set sd:sd_io_time = 0x78
```

3. Set `sd_max_throttle` to 20. This setting prevents the host from over-sending tag queuing commands which may cause `scsi cmd timeout` and `scsi bus reset`. The edited entry should read:

```
set sd:sd_max_throttle = 20
```

Format, Partition, Label

Fujitsu DS-90 or GP7000D Host

After Symmetrix VTOC, the drive will have not a VTOC table. The `fmthard` utility can be used to write the VTOC information onto the disk. The VTOC information defines the number of partition, the partition size, the starting and ending blocks. The sample VTOC information file called **/emc/emcdisk.part** is provided in step 3 below.

Use the `fmthard` command to partition and label the new devices. The devices will appear under `/dev`.

To partition and label new devices:

1. Check the **/dev/rdisk/** directory for the "hd10" entry. (Assume you are already logged in as **root**.)

```
ls -lL /dev/rdisk/hd#
ls -lL /dev/dsk/hd*
```

If this is the first time the disk is used, the disks will have no VTOC data.

2. Check if VTOC data is on the disk. (Assume that `hd12` in all commands that follow.)

```
prtvtoc /dev/rdisk/hd12
```

3. Create VTOC tables using the partition definition file for each disk.

```
fmthard -s /emc/emcdisk.part hd12
```

where `/emc/emcdisk.part` contains the following:

```
=====
0 05      02      0          1024      1023
1 12      00      1024      100000     1001023
2 12      12      1001024  100000     2002023
3 12      12      2002024  4194240    6196263
4 12      12      6196264  4194240    10390504
=====
```

Note: Do not modify Partition 0 as it's used for VTOC table.

Column 1: Partition # - UXP/DS OS supports up to 15 partitions

Column 2: tag #
 - 5 = Boot
 - 12 = usrptn

Column 3: flag = 00 = read/write, mountable
 flag = 01 = unmountable
 flag = 02 = read-only

4. Verify that the new VTOC table is on the disk.

```
prtvtoctoc /dev/rdisk/hd12
```

You should see the new partitions.

5. Verify that device attributes were created for each device by the kernel rebuild.

```
devattr -v hd12
```

This will list a number of device attributes. But there will be no attribute **dpartlist** on a new disk. Each device (e.g. hd12) and its attributes are listed in **/etc/device.tab**.

```
putdev -m hd12 dpartlist='hd1201,hd1202,hd1203,hd1204'
```

6. Use the **disks** command to create **/dev/rdisk** and **/dev/dsk** device paths for the partition of each disk.

```
/sbin/disks hd12
```

7. Verify that device paths for the disks partition are created.

```
ls -lL /dev/rdisk/hd12*  
ls -lL /dev/dskhd12*
```

Create a New Filesystem for Each Device

Once you have labeled and partitioned each Symmetrix disk device, create a new filesystem for each Symmetrix disk. To do this, use the **mkfs** command in a statement similar to the following:

```
prtvtoctoc /dev/rdisk/hd12 {determine the size of your partition}  
mkfs -F ufs /dev/rdisk/hd1201 4194240 {in sectors}
```

The “Construct new filesystem?” prompt appears. Press the RETURN key at the prompt.

The actions above created a new filesystem on partition 01 for the Symmetrix disk configured onto the Fujitsu host system under the device name hd1201.

Create the Mount Directory

Repeat these steps for each new device and its partitions. Refer to the kernel configuration files for all defined Symmetrix names.

Once the filesystems for each Symmetrix device are created, create a mount directory for each device. To do this, type a statement similar to the following for each Symmetrix device:

```
mkdir /fs/hd1201
```

Mount the Filesystem

To mount each filesystem, type a statement similar to the following:

```
mount -F ufs /dev/dsk/hd1201 /fs/hd1201
```

Repeat this action for each Symmetrix device.



CAUTION

You may wish to place the Symmetrix devices in the mount table. This requires /etc/vfstab to be edited. /etc/vfstab is syntax-sensitive, and if not edited properly may prevent the system from booting.

Fujitsu PRIMEPOWER (GP7000F) Host

Solaris 7 and Later

Use the `format` command to partition and label the new devices. The devices will appear under `/dev/dsk`.

To partition and label new devices:

1. At the root prompt, type `format`
The host searches for all disks. It generates a display placing all unlabeled disks at the beginning of the listing.
2. At the `Specify Disk` prompt, enter the number of the first EMC drive. The Format menu appears.
3. At the `format>` prompt, type `label`
4. At the `Disk not labeled. Label it now?` prompt, type `y`
5. Type `disk` at the prompt to display a listing of the disks.
6. Repeat steps 2 through 5 for all Symmetrix disks.

If you wish to use raw devices, there is no need to create filesystems for each device. Otherwise, create new filesystems for each device as described in *Create and Mount a File System* on page 15-42.

Format, Partition, Label using sysadm

The DS-90 and GP7000D hosts provide a character-based menu-driven system administration tool, **sysadm**. You can use the **sysadm** tool to format, partition, and label disks instead of using the commands described in the previous section.

Follow these steps to format, partition, and label disks using the **sysadm** utility.

1. Log into the system as the root user and type:

```
sysadm
```

2. At the **UXP/DS System Administration** window, choose **Partition**.
3. At the **Disk Partition** window, press the **F2** key (Choices) to list all the disk drives seen by the host. The **Device Drive Selection** window appears.
4. At the **Device Drive Selection** window, choose the disk drive to partition.



CAUTION

Take care not to select the internal disk drive for format or partitioning as the system may crash.

5. At the **Disk Partition** window, press the **F3** key (Save). The **Partition Selection** window appears.
6. At the **Partition Selection** window, enter the partition information you wish to change.
7. Press the **F8** key (Format) to perform the partitioning. The **sysadm** utility saves the modified partition table to disk.
8. Repeat steps 2 through 7 to format and partition additional disks.
9. Exit from the **sysadm** utility.

Adding Devices Online

Whenever devices are added online to the Symmetrix unit or device channel addresses are changed, you will need to perform the actions described below in order to introduce the new devices to the system.

Fujitsu DS-90 or GP7000D Host

The Fujitsu DS-90 and GP7000D hosts currently do not support adding devices online since the kernel has to be rebuilt and rebooted for new devices to be recognized.

The EMC PSE Configuration Group can add new devices to the Symmetrix configuration while Symmetrix remains online to the host. Once this is complete, stop all host activity, configure the new Symmetrix devices to the host, and follow the instructions in *Incorporating Symmetrix* on page 15-8 and *Format, Partition, Label* on page 15-12 to introduce the new devices.

Fujitsu PRIMEPOWER (GP7000F) Host

To add new Symmetrix devices while online in the Solaris environment, perform the following steps:

1. Use the **online upgrade** feature at the Symmetrix service processor to add/map new drives to the SCSI host channel.
2. Execute the following two utilities:

```
drvconfig  
disks
```

3. Follow the instructions in *Format, Partition, Label* for PRIMEPOWER on page 15-14 to introduce new devices to the host environment.

Solstice DiskSuite

Fujitsu PRIMEPOWER (GP7000F) running the Solaris operating system supports Solstice DiskSuite.

Solstice DiskSuite is a tool for disk and file management. This tool can be used to create and manage logical disks, mirrored or striped volumes, and file systems. DiskSuite supports large file systems, file system expansion and volume manager level intent logging for fast file system recovery.

The following paragraphs provide instructions for using Solstice DiskSuite version 4.0/4.1 (SDS) to create concatenated, striped, and mirrored logical volumes for use with either standard UNIX file systems (UFS) or UFS with intent logging.

See *Intent Logging* on page 15-42 for information on the SDS UFS logging feature

For additional information on using DiskSuite in the Solaris environment refer to *Solstice Disk 4.0/4.1 Administrator's Guide*.

For readability purposes, many of the example lines in this section wrap as indicated by "\" at the end of a line. Your entries should be typed on a single line.

Creating a Concatenated Volume

Use this procedure to add a concatenated metadvice (volume). The example below uses four disks: c1t0d0, c1t0d1, c1t0d2, and c1t0d3. The metadvice is configured to be contained in slice s0 on all devices.

To create a concatenated metadvice comprising slice s0 on four disks, perform the following steps:

1. Edit the **/etc/opt/SUNWmd/md.tab** and add the entry below to define the metadvice:

```
/dev/md/dsk/d0 4 1 /dev/dsk/c1t0d0s0 1 /dev/dsk/c1t0d1s0
1 /dev/dsk/c1t0d2s0 1 /dev/dsk/c1t0d3s0
```

where:

- d0** = name of concatenated metadvice
- 4** = number of device rows to be concatenated
- 1** = number of components in each row

Concatenated Volumes for UFS Logging

2. To create the concatenated metadvice, type

```
metainit d0
```

3. To verify configuration, type

```
metastat  
(or metastat d0)
```

Use this procedure to add a concatenated metadvice (volume) for use with UFS logging. See *Intent Logging* on page 15-42 for a description of the UFS logging feature.

To use UFS logging, you must create a metatrans device. A metatrans device consist of two components (metadevice):

- ◆ Master Metadvice - user data space
- ◆ Logging Metadvice - logging space

You must create one or more slice on one or more disk to be the logging slice. The general guideline of space to allocate for the UFS logging is 10 MB per 1 GB of user space up to a maximum of 64 MB.

The example below uses four disks: c1t0d0, c1t0d1, c1t0d2, and c1t0d3.

- ◆ All disks are on same controller c1
- ◆ Slice s0 is configured to contain the master user data
- ◆ Slice s3 is configured to contain the UFS logging space (16MB/disk)

This example spreads the logging across four disks. You may want to put logging on a single disk to minimize risk of failure when any one of the four volumes fails. Trade-off is a penalty in wasted space on other disks: if disk1 logging is 64 MB, then disk 2, 3, and 4 will each waste 64 MB.

To create a concatenated metadvice for a UFS logging file system comprising slice s0 and s3 on four disks, perform the following steps:

1. Edit the **/etc/opt/SUNWmd/md.tab** and add the entries below to define the metadevices:

```
/dev/md/dsk/d1  4  1  /dev/dsk/c1t0d0s0  1  /dev/dsk/c1t0d1s0  \
                 1  /dev/dsk/c1t0d2s0  1  /dev/dsk/c1t0d3s0
/dev/md/dsk/d2  4  1  /dev/dsk/c1t0d0s3  1  /dev/dsk/c1t0d1s3  \
                 1  /dev/dsk/c1t0d2s3  1  /dev/dsk/c1t0d3s3
/dev/md/dsk/d0 -t  /dev/md/dsk/d1  /dev/md/dsk/d2
```

where:

d0 = name of metatrans device consisting of:
d1 = name of master concatenated metadevice
d2 = name of logging concatenated metadevice
4 = number of device rows to be concatenated
1 = number of components in each row
-t = option to define a metatrans device.

2. To create the metatrans device, type

```
metainit d2
(create the logging device)
```

```
metainit d1
(create the master device)
```

```
metainit d0
(create the metatrans device)
```

3. To verify configuration, type

```
# metastat
(or metastat d0)
```

Creating a Striped Volume

Use this procedure to add a striped metadevice (volume). The example below uses four disks: c1t0d0 (disk1), c2t0d0 (disk2), c1t0d1 (disk3), and c2t0d1 (disk4).

- ◆ The metadevice is configured to be contained in slice s0 on all devices.
- ◆ Stripe width is 64 KB.
- ◆ Striping is across controllers for performance. Trade-off consideration is that a single controller failure will disable the whole volume.
- ◆ Stripe sequence is disk1, disk2, disk3, disk4, then back to disk1 and so on.

To create a striped metadvice comprising slice s0 on four disks, perform the following steps:

1. Edit the `/etc/opt/SUNWmd/md.tab` and add the entry below to define the striped metadvice:

```
/dev/md/dsk/d0 1 4 /dev/dsk/c1t0d0s0 /dev/dsk/c2t0d0s0 \
/dev/dsk/c1t0d1s0 /dev/dsk/c2t0d1s0 -i 64k
```

where:

d0 = name of striped metadvice

1 = number of device rows to be striped

4 = number of components in each row

-i 64k = Specify the stripe width size (8k, 16k [default], 32k, or 64k)

2. To create the striped metadvice, type

```
metainit d0
```

3. To verify configuration

```
metastat
```

(or `metastat d0`)

Creating a Striped Volume for UFS Logging

Use this procedure to add a striped metadvice (volume) for use with UFS logging. See *Intent Logging* on page 15-42 for a description of the UFS logging feature.

To use UFS logging, you must create a metatrans device. A metatrans device consist of two components (metadevices):

- ◆ Master Metadvice - user data space
- ◆ Logging Metadvice - logging space

You must create one or more slice on one or more disk to be the logging slice. The general guideline of space to allocate for the UFS logging is 10 MB per 1 GB of user space up to a maximum of 64 MB.

Logging metadevices should be placed on disks separate from the data disks to avoid hot spots resulting from constant access to the logging area.

The example below uses four disks: c1t0d0 (disk1), c2t0d0 (disk2), c1t0d1 (disk3), and c2t0d1 (disk4).

- ◆ Disks are on two controllers: c1 and c2

- ◆ Slice s0 is configured to contain the master user data on all devices.
- ◆ Slice s3 is configured to contain the UFS logging space (16MB/disk)
- ◆ Stripe width is 64 KB.
- ◆ Striping is across controller for performance. The trade-off is that a single controller failure will disable the whole volume. Stripe sequence is disk1, disk2, disk3, then back to disk1 and so on.

To create a striped metadvice for UFS logging comprising slices s0 and s3 on four disks, perform the following steps:

1. Edit the **/etc/opt/SUNWmd/md.tab** and add the entries below to define the striped metatrans device:

```
/dev/md/dsk/d1  1  4  /dev/dsk/c1t0d0s0  /dev/dsk/c2t0d0s0 \
                  /dev/dsk/c1t0d1s0  /dev/dsk/c2t0d1s0 -i 64k
/dev/md/dsk/d2  1  4  /dev/dsk/c1t0d0s3  /dev/dsk/c2t0d0s3 \
                  /dev/dsk/c1t0d1s3  /dev/dsk/c2t0d1s3 -i 64k
/dev/md/dsk/d0 -t /dev/md/dsk/d1 /dev/md/dsk/d2
```

where:

d0 = name of metatrans device consisting of:

d1 = name of master striped metadvice

d2 = name of logging striped metadvice

1 = number of device rows to be striped

4 = number of components in each row

-i 64k = Specify the stripe width (8k, 16k [default], 32k, 64k)

-t = option to define a metatrans device

2. To create the striped metatrans device, type

```
metainit d2
(create the striped logging device)
```

```
metainit d1
(create the striped master device)
```

```
metainit d0
(create the striped metatrans device)
```

3. To verify configuration, type

```
metastat
(or metastat d0)
```

Creating a Mirrored Striped Volume

Use this procedure to add a mirrored and striped metadvice (volume). The example below uses eight disks:

Submirror 1 - c1t0d0, c1t0d1, c1t0d2, and c1t0d3

Submirror 2 - c2t0d0, c2t0d1, c2t0d2, and c2t0d3

- ◆ Slice s0 is configured to contain data on all devices.
- ◆ Stripe width is 64 KB.
- ◆ Metamirror consists of two submirrors.
Each submirror comprises four disks on the same controller: c1 for submirror 1 and c2 for submirror 2
- ◆ Metamirror is created as a 1-way mirror, then a second mirror is added using the **metattach** command
- ◆ Stripe sequence is disk1, disk2, disk3, disk4, then back to disk1 and so on.

To create a mirrored and striped metamirror device comprising slice s0 on eight disks, perform the following steps:

1. Edit the **/etc/opt/SUNWmd/md.tab** and add the entries below to define the metamirror:

```
/dev/md/dsk/d1  1  4 /dev/dsk/c1t0d0s0 /dev/dsk/c1t0d1s0 \
                  /dev/dsk/c1t0d2s0 /dev/dsk/c1t0d3s0 -i 64k
/dev/md/dsk/d2  1  4 /dev/dsk/c2t0d0s0 /dev/dsk/c2t0d1s0 \
                  /dev/dsk/c2t0d2s0 /dev/dsk/c2t0d3s0 -i 64k
/dev/md/dsk/d0 -m /dev/md/dsk/d1 1
```

where:

- d0** = name of metamirror consisting of:
- d1** = name of striped metadvice (submirror 1)
- d2** = name of striped metadvice (submirror 2)
- 1** = number of device rows to be striped
- 4** = number of components in each row.
- i 64k** = Specify the stripe width (8k, 16k [default], 32k, 64k)
- m** = option to define a metamirror device.
- 1** = Define as a 1-way mirror.

2. To create the striped metatrans device, type

```
metainit d1
(create the striped (submirror 1) metadvice)

metainit d2
(create the striped (submirror 2) metadvice)

metainit d0
```


(create the 1-way Metamirror)

3. To add the second submirror to the metamirror, type

```
metattach /dev/md/dsk/d0 /dev/md/dsk/d2
```

4. To verify configuration, type

```
metastat
```

(or `metastat d0`)

Mirrored Striped Volume for UFS Logging

Use this procedure to add a mirrored and striped metadvice (volume) for use with UFS logging. Refer to *Intent Logging* on page 42 for a description of the UFS logging feature.

To use UFS logging, you must create a metatrans device. A metatrans device consist of two components (metadevice):

- ◆ Master Metadvice - user data space
- ◆ Logging Metadvice - logging space

You must create one or more slice on one or more disk to be the logging slice. The general guideline of space to allocate for the UFS logging is 10 MB per 1 GB of user space up to a maximum of 64 MB.

Logging metadevice should be placed on disks separate from the data disks to avoid hot spots resulting from constant access to the logging area.

The example below uses eight disks:

Submirror 1 - c1t0d0 (disk1), c1t0d1 (disk2), c1t0d2 (disk3), and c1t0d3 (disk4)

Submirror 2 - c2t0d0 (disk1), c2t0d1 (disk2), c2t0d2 (disk3), and c2t0d3 (disk4)

- The metadvice is configured to contain master user data on slice s0 on all devices.
- Slice s3 is configured to contain the UFS logging Space (16MB/disk)
- Stripe width is 64 KB
- Metamirror consists of two submirrors
- Each submirror comprises four disks on two controllers: c1 for submirror 1 and c2 for submirror 2

- Metamirror is created as a 1-way mirror, then a second mirror is added using the `metattach` command
- Stripe sequence is disk1, disk2, disk3, disk4, then back to disk1 and so on.

To create a mirrored and striped metamirror device for use with UFS logging, and comprising slice s0 and s3 on eight disks, perform the following steps:

1. Edit the `/etc/opt/SUNWmd/md.tab` and add the entries below to define the metamirror:

```
/dev/md/dsk/d1  1  4 /dev/dsk/c1t0d0s0 /dev/dsk/c1t0d1s0 \
                  /dev/dsk/c1t0d2s0 /dev/dsk/c1t0d3s0 -i 64k
/dev/md/dsk/d2  1  4 /dev/dsk/c2t0d0s0 /dev/dsk/c2t0d1s0 \
                  /dev/dsk/c2t0d2s0 /dev/dsk/c2t0d3s0 -i 64k
/dev/md/dsk/d3  1  4 /dev/dsk/c1t0d0s3 /dev/dsk/c1t0d1s3 \
                  /dev/dsk/c1t0d2s3 /dev/dsk/c1t0d3s3 -i 64k
/dev/md/dsk/d4  1  4 /dev/dsk/c2t0d0s3 /dev/dsk/c2t0d1s3
                  /dev/dsk/c2t0d2s3 /dev/dsk/c2t0d3s3 -i 64k
/dev/md/dsk/d0 -m /dev/md/dsk/d1 1
/dev/md/dsk/d5 -m /dev/md/dsk/d3 1
/dev/md/dsk/d6 -t /dev/md/dsk/d0 /dev/md/dsk/d5
```

where:

d6 = name of mirrored metatrans device consisting of:

d1 = name of striped master metadvice (master submirror 1)

d2 = name of striped master metadvice (master submirror 2)

d3 = name of striped UFS log metadvice (log submirror 1)

d4 = name of striped UFS log metadvice (log submirror 2)

d0 = name of master metamirror

d5 = name of log metamirror

1 = number of device rows to be striped

4 = number of components in each row

-i 64k = Specify the stripe width (8k, 16k [default], 32k, 64k)

-m = option to define a metamirror device

-t = option to define a metatrans device

2. To create the mirrored metatrans device, type

```
metainit d1
(Create the striped (submirror 1) master)
```

```
metainit d2
(Create the striped (submirror 2) master)
```

```
metainit d3
(Create the striped (submirror 1) logging)

metainit d4
(Create the striped (submirror 2) logging)

metainit d0
(Create the 1-way master metamirror d0)

metainit d5
(Create the 1-way logging metamirror d5)

metattach d0 d2
(Attach second submirror d2 to master metamirror d0)

metattach d5 d4
(Attach second submirror d4 to log metamirror d5)

metainit d6
(Create the metatrans device d6 consisting of master device d0
and logging device d5)
```

3. To verify configuration, type

```
metastat
(or metastat d0)
```

DiskSuite State Database Replicas

The *state database* stores all configuration and status information for DiskSuite objects. Without the state database DiskSuite is unable to access any devices and all data could be lost. Replicas of the database are created so that DiskSuite can compare copies to verify the current configuration and running state of all objects.

It is recommended that at least three replicas be created. If one replica is not available, the remaining two can still be compared to verify configuration and state information. Three replicas can be stored on a system boot disk, however, this creates a single point of failure at the boot disk. Additional replicas should be created on other system disks including Symmetrix devices to protect against the loss of the boot disk. The following considerations apply when planning locations for state database replicas.

- ◆ Always create at least three replicas. DiskSuite will not function if it has only one state database.

- ◆ Replicas can be stored on any unused partition or on partitions that are also part of a metadvice or logging device with the exceptions of **root**, **swap**, **/usr** or file system.
- ◆ Replicas should be spread evenly across host controllers. In Symmetrix, replica devices should be behind separate even and odd SCSI adapters.
- ◆ Store replicas on least two disks on each controller.
- ◆ Each disk should contain only one replica.

Useful DiskSuite Commands

This section lists DiskSuite functions and utilities you may use to define and manage Symmetrix devices. The use of these functions and utilities is optional. They are listed for reference only.

- ◆ `metaclear devicename`

Delete a metadvice. When removing a metadvice hierarchy, use **metaclear** in top down sequence, for example

```
metaclear meta trans
metaclear master_mirror
metaclear log_mirror
metaclear submirrora1
metaclear submirrora2
metaclear submirrorb1
metaclear submirrorb2
```

- ◆ `metadb -i`

Check status of metastate database replicas.

- ◆ `metareplace`

Fix a defective mirrored or RAID device.

- ◆ `metastat`

Display status of mirrored or RAID device.

- ◆ `metahs`

Command to add a hot spare.

Veritas Volume Manager

Fujitsu PRIMEPOWER (GP7000F) running the Solaris operating system supports Veritas Volume Manager and the Veritas file system.

Veritas VxVM and VxFS are tools for disk and file management. VxVM can be used to create logical disks, mirrored and striped volumes. VxFS supports large file systems, file system expansion and a journaling file system.

The following paragraphs provide instructions for using VxVM to create concatenated, striped, and mirrored logical volumes for use with either standard UNIX file systems (UFS) or VxFS journaled file systems.

See *Intent Logging* on page 15-42 for a description of journaling file systems.

For additional information on using Veritas in the Solaris environment refer to *Veritas Volume Manager (VxVM) System Administrator's Guide*.

Initializing Disks

All disks must be initialized before they can be used by VxVM. The initialization process stores configuration and identification data on the disk for use by VxVM.

Disks must be formatted before being initialized by VxVM. Formatting is done outside the volume manager using the **format** command. See *Format, Partition, Label* on page 15-12.

Use the **vxdiskadm** menu-driven interface to initialize disks before attempting to create VxVM volumes. To start **vxdiskadm**, log in as **root** and type

```
vxdiskadm
```

Select **1 Add or Initialize one or more disk**, and follow the system prompts to select devices to be initialized.

Creating a Concatenated Volume

Use this procedure to add a 2 GB concatenated volume of four disks. The configuration uses 500 MB on each 2 GB device.

There are two methods for creating a concatenated volume presented below. The manual `vxmake` method with more control over volume subdisk and plex creation, and the more convenient `vxassist` method with less control (automatic creation and allocation of volume components).

The disk must be initialized before attempting to create volumes (refer to *Initializing Disks* on page 15-27). The initialization process creates slice 02 on the disk. When VxVM puts the disk under its control it must first detect slice 02. VxVM can then create slice 03 as a private region and slice 04 a public region.

Creating a Concatenated Volume using vxmake

Use the manual `vxmake` method to create a concatenated volume as follows:

- To verify that the disk are online and not being used, type

```
vxdisk list
```

Check the displayed list for available disks. For example:

DEVICE	TYPE	DISK	GROUP	STATUS
c0t1d0s2	sliced	-	-	online
c0t3d0s2	sliced	-	-	error
c1t0d0s2	sliced	-	-	error
c1t0d1s2	sliced	-	-	error
c1t0d2s2	sliced	-	-	online
c1t0d3s2	sliced	-	-	online
c1t0d4s2	sliced	-	-	online
c1t0d5s2	sliced	-	-	online
c1t0d6s2	sliced	-	-	online
c1t0d7s2	sliced	-	-	online
* c1t1d0s2	sliced	-	-	online
* c1t1d1s2	sliced	-	-	online
* c1t1d2s2	sliced	-	-	online
* c1t1d3s2	sliced	-	-	online
c1t1d4s2	sliced	-	-	online
c1t1d5s2	sliced	-	-	online
c1t1d6s2	sliced	-	-	online
c1t1d7s2	sliced	-	-	online

* Disks used in this example

2. To create disk group named **cust** and add four disks to disk group **cust**, type

```
vxvg init cust cust-00=c1t1d0
vxvg -g cust adddisk cust-01=c1t1d1
vxvg -g cust adddisk cust-02=c1t1d2
vxvg -g cust adddisk cust-03=c1t1d3
```

3. To verify subdisk creation, type

```
vxdisk list
```

Check the displayed list for subdisk and group. For example:

DEVICE	TYPE	DISK	GROUP	STATUS
		.		
		.		
		.		
c1t1d0s2	sliced	cust-00	cust	online
c1t1d1s2	sliced	cust-01	cust	online
c1t1d2s2	sliced	cust-02	cust	online
c1t1d3s2	sliced	cust-03	cust	online

4. To create 500 MB subdisks on each disks, type

```
vxmake -g cust sd cust-sd0 cust-00,0,500m
vxmake -g cust sd cust-sd1 cust-01,0,500m
vxmake -g cust sd cust-sd2 cust-02,0,500m
vxmake -g cust sd cust-sd3 cust-03,0,500m
```

5. To verify configuration database information, type

```
vxprint -st
```

Check the display for correct configuration data.

For example:

Disk group: cust

SD NAME	PLEX	DISK	DISKOFFS	LENGTH	[COL/]OFF	DEVICE	MODE
sd cust-sd0	-	cust-00	0	1024000	-	c1t1d0	ENA
sd cust-sd1	-	cust-01	0	1024000	-	c1t1d1	ENA
sd cust-sd2	-	cust-02	0	1024000	-	c1t1d2	ENA
sd cust-sd3	-	cust-03	0	1024000	-	c1t1d3	ENA

6. To create striped plex named **cust-plex** out of the subdisks, type

```
vxmake -g cust plex cust-plex sd=cust-sd0,cust-sd1, \ cust-sd2,cust-sd3
layout=concat stwidth=64k
```

7. To verify configuration database information, type

```
vxprint -pt
```

Check display for correct configuration data.

For example:

Disk group: cust

PL NAME	VOLUME	KSTATE	STATE	LENGTH	LAYOUT	NCOL/WID	MODE
pl cust-plex	cust-vol	ENABLED	ACTIVE	4096000	STRIPE	4/128	

8. To create volume out of the concatenated plex **cust-plex**, type

```
vxmake -g cust -U fsgen vol cust-vol \ plex=cust-plex
```

9. To verify configuration database information, type

```
vxprint -vt
```

Check display for correct configuration data.

For example:

Disk group: cust

V NAME	USETYPE	KSTATE	STATE	LENGTH	READPOL	PREFPLEX
v cust-vol	fsgen	DISABLED	EMPTY	4096000	ROUND	-

10. To initialize the volume to the CLEAN state, type

```
vxvol -g cust init clean cust-vol
```

11. To start the concatenated volume, type

```
vxvol -g cust start cust-vol
```

12. To verify configuration database information, type

```
vxprint -vt
```

Check display for correct configuration data.

For example:

Disk group: cust

V NAME	USETYPE	KSTATE	STATE	LENGTH	READPOL	PREFPLEX
v cust-vol	fsgen	ENABLED	ACTIVE	4096000	ROUND	-

Creating a Concatenated Volume using vxassist

Creating volumes with **vxassist** is the recommended way of creating volumes since the user does not deal with the complicated commands as seen in the manual **vxmake** method. This example creates 2 GB concatenated volumes **cust-vol1** and **cust-vol2** on disk group **cust**. **vxassist** creates the subdisks and plexes, and initializes and starts the volume automatically.

1. To create disk group named **cust** and add four disks to disk group **cust**, type

```
vx dg init cust cust-00=c1t1d0
vx dg -g cust adddisk cust-01=c1t1d1
vx dg -g cust adddisk cust-02=c1t1d2
vx dg -g cust adddisk cust-03=c1t1d3
```

2. To create concatenated volume named **cust-vol1** with 2.0 GB size using disk in disk group **cust**, type

```
vxassist -g cust -fsgen make cust-vol1 2g \ layout=span
```

3. To verify configuration database information, type

```
vxprint -vt cust-vol1
```

Check display for correct configuration data.

For example:

Disk group: cust

V NAME	USETYPE	KSTATE	STATE	LENGTH	READPOL	PREFPLEX
v cust-vol	fsgen	ENABLED	ACTIVE	4194394	SELECT	-

Creating a Striped Volume

Use this procedure to add a 2 GB striped volume of four disks. The configuration uses 500 MB on each 2 GB device.

There are two methods for creating a striped volume presented below. The manual **vxmake** method with more control over volume subdisk and plex creation, and the more convenient **vxassist** method with less control (automatic creation and allocation of volume components).

Creating a Striped Volume with vxmake

Use the manual **vxmake** method to create a striped volume as follows:

1. To verify that the disk are online and not being used, type

```
vxdisk list
```

Check the displayed list for available disks. For example:

	DEVICE	TYPE	DISK	GROUP	
	c0t1d0s2	sliced	-	-	online
	c0t3d0s2	sliced	-	-	error
*	c1t0d0s2	sliced	-	-	online
*	c1t0d1s2	sliced	-	-	online
	c1t0d2s2	sliced	-	-	online
	c1t0d3s2	sliced	-	-	online
	c1t0d4s2	sliced	-	-	online
	c1t0d5s2	sliced	-	-	online
	c1t0d6s2	sliced	-	-	online
	c1t0d7s2	sliced	-	-	online
*	c2t0d0s2	sliced	-	-	online
*	c2t0d1s2	sliced	-	-	online
	c2t0d2s2	sliced	-	-	online
	c2t0d3s2	sliced	-	-	online
	c2t0d4s2	sliced	-	-	online
	c2t0d5s2	sliced	-	-	online
	c2t0d6s2	sliced	-	-	online
	c2t0d7s2	sliced	-	-	online

* Disks used in this example

2. To create disk group named **cust** and add four disks to disk group **cust**, type

```
vx dg init cust cust-00=c1t0d0
vx dg -g cust adddisk cust-01=c2t0d0
vx dg -g cust adddisk cust-02=c1t0d1
vx dg -g cust adddisk cust-03=c2t0d1
```

3. To verify subdisk creation, type

```
vxdisk list
```

Check the displayed list for subdisk and group. For example:

DEVICE	TYPE	DISK	GROUP	STATUS
			.	
			.	
			.	
c1t0d0s2	sliced	cust	-00	customonline
c2t0d0s2	sliced	cust	-01	customonline
c1t0d1s2	sliced	cust	-02	customonline
c2t0d1s2	sliced	cust	-03	customonline

4. To create 500 MB subdisks on each disks, type

```
vxmake -g cust sd cust-sd0 cust-00,0,500m
vxmake -g cust sd cust-sd1 cust-01,0,500m
vxmake -g cust sd cust-sd2 cust-02,0,500m
vxmake -g cust sd cust-sd3 cust-03,0,500m
```

5. To verify configuration database information, type

```
vxprint -st
```

Check display for correct configuration data.

For example:

Disk group: cust

SD	NAMEPLEXDISK	DISKOFFS	LENGTH	[COL/]OFF	DEVICE	MODE
sd	cust-sd0-cust-000		1024000	-	c1t0d0	ENA
sd	cust-sd1-cust-010		1024000	-	c2t0d0	ENA
sd	cust-sd2-cust-020		1024000	-	c1t0d1	ENA
sd	cust-sd3-cust-030		1024000	-	c2t0d1	ENA

6. To create striped plex named **cust-plex** out of the subdisks, type

```
vmake -g cust plex cust-plex sd=cust-sd0,cust-sd1,cust-sd2,cust-sd3 \
                                layout=stripe stwidth=64k
```

Stripe width values are 1Kbytes to 256 Kbytes. The default value is 64 Kbytes.

7. To verify configuration database information, type

```
vxprint -pt
```

Check display for correct configuration data.

For example:

Disk group: cust

PL NAME	VOLUME	KSTATE	STATE	LENGTH	LAYOUT	NCOL/WID	MODE
pl cust-plex	cust-vol	ENABLED	ACTIVE	4096000	STRIPE	4/128	

8. To create volume **cust-vol** out of the plex **cust-plex**, type

```
vxmake -g cust -Ufsgen vol cust-vol plex=cust-plex
```

9. To verify configuration database information, type

```
vxprint -vt
```

Check display for correct configuration data.

For example:

Disk group: cust

V NAME	USETYPE	KSTATE	STATE	LENGTH	READPOL	PREFPLEX
v cust-vol	fsgen	DISABLED	EMPTY	4096000	ROUND	-

10. To initialize the volume to the CLEAN state, type

```
vxvol -g cust init clean cust-vol
```

11. To start the concatenated volume, type

```
vxvol -g cust start cust-vol
```

12. To verify configuration database information, type

```
vxprint -vt
```

Check display for correct configuration data.

For example:

Disk group: cust

V NAME	USETYPE	KSTATE	STATE	LENGTH	READPOL	PREFPLEX
v cust-vol	fsgen	ENABLED	ACTIVE	4096000	ROUND	-

Creating a Striped Volume with vxassist

Creating volumes with **vxassist** is the recommended way of creating volumes since the user does not deal with the complicated commands as seen in the manual **vxmake** method. This example, creates a 2 GB striped volume named **cust-vol1**. **vxassist** creates the subdisk and plex, and initializes and starts the volume automatically.

1. To create disk group named **cust** and add four disks to disk group **cust**, type

```
vx dg init cust cust-00=c1t0d0
vx dg -g cust adddisk cust-01=c2t0d0
vx dg -g cust adddisk cust-02=c1t0d1
vx dg -g cust adddisk cust-03=c2t0d1
```

2. To create striped volume named **cust-vol1** with 2.0 GB size using disks in disk group **cust**, type

```
vxassist -g cust -Ufsgen make cust-vol1 2g layout=stripe stwidth=64k \
                                         cust-00 cust-01 cust-02 cust-03
```

Stripe width values are 1Kbytes to 256 Kbytes. The default value is 64 Kbytes.

3. To verify configuration database information, type

```
vxprint -vt cust-vol1
```

Check display for correct configuration data. For example:

Disk group: cust

V NAME	USETYPE	KSTATE	STATE	LENGTH	READPOL	PREFPLEX
v cust-vol1	fsgen	ENABLED	ACTIVE	4194394	ROUND	-

Creating a Striped Mirrored Volume

Use this procedure to add a 2 GB striped volume of four disks with a mirror on another four disks. The configuration uses 500 MB on each 2 GB device.

There are two methods for creating a striped volume presented below. The manual **vxmake** method with more control over volume subdisk and plex creation, and the more convenient **vxassist** method with less control (automatic creation and allocation of volume components).

Manual vxmake Method

Use the manual **vxmake** method to create a concatenated volume as follows:

1. To verify that the disk are online and not being used, type

```
vxdisk list
```

Check the displayed list for available disks.
For example:

DEVICE	TYPE	DISK	GROUP	STATUS
c0t1d0s2	sliced	-	-	online
c0t3d0s2	sliced	-	-	error
c1t0d0s2	sliced	-	-	error
c1t0d1s2	sliced	-	-	error
c1t0d2s2	sliced	-	-	online
c1t0d3s2	sliced	-	-	online
c1t0d4s2	sliced	-	-	online
c1t0d5s2	sliced	-	-	online
c1t0d6s2	sliced	-	-	online
c1t0d7s2	sliced	-	-	online
* c1t1d0s2	sliced	-	-	online
* c1t1d1s2	sliced	-	-	online
* c1t1d2s2	sliced	-	-	online
* c1t1d3s2	sliced	-	-	online
c1t1d4s2	sliced	-	-	online
c1t1d5s2	sliced	-	-	online
c1t1d6s2	sliced	-	-	online
c1t1d7s2	sliced	-	-	online
* c2t1d0s2	sliced	-	-	online
* c2t1d1s2	sliced	-	-	online
* c2t1d2s2	sliced	-	-	online
* c2t1d3s2	sliced	-	-	online

* Disks used in this example

2. To create Disk group named **cust** and add eight disks to disk group **cust**, type

```

vxdg init cust cust-00=c1t1d0
vxdg -g cust adddisk cust-01=c1t1d1
vxdg -g cust adddisk cust-02=c1t1d2
vxdg -g cust adddisk cust-03=c1t1d3
vxdg -g cust adddisk cust-04=c2t1d0
vxdg -g cust adddisk cust-05=c2t1d1
vxdg -g cust adddisk cust-06=c2t1d2
vxdg -g cust adddisk cust-07=c2t1d3

```

3. To verify subdisk creation, type

```
vxdisk list
```

Check the displayed list for subdisk and group. For example:

DEVICE	TYPE	DISK	GROUP	STATUS
		.		
		.		
		.		
c1t1d0s2	sliced	cust-00	cust	online
c1t1d1s2	sliced	cust-01	cust	online
c1t1d2s2	sliced	cust-02	cust	online
c1t1d3s2	sliced	cust-03	cust	online
c1t1d4s2	sliced	-	-	online
c2t1d0s2	sliced	cust-04	cust	online
c2t1d1s2	sliced	cust-05	cust	online
c2t1d2s2	sliced	cust-06	cust	online
c2t1d3s2	sliced	cust-07	cust	online

4. To create 500 MB subdisks on each disks, type

```

vxmake -g cust sd cust-sd0 cust-00,0,500m
vxmake -g cust sd cust-sd1 cust-01,0,500m
vxmake -g cust sd cust-sd2 cust-02,0,500m
vxmake -g cust sd cust-sd3 cust-03,0,500m
vxmake -g cust sd cust-sd4 cust-04,0,500m
vxmake -g cust sd cust-sd5 cust-05,0,500m
vxmake -g cust sd cust-sd6 cust-06,0,500m
vxmake -g cust sd cust-sd7 cust-07,0,500m

```

5. To verify configuration database information, type

```
vxprint -st
```

Check display for correct configuration data.
For example:

Disk group: cust

SD NAME	PLEX	DISK	DISKOFFS	LENGTH	[COL/]OFF	DEVICE	MODE
sd cust-sd0	-	cust-00	0	1024000	-	c1t1d0	ENA
sd cust-sd1	-	cust-01	0	1024000	-	c1t1d1	ENA
sd cust-sd2	-	cust-02	0	1024000	-	c1t1d2	ENA
sd cust-sd3	-	cust-03	0	1024000	-	c1t1d3	ENA
sd cust-sd4	-	cust-00	0	1024000	-	c2t1d0	ENA
sd cust-sd5	-	cust-01	0	1024000	-	c2t1d1	ENA
sd cust-sd6	-	cust-02	0	1024000	-	c2t1d2	ENA
sd cust-sd7	-	cust-03	0	1024000	-	c2t1d3	ENA

6. To create striped plexes **cust-plex1** and **cust-plex2** out of the subdisks, type

```
cust-plex1 make up submirror 1, cust-plex2 make up submirror 2

vxmake -g cust plex cust-plex1 sd=cust-sd0,cust-sd1,cust-sd2,cust-sd3\
  layout=stripe stwidth=64k
vxmake -g cust plex cust-plex2 sd=cust-sd4,cust-sd5,cust-sd6,cust-sd7\
  layout=stripe stwidth=64k
```

Stripe width values are 1Kbytes to 256 Kbytes. The default value is 64 Kbytes.

7. To verify configuration database information, type

```
vxprint -pt
```

Check display for correct configuration data.
For example:

Disk group: cust

PL NAME	VOLUME	KSTATE	STATE	LENGTH	LAYOUT	NCOL/WID	MODE
pl cust-plex1-vol		DISABLED	-	4096000	STRIPE	4/128RW	

8. To create mirror volume out of the **cust-plex1** and then attach second **cust-plex2**

```
vxmake -g cust -U fsgen vol cust-mirvol plex=cust-plex1
vxplex -g cust att cust-mirvol cust-plex2
```

9. To verify configuration database information, type

```
vxprint -vt
```

Check display for correct configuration data.

For example:

Disk group: cust

```
V NAME          USETYPE KSTATE  STATE  LENGTH READPOL PREFPLEX
v cust-mirvol fsgen     DISABLED EMPTY  4096000 ROUND -
```

10. To initialize the volume to the CLEAN state, type

```
vxvol -g cust init clean cust-mirvol
```

11. To start the concatenated volume, type

```
vxvol -g cust start cust-mirvol
```

12. To verify configuration database information, type

```
vxprint -vt
```

Check display for correct configuration data.

For example:

Disk group: cust

```
V NAME          USETYPE KSTATE  STATE  LENGTH READPOL PREFPLEX
v cust-mirvol fsgen     ENABLED  ACTIVE  4096000 ROUND -
```

13. To add two copies of dirty region logging (DRL) to the mirror volume, type

```
vxassist -g cust addlog cust-mirvol
(Add first log)
```

```
vxassist -g cust addlog cust-mirvol
(Add second log)
```

Dirty region logging (DRL) is a bit map mechanism used to facilitate mirror resynchronization after a system failure. Vxassist is the easy way to add a DRL log to the volume. Vxassist creates a subdisk in the diskgroup; its size based on the volume size (1 KB for every 2 GB). The log for this volume may be 4 sectors or 2 Kbytes (1 KB for each 2 GB mirror). When possible, the logging subdisk is created on a disk separate from the data volume. Vxassist creates a plex out of the subdisk and attaches the log plex to the mirrored volume.

Creating Mirrored Volume with vxassist

Creating volumes with **vxassist** is the recommended way of creating volumes since the user does not deal with the complicated commands as seen in the manual **vxmake** method. This example creates a mirrored volume named **cust-mirvol1** with four disks on controller c1 and four disks on controller c2. **vxassist** creates the subdisk and plexes, and initializes and starts the volume automatically.

1. To create disk group named **cust** and add eight disks to disk group **cust**, type

```
vxdbg init cust cust-00=c1t1d0
vxdbg -g cust adddisk cust-01=c1t1d1
vxdbg -g cust adddisk cust-02=c1t1d2
vxdbg -g cust adddisk cust-03=c1t1d3
vxdbg -g cust adddisk cust-04=c2t1d0
vxdbg -g cust adddisk cust-05=c2t1d1
vxdbg -g cust adddisk cust-06=c2t1d2
vxdbg -g cust adddisk cust-07=c2t1d3
```

2. To create mirrored volume named **cust-mirvol** with 2.0 GB size using disks **cust-00**, **cust-01**, **cust-02** and **cust-03** (controller c1) in disk group **cust**, type

```
vxassist -g cust -Ufsgen make cust-mirvol1 2g layout=stripe stwidth=64k \
                                         cust-00 cust-01 cust-02 cust-03
```

Stripe width values are 1Kbytes to 256 Kbytes. The default value is 64 Kbytes.

3. To verify configuration database information, type

```
vxprint -vt cust-mirvol
```

Check display for correct configuration data.

For example:

Disk group: cust

```
V NAME          USETYPE KSTATE  STATE  LENGTH READPOL PREFPLEX
v cust-mirvol fsgen    ENABLED  ACTIVE 4194394 ROUND -
```

4. To create submirror for mirrored volume **cust-mirvol** of 2.0 GB size with disks **cust-04**, **cust-05**, **cust-06** and **cust-07** on second controller c2, type

```
vxassist -g cust mirror cust-mirvol layout=stripe cust-04 cust-05\
cust-06 cust-07 &
```

The mirror volume automatically synchronizes the second submirror. This can take a while.

5. To add two copies of dirty region logging (DRL) to the mirror volume, type

```
vxassist -g cust addlog cust-mirvol1
(Add first log)
```

```
vxassist -g cust addlog cust-mirvol1
(Add second log)
```

Dirty region logging (DRL) is a bit map mechanism used to facilitate mirror resynchronization after a system failure. Vxassist is the easy way to add a DRL log to the volume. Vxassist creates a subdisk in the diskgroup; its size based on the volume size (1 KB for every 2 GB). The log for this volume may be 4 sectors or 2 Kbytes (1 KB for each 2 GB mirror). When possible, the logging subdisk is created on a disk separate from the data volume. Vxassist creates a plex out of the subdisk and attaches the log plex to the mirrored volume.

Create and Mount a File System

Volumes created and managed by Solstice DiskSuite or VxVM volume manager may be used as raw devices, with the standard UNIX file system, or with the VxFS journaled file system.

Intent Logging

Intent logging records pending changes to the file system structure in an intent log. The intent log is replayed during system failure recovery to complete or abandon changes to the structure that were pending at the time of system failure. The file system can then be mounted without completing a full structural check (fsck). An intent logging system can significantly reduce recovery time following a system failure.

Intent logging may be preformed at the file system level or at the volume manager level. Intent logging at the file system level, generally known as a journaling file system (JFS), is usually more effective than volume manager intent logging.

VxFS and VxVM Intent Logging

The Veritas File System (VxFS) includes a journaling file system that provides intent logging at the file system level. Pending changes to the file system, written to an intent log, are scanned during recovery from a system failure. Changes that were active at the time of failure are completed, and the file system is mounted without requiring a fsck of the entire file system. File system recovery is done in a few seconds; much faster than a standard UFS recovery that requires a complete fsck.

VxVM Dirty Region Logging (DRL) provides intent logging at the volume manager level to reduce the time required to resynchronize mirrored volumes after a system failure. DRL applies to VxVM mirrored volumes only. Striped or concatenated volumes do not use intent logging, but may rely on VxFS for fast recovery.

SDS Intent Logging

Solstice DiskSuite (SDS) uses a standard UNIX file system (UFS) that does not provide intent logging at the file system level. However, the VxFS Veritas file system may be used in the SDS environment to provide a journaling file system for SDS volumes.

At the volume manger level, SDS uses a method called the *UFS logging feature* to provide intent logging for all volume types (striped, concatenated, mirrored, and RAID). UFS logging is not a journaling file system. It uses the standard UFS and does intent logging at the volume manager level.

Procedures for creating volumes with UFS logging are provided in *Solstice DiskSuite* on page 15-17.

Procedures for creating mirrored volumes with DRL are provided in *Veritas Volume Manager* on page 15-27.

The following paragraphs outline the steps required to add, create and mount a standard UNIX file system or a VxFS journaled file system for volumes and raw devices.



CAUTION

You may wish to place the Symmetrix devices in the mount table. This requires `/etc/vfstab` to be edited. `/etc/vfstab` is syntax-sensitive, and if not edited properly may prevent the system from booting.

UFS File System on Raw Device

Create New File System

To create standard UNIX file systems under Solaris OS, log in as **root** and proceed as follows for each new device.

Once you have formatted, partitioned, and labeled each Symmetrix disk device, create a new file system for each Symmetrix disk. To do this, use the **newfs** command in a statement similar to the following:

```
newfs -v /dev/rdisk/c1t0d0s0
```

The **Construct new filesystem?** prompt appears. Type **y** at the prompt.

The actions above created a new file system for the Symmetrix disk connected to SCSI controller 1, target ID 0, lun 0, partition 0.

Create Mount Directory

Once the file systems for each Symmetrix device are created, create a mount directory for each device. To do this, type a statement similar to the following for each Symmetrix device:

```
mkdir /fs/c1t0d0s0
```

where **/fs/c1t0d0s0** is the complete path for the new file system directory.

The **/fs** directory must exist prior to creating the mount directories.

Mount the Filesystem

To mount each filesystem, type a statement similar to the following:

```
mount /dev/dsk/c1t0d0s0 /fs/c1t0d0s0
```

VxFS File System on Raw Device

A VxFS journaling file system is created using the **mkfs** command with arguments provided for block size, log size, device name and size. To create a VxFS file system first determine the size in sectors of the volume. The size of the volume is displayed under the **Sector Count** field of the **prtvtoc** output.

Once you have formatted, partitioned, and labeled each Symmetrix disk device, create a new file system for each Symmetrix disk. To create VxVM journaling file systems under VxFS and Solaris OS, log in as **root** and proceed as follows for each new device.

Create New File System

1. To display sector count information, type

```
prtvtoc /dev/rdisk/clt0d0s0
```

The size of the disk is displayed under the **Sector Count** field.
(Assume 4099000 for this example.)

2. To create the VxFS file system for the volume, type

```
mkfs -F vxfs -o bsize=4096,logsize=256 \ /dev/rdisk/clt0d0s0 4099000
```

where:

bsize = block size in bytes (1k, 2k, 4k, or 8k - 1k default for filesystems < 4 GB, 4k default for filesystems > 4 GB)

logsize = size of VxFS file system logging in blocks (256 blocks default, 32 to 1024 blocks)

4099000 = file system size in sectors (from **prtvtoc** command)

Create Mount Directory

Once the file systems for each Symmetrix device are created, create a mount directory for each device. To do this, type a statement similar to the following for each Symmetrix device:

```
mkdir /fs/clt0d0s0
```

where **/fs/clt0d0s0** is the complete path for the new file system directory.

Mount the Filesystem

To mount each filesystem, type a statement similar to the following:

```
mount -F vxfs /dev/dsk/clt0d0s0 /fs/clt0d0s0
```

UFS File System on SDS Device

Create New File System

To create standard UNIX file systems under Solstice DiskSuite, log in as **root** and proceed as follows for each new device.

Once you have formatted, partitioned, and labeled each Symmetrix disk device, and created SDS volumes, create a new file system on each volume. To do this, use the **newfs** command in a statement similar to the following:

```
newfs -v /dev/md/rdisk/d0
```

The Construct new filesystem? prompt appears. Type **Y** at the prompt.

The actions above created a new file system for the Symmetrix disk defined as SDS metadvice **d0**.

Create Mount Directory

Once the file systems for each device are created, create a mount directory for each device. To do this, type a statement similar to the following for each device:

```
mkdir /fs/d0
```

where **/fs/d0** is the complete path for the new file system directory.

You can assume that the **/fs** directory existed prior to creating the **d0** directory.

Mount the Filesystem

To mount each filesystem, type a statement similar to the following:

```
mount /dev/md/dsk/d0 /fs/d0
```

VxFS File System on SDS Device

A VxFS journaling file system is created using the **mkfs** command with arguments provided for block size, log size, device name and size. To create a VxFS file system first determine the size in sectors of the volume. The size of the volume is displayed under the **Sector Count** field of the **prtvtoc** output.

To create VxFS journaling file systems under Solstice DiskSuite, log in as **root** and proceed as follows for each new device.

Create New File System

1. To display sector count information, type

```
prtvtoc /dev/md/rdisk/d0
```

The size of the disk is displayed under the **Sector Count** field.
(Assume 4099000 for this example.)

2. To create the VxFS file system for the volume, type

```
mkfs -F vxfs -o bsize=1024 logsize=512\ /dev/md/rdisk/d0 4099000
```

Where:

bsize = block size in bytes (1k, 2k, 4k, or 8k - 1k default for filesystems < 4 GB, 4k default for filesystems > 4 GB)

logsize = size of VxFS file system logging in blocks (256 blocks default, 32 to 1024 blocks)

4099000 = file system size in sectors (from **prtvtoc** command)

Create Mount Directory

Once the file systems for each device are created, create a mount directory for each device. To do this, type a statement similar to the following for each device:

```
mkdir /fs/d0
```

where **/fs/d0** is the complete path for the new file system directory.

Mount the Filesystem

To mount each filesystem, type a statement similar to the following:

```
mount -F vxfs /dev/md/dsk/d0 /fs/d0
```


UFS File System on VxVM Device

Create New File System

To create standard UNIX file systems under Veritas Volume Manager, log in as **root** and proceed as follows for each new device.

Once you have formatted, partitioned, and labeled each Symmetrix disk device, and created VxVM volumes, create a new file system on each volume. To do this, use the **newfs** command in a statement similar to the following:

```
newfs -v /dev/vx/rdisk/dskgrp/vol1
```

The **Construct new filesystem?** prompt appears. Type **y** at the prompt.

The actions above created a new file system for the Symmetrix disk defined as **vol1** in the diskgroup named **dskgrp**.

Create Mount Directory

Once the file systems for each device are created, create a mount directory for each device. To do this, type a statement similar to the following for each device:

```
mkdir /fs/vol1
```

where **/fs/vol1** is the complete path for the new file system directory.

Mount the Filesystem

To mount each filesystem, type a statement similar to the following:

```
mount -F vsfs /dev/vx/dsk/dskgrp/vol1 /fs/vol1
```

VxFS File System on VxVM Device

A VxFS journaling file system is created using the **mkfs** command with arguments provided for block size, log size, device name and size. To create a VxFS file system first determine the size in sectors of the volume. The size of the volume is displayed under the **Sector Count** field of the **prvtoc** output.

To create VxFS journaling file systems under Veritas Volume Manager, log in as **root** and proceed as follows for each new device.

Create New File System

1. To display sector count information, type

```
prvtoc /dev/vx/rdisk/dskgrp/vol1
```

The size of the disk is displayed under the **Sector Count** field.
(Assume 4099000 for this example.)

2. To create the VxFS file system for the volume, type

```
mkfs -F vxfs -o bsize=4096 logsize=512 /dev/vx/rdisk/dskgrp/vol1 4099000
```

Where:

bsize = block size in bytes (1k, 2k, 4k, or 8k - 1k default for filesystems < 4 GB, 4k default for filesystems > 4 GB)

logsize = size of VxFS file system logging in blocks (256 blocks default, 32 to 1024 blocks)

4099000 = file system size in sectors (from **prvtoc** command)

Create Mount Directory

Once the file systems for each device are created, create a mount directory for each device. To do this, type a statement similar to the following for each device:

```
mkdir /fs/vol1
```

Where **/fs/vol1** is the complete path for the new file system directory.

Mount the Filesystem

To mount each filesystem, type a statement similar to the following:

```
mount -F vxfs /dev/vx/dsk/dskgrp/vol1 /fs/vol1
```

File System Expansion

DiskSuite volume manager allows you to increase the available storage space of an existing volume by concatenating additional volumes to the metadvice using **metattach**. A file system can then be expanded to fill all or part of the additional space using the **growfs** command. The file system can remain mounted, but will be locked (**lockfs**) during the expansion. For detailed information on expanding a file system under DiskSuite refer to *Solstice DiskSuite 4.0 /4.1 Administrator's Guide*.

VxVM allows *growing* of a mounted file system using **vxassist** commands. The following steps outline the procedure.

1. Log in as **root**.
2. To determine how large the volume can grow, type

```
vxassist maxgrow vol
```

Where **vol** is the volume name

Volume vol can be extended by 12533760 to 16629760 (8120Mb)

3. Determine the size of the current volume, type

```
vxprint -vt
```

Disk group: cust

V	NAME	USETYPE	KSTATE	STATE	LENGTH	READPOL	PREFPLEX
v	vol	fsgen	ENABLED	ACTIVE	4096000	ROUND	-

The length unit is in sectors (1 sectors = 512 bytes). Therefore **vol** is approximately 2 GB.

4. To grow the volume **vol** to 4 GB, type

```
vxassist growto cust-mirvol 4g
vxprint -vt
```

Disk group: cust

V	NAME	USETYPE	KSTATE	STATE	LENGTH	READPOL	PREFPLEX
v	cust-mirvol	fsgen	ENABLED	ACTIVE	8388608	ROUND	-

System and Error Messages

The Fujitsu DS-90 and GP7000D hosts log system and error messages including SCSI error messages to a file called `/var/adm/syslog.conf`. The hosts also display these messages at the system console.

The Fujitsu PRIMEPOWER (GP7000F) host logs system and error messages, including SCSI and Fibre Channel error messages, to a file called `/var/adm/messages` and also displays these messages at the system console.

This chapter describes how to incorporate Symmetrix in the Fujitsu Siemens PRIMEPOWER (formerly GP7000F) environment.

◆ Overview	16-2
◆ Useful Solaris Utilities and Functions	16-4
◆ Hardware Mapping	16-5
◆ Symmetrix Configuration	16-6
◆ Incorporating Symmetrix	16-8
◆ Format, Partition, Label	16-11
◆ Adding Devices Online	16-12
◆ Solstice DiskSuite	16-13
◆ Veritas Volume Manager	16-23
◆ Create and Mount a File System	16-38
◆ System and Error Messages	16-46

Overview

This chapter covers the Fujitsu Siemens PRIMEPOWER (GP7000F) host in the Solaris environment. See Chapter 16, *Fujitsu Siemens PRIMEPOWER (GP7000F) Environment*, for the GP7000D and earlier hosts. Refer to EMC's *Symmetrix Fibre Channel Product Guide* P/N 200-999-642 for a description of the the Fujitsu Siemens PRIMEPOWER host in the Fibre Channel environment.

When using Symmetrix in the Fujitsu Siemens environment, please note the following:

- ◆ For Symmetrix models 57/3700, 54/34xx, and 53/33xx, EMC recommends that you use the Fast-Wide Differential SCSI director P/N 201-207-917 (60 MHz processor), P/N 201-207-927 (66 MHz processor), Ultra-Wide Differential SCSI director P/N 201-277-917 (60 MHz processor), P/N 201-277-927 (66 MHz processor), or Fibre Channel director P/N 200-520-910 or P/N 200-520-920.
- ◆ Symmetrix attaches to a SCSI Controller (Model X6541A Fujitsu Siemens P/N GP70F-CS02) or to an Emulex LP8000 (Fujitsu Siemens P/N GP70F-CF10).
- ◆ The Fujitsu Siemens PRIMEPOWER host must have Ultra Wide SCSI Differential or Fibre Channel cards installed in order to connect to the Symmetrix unit.
- ◆ When using Symmetrix with the Fujitsu Siemens PRIMEPOWER host in the Solaris environment, the host must be running Solaris 2.6 or greater.

Table 16-1 below summarizes the operating system features in the Fujitsu Siemens PRIMEPOWER Solaris environment.

Table 16-1 Fujitsu Siemens PRIMEPOWER Solaris Operating System Features

O/S	Solaris 7 or greater
Largest recognizable drive size	N/A (unlimited)
Largest defined partition	N/A (unlimited)
Largest data file	2 GB (4,194,304 blks)
Largest filesystem	N/A (unlimited)
Supports multiple Target IDs	Yes
Supports multiple LUNs per target ID	Yes
Partition representing the entire disk	s2
To see the drive configuration, use this utility:	format
To configure drives use this utility:	format
Special drive definition files needed	No

Useful Solaris Utilities and Functions

This section lists Solaris functions and utilities you may use to define and manage Symmetrix devices. The use of these functions and utilities is optional. They are listed for reference only.

◆ `format`

The Solaris disk format utility. Allows you to format, partition, and label disk drives.

◆ `newfs`

Creates a filesystem.

◆ `shutdown`

Gracefully shutdown the system.

shutdown is the preferred command for system shutdown. The **halt** command is not recommend.

Hardware Mapping

Use the following procedure to determine the target IDs available to the Symmetrix unit.

If the host is not powered:

1. Power on the Fujitsu Siemens PRIMEPOWER host system and wait for it to perform its self-tests.
2. At the **Open Firmware** prompt (**OK>**), type `printenv`
3. Review the listing that appears and determine the value of **scsi_initiator_id**. This target ID is reserved for the system and may not be used by any other device.

If the host is powered up:

1. Type `eeeprom` at the system prompt.
2. Review the listing that appears and determine the value of **scsi_initiator_id**. This target ID is reserved for the system and may not be used by any other device.

Do NOT use the **scsi_initiator_id** value as a target ID value for Symmetrix devices.

Symmetrix Configuration

In the Solaris host environment, you can configure the Symmetrix disk devices into logical volumes. EMC recommends the following logical volume ratios as indicated in the table below:

Table 16-2 Disk Device Configuration

		Symmetrix Logical Volume to Physical Splits ^a							
		1:1	2:1	3:1	4:1	5:1	6:1	7:1	8:1 ^b
Physical Devices	4 GB	x	x						
	9 GB	x	x	x	x				
	18 GB		x	x	x	x	x	x	x
	23 GB		*	x	x	x	x	x	x
	36 GB			*	x	x	x	x	x
	47 GB			*	*	x	x	x	x
	50 GB ^c				x	x	x	x	x
	73 GB ^{c, d}					x	x	x	x

- a. * indicates ratios available only in microcode level 5265 or higher.
- b. With microcode revision level 5263, maximum is 8:1. With 5264, maximum is 16:1. With 5265, maximum is 32:1.
- c. The 50 and 73 GB drives are supported only by microcode revision levels 5x66 and higher.
- d. The 73 GB drive does not support the RAID-S configuration on Symmetrix.

For specific cylinder counts and logical volume sizes, refer to Table 2-1 on page 2-5 and the following tables in that chapter.

The EMC Customer Engineer (CE) should contact the EMC Configuration Specialist for updated online information. This information is necessary to configure Symmetrix to support the customer's host environment.

PRIMEPOWER (GP7000F) Servers

Perform this procedure prior to booting the operating system when running Solaris on Fujitsu Siemens PRIMEPOWER servers.

1. Power on the server and wait for it to perform its self-tests.
2. When the system banner appears on the screen display, press <BREAK> if you are using an ASCII terminal or <STOP> <A> if you are using a workstation terminal.
3. At the **OK>** prompt, type `probe-scsi-all`

The host looks for all attached SCSI devices and displays information similar to the following (Units 0-3 are LUNs):

```
Target 0
Unit 0 Disk EMC SYMMETRIX 50607801E153
      Copyright (c) 1994
      EMC Corp. All rights reserved.
Unit 1 Disk EMC SYMMETRIX 50607801D153
      Copyright (c) 1994
      EMC Corp. All rights reserved.
Unit 2 Disk EMC SYMMETRIX 50607801F160
      Copyright (c) 1994
      EMC Corp. All rights reserved.
Unit 3 Disk EMC SYMMETRIX 506078020161
      Copyright (c) 1994
      EMC Corp. All rights reserved.
```

This disk value may change based on microcode rev in use

The disk value returned (e.g., 50607810E153) provides additional detail. The syntax of this number is MMMMWXXYYZ, where:

MMMM = Symmetrix microcode revision level

WW = Last two digits of the Symmetrix serial number

XXX = Symmetrix device number

YY = Symmetrix SCSI Director number

Z = Symmetrix SCSI Direct port number

4. Make sure the target ID and LUNs listed agree with the device information configured in the Symmetrix unit.

If this information does not agree, check the Symmetrix-host port connection and the AutoInstall configuration.

5. At the **OK>** prompt, type `boot -r`

The host will boot to a login prompt. The host will also issue a warning message for each Symmetrix logical unit attached but not yet labeled.

Incorporating Symmetrix

The following sections specify the actions necessary to incorporate Symmetrix volumes into a PRIMEPOWER system.

To Recognize LUNs

Symmetrix devices are addressed on the SCSI bus using SCSI target IDs and SCSI LUNs. Each target ID can have up to 8 LUNs associated with it. By default, Solaris only searches for target IDs on the SCSI bus. If LUN addressing is required in addition to target ID addressing, you will need to modify the file `/kernel/drv/sd.conf` to have the following definitions for each target ID which needs to support multiple LUNs.

```
name="sd" class="scsi"
  target=0 lun=0;
name="sd" class="scsi"
  target=0 lun=1;
name="sd" class="scsi"
  target=0 lun=2;
name="sd" class="scsi"
  target=0 lun=3;
name="sd" class="scsi"
  target=0 lun=4;
name="sd" class="scsi"
  target=0 lun=5;
name="sd" class="scsi"
  target=0 lun=6;
name="sd" class="scsi"
  target=0 lun=7;
```

Each target ID can have up to 8 LUNs.

This addition to `sd.conf` is necessary when using LUN addressing. Make sure that for target IDs that do not need LUN support, only the `target=x lun=0` line is specified. This decreases the time needed to reboot the system.

To enable Solaris to support the Symmetrix system, EMC provides Symmetrix device definition files. The device definition files are available on EMC's FTP server, <ftp.emc.com>, in `/pub/symm3000/solaris`.

Appendix A, *Symmetrix Device Definition Files* has instructions on how to transfer these files to your host.

System Specification File Modification

There are three parameters in the file, `/etc/system`, that require modification when operating in a Solaris environment. These parameters are

1. `scsi_options`
2. `sd_io_time`
3. `sd_max_throttle`

This section describes the necessary modifications to the three parameters.

1. For hosts configured with the PCI Ultra-Wide Differential SCSI Controller (X6541A) running Solaris 7 or higher, set the following in the file:

```
set scsi_options = 0x7F8
```

The bit meaning for this parameter is as follows:

Table 16-3 Bit Settings

Bit	Mask	Meaning
3	0x8	Disconnect enable
4	0x10	Link enable
5	0x20	Sync xfer enable
6	0x40	Parity support enable
7	0x80	Tag commands enable
8	0x100	Fast SCSI enable
9	0x200	Wide SCSI enable
10	0x400	Ultra SCSI enable

Since the `/etc/system` file can also affect other disks in the system, consider the effects of implementing the `scsi_options=0x7f8` mask on the other disks. If this mask is not set, Symmetrix operations will not benefit from these options (e.g., Symmetrix will not use Tag Command Queuing).

2. Set `sd_io_time` to 120 seconds. This setting will prevent the host from issuing warning messages while non-disruptive operations are performed on Symmetrix. The entry should read:

```
set sd:sd_io_time = 0x78
```

3. Set `sd_max_throttle` to 20. This setting prevents the host from over-sending tag queuing commands which may cause `scsi cmd timeout` and `scsi bus reset`. The edited entry should read:

```
set sd:sd_max_throttle = 20
```

Format, Partition, Label

Use the `format` command to partition and label the new devices. The devices will appear under `/dev/dsk`.

Solaris 7 and Later

To partition and label new devices:

1. At the **root** prompt, type `format`
The host searches for all disks. It generates a display placing all unlabeled disks at the beginning of the listing.
2. At the **Specify Disk** prompt, enter the number of the first EMC drive. The Format menu appears.
3. At the **format>** prompt, type `label`
4. At the **Disk not labeled. Label it now?** prompt, type `y`
5. Type `disk` at the prompt to display a listing of the disks.
6. Repeat steps 2 through 5 for all Symmetrix disks.

If you wish to use raw devices, there is no need to create filesystems for each device. Otherwise, create new filesystems for each device as described in *Create and Mount a File System* on page 16-38.

Adding Devices Online

Whenever devices are added online to the Symmetrix unit or device channel addresses are changed, you will need to perform the actions described below in order to introduce the new devices to the system.

To add new Symmetrix devices while online in the Solaris environment, perform the following steps:

1. Use the **online upgrade** feature at the Symmetrix service processor to add/map new drives to the SCSI host channel.
2. Execute the following two utilities:

```
drvconfig  
disks
```
3. Follow the instructions in *Format, Partition, Label* on page 16-11 to introduce new devices to the host environment.

Solstice DiskSuite

Solstice DiskSuite is a tool for disk and file management. This tool can be used to create and manage logical disks, mirrored or striped volumes, and file systems. DiskSuite supports large file systems, file system expansion and volume manager level intent logging for fast file system recovery.

The following paragraphs provide instructions for using Solstice DiskSuite version 4.0/4.1 (SDS) to create concatenated, striped, and mirrored logical volumes for use with either standard UNIX file systems (UFS) or UFS with intent logging.

See *Intent Logging* on page 16-38 for information on the SDS UFS logging feature.

For additional information on using DiskSuite in the Solaris environment refer to *Solstice Disk 4.0/4.1 Administrator's Guide*.

For readability purposes, many of the example lines in this section wrap as indicated by "\" at the end of a line. Your entries should be typed on a single line.

Creating a Concatenated Volume

Use this procedure to add a concatenated metadevice (volume). The example below uses four disks: c1t0d0, c1t0d1, c1t0d2, and c1t0d3. The metadevice is configured to be contained in slice s0 on all devices.

To create a concatenated metadevice comprising slice s0 on four disks, perform the following steps:

1. Edit the `/etc/opt/SUNWmd/md.tab` and add the entry below to define the metadevice:

```
/dev/md/dsk/d0 4 1 /dev/dsk/c1t0d0s0 1 /dev/dsk/c1t0d1s0
1 /dev/dsk/c1t0d2s0 1 /dev/dsk/c1t0d3s0
```

where:

d0 = name of concatenated metadevice
4 = number of device rows to be concatenated
1 = number of components in each row

2. To create the concatenated metadevice, type

```
metainit d0
```

Concatenated Volumes for UFS Logging

3. To verify configuration, type

```
metastat
(or metastat d0)
```

Use this procedure to add a concatenated metadvice (volume) for use with UFS logging. See *Intent Logging* on page 16-38 for a description of the UFS logging feature.

To use UFS logging, you must create a metatrans device. A metatrans device consist of two components (metadevices):

- ◆ Master Metadvice - user data space
- ◆ Logging Metadvice - logging space

You must create one or more slice on one or more disk to be the logging slice. The general guideline of space to allocate for the UFS logging is 10 MB per 1 GB of user space up to a maximum of 64 MB.

The example below uses four disks: c1t0d0, c1t0d1, c1t0d2, and c1t0d3.

- ◆ All disks are on same controller c1
- ◆ Slice s0 is configured to contain the master user data
- ◆ Slice s3 is configured to contain the UFS logging space (16MB/disk)

This example spreads the logging across four disks. You may want to put logging on a single disk to minimize risk of failure when any one of the four volumes fails. Trade-off is a penalty in wasted space on other disks: if disk1 logging is 64 MB, then disk 2, 3, and 4 will each waste 64 MB.

To create a concatenated metadvice for a UFS logging file system comprising slice s0 and s3 on four disks, perform the following steps:

1. Edit the **/etc/opt/SUNWmd/md.tab** and add the entries below to define the metadevices:

```
dev/md/dsk/d1  4  1 /dev/dsk/c1t0d0s0 1 /dev/dsk/c1t0d1s0 \
                1 /dev/dsk/c1t0d2s0 1 /dev/dsk/c1t0d3s0
/dev/md/dsk/d2  4  1 /dev/dsk/c1t0d0s3 1 /dev/dsk/c1t0d1s3 \
                1 /dev/dsk/c1t0d2s3 1 /dev/dsk/c1t0d3s3
/dev/md/dsk/d0 -t /dev/md/dsk/d1 /dev/md/dsk/d2
```

where:

d0 = name of metatrans device consisting of:
d1 = name of master concatenated metadvice
d2 = name of logging concatenated metadvice
4 = number of device rows to be concatenated
1 = number of components in each row
-t = option to define a metatrans device.

2. To create the metatrans device, type

```
metainit d2
(create the logging device)

metainit d1
(create the master device)

metainit d0
(create the metatrans device)
```

3. To verify configuration, type

```
# metastat
(or metastat d0)
```

Creating a Striped Volume

Use this procedure to add a striped metadvice (volume). The example below uses four disks: c1t0d0 (disk1), c2t0d0 (disk2), c1t0d1 (disk3), and c2t0d1 (disk4).

- ◆ The metadvice is configured to be contained in slice s0 on all devices.
- ◆ Stripe width is 64 KB.
- ◆ Striping is across controllers for performance. Trade-off consideration is that a single controller failure will disable the whole volume.
- ◆ Stripe sequence is disk1, disk2, disk3, disk4, then back to disk1 and so on.

To create a striped metadvice comprising slice s0 on four disks, perform the following steps:

1. Edit the **/etc/opt/SUNWmd/md.tab** and add the entry below to define the striped metadvice:

```
/dev/md/dsk/d0 1 4 /dev/dsk/c1t0d0s0 /dev/dsk/c2t0d0s0 \
/dev/dsk/c1t0d1s0 /dev/dsk/c2t0d1s0 -i 64k
```

Where:

d0 = name of striped metadvice

1 = number of device rows to be striped

4 = number of components in each row

-i 64k = Specify the stripe width size (8k, 16k [default], 32k, or 64k)

2. To create the striped metadvice, type

```
metainit d0
```

3. To verify configuration

```
metastat
```

(or `metastat d0`)

Creating a Striped Volume for UFS Logging

Use this procedure to add a striped metadvice (volume) for use with UFS logging. See *Intent Logging* on page 16-38 for a description of the UFS logging feature.

To use UFS logging, you must create a metatrans device. A metatrans device consist of two components (metadevices):

- ◆ Master Metadvice - user data space
- ◆ Logging Metadvice - logging space

You must create one or more slice on one or more disk to be the logging slice. The general guideline of space to allocate for the UFS logging is 10 MB per 1 GB of user space up to a maximum of 64 MB.

Logging metadevices should be placed on disks separate from the data disks to avoid hot spots resulting from constant access to the logging area.

The example below uses four disks: c1t0d0 (disk1), c2t0d0 (disk2), c1t0d1 (disk3), and c2t0d1 (disk4).

- ◆ Disks are on two controllers: c1 and c2
- ◆ Slice s0 is configured to contain the master user data on all devices.
- ◆ Slice s3 is configured to contain the UFS logging space (16MB/disk)
- ◆ Stripe width is 64 KB.
- ◆ Striping is across controller for performance. The trade-off is that a single controller failure will disable the whole volume. Stripe sequence is disk1, disk2, disk3, then back to disk1 and so on.

To create a striped metadvice for UFS logging comprising slices s0 and s3 on four disks, perform the following steps:

1. Edit the `/etc/opt/SUNWmd/md.tab` and add the entries below to define the striped metatrans device:

```
/dev/md/dsk/d1  1  4 /dev/dsk/c1t0d0s0 /dev/dsk/c2t0d0s0 \
                  /dev/dsk/c1t0d1s0 /dev/dsk/c2t0d1s0 -i 64k
/dev/md/dsk/d2  1  4 /dev/dsk/c1t0d0s3 /dev/dsk/c2t0d0s3 \
                  /dev/dsk/c1t0d1s3 /dev/dsk/c2t0d1s3 -i 64k
/dev/md/dsk/d0 -t /dev/md/dsk/d1 /dev/md/dsk/d2
```

Where:

d0 = name of metatrans device consisting of:

d1 = name of master striped metadvice

d2 = name of logging striped metadvice

1 = number of device rows to be striped

4 = number of components in each row

-i 64k = Specify the stripe width (8k, 16k [default], 32k, 64k)

-t = option to define a metatrans device

2. To create the striped metatrans device, type

```
metainit d2
(create the striped logging device)
```

```
metainit d1
(create the striped master device)
```

```
metainit d0
(create the striped metatrans device)
```

3. To verify configuration, type

```
metastat
(or metastat d0)
```

Creating a Mirrored Striped Volume

Use this procedure to add a mirrored and striped metadvice (volume). The example below uses eight disks:

Submirror 1 - c1t0d0, c1t0d1, c1t0d2, and c1t0d3

Submirror 2 - c2t0d0, c2t0d1, c2t0d2, and c2t0d3

- ◆ Slice s0 is configured to contain data on all devices.
- ◆ Stripe width is 64 KB.
- ◆ Metamirror consists of two submirrors.
Each submirror comprises four disks on the same controller: c1 for submirror 1 and c2 for submirror 2

- ◆ Metamirror is created as a 1-way mirror, then a second mirror is added using the **metattach** command
- ◆ Stripe sequence is disk1, disk2, disk3, disk4, then back to disk1 and so on.

To create a mirrored and striped metamirror device comprising slice s0 on eight disks, perform the following steps:

1. Edit the **/etc/opt/SUNWmd/md.tab** and add the entries below to define the metamirror:

```
/dev/md/dsk/d1  1  4  /dev/dsk/c1t0d0s0 /dev/dsk/c1t0d1s0 \
                  /dev/dsk/c1t0d2s0 /dev/dsk/c1t0d3s0 -i 64k
/dev/md/dsk/d2  1  4  /dev/dsk/c2t0d0s0 /dev/dsk/c2t0d1s0 \
                  /dev/dsk/c2t0d2s0 /dev/dsk/c2t0d3s0 -i 64k
/dev/md/dsk/d0 -m /dev/md/dsk/d1 1
```

Where:

d0 = name of metamirror consisting of:

d1 = name of striped metadvice (submirror 1)

d2 = name of striped metadvice (submirror 2)

1 = number of device rows to be striped

4 = number of components in each row.

-i 64k = Specify the stripe width (8k, 16k [default], 32k, 64k)

-m = option to define a metamirror device.

1 = Define as a 1-way mirror.

2. To create the striped metatrans device, type

```
metainit d1
(create the striped (submirror 1) metadvice)
```

```
metainit d2
(create the striped (submirror 2) metadvice)
```

```
metainit d0
(create the 1-way Metamirror)
```

3. To add the second submirror to the metamirror, type

```
metattach /dev/md/dsk/d0 /dev/md/dsk/d2
```

4. To verify configuration, type

```
metastat
(or metastat d0)
```

Mirrored Striped Volume for UFS Logging

Use this procedure to add a mirrored and striped metadvice (volume) for use with UFS logging. Refer to *Intent Logging* on page 38 for a description of the UFS logging feature.

To use UFS logging, you must create a metatrans device. A metatrans device consist of two components (metadevices):

- ◆ Master Metadvice - user data space
- ◆ Logging Metadvice - logging space

You must create one or more slice on one or more disk to be the logging slice. The general guideline of space to allocate for the UFS logging is 10 MB per 1 GB of user space up to a maximum of 64 MB.

Logging metadevices should be placed on disks separate from the data disks to avoid hot spots resulting from constant access to the logging area.

The example below uses eight disks:

Submirror 1 - c1t0d0 (disk1), c1t0d1 (disk2), c1t0d2 (disk3), and c1t0d3 (disk4)

Submirror 2 - c2t0d0 (disk1), c2t0d1 (disk2), c2t0d2 (disk3), and c2t0d3 (disk4)

- The metadvice is configured to contain master user data on slice s0 on all devices.
- Slice s3 is configured to contain the UFS logging Space (16MB/disk)
- Stripe width is 64 KB
- Metamirror consists of two submirrors
- Each submirror comprises four disks on two controllers: c1 for submirror 1 and c2 for submirror 2
- Metamirror is created as a 1-way mirror, then a second mirror is added using the `metattach` command
- Stripe sequence is disk1, disk2, disk3, disk4, then back to disk1 and so on.

To create a mirrored and striped metamirror device for use with UFS logging, and comprising slice s0 and s3 on eight disks, perform the following steps:

1. Edit the **/etc/opt/SUNWmd/md.tab** and add the entries below to define the metamirror:

```
/dev/md/dsk/d1  1  4 /dev/dsk/c1t0d0s0 /dev/dsk/c1t0d1s0 \
                  /dev/dsk/c1t0d2s0 /dev/dsk/c1t0d3s0 -i 64k
/dev/md/dsk/d2  1  4 /dev/dsk/c2t0d0s0 /dev/dsk/c2t0d1s0 \
                  /dev/dsk/c2t0d2s0 /dev/dsk/c2t0d3s0 -i 64k
/dev/md/dsk/d3  1  4 /dev/dsk/c1t0d0s3 /dev/dsk/c1t0d1s3 \
                  /dev/dsk/c1t0d2s3 /dev/dsk/c1t0d3s3 -i 64k
/dev/md/dsk/d4  1  4 /dev/dsk/c2t0d0s3 /dev/dsk/c2t0d1s3
                  /dev/dsk/c2t0d2s3 /dev/dsk/c2t0d3s3 -i 64k
/dev/md/dsk/d0 -m /dev/md/dsk/d1  1
/dev/md/dsk/d5 -m /dev/md/dsk/d3  1
/dev/md/dsk/d6 -t /dev/md/dsk/d0 /dev/md/dsk/d5
```

Where:

d6 = name of mirrored metatrans device consisting of:

d1 = name of striped master metadvice (master submirror 1)

d2 = name of striped master metadvice (master submirror 2)

d3 = name of striped UFS log metadvice (log submirror 1)

d4 = name of striped UFS log metadvice (log submirror 2)

d0 = name of master metamirror

d5 = name of log metamirror

1 = number of device rows to be striped

4 = number of components in each row

-i 64k = Specify the stripe width (8k, 16k [default], 32k, 64k)

-m = option to define a metamirror device

-t = option to define a metatrans device

2. To create the mirrored metatrans device, type

```
metainit d1
(Create the striped (submirror 1) master)

metainit d2
(Create the striped (submirror 2) master)

metainit d3
(Create the striped (submirror 1) logging)

metainit d4
(Create the striped (submirror 2) logging)
```



```
metainit d0
(Create the 1-way master metamirror d0)

metainit d5
(Create the 1-way logging metamirror d5)

metattach d0 d2
(Attach second submirror d2 to master metamirror d0)

metattach d5 d4
(Attach second submirror d4 to log metamirror d5)

metainit d6
(Create the metatrans device d6 consisting of master device d0
and logging device d5)
```

3. To verify configuration, type

```
metastat
(or metastat d0)
```

DiskSuite State Database Replicas

The *state database* stores all configuration and status information for DiskSuite objects. Without the state database DiskSuite is unable to access any devices and all data could be lost. Replicas of the database are created so that DiskSuite can compare copies to verify the current configuration and running state of all objects.

It is recommended that at least three replicas be created. If one replica is not available, the remaining two can still be compared to verify configuration and state information. Three replicas can be stored on a system boot disk, however, this creates a single point of failure at the boot disk. Additional replicas should be created on other system disks including Symmetrix devices to protect against the loss of the boot disk. The following considerations apply when planning locations for state database replicas.

- ◆ Always create at least three replicas. DiskSuite will not function if it has only one state database.
- ◆ Replicas can be stored on any unused partition or on partitions that are also part of a metadevice or logging device with the exceptions of **root**, **swap**, **/usr** or file system.

- ◆ Replicas should be spread evenly across host controllers. In Symmetrix, replica devices should be behind separate even and odd SCSI adapters.
- ◆ Store replicas on least two disks on each controller.
- ◆ Each disk should contain only one replica.

Useful DiskSuite Commands

This section lists DiskSuite functions and utilities you may use to define and manage Symmetrix devices. The use of these functions and utilities is optional. They are listed for reference only.

- ◆ `metaclear devicename`

Delete a metadevice. When removing a metadevice hierarchy, use **metaclear** in top down sequence, for example

```
metaclear meta trans
metaclear master_mirror
metaclear log_mirror
metaclear submirrora1
metaclear submirrora2
metaclear submirrorb1
metaclear submirrorb2
```

- ◆ `metadb -i`

Check status of metastate database replicas.

- ◆ `metareplace`

Fix a defective mirrored or RAID device.

- ◆ `metastat`

Display status of mirrored or RAID device.

- ◆ `metahs`

Command to add a hot spare.

Veritas Volume Manager

Veritas VxVM and VxFS are tools for disk and file management. VxVM can be used to create logical disks, mirrored and striped volumes. VxFS supports large file systems, file system expansion and a journaling file system.

The following paragraphs provide instructions for using VxVM to create concatenated, striped, and mirrored logical volumes for use with either standard UNIX file systems (UFS) or VxFS journaled file systems.

See *Intent Logging* on page 16-38 for a description of journaling file systems.

For additional information on using Veritas in the Solaris environment refer to *Veritas Volume Manager (VxVM) System Administrator's Guide*.

Initializing Disks

All disks must be initialized before they can be used by VxVM. The initialization process stores configuration and identification data on the disk for use by VxVM.

Disks must be formatted before being initialized by VxVM. Formatting is done outside the volume manager using the **format** command. See *Format, Partition, Label* on page 16-11.

Use the **vxdiskadm** menu-driven interface to initialize disks before attempting to create VxVM volumes. To start **vxdiskadm**, log in as **root** and type

```
vxdiskadm
```

Select **1 Add or Initialize one or more disk**, and follow the system prompts to select devices to be initialized.

Creating a Concatenated Volume

Use this procedure to add a 2 GB concatenated volume of four disks. The configuration uses 500 MB on each 2 GB device.

There are two methods for creating a concatenated volume presented below. The manual **vxmake** method with more control over volume subdisk and plex creation, and the more convenient **vxassist** method with less control (automatic creation and allocation of volume components).

The disk must be initialized before attempting to create volumes (refer to *Initializing Disks* on page 16-23). The initialization process creates slice 02 on the disk. When VxVM puts the disk under its control it must first detect slice 02. VxVM can then create slice 03 as a private region and slice 04 a public region.

Creating a Concatenated Volume using vxmake

Use the manual **vxmake** method to create a concatenated volume as follows:

- 1. To verify that the disk are online and not being used, type

```
vxdisk list
```

Check the displayed list for available disks. For example:

DEVICE	TYPE	DISK	GROUP	STATUS
c0t1d0s2	sliced	-	-	online
c0t3d0s2	sliced	-	-	error
c1t0d0s2	sliced	-	-	error
c1t0d1s2	sliced	-	-	error
c1t0d2s2	sliced	-	-	online
c1t0d3s2	sliced	-	-	online
c1t0d4s2	sliced	-	-	online
c1t0d5s2	sliced	-	-	online
c1t0d6s2	sliced	-	-	online
c1t0d7s2	sliced	-	-	online
* c1t1d0s2	sliced	-	-	online
* c1t1d1s2	sliced	-	-	online
* c1t1d2s2	sliced	-	-	online
* c1t1d3s2	sliced	-	-	online
c1t1d4s2	sliced	-	-	online
c1t1d5s2	sliced	-	-	online
c1t1d6s2	sliced	-	-	online
c1t1d7s2	sliced	-	-	online

* Disks used in this example

2. To create disk group named **cust** and add four disks to disk group **cust**, type

```
vxvg init cust cust-00=c1t1d0
vxvg -g cust adddisk cust-01=c1t1d1
vxvg -g cust adddisk cust-02=c1t1d2
vxvg -g cust adddisk cust-03=c1t1d3
```

3. To verify subdisk creation, type

```
vxdisk list
```

Check the displayed list for subdisk and group. For example:

DEVICE	TYPE	DISK	GROUP	STATUS
		.		
		.		
		.		
c1t1d0s2	sliced	cust-00	cust	online
c1t1d1s2	sliced	cust-01	cust	online
c1t1d2s2	sliced	cust-02	cust	online
c1t1d3s2	sliced	cust-03	cust	online

4. To create 500 MB subdisks on each disks, type

```
vxmake -g cust sd cust-sd0 cust-00,0,500m
vxmake -g cust sd cust-sd1 cust-01,0,500m
vxmake -g cust sd cust-sd2 cust-02,0,500m
vxmake -g cust sd cust-sd3 cust-03,0,500m
```

5. To verify configuration database information, type

```
vxprint -st
```

Check display for correct configuration data.

For example:

Disk group: cust

SD NAME	PLEX	DISK	DISKOFFS	LENGTH	[COL/]OFF	DEVICE	MODE
sd cust-sd0	-	cust-00	0	1024000	-	c1t1d0	ENA
sd cust-sd1	-	cust-01	0	1024000	-	c1t1d1	ENA
sd cust-sd2	-	cust-02	0	1024000	-	c1t1d2	ENA
sd cust-sd3	-	cust-03	0	1024000	-	c1t1d3	ENA

6. To create striped plex named **cust-plex** out of the subdisks, type

```
vxmake -g cust plex cust-plex sd=cust-sd0,cust-sd1, \ cust-sd2,cust-sd3
layout=concat stwidth=64k
```

7. To verify configuration database information, type

```
vxprint -pt
```

Check display for correct configuration data.

For example:

Disk group: cust

PL NAME	VOLUME	KSTATE	STATE	LENGTH	LAYOUT	NCOL/WID	MODE
pl cust-plex	cust-vol	ENABLED	ACTIVE	4096000	STRIPE	4/128	

8. To create volume out of the concatenated plex **cust-plex**, type

```
vxmake -g cust -U fsgen vol cust-vol \ plex=cust-plex
```

9. To verify configuration database information, type

```
vxprint -vt
```

Check display for correct configuration data.

For example:

Disk group: cust

V NAME	USETYPE	KSTATE	STATE	LENGTH	READPOL	PREFPLEX
v cust-vol	fsgen	DISABLED	EMPTY	4096000	ROUND	-

10. To initialize the volume to the CLEAN state, type

```
vxvol -g cust init clean cust-vol
```

11. To start the concatenated volume, type

```
vxvol -g cust start cust-vol
```

12. To verify configuration database information, type

```
vxprint -vt
```

Check display for correct configuration data.

For example:

Disk group: cust

V NAME	USETYPE	KSTATE	STATE	LENGTH	READPOL	PREFPLEX
v cust-vol	fsgen	ENABLED	ACTIVE	4096000	ROUND	-

Creating a Concatenated Volume using vxassist

Creating volumes with **vxassist** is the recommended way of creating volumes since the user does not deal with the complicated commands as seen in the manual **vxmake** method. This example creates 2 GB concatenated volumes **cust-vol1** and **cust-vol2** on disk group **cust**. **vxassist** creates the subdisks and plexes, and initializes and starts the volume automatically.

1. To create disk group named **cust** and add four disks to disk group **cust**, type

```
vx dg init cust cust-00=c1t1d0
vx dg -g cust adddisk cust-01=c1t1d1
vx dg -g cust adddisk cust-02=c1t1d2
vx dg -g cust adddisk cust-03=c1t1d3
```

2. To create concatenated volume named **cust-vol1** with 2.0 GB size using disk in disk group **cust**, type

```
vxassist -g cust -fsgen make cust-vol1 2g \ layout=span
```

3. To verify configuration database information, type

```
vxprint -vt cust-vol1
```

Check display for correct configuration data.

For example:

Disk group: cust

V NAME	USETYPE	KSTATE	STATE	LENGTH	READPOL	PREFPLEX
v cust-vol	fsgen	ENABLED	ACTIVE	4194394	SELECT	-

Creating a Striped Volume

Use this procedure to add a 2 GB striped volume of four disks. The configuration uses 500 MB on each 2 GB device.

There are two methods for creating a striped volume presented below. The manual **vxmake** method with more control over volume subdisk and plex creation, and the more convenient **vxassist** method with less control (automatic creation and allocation of volume components).

Creating a Striped Volume with vxmake

Use the manual **vxmake** method to create a striped volume as follows:

- 1. To verify that the disk are online and not being used, type

```
vxdisk list
```

Check the displayed list for available disks. For example:

	DEVICE	TYPE	DISK	GROUP	
	c0t1d0s2	sliced	-	-	online
	c0t3d0s2	sliced	-	-	error
*	c1t0d0s2	sliced	-	-	online
*	c1t0d1s2	sliced	-	-	online
	c1t0d2s2	sliced	-	-	online
	c1t0d3s2	sliced	-	-	online
	c1t0d4s2	sliced	-	-	online
	c1t0d5s2	sliced	-	-	online
	c1t0d6s2	sliced	-	-	online
	c1t0d7s2	sliced	-	-	online
*	c2t0d0s2	sliced	-	-	online
*	c2t0d1s2	sliced	-	-	online
	c2t0d2s2	sliced	-	-	online
	c2t0d3s2	sliced	-	-	online
	c2t0d4s2	sliced	-	-	online
	c2t0d5s2	sliced	-	-	online
	c2t0d6s2	sliced	-	-	online
	c2t0d7s2	sliced	-	-	online

* Disks used in this example

- 2. To create disk group named **cust** and add four disks to disk group **cust**, type

```
vx dg init cust cust-00=c1t0d0
vx dg -g cust adddisk cust-01=c2t0d0
vx dg -g cust adddisk cust-02=c1t0d1
vx dg -g cust adddisk cust-03=c2t0d1
```


3. To verify subdisk creation, type

```
vxdisk list
```

Check the displayed list for subdisk and group. For example:

DEVICE	TYPE	DISK	GROUP	STATUS
			.	
			.	
			.	
c1t0d0s2	sliced	cust	-00	customonline
c2t0d0s2	sliced	cust	-01	customonline
c1t0d1s2	sliced	cust	-02	customonline
c2t0d1s2	sliced	cust	-03	customonline

4. To create 500 MB subdisks on each disks, type

```
vxmake -g cust sd cust-sd0 cust-00,0,500m
vxmake -g cust sd cust-sd1 cust-01,0,500m
vxmake -g cust sd cust-sd2 cust-02,0,500m
vxmake -g cust sd cust-sd3 cust-03,0,500m
```

5. To verify configuration database information, type

```
vxprint -st
```

Check display for correct configuration data.

For example:

Disk group: cust

SD	NAMEPLEXDISK	DISKOFFS	LENGTH	[COL/]OFF	DEVICE	MODE
sd	cust-sd0-cust-000		1024000	-	c1t0d0	ENA
sd	cust-sd1-cust-010		1024000	-	c2t0d0	ENA
sd	cust-sd2-cust-020		1024000	-	c1t0d1	ENA
sd	cust-sd3-cust-030		1024000	-	c2t0d1	ENA

6. To create striped plex named **cust-plex** out of the subdisks, type

```
vmake -g cust plex cust-plex sd=cust-sd0,cust-sd1,cust-sd2,cust-sd3 \
layout=stripe stwidth=64k
```

Stripe width values are 1Kbytes to 256 Kbytes. The default value is 64 Kbytes.

7. To verify configuration database information, type

```
vxprint -pt
```

Check display for correct configuration data.

For example:

Disk group: cust

PL NAME	VOLUME	KSTATE	STATE	LENGTH	LAYOUT	NCOL/WID	MODE
pl cust-plex	cust-vol	ENABLED	ACTIVE	4096000	STRIPE	4/128	

8. To create volume **cust-vol** out of the plex **cust-plex**, type

```
vxmake -g cust -Ufsgen vol cust-vol plex=cust-plex
```

9. To verify configuration database information, type

```
vxprint -vt
```

Check display for correct configuration data.

For example:

Disk group: cust

V NAME	USETYPE	KSTATE	STATE	LENGTH	READPOL	PREFPLEX
v cust-vol	fsgen	DISABLED	EMPTY	4096000	ROUND	-

10. To initialize the volume to the CLEAN state, type

```
vxvol -g cust init clean cust-vol
```

11. To start the concatenated volume, type

```
vxvol -g cust start cust-vol
```

12. To verify configuration database information, type

```
vxprint -vt
```

Check display for correct configuration data.

For example:

Disk group: cust

V NAME	USETYPE	KSTATE	STATE	LENGTH	READPOL	PREFPLEX
v cust-vol	fsgen	ENABLED	ACTIVE	4096000	ROUND	-

Creating a Striped Volume with vxassist

Creating volumes with **vxassist** is the recommended way of creating volumes since the user does not deal with the complicated commands as seen in the manual **vxmake** method. This example, creates a 2 GB striped volume named **cust-vol1**. **vxassist** creates the subdisk and plex, and initializes and starts the volume automatically.

1. To create disk group named **cust** and add four disks to disk group **cust**, type

```
vx dg init cust cust-00=c1t0d0
vx dg -g cust adddisk cust-01=c2t0d0
vx dg -g cust adddisk cust-02=c1t0d1
vx dg -g cust adddisk cust-03=c2t0d1
```

2. To create striped volume named **cust-vol1** with 2.0 GB size using disks in disk group **cust**, type

```
vxassist -g cust -Ufsgen make cust-vol1 2g layout=stripe stwidth=64k \ cust-00
cust-01 cust-02 cust-03
```

Stripe width values are 1Kbytes to 256 Kbytes. The default value is 64 Kbytes.

3. To verify configuration database information, type

```
vxprint -vt cust-vol1
```

Check display for correct configuration data. For example:

Disk group: cust

V NAME	USETYPE	KSTATE	STATE	LENGTH	READPOL	PREFPLEX
v cust-vol1	fsgen	ENABLED	ACTIVE	4194394	ROUND	-

Creating a Striped Mirrored Volume

Use this procedure to add a 2 GB striped volume of four disks with a mirror on another four disks. The configuration uses 500 MB on each 2 GB device.

There are two methods for creating a striped volume presented below. The manual **vxmake** method with more control over volume subdisk and plex creation, and the more convenient **vxassist** method with less control (automatic creation and allocation of volume components).

Manual vxmake Method

Use the manual **vxmake** method to create a concatenated volume as follows:

1. To verify that the disk are online and not being used, type

```
vxdisk list
```

Check the displayed list for available disks.
For example:

DEVICE	TYPE	DISK	GROUP	STATUS
c0t1d0s2	sliced	-	-	online
c0t3d0s2	sliced	-	-	error
c1t0d0s2	sliced	-	-	error
c1t0d1s2	sliced	-	-	error
c1t0d2s2	sliced	-	-	online
c1t0d3s2	sliced	-	-	online
c1t0d4s2	sliced	-	-	online
c1t0d5s2	sliced	-	-	online
c1t0d6s2	sliced	-	-	online
c1t0d7s2	sliced	-	-	online
* c1t1d0s2	sliced	-	-	online
* c1t1d1s2	sliced	-	-	online
* c1t1d2s2	sliced	-	-	online
* c1t1d3s2	sliced	-	-	online
c1t1d4s2	sliced	-	-	online
c1t1d5s2	sliced	-	-	online
c1t1d6s2	sliced	-	-	online
c1t1d7s2	sliced	-	-	online
* c2t1d0s2	sliced	-	-	online
* c2t1d1s2	sliced	-	-	online
* c2t1d2s2	sliced	-	-	online
* c2t1d3s2	sliced	-	-	online

* Disks used in this example

2. To create Disk group named **cust** and add eight disks to disk group **cust**, type

```

vxdg init cust cust-00=c1t1d0
vxdg -g cust adddisk cust-01=c1t1d1
vxdg -g cust adddisk cust-02=c1t1d2
vxdg -g cust adddisk cust-03=c1t1d3
vxdg -g cust adddisk cust-04=c2t1d0
vxdg -g cust adddisk cust-05=c2t1d1
vxdg -g cust adddisk cust-06=c2t1d2
vxdg -g cust adddisk cust-07=c2t1d3

```

3. To verify subdisk creation, type

```
vxdisk list
```

Check the displayed list for subdisk and group. For example:

DEVICE	TYPE	DISK	GROUP	STATUS
		.		
		.		
		.		
c1t1d0s2	sliced	cust-00	cust	online
c1t1d1s2	sliced	cust-01	cust	online
c1t1d2s2	sliced	cust-02	cust	online
c1t1d3s2	sliced	cust-03	cust	online
c1t1d4s2	sliced	-	-	online
c2t1d0s2	sliced	cust-04	cust	online
c2t1d1s2	sliced	cust-05	cust	online
c2t1d2s2	sliced	cust-06	cust	online
c2t1d3s2	sliced	cust-07	cust	online

4. To create 500 MB subdisks on each disks, type

```

vxmake -g cust sd cust-sd0 cust-00,0,500m
vxmake -g cust sd cust-sd1 cust-01,0,500m
vxmake -g cust sd cust-sd2 cust-02,0,500m
vxmake -g cust sd cust-sd3 cust-03,0,500m
vxmake -g cust sd cust-sd4 cust-04,0,500m
vxmake -g cust sd cust-sd5 cust-05,0,500m
vxmake -g cust sd cust-sd6 cust-06,0,500m
vxmake -g cust sd cust-sd7 cust-07,0,500m

```

5. To verify configuration database information, type

```
vxprint -st
```

Check display for correct configuration data.

For example:

Disk group: cust

SD NAME	PLEX	DISK	DISKOFFS	LENGTH	[COL/]OFF	DEVICE	MODE
sd cust-sd0	-	cust-00	0	1024000	-	c1t1d0	ENA
sd cust-sd1	-	cust-01	0	1024000	-	c1t1d1	ENA
sd cust-sd2	-	cust-02	0	1024000	-	c1t1d2	ENA
sd cust-sd3	-	cust-03	0	1024000	-	c1t1d3	ENA
sd cust-sd4	-	cust-00	0	1024000	-	c2t1d0	ENA
sd cust-sd5	-	cust-01	0	1024000	-	c2t1d1	ENA
sd cust-sd6	-	cust-02	0	1024000	-	c2t1d2	ENA
sd cust-sd7	-	cust-03	0	1024000	-	c2t1d3	ENA

6. To create striped plexes **cust-plex1** and **cust-plex2** out of the subdisks, type

```
cust-plex1 make up submirror 1, cust-plex2 make up submirror 2
```

```
vxmake -g cust plex cust-plex1 sd=cust-sd0,cust-sd1,cust-sd2,cust-sd3\  
layout=stripe stwidth=64k
```

```
vxmake -g cust plex cust-plex2 sd=cust-sd4,cust-sd5,cust-sd6,cust-sd7\  
layout=stripe stwidth=64k
```

Stripe width values are 1Kbytes to 256 Kbytes. The default value is 64 Kbytes.

7. To verify configuration database information, type

```
vxprint -pt
```

Check display for correct configuration data.

For example:

Disk group: cust

PL NAME	VOLUME	KSTATE	STATE	LENGTH	LAYOUT	NCOL/WID	MODE
pl cust-plex1-vol		DISABLED	-	4096000	STRIPE	4/128RW	

8. To create mirror volume out of the **cust-plex1** and then attach second **cust-plex2**

```
vxmake -g cust -U fsgen vol cust-mirvol plex=cust-plex1
vxplex -g cust att cust-mirvol cust-plex2
```

9. To verify configuration database information, type

```
vxprint -vt
```

Check display for correct configuration data.

For example:

Disk group: cust

```
V NAME          USETYPE KSTATE  STATE  LENGTH READPOL PREFPLEX
v cust-mirvol fsgen     DISABLED EMPTY  4096000 ROUND -
```

10. To initialize the volume to the CLEAN state, type

```
vxvol -g cust init clean cust-mirvol
```

11. To start the concatenated volume, type

```
vxvol -g cust start cust-mirvol
```

12. To verify configuration database information, type

```
vxprint -vt
```

Check display for correct configuration data.

For example:

Disk group: cust

```
V NAME          USETYPE KSTATE  STATE  LENGTH READPOL PREFPLEX
v cust-mirvol fsgen     ENABLED  ACTIVE  4096000 ROUND -
```

13. To add two copies of dirty region logging (DRL) to the mirror volume, type

```
vxassist -g cust addlog cust-mirvol
(Add first log)
```

```
vxassist -g cust addlog cust-mirvol
(Add second log)
```

Dirty region logging (DRL is a bit map mechanism used to facilitate mirror resynchronization after a system failure. Vxassist is the easy way to add a DRL log to the volume. Vxassist creates a subdisk in the diskgroup; its size based on the volume size (1 KB for every 2 GB). The log for this volume may be 4 sectors or 2 Kbytes (1 KB for each 2 GB mirror). When possible, the logging subdisk is created on a disk separate from the data volume. Vxassist creates a plex out of the subdisk and attaches the log plex to the mirrored volume.

Creating Mirrored Volume with vxassist

Creating volumes with **vxassist** is the recommended way of creating volumes since the user does not deal with the complicated commands as seen in the manual **vxmake** method. This example creates a mirrored volume named **cust-mirvol1** with four disks on controller c1 and four disks on controller c2. **vxassist** creates the subdisk and plexes, and initializes and starts the volume automatically.

1. To create disk group named **cust** and add eight disks to disk group **cust**, type

```
vxdbg init cust cust-00=c1t1d0
vxdbg -g cust adddisk cust-01=c1t1d1
vxdbg -g cust adddisk cust-02=c1t1d2
vxdbg -g cust adddisk cust-03=c1t1d3
vxdbg -g cust adddisk cust-04=c2t1d0
vxdbg -g cust adddisk cust-05=c2t1d1
vxdbg -g cust adddisk cust-06=c2t1d2
vxdbg -g cust adddisk cust-07=c2t1d3
```

2. To create mirrored volume named **cust-mirvol** with 2.0 GB size using disks **cust-00**, **cust-01**, **cust-02** and **cust-03** (controller c1) in disk group **cust**, type

```
vxassist -g cust -Ufsgen make cust-mirvol1
2g layout=stripe stwidth=64k \
cust-00 cust-01 cust-02 cust-03
```

Stripe width values are 1Kbytes to 256 Kbytes. The default value is 64 Kbytes.

3. To verify configuration database information, type

```
vxprint -vt cust-mirvol
```

Check display for correct configuration data.

For example:

Disk group: cust

```
V NAME          USETYPE KSTATE  STATE  LENGTH READPOL PREFPLEX
v cust-mirvol fsgen    ENABLED  ACTIVE 4194394 ROUND  -
```

4. To create submirror for mirrored volume **cust-mirvol** of 2.0 GB size with disks **cust-04**, **cust-05**, **cust-06** and **cust-07** on second controller c2, type

```
vxassist -g cust mirror cust-mirvol layout=stripe cust-04 cust-05\
cust-06 cust-07 &
```

The mirror volume automatically synchronizes the second submirror. This can take a while.

5. To add two copies of dirty region logging (DRL) to the mirror volume, type

```
vxassist -g cust addlog cust-mirvol1
(Add first log)
```

```
vxassist -g cust addlog cust-mirvol1
(Add second log)
```

Dirty region logging (DRL) is a bit map mechanism used to facilitate mirror resynchronization after a system failure. Vxassist is the easy way to add a DRL log to the volume. Vxassist creates a subdisk in the diskgroup; its size based on the volume size (1 KB for every 2 GB). The log for this volume may be 4 sectors or 2 Kbytes (1 KB for each 2 GB mirror). When possible, the logging subdisk is created on a disk separate from the data volume. Vxassist creates a plex out of the subdisk and attaches the log plex to the mirrored volume.

Create and Mount a File System

Volumes created and managed by Solstice DiskSuite or VxVM volume manager may be used as raw devices, with the standard UNIX file system, or with the VxFS journaled file system.

Intent Logging

Intent logging records pending changes to the file system structure in an intent log. The intent log is replayed during system failure recovery to complete or abandon changes to the structure that were pending at the time of system failure. The file system can then be mounted without completing a full structural check (fsck). An intent logging system can significantly reduce recovery time following a system failure.

Intent logging may be preformed at the file system level or at the volume manager level. Intent logging at the file system level, generally known as a journaling file system (JFS), is usually more effective than volume manager intent logging.

VxFS and VxVM Intent Logging

The Veritas File System (VxFS) includes a journaling file system that provides intent logging at the file system level. Pending changes to the file system, written to an intent log, are scanned during recovery from a system failure. Changes that were active at the time of failure are completed, and the file system is mounted without requiring a fsck of the entire file system. File system recovery is done in a few seconds; much faster than a standard UFS recovery that requires a complete fsck.

VxVM Dirty Region Logging (DRL) provides intent logging at the volume manager level to reduce the time required to resynchronize mirrored volumes after a system failure. DRL applies to VxVM mirrored volumes only. Striped or concatenated volumes do not use intent logging, but may rely on VxFS for fast recovery.

SDS Intent Logging

Solstice DiskSuite (SDS) uses a standard UNIX file system (UFS) that does not provide intent logging at the file system level. However, the VxFS Veritas file system may be used in the SDS environment to provide a journaling file system for SDS volumes.

At the volume manger level, SDS uses a method called the *UFS logging feature* to provide intent logging for all volume types (striped, concatenated, mirrored, and RAID). UFS logging is not a journaling file system. It uses the standard UFS and does intent logging at the volume manager level.

Procedures for creating volumes with UFS logging are provided in *Solstice DiskSuite* on page 16-13.

Procedures for creating mirrored volumes with DRL are provided in *Veritas Volume Manager* on page 16-23.

The following paragraphs outline the steps required to add, create and mount a standard UNIX file system or a VxFS journaled file system for volumes and raw devices.



CAUTION

You may wish to place the Symmetrix devices in the mount table. This requires `/etc/vfstab` to be edited. `/etc/vfstab` is syntax-sensitive, and if not edited properly may prevent the system from booting.

UFS File System on Raw Device

Create New File System

To create standard UNIX file systems under Solaris OS, log in as **root** and proceed as follows for each new device.

Once you have formatted, partitioned, and labeled each Symmetrix disk device, create a new file system for each Symmetrix disk. To do this, use the **newfs** command in a statement similar to the following:

```
newfs -v /dev/rdisk/c1t0d0s0
```

The **Construct new filesystem?** prompt appears. Type **y** at the prompt.

The actions above created a new file system for the Symmetrix disk connected to SCSI controller 1, target ID 0, lun 0, partition 0.

Create Mount Directory

Once the file systems for each Symmetrix device are created, create a mount directory for each device. To do this, type a statement similar to the following for each Symmetrix device:

```
mkdir /fs/c1t0d0s0
```

where **/fs/c1t0d0s0** is the complete path for the new file system directory.

The **/fs** directory must exist prior to creating the mount directories.

Mount the Filesystem

To mount each filesystem, type a statement similar to the following:

```
mount /dev/dsk/c1t0d0s0 /fs/c1t0d0s0
```

VxFS File System on Raw Device

A VxFS journaling file system is created using the **mkfs** command with arguments provided for block size, log size, device name and size. To create a VxFS file system first determine the size in sectors of the volume. The size of the volume is displayed under the **Sector Count** field of the **prtvtoc** output.

Once you have formatted, partitioned, and labeled each Symmetrix disk device, create a new file system for each Symmetrix disk. To create VxVM journaling file systems under VxFS and Solaris OS, log in as **root** and proceed as follows for each new device.

Create New File System

1. To display sector count information, type

```
prtvtoc /dev/rdsk/clt0d0s0
```

The size of the disk is displayed under the **Sector Count** field.
(Assume 4099000 for this example.)

2. To create the VxFS file system for the volume, type

```
mkfs -F vxfs -o bsize=4096,logsize=256 \ /dev/rdsk/clt0d0s0 4099000
```

Where:

bsize = block size in bytes (1k, 2k, 4k, or 8k - 1k default for filesystems < 4 GB, 4k default for filesystems > 4 GB)

logsize = size of VxFS file system logging in blocks (256 blocks default, 32 to 1024 blocks)

4099000 = file system size in sectors (from **prtvtoc** command)

Create Mount Directory

Once the file systems for each Symmetrix device are created, create a mount directory for each device. To do this, type a statement similar to the following for each Symmetrix device:

```
mkdir /fs/clt0d0s0
```

where **/fs/clt0d0s0** is the complete path for the new file system directory.

Mount the Filesystem

To mount each filesystem, type a statement similar to the following:

```
mount -F vxfs /dev/dsk/clt0d0s0 /fs/clt0d0s0
```

UFS File System on SDS Device

Create New File System

To create standard UNIX file systems under Solstice DiskSuite, log in as **root** and proceed as follows for each new device.

Once you have formatted, partitioned, and labeled each Symmetrix disk device, and created SDS volumes, create a new file system on each volume. To do this, use the **newfs** command in a statement similar to the following:

```
newfs -v /dev/md/rdisk/d0
```

The **Construct new filesystem?** prompt appears. Type **y** at the prompt.

The actions above created a new file system for the Symmetrix disk defined as SDS metadvice **d0**.

Create Mount Directory

Once the file systems for each device are created, create a mount directory for each device. To do this, type a statement similar to the following for each device:

```
mkdir /fs/d0
```

where **/fs/d0** is the complete path for the new file system directory.

You can assume that the **/fs** directory existed prior to creating the **d0** directory.

Mount the Filesystem

To mount each filesystem, type a statement similar to the following:

```
mount /dev/md/dsk/d0 /fs/d0
```

VxFS File System on SDS Device

A VxFS journaling file system is created using the **mkfs** command with arguments provided for block size, log size, device name and size. To create a VxFS file system first determine the size in sectors of the volume. The size of the volume is displayed under the **Sector Count** field of the **prtvtoc** output.

To create VxFS journaling file systems under Solstice DiskSuite, log in as **root** and proceed as follows for each new device.

Create New File System

1. To display sector count information, type

```
prtvtoc /dev/md/rdisk/d0
```

The size of the disk is displayed under the Sector Count field.
(Assume 4099000 for this example.)

2. To create the VxFS file system for the volume, type

```
mkfs -F vxfs -o bsize=1024 logsize=512\ /dev/md/rdisk/d0 4099000
```

Where:

bsize = block size in bytes (1k, 2k, 4k, or 8k - 1k default for filesystems < 4 GB, 4k default for filesystems > 4 GB)

logsize = size of VxFS file system logging in blocks (256 blocks default, 32 to 1024 blocks)

4099000 = file system size in sectors (from **prtvtoc** command)

Create Mount Directory

Once the file systems for each device are created, create a mount directory for each device. To do this, type a statement similar to the following for each device:

```
mkdir /fs/d0
```

where **/fs/d0** is the complete path for the new file system directory.

Mount the Filesystem

To mount each filesystem, type a statement similar to the following:

```
mount -F vxfs /dev/md/dsk/d0 /fs/d0
```

UFS File System on VxVM Device

To create standard UNIX file systems under Veritas Volume Manager, log in as **root** and proceed as follows for each new device.

Create New File System

Once you have formatted, partitioned, and labeled each Symmetrix disk device, and created VxVM volumes, create a new file system on each volume. To do this, use the **newfs** command in a statement similar to the following:

```
newfs -v /dev/vx/rdisk/dskgrp/vol11
```

The **Construct new filesystem?** prompt appears. Type **y** at the prompt.

The actions above created a new file system for the Symmetrix disk defined as **vol1** in the diskgroup named **dskgrp**.

Create Mount Directory

Once the file systems for each device are created, create a mount directory for each device. To do this, type a statement similar to the following for each device:

```
mkdir /fs/vol1
```

where **/fs/vol1** is the complete path for the new file system directory.

Mount the Filesystem

To mount each filesystem, type a statement similar to the following:

```
mount -F vsfs /dev/vx/dsk/dskgrp/vol1 /fs/vol1
```

VxFS File System on VxVM Device

A VxFS journaling file system is created using the **mkfs** command with arguments provided for block size, log size, device name and size. To create a VxFS file system first determine the size in sectors of the volume. The size of the volume is displayed under the **Sector Count** field of the **prtvoc** output.

To create VxFS journaling file systems under Veritas Volume Manager, log in as **root** and proceed as follows for each new device.

Create New File System

1. To display sector count information, type

```
prtvoc /dev/vx/rdsk/dskgrp/vol1
```

The size of the disk is displayed under the **Sector Count** field. (Assume 4099000 for this example.)

2. To create the VxFS file system for the volume, type

```
mkfs -F vxfs -o bsize=4096 logsize=512 /dev/vx/rdsk/dskgrp/vol1 4099000
```

Where:

bsize = block size in bytes (1k, 2k, 4k, or 8k - 1k default for

filesystems < 4 GB, 4k default for filesystems > 4 GB
logsize = size of VxFS file system logging in blocks (256 blocks default, 32 to 1024 blocks)
4099000 = file system size in sectors (from **prvtoc** command)

Create Mount Directory

Once the file systems for each device are created, create a mount directory for each device. To do this, type a statement similar to the following for each device:

```
mkdir /fs/vol1
```

Where **/fs/vol1** is the complete path for the new file system directory.

Mount the Filesystem

To mount each filesystem, type a statement similar to the following:

```
mount -F vxfs /dev/vx/dsk/dskgrp/vol1 /fs/vol1
```

File System Expansion

DiskSuite volume manager allows you to increase the available storage space of an existing volume by concatenating additional volumes to the metadvice using **metattach**. A file system can then be expanded to fill all or part of the additional space using the **growfs** command. The file system can remain mounted, but will be locked (**lockfs**) during the expansion. For detailed information on expanding a file system under DiskSuite refer to *Solstice DiskSuite 4.0 /4.1 Administrator's Guide*.

VxVM allows *growing* of a mounted file system using **vxassist** commands. The following steps outline the procedure.

1. Log in as **root**.
2. To determine how large the volume can grow, type

```
vxassist maxgrow vol
```

Where **vol** is the volume name

Volume **vol** can be extended by 12533760 to 16629760 (8120Mb)

3. Determine the size of the current volume, type

```
vxprint -vt
```

```
Disk group: cust
```

V	NAME	USETYPE	KSTATE	STATE	LENGTH	READPOL	PREFPLEX
v	vol	fsgen	ENABLED	ACTIVE	4096000	ROUND	-

The length unit is in sectors (1 sectors = 512 bytes). Therefore **vol** is approximately 2 GB.

4. To grow the volume **vol** to 4 GB, type

```
vxassist growto cust-mirvol 4g
vxprint -vt
```

```
Disk group: cust
```

V	NAME	USETYPE	KSTATE	STATE	LENGTH	READPOL	PREFPLEX
v	cust-mirvol	fsgen	ENABLED	ACTIVE	8388608	ROUND	-

System and Error Messages

The Fujitsu Siemens PRIMEPOWER host running the Solaris OS logs system and error messages including SCSI and Fibre Channel error messages to a file called **/var/adm/messages** and also displays these messages at the system console.

Stratus Fault Tolerant Environment

This chapter describes how to incorporate Symmetrix in the Stratus® Fault Tolerant environment running the HP-UX operating system.

◆ Overview	17-2
◆ Useful Stratus Utilities and Functions	17-4
◆ Symmetrix Configuration	17-7
◆ Continuum 400 Series Host Cabling	17-9
◆ Recognizing New Devices and LUNs	17-14
◆ Configuring Logical Volumes	17-15
◆ Using SAM	17-17
◆ Adding Devices Online	17-18
◆ System and Error Messages	17-19

Overview

When using Symmetrix in a Stratus Fault Tolerant environment on Stratus Fault Tolerant Hosts, please note the following:

- ◆ Symmetrix attaches to the STRATUS Fast-Wide Differential SCSI Controller Card (P/N AA-U50300) (PCI based) on Stratus Continuum 400 Series Hosts (AC POLO 8000 Models)
- ◆ The STRATUS Fast-Wide Differential (FWD) Controller U50300 SCSI Host Adapter Card must be running with Firmware version S06 (μ5030fw0S06raw) or higher.
- ◆ The Stratus host requires the HP-UX Version 10.20.1 Operating System or higher.
- ◆ Stratus with HP_UX supports target IDs 0 through D. Target IDs E and F are reserved for the Stratus U503 SCSI Controller.
- ◆ Stratus U503 SCSI Controller supports 8 LUNs (0-7).
- ◆ Stratus U503 SCSI Controller supports a maximum of 32 devices.
- ◆ Stratus with HP-UX does not support target and LUN skipping. LUNs must be consecutive starting with LUN 0.

Table 17-1 below summarizes the HP-UX operating system features in the Stratus environment.

Table 17-1 Stratus with HP-UX v10.20.1 OS Features

O/S	Stratus with HP-UX
Largest recognizable drive size	4 GB or larger using SAM LVM
Largest defined partition	4 GB or larger using SAM LVM
Largest filesystem	4 GB to 128 GB ^a
Largest data file	2 GB to 128 GB ^b
Supports multiple target IDs	Yes
Supports multiple LUNs per target ID	Yes
Partition representing the entire disk	s2

Table 17-1 **Stratus with HP-UX v10.20.1 OS Features** (*continued*)

O/S	Stratus with HP-UX
To see the drive configuration, use this utility:	<code>ioscan</code> or <code>SAM</code>
To configure drives use this utility:	<code>SAM</code> or line commands
Special drive definition files needed	No

- a. HP-UX 10.20 supports a maximum filesystem size of 128 GB
- b. HP-UX 10.20 supports a maximum file size of 128 GB

Useful Stratus Utilities and Functions

This section lists Stratus and HP-UX functions and utilities you may use to define and manage Symmetrix devices. The use of these functions and utilities is optional. They are listed for reference only. See Stratus man pages for more complete information.

- ◆ `addhardware`
Finds new hardware on the SCSI bus or ports in the PCI bays, runs `ioscan` and `ioinit -i`, then updates the *ioconfig* file, copying the file to the booting flash card, making the new hardware available for configuration.
- ◆ `ftsmaint`
Provides a set of hardware control commands to view or modify the configuration of devices on your system.
- ◆ `showboot`
Displays the device name of the flash card used to boot the system and optionally shows the name of the boot kernel.
- ◆ `disksecn`
Manually creates a device ID.
- ◆ `flifcp`
Copies an LIF file in a flash image to a file on a disk, or copies a file on a disk to an LIF file in a flash image. The flash image can be on a flash card or on a disk.
- ◆ `flifrm`
Removes one LIF file from a flash image.
- ◆ `ioscan`
An online hardware mapping program that provides a listing of devices connected to the host. It scans the system hardware for usable I/O systems devices or kernel I/O system data structures as appropriate and lists the results. See the following table of options.

Useful options for this function are:

-C disk	Restrict the output to all devices belonging to the disk class.
-f	Generate a full listing.
-F	Generate a compact listing.
-n	List device file names in the output.
-M -H <hardware path>	Force the software driver (specified hardware path) into the kernel.

◆ insf

Installs special files and assigns logical unit numbers to devices in the system. Used for introducing newly attached devices to the system without rebooting.

◆ mknod

Creates a directory, special, or ordinary file.

◆ pvcreate

Creates a physical volume.

◆ pvchange

Changes characteristics of a physical volume in a volume group. It also redirects I/O to different physical volume paths.

◆ pvdisplay

Displays information about physical drives within a volume group.

◆ vgscan

Scans physical volumes for volume groups.

◆ vgcreate

Create volume group on a set of physical volumes.

◆ vgchange

Sets volume group availability.

◆ vgdisplay

Displays information about volume groups.

- ◆ `vgexport`
Exports a volume group and its associated logical volumes.
- ◆ `vgimport`
Imports a volume group onto a system.
- ◆ `vgremove`
Removes a volume group definition from the system.
- ◆ `lvcreate`
Stripes; creates logical volume in volume group.
- ◆ `lvchange`
Changes logical volume characteristics.
- ◆ `lvdisplay`
Displays information about logical volumes.
- ◆ `lvextend`
Increases space for logical volumes, but not size of filesystems.
- ◆ `lvreduce`
Decreases number of physical extents allocated to logical volume.
- ◆ `lvremove`
Removes one or more logical volumes.
- ◆ `newfs`
Creates a new filesystem.
- ◆ `bdf`
Shows all mounted filesystems and available free space.
- ◆ `lsdev`
Lists device drivers in the system.

Symmetrix Configuration

In the Stratus Fault Tolerant Continuum 400 Series host environment, you can configure the Symmetrix disk devices into logical volumes. EMC recommends the following logical volume ratios as indicated in the table below:

Table 17-2 Disk Device Configuration

		Symmetrix Logical Volume to Physical Splits ^a							
		1:1	2:1	3:1	4:1	5:1	6:1	7:1	8:1 ^b
Physical Devices	4 GB	x	x						
	9 GB	x	x	x	x				
	18 GB		x	x	x	x	x	x	x
	23 GB		*	x	x	x	x	x	x
	36 GB			*	x	x	x	x	x
	47 GB			*	*	x	x	x	x
	50 GB ^c				x	x	x	x	x
	73 GB ^{c, d}					x	x	x	x

- a. * indicates ratios available only in microcode level 5265 or higher.
- b. With microcode revision level 5263, maximum is 8:1. With 5264, maximum is 16:1. With 5265, maximum is 32:1.
- c. The 50 and 73 GB drives are supported only by microcode revision levels 5x66 and higher.
- d. The 73 GB drive does not support the RAID-S configuration on Symmetrix.

For specific cylinder counts and logical volume sizes, refer to Table 2-1 on page 2-5 and the following tables in that chapter.

The EMC Customer Engineer (CE) specifies the volume configuration and size during Symmetrix configuration. The CE also assigns target IDs and LUNs to these volumes at that time.

The CE should contact the EMC Configuration Specialist for updated online information. This information is necessary to configure Symmetrix to support the customer's host environment.

Symmetrix Configuration for the Stratus Host

For Stratus High Availability and fault tolerant configuration, Symmetrix devices must be configured as follows:

- ◆ Shared device across two SA ports
- ◆ The parameter " Enable propagate reset " must be set to NO as below:
`Enable Propagate Reset? N`
- ◆ The following bits must be set in Edit Director Table for all shared SA ports :

Table 17-3 Director Bits to be Set

Director Bit	Description
T	Tag Command Queuing
L	Command Links
Y	Synchronous Negotiation
W	Wide Negotiation
B	Return BUSY for Abort
C	Common Serial Number
D	Disable Queue Reset on Unit Attention

- ◆ Stratus only supports Target ID 0-D. TID E, F are reserved for Stratus. LUN 0-7 is supported, but target and LUN skipping are not supported.
- ◆ Maximum devices supported per SCSI bus is 32 devices

Continuum 400 Series Host Cabling

The Stratus Continuum 400 Series Host is a PCI-based, fault tolerant system that provides fault detection, failover and recovery of all hardware components. Every hardware component is backed up by another redundant component. The Stratus Host has two PCI card cages that house a bootcode flash card, ethernet adapter and SCSI controllers. One cage is active and the other is idle. Failure in any component housed in card cage 1 will result in failover to the other component housed in idle card cage 2. The Operating System, together with the redundant hardware, provide transparent I/O protection against failure in the Host Based Adapter (HBA), Cable disconnect, or SA failure along any one path.

To support this high availability fault tolerance, the Stratus host must be cabled to the Symmetrix as shown in Figure 17-1 or Figure 17-2 (on the following pages). The Symmetrix devices must be configured as shared devices across two SAs.

HP-UX 11.00 is supported only on the AC POLO 8000 and is not supported on the AC POLO 7100.

The Stratus Continuum 400 Series host model AC POLO 7100 supports maximum of one U503 SCSI controller per PCI card cage as shown in Figure 17-1. The Stratus Continuum 400 Series host model AC POLO 8000 supports a maximum of two (pairs) of U503 SCSI controllers as shown in Figure 17-2.

The U503 SCSI controller connector is 68 pin-female Very-High-Density Cable Interconnect (VHDCI) 0.8 mm type connector. Stratus provides this cable (P/N AW-B24200-12M) with the VHDCI-0.8mm connector at one end and a 68-pin SCSI III connector at the other end for connecting to the SYMMETRIX.

EMC also provides the same cable with P/N 038-001-537, 038-001-538, and 038-001-539 for 6 meters, 12 meters, and 20 meters respectively.

Figure 17-1 below shows the Stratus Continuum 400 Series -- AC POLO 7100 model host cabling configuration for high availability.




Figure 17-1 Host Cabling Configuration (7100)

Figure 17-2 below shows the Stratus Continuum 400 Series -- AC POLO 8000 model host cabling configuration for high availability.




Figure 17-2 Host Cabling Configuration (8000)

Once the Symmetrix has been configured as specified in *Symmetrix Configuration* on page 17-7, and the Stratus host has been cabled to the Symmetrix as shown in Figure 17-1 or Figure 17-2, you can power up the Stratus Host.

Editing the /stand/conf File

Set the dpt1port entry

After installing the U503 card(s) into the system, edit the `/stand/conf` file to include the `lsm` entries for the cards and the `dpt1port` entry. Set the `dpt1port` entry to the bay and slot associated with the U503 cards. See the following examples and text.

For HP-UX 10.20:

```
# cat /stand/conf
# $Stratus_Id: @(#) /vobs/ux/ux10/Stratus/INSTALL/conf@@/main/sra10.20/6 1998-01
-20 12:30 misha $
# Copyright 1997 Stratus Computer Inc. All rights reserved.
rootdev=disc(14/0/0.0.0;0)/stand/vmunix
consdev=(15/2/0;0)
kbddev=(;)
dumpdev=(;)
swapdev=(;)
kernel=HP_UX
lsm0=2/7/1,3/7/1:id0=15,id1=14,tm0=0,tp0=1,tm1=0,tp1=1,rt=1,bt=1
lsm1=2/7/2,3/7/2:id0=15,id1=14,tm0=0,tp0=1,tm1=0,tp1=1
lsm2=2/7/0:id0=7,tm0=1,tp0=1
lsm3=3/7/0:id0=7,tm0=1,tp0=1
lsm4=2/3/0,3/3/0:id0=15,id1=14,tm0=1,tp0=1,tm1=1,tp1=1,rt=1,bt=1,fl0=0x11,fl1=0x
11
dptlport=2/3,3/3
```

For HP-UX 11.01:

```
# $Stratus_Id: @(#) /vobs/ux/ux11/Stratus/INSTALL/conf.jetta@@/main/sra11.00/2 1
998-10-22 11:27 utpal $
# Copyright 1998 Stratus Computer Inc. All rights reserved.
rootdev=disc(14/0/1.0.0;0)/stand/vmunix
consdev=(15/2/0;0)
kbddev=(;)
dumpdev=(;)
swapdev=(;)
kernel=BOOT
lsm0=0/2/7/1,0/3/7/1:id0=15,id1=14,tm0=0,tp0=1,tm1=0,tp1=1,rt=1,bt=1,rt=1,bt=
1,r
t=1,bt=1,rt=1,bt=1,rt=1,bt=1,rt=1,bt=1
lsm1=0/2/7/2,0/3/7/2:id0=15,id1=14,tm0=0,tp0=1,tm1=0,tp1=1,rt=1,bt=1,rt=1,bt=
1
lsm2=0/2/7/0:id0=7,tm0=1,tp0=1
lsm3=0/3/7/0:id0=7,tm0=1,tp0=1
lsm4=0/2/3/0,0/3/3/0:id0=15,id1=14,tm0=1,tp0=1,tm1=1,tp1=1,rt=1,bt=1,fl0=0x11
,fl
1=0x11,rt=1,bt=1
lsm5=0/2/6/0,0/3/6/0:id0=15,id1=14,tm0=1,tp0=1,tm1=1,tp1=1,rt=1,bt=1,fl0=0x11
,fl
1=0x11
dptlport=2/3,2/6,3/3,3/6
```

This **/stand/conf** file *must* be copied to the booting flashcard, usually in bay2. Use the **flifrm** and **flifcp** commands to update the flash card. After this has been performed and verified as operational, update the flash card in the other bay to maintain consistency between the active and standby flashcards.

**CAUTION**

Be sure that the latest kernel patch and any dependencies have been installed on the Stratus/Lucent host.

All patches should be obtained from Stratus (www.stratus.com).
Do not install any patches from HP.

Recognizing New Devices and LUNs

The HP-UX operating system will automatically detect all Symmetrix devices after you have done the following:

- 1. Symmetrix has been configured to support the Stratus Host as specified in section 15.3 of this chapter
- 2. The Stratus Host has been cabled to Symmetrix in a high availability configuration shown in Figure 17-1 and Figure 17-2.
- 3. The Host has been re-booted

To verify that the host has identified the connected devices, type the following:

```
ioscan -fn -C disk [> filename]
```

where filename is the name of file in which to store this information. A listing of all devices found will be stored in the given filename. See the example in the following display.

Label1

```
# ioscan -fn -C disk | more
Class I  H/W Path      Driver  S/W State  H/W Type      Description
=====
.
disk   5   14/0/4.0.0    sdisk   CLAIMED    DEVICE         EMC   SYMMETRIX
        /dev/dsk/c4t0d0 /dev/rdisk/c4t0d0
disk   6   14/0/4.0.1    sdisk   CLAIMED    DEVICE         EMC   SYMMETRIX
        /dev/dsk/c4t0d1 /dev/rdisk/c4t0d1
disk   7   14/0/4.0.2    sdisk   CLAIMED    DEVICE         EMC   SYMMETRIX
        /dev/dsk/c4t0d2 /dev/rdisk/c4t0d2
.
.
.
disk  13   14/0/4.1.0    sdisk   CLAIMED    DEVICE         EMC   SYMMETRIX
        /dev/dsk/c4t1d0 /dev/rdisk/c4t1d0
disk  14   14/0/4.1.1    sdisk   CLAIMED    DEVICE         EMC   SYMMETRIX
        /dev/dsk/c4t1d1 /dev/rdisk/c4t1d1
.
.
```

where in /dev/rdisk/cNtYdZ
N = Logical controller Number
Y = target ID number
Z = LUN number

Note the card instance number next to the H/W Path for the disk devices. Also note that the H/W Path will differ from one machine model to another. Slice (partition) is not supported by HP-UX v10.0 or higher.

Configuring Logical Volumes

Once the Stratus Host has seen all the Symmetrix disks, the HP-UX Logical Volume Manager (LVM) can be used to configure the Symmetrix disks. HP-UX v10.0 and higher does not support hard partitioning or the use of the physical disk in its raw non-LVM form. Using LVM is recommended on Stratus Continuum 400 Series hosts.

A logical volume is a flexible model for storage management. One or more physical devices are introduced into a volume group. Logical volumes are created from storage available for the volume group. Additional functionality such as striping, volumes that span multiple physical devices, etc. are achieved by volume management. You need to understand LVMs before implementing logical volumes on your host. For more information on logical volumes, refer to Chapter 9 in the *HP Concepts for the System Administrator* (P/N B2355-90029).

You must configure logical volumes to store data when using HP 10.x.

Creating Physical Volumes

After you have verified that the host has recognized each Symmetrix device, you need to create a physical volume for each device.

To create a physical volume for each Symmetrix disk device (logical volume), enter a statement similar to the following for each physical volume to create:

```
pvccreate /dev/rdisk/c3t0d0
```

This statement creates a physical volume, `"/dev/rdisk/c3t0d0"`.

You can also create physical volumes using the SAM utility.

Volume Groups & Logical Volumes

At this time you can use physical volumes or create volume groups and logical volumes. Volume groups and logical volumes can be created using HP-UX stand-alone commands or SAM. To create volume groups and logical volumes, refer to the HP documentation set.

Creating the Mount Directory

If you want to use filesystems on the volumes, you must create a mount directory for each of these volumes. To do this, type a statement similar to the following for each physical volume:

```
mkdir /symm
```

This example creates a mount directory called “symm”.

Creating a New Filesystem

Now that the physical volumes and their respective mount directories are available to the operating system, a new filesystem must be created for each Symmetrix disk. To do this, use the **newfs** command in a statement similar to the following:

```
newfs /dev/vg-group/volx
```

Mount the Directory

For HP-UX 10.x and volume groups, mount the directory by adding the path to **/etc/fstab**. To do this, type the following statement:

```
vi /etc/fstab
```

Specify each device in a statement similar to the following:

```
/dev/vg-group/volx
```

Mount the filesystem by typing a statement similar to the following in either operating system:

```
/etc/mount /symm
```

Use the **bdf** command to show all mounted filesystems and the available free space. For example,

```
bdf
```

Filesystem	kbytes	used	available	%used	mounted on
/dev/vg00/lvol1	104299	52936	40933	56%	/
/dev/vg00/lvol5	349530	125692	188884	40%	/users
/dev/vg00/lvol4	124507	836	111220	1%	/tmp
/dev/vg00/lvol3	397867	137494	220586	38%	/usr
/dev/dsk/c10d0s22741142		104041	2362986	4%	/symm5
/dev/dsk/c11d0s22741142		0	2741142	0%	/symm
/dev/dsk/c12d0s22741142		104041	2362986	4%	/symm3
/dev/dsk/c13d0s22741142		104041	2362986	4%	/symm1
/dev/dsk/c29d0s22741142		104041	2362986	4%	/symm4
/dev/dsk/c30d0s22741142		104041	2362986	4%	/symm2

shows **/dev/dsk/c11d0s2** mounted.

Using SAM

HP-UX has a menu-driven utility for managing system resources, specifically disk drives. This utility is called SAM and users familiar with it might find it more useful to use SAM for configuring new Symmetrix disks.

Adding Devices Online

Whenever devices are added online to the Symmetrix unit or device channel addresses are changed, you will need to perform the actions described below in order to introduce the new devices to the system.

You can define newly connected physical volumes to the system without rebooting the host. This requires use of the **insf** command in a statement similar to the following:

```
insf -H 2/4.0.0
```

where **2/4.0.0** is a hardware path to which the new device has been connected.

Alternatively, you can use the following command sequence to look for all new hardware devices:

```
ioscan -fn  
insf -e
```

System and Error Messages

HP-UX 10.x logs system and error messages including SCSI error messages to a file called `/var/adm/syslog/syslog.log`.

Diagnostic Patches

Installing the appropriate patches for your operating system level causes Symmetrix drives to be recognized properly by HP hosts.

Symmetrix definitions for STM and Sysdiag should already exist for operating system 10.20.

These patches have an effect on the output of the STM and sysdiag utilities when run on Symmetrix devices. The effects of these patches are as follows:

- ◆ SCSTIDSK2 - HP's SCSI disk diagnostic utility (which includes commands like INQUIRY, CAPACITY, etc.) now works with Symmetrix devices, but reports the following message:

`SCSTIDSK2 is not designed to test the selected device (which is "EMCSYMMETRIX")`

- ◆ The IOMAP utility recognizes Symmetrix devices as EMCSYMMETRIX DISK. (If the patch is not installed, Symmetrix devices are reported as "Device Unknown".)
- ◆ In the System Error Log Report, any entries in the log for Symmetrix devices now report the product name as "EMCSYMMETRIX" and the device class as "Hard Disk".

`logtool` decodes the hex bytes presented in the environmental data to text.

HP's patch for operating system 10.20 is PHSS_8590.

Lucent Environment

This chapter lists the Lucent (formerly Ascend) DNCP-400 and identifies it with the corresponding Stratus model.

- ◆ DNCP-400 Hardware System..... 18-2

DNCP-400 Hardware System

The Lucent (formerly Ascend) Distributed Network Computing Platform (DNCP-400) corresponds to the Stratus Continuum 400 model.

For more information about the Stratus/Lucent environment, see the Stratus chapter.

This chapter describes how to incorporate Symmetrix into the Santa Cruz Operation UnixWare environment.

- ◆ Overview 19-2
- ◆ Host Bus Adapters 19-3
- ◆ Useful Functions and Utilities 19-7
- ◆ Symmetrix Configuration 19-9
- ◆ Configuring a New Device 19-10
- ◆ Adding Devices Online 19-15
- ◆ System and Console Error Messages 19-16

Overview

When using the Symmetrix in a SCO Unixware environment:

- ◆ The SCO host must have either Wide or Ultra Wide SCSI Differential cards installed in order to connect to the Symmetrix.
- ◆ When using Symmetrix with a SCO UnixWare host, the host must be running SCO Unixware 7.01 or greater.
- ◆ For a list of required and qualified optional patches, see your EMC Customer Engineer.
- ◆ Targets 0 through f are valid for Symmetrix devices, with the conventional exception of 7 which is reserved for Initiator ID use.
- ◆ The Multipath I/O feature is not qualified for UnixWare version 7.0.0, 7.0.1 or 7.1. Disable MPIO before connecting to Symmetrix.

Table 19-1 SCO UnixWare 7 Summary of Features

O/S	SCO UnixWare 7 (7.00 or greater)
Largest recognizable drive size	1 TB
Largest defined partition	1 TB
Largest data file	1 TB possible -- defaults to 2 GB
Largest filesystem	1 TB
Supports multiple target IDs	Yes
Supports multiple LUNs per target ID	Yes
Partition representing the entire disk	s0
To see the drive configuration, use these utilities:	fdisk, prtvtoc
To configure drives use these utilities:	fdisk, disksetup

Host Bus Adapters

Consult your EMC sales representative for the most current listing of supported hosts and supported HBAs.

Table 19-2 QLogic Host Bus Adapter QLA2200F-EMC

Hosts Qualified by EMC:	For qualified hosts and BIOS, see your EMC representative for information in the latest revision of the Open Systems Support Matrix.
Targets supported	1 target supported per HBA (FC-AL)
LUNs supported	0-127
Special OS Patches Needed	Contact your EMC customer representative for the latest required patches.
Firmware	Version 1.59
Driver	3.10
HBA Driver configuration	Press <Alt> Q at boot time
Notes	For complete setup procedures, refer to the latest revision of the Symmetrix Fibre Channel Product Guide, Part Number 200-999-642.

Table 19-3 QLogic Host Bus Adapter QLA1240D

Hosts Qualified by EMC:	For qualified hosts and BIOS, see your EMC representative for information in the latest revision of the Open Systems Support Matrix.
Operating System Level	7.0.1
Targets supported	0-F
LUNs supported	0-F
Special Patches Needed	None
Firmware	1.12
Driver	Driver Version 3.02 ^a
HBA Driver configuration	Disable BIOS, Auto-Configure is recommended
Notes	Supports LUN Skipping. Supports LUNs out of sequence. LUNs series must begin with 0.

a. This driver is available from QLogic's web page.

Table 19-4 QLogic Host Bus Adapter QLA1041D

Hosts Qualified by EMC:	For qualified hosts and BIOS, see your EMC representative for information in the latest revision of the Open Systems Support Matrix.
Operating System Level	7.0.1
Targets supported	0-F
LUNs supported	0-F
Special Patches Needed	None
Firmware	2.16
Driver	Driver Version 2.07 ^a
HBA Driver configuration	Disable BIOS Auto Configure is recommended
Notes	Supports LUN Skipping. Supports LUNs out of sequence. LUNs series must begin with 0.

a. This driver is available from QLogic's web page.

Table 19-5 Adaptec Host Bus Adapter AHA2944UW^a

Hosts Qualified by EMC:	For qualified hosts and BIOS, see your EMC representative for information in the latest revision of the Open Systems Support Matrix.
Operating System Levels	7.0.0 7.0.1
Targets supported	0-F
LUNs supported	0-7
Special Patches Needed	Contact your EMC customer representative for the latest required patches.
Firmware	1.34
Driver	7.0.1 - Version 7.0.1a
HBA Driver configuration	Max Sync Transfer Rate (Mb/Sec): 40.0 Initiate Wide Negotiation: Yes Disable BIOS
Notes	Will only run Ultra speed Narrow. No LUN Skipping. No LUNs out of sequence. LUNs series must begin with 0.

a. For more information regarding SCSI parameter settings for this card, refer to Appendix C.

Useful Functions and Utilities

This section lists SCO UnixWare 7 functions and utilities you may use to define and manage Symmetrix devices. The use of these functions and utilities is optional. They are listed for reference only.

- ◆ `prtvtoc`

This command displays disk information from the VTOC table (volume table of contents).

- ◆ `fdisk`

This command creates and modifies the partition table.

- ◆ `disksetup`

This command performs low-level activities required to install the primary drive or additional drives. `disksetup` calls the `mkfs` and `mount` commands.

- ◆ `fsck`

This command audits and interactively compares inconsistent conditions for file systems.

- ◆ `mkdir`

This command creates a specified directory.

- ◆ `sysdef`

This command analyzes the default bootable kernel and extracts configuration information.

- ◆ `mkfs`

This command makes a file system on a raw device file.

◆ `mount`

This command attaches a file system to a specified directory called the mount point.

◆ `umount`

This command un-mounts a mounted file system.

◆ `shutdown`

This command shuts down the system gracefully.

◆ `df` or `dfspace`

These commands summarize free disk space.

◆ `scoadmin`

Similar to the GUI system administration program. Used in the ASCII environment. Type `scoadmin` at any prompt.

Graphic User Interface system administration is available at `<ctr><alt><f1>`. (Many command line administrative functions can be done through the GUI interface.)

Symmetrix Configuration

In the SCO UnixWare 7 host environment, you can configure the Symmetrix disk devices into logical volumes. EMC recommends the following logical volume ratios as indicated in the table below:

Table 19-6 Disk Device Configuration

		Symmetrix Logical Volume to Physical Splits ^a							
		1:1	2:1	3:1	4:1	5:1	6:1	7:1	8:1 ^b
Physical Devices	4 GB	x	x						
	9 GB	x	x	x	x				
	18 GB		x	x	x	x	x	x	x
	23 GB		*	x	x	x	x	x	x
	36 GB			*	x	x	x	x	x
	47 GB			*	*	x	x	x	x
	50 GB ^c				x	x	x	x	x
	73 GB ^{c, d}					x	x	x	x

- * indicates ratios available only in microcode level 5265 or higher.
- With microcode revision level 5263, maximum is 8:1. With 5264, maximum is 16:1. With 5265, maximum is 32:1.
- The 50 and 73 GB drives are supported only by microcode revision levels 5x66 and higher.
- The 73 GB drive does not support the RAID-S configuration on Symmetrix.

For specific cylinder counts and logical volume sizes, refer to Table 2-1 on page 2-5 and the following tables in that chapter.

The EMC Customer Engineer (CE) should contact the EMC Configuration Specialist for updated online information. This information is necessary to configure Symmetrix to support the customer's host environment. After the EMC CE has assigned target IDs and LUNs and brought the Symmetrix channel and disk directors online, you must reboot the host computer. During reboot use the HBA utility to verify SCSI settings (see the HBA section of this chapter and Appendix C). After reboot is complete, Symmetrix raw devices should be seen in the `/dev/rdisk/` directory.

Configuring a New Device

Disk Device File Naming Conventions

Raw disk device filenames take two forms:

`/dev/rdisk/cCbBtTdDsS`

or

`/dev/rdisk/cCbBtTdDpP`

where:

cC

C is the occurrence of the host adapter board. C can range from 0 to 31.

bB

B is the HBA bus number, ranging from 0 to 7 (dependent on the add-on adapter).

tT

T is the SCSI target controller number (ID).

dD

D is the logical unit number of the disk device.

sS

S is the disk slice number, ranging from 0 to b7 (hex).

pP

P is the fdisk partition number. P can range from 0 to 4. The p0 node corresponds to the whole disk, the p1 to p4 nodes correspond to partitions displayed by the fdisk command.

Example: `c0b0t0d0s0`

Format and Partition a Disk

To format and partition a Symmetrix device for the use of SCO UNIX, log in as root and proceed as follows for each new device.

Once the Symmetrix and the host have been integrated there will be raw device files in the **/dev/rdisk/** directory representing the target and LUNs specified by the EMC customer engineer.

Use the `fdisk` command to partition each Symmetrix raw device in a statement similar to the following:

```
fdisk /dev/rdisk/c0b0t0d0s0
```

c0b0t0d0s0 represents any raw Symmetrix device visible in the `/dev/rdisk/` directory.

- ◆ If there is no partition table on the Symmetrix device, you have the option to create a default partition or to specify the initial table values. Select "Y" and <enter> for the default option. The default option will allocate 100% of the disk for the UnixWare system and make the UNIX system partition active.
- ◆ If there is an existing partition table, the command program will display a partition table as shown below and a menu with options allowing you to modify the table.

Table 19-7 Partition Table

Partition	Status	Type	Start	End	Length	%	MB
1	Active	UNIX System	0	521	522	100	4094.7

SELECT ONE OF THE FOLLOWING:
0. Overwrite system master boot code
1. Create a partition
2. Change Active (Boot from) partition
3. Delete a partition
4. Exit (Update disk configuration and exit)
5. Cancel (Exit without updating disk configuration)
Enter Selection:

- ◆ **Warning: Changing the partition table will result in loss of any previous data.** To exit without making any changes to the partition table, type 5 and press <enter>.
- ◆ If the type of partition is **UNIX System** and the status of the UNIX System partition is **Active**, type 4 and press <enter>. This will exit and update the disk configuration.
- ◆ To activate a UNIX System partition, type 2 and press <enter>. You will be prompted to enter the partition number to activate or to boot from. Select the number of the UNIX System partition.
- ◆ To create a UNIX System partition, type 1 and press <enter>. You will be prompted with a list of partition types. Type 1 and press <enter> for a UNIX System partition.

- ◆ You will then be prompted for the percentage of the disk you would like use for the partition. Type 100 and press <enter>. If 100% of the disk is not available, use as much of the disk that is available. The UNIX System partition requires at least 4% of the disk.
- ◆ You will then be prompted with the option to activate the partition. Type y and press <enter>. The UNIX System partitions must be in the **Active** state to be accessible.
- ◆ If at the main menu there is a UNIX System partition that has a status of **Active**, select type 4 and press <enter>. This will exit partitioning and update disk configuration.

Setting Up a Raw Device

To create a raw device under SCO UNIX, log in as root and proceed as follows for each new device.

Once you have run the `fdisk` command on each Symmetrix disk device, use the `disksetup` command in a statement similar to the following:

```
disksetup -I /dev/rdisk/c0b0t0d0s0
```

NOTE: `c0b0t0d0s0` represents any raw Symmetrix device visible in the `/dev/rdisk/` directory. When using the `disksetup` command you must use only slice 0 of the raw device.

- ◆ You will have the option to do a surface analysis of the disk. A surface analysis is recommended but not required. Type y or n and press <enter>.
- ◆ You will then be prompted for how many slices/filesystems you want created on the disk (1 - 13). Type 1 and press <enter>.
- ◆ For the created slice, you will be prompted to enter the absolute pathname to serve as a mount point. Enter the path `/mnt/` and press <enter>. No file system will be mounted.
- ◆ For the created slice, you will be prompted to enter the filesystem type. Type na and press <enter> to not create a file system.

- ◆ For the created slice, you will be given the option to automatically remount the file system during a reboot. Type `no` and press `<enter>`.
- ◆ You will now specify the size in cylinders of each slice. The maximum number of cylinders will be provided in a display similar to the following:

```
(One megabyte of disk space is approximately 1 cylinders.)
There are now 10780 cylinders available on your disk.
How many cylinders would you like for /mnt/ (0 - 10780)?
Press <enter> for 0 cylinders:
```

Type in the maximum number of cylinders and press `<enter>`.
In this example you would type `10780` and then press `<enter>`.

- ◆ After you have entered the cylinders allocated, the allocations you specified will be displayed. You will be asked whether the allocations are acceptable to you. Type `y` and press `<enter>`.
- ◆ Use only slice 1 for I/O with a raw device.
- ◆ You can use the `prtvtoc` command to verify that slice 1 is valid and unmountable.

Create and Mount a File System

To create file systems under SCO UNIX, log in as root and proceed as follows for each new device.

Before using the `disksetup` command you must predetermine directories that will serve as mount points for each slice of each device created. (See the `mkdir` command in Section)

Once you have run the `fdisk` command on each Symmetrix disk device and created mount points, use the `disksetup` command in a statement similar to the following:

```
disksetup -I [e] /dev/rdisk/c0b0t0d0s0
```

`c0b0t0d0s0` represents any raw Symmetrix device visible in the `/dev/rdisk/` directory. **When using the `disksetup` command you must only slice 0 of the raw device.**

The `Disksetup -I` option defaults to a maximum of 13 slices per device. With the `Disksetup -Ie` option, the maximum number of slices that can be created is 184.

- ◆ You will have the option to do a surface analysis of the disk. A surface analysis is recommended but not required. Type `y` or `n` and press `<enter>`.
- ◆ You will then be prompted for how many slices/filesystems you want created on the disk (1 - 13). Type the number of slices needed and press `<enter>`. If the extended slice option was used in the disk setup command, you may enter up to 184 slices.
- ◆ For each slice created, you will be prompted to enter the absolute pathname which will serve as a mount point for the file system you create. Enter the full predetermined path of the mount point and press `<enter>`.
- ◆ For each slice created, you will be prompted to enter the filesystem type (`vxfs`, `ufs`, `s5`, or `sfs`). A `vxfs` file system is the default. Press `<enter>` to select the default `vxfs` file system.
- ◆ For each slice created, you will be prompted to enter the block size (1024, 2048, 4096 or 8192). The default block size is 1024. Press `<enter>` to select the default.
- ◆ For each slice created, you will be given the option to automatically remount the file system during a reboot. The default option is yes. Press `<enter>` for auto remount.
- ◆ After all mount points, block sizes and auto remount options have been determined, you will be prompted to specify the size in cylinders of each slice. You will be given the number of cylinders available and which mount point you will be allocating cylinders for. Type in cylinder allocations and press `<enter>`.
- ◆ After you have entered the cylinders allocated for each mount point, the allocations you specified will be displayed and you will be asked whether the allocations are acceptable to you. Type `y` and press `<enter>` if the allocations are correct. Press `n` and `<enter>` to start again.
- ◆ If all of the allocations were correct, you will be prompted with a default number of inodes for each file system. Press `<enter>` to accept default number of inodes.
- ◆ To verify that all file systems were created at any command prompt, type `df`. This will display the all mount points and associated raw devices, sizes in blocks, and the number of files.

Adding Devices Online

Whenever devices are added online to the Symmetrix unit or device channel addresses are changed, you will need to perform the actions described below in order to introduce the new devices to the system.

SCO UnixWare 7 currently does not support adding devices online.

The EMC PSE Configuration Group can add new devices to the Symmetrix configuration while Symmetrix remains online to the host. Once this is complete, stop all host activity and reboot the host. The new Symmetrix devices should be seen in the **/dev/rdisk/** directory.

Follow the instructions in *Configuring a New Device* on page 19-10 to introduce the new devices.

System and Console Error Messages

Console error messages can be viewed in the `/var/adm/syslog` file.

System I/O errors that may appear on the application screen are as follows:

```
System error 2 No such file or directory
System error 5 I/O Error
System error 6 No such device or address
System error 11Resource temporarily univaible
System error 16Mounted Device Busy
System error 19No such device
System error 22 Invalid argument
```

For a complete list of error codes see the file `/usr/include/sys/errno.h`

This chapter describes how to incorporate Symmetrix in the Linux environment.

◆ Host Connectivity	20-2
◆ Useful Linux Utilities and Functions	20-3
◆ Symmetrix Configuration	20-5
◆ Host Configuration with Adaptec AHA-2944UW HBA	20-7
◆ Kernel Configuration & Compilation with Adaptec HBA	20-10
◆ Configuring and Compiling the Kernel.....	20-16
◆ Partitioning and Labeling a New Device	20-24
◆ Creating and Mounting a File System	20-29
◆ System and Error Messages.....	20-30

Host Connectivity

As of the current revision of this manual, Symmetrix is qualified for use with Linux hosts on SCSI. Contact your EMC representative for the latest information on supported hosts and HBAs, as updates are frequent. See the following table.

Table 20-1 HBAs for Linux Hosts on SCSI

Host Type	IA32 x86 platforms
Qualified OS versions	See your EMC representative for online information in the most recent EMC Symmetrix Open Systems Support Matrix.
Host Bus Adapter	Adaptec AHA-2944UW

Useful Linux Utilities and Functions

You may optionally use these Linux functions and utilities to define and manage Symmetrix devices, partition tables, and file systems. These commands are listed for reference only.

- ◆ `fdisk`
Command used to create and manipulate Linux partition tables.
- ◆ `cfdisk`
Curses-based disk partition table manipulator for Linux.
- ◆ `fsck`
Command used to check and repair a Linux file system.
- ◆ `mkfs`
Command used to create a Linux file system on a device.
- ◆ `mount`
Command used to attach the file system on a device to the file tree.
- ◆ `umount`
Command used to detach a file system.
- ◆ `shutdown`
Command used to shut down the system gracefully.
- ◆ `reboot`
Command used to stop and restart the operating system.

- ◆ `modprobe`

Utility used to load or remove a set of modules that can be either a single module or a stack of dependent modules.

- ◆ `insmod`

Utility used to dynamically load a single module into the running kernel.

- ◆ `rmmod`

Utility used to unload loadable modules from the running kernel if they are not in use and if other modules are not dependent upon those being removed.

Symmetrix Configuration

In the Linux host environment, you can configure the Symmetrix disk devices into logical volumes. EMC recommends the following logical volume ratios as indicated in the table below:

Table 20-2 Disk Device Configuration

		Symmetrix Logical Volume to Physical Splits ^a							
		1:1	2:1	3:1	4:1	5:1	6:1	7:1	8:1 ^b
Physical Devices	4 GB	x	x						
	9 GB	x	x	x	x				
	18 GB		x	x	x	x	x	x	x
	23 GB		*	x	x	x	x	x	x
	36 GB			*	x	x	x	x	x
	47 GB			*	*	x	x	x	x
	50 GB ^c				x	x	x	x	x
	73 GB ^{c, d}					x	x	x	x

- a. * indicates ratios available only in microcode level 5265 or higher.
- b. With microcode revision level 5263, maximum is 8:1. With 5264, maximum is 16:1. With 5265, maximum is 32:1.
- c. The 50 and 73 GB drives are supported only by microcode revision levels 5x66 and higher.
- d. The 73 GB drive does not support the RAID-S configuration on Symmetrix.

For specific cylinder counts and logical volume sizes, refer to Table 2-1 on page 2-5 and the following tables in that chapter.

Set the director bits listed in the following table to configure Symmetrix for Linux hosts with SCSI connectivity.

Table 20-3 HBAs for Linux Hosts on SCSI

Set Director Bits
TAGD
LINK
WIDE
SYNC

Linux supports the use of SCSI targets 0-F, excluding 7. 7 is reserved for the HBA.

LUNs 0-7 are supported. There is a limitation in Linux that requires the use of LUN 0. Skipping LUNs after zero is acceptable.

Linux v2.2.x supports accessing up to 128 devices in a range of 256. This limitation is per host, not per adapter.

To verify that the Linux host can see all of the target Symmetrix devices, configure the host according to the section, *Host Configuration with Adaptec AHA-2944UW HBA* on page 20-7, and use the command `#more /proc/scsi/scsi`. This command provides a listing of the attached Symmetrix devices.

Host Configuration with Adaptec AHA-2944UW HBA

Installing the HBA

The following procedures must be done with the Adaptec AHA-2944UW installed in the host.

Configuring BIOS Settings

Verifying the Correct BIOS Version

The HBA BIOS version can be determined at boot time from the Adaptec AHA-2944UW banner or from the SCSISelect options menu.

- ◆ To determine the firmware version from the Adaptec AHA-2944UW banner:
 - a. Boot the host and watch for the banner like the one shown below:

```
Adaptec AHA-2944UW Differential SCSI BIOS v2.20.0
© 1999 Adaptec, Inc. All Rights Reserved
Press <Ctrl> <A> for SCSISelect™ Utility
```

Figure 20-1 Adaptec AHA-2944UW Banner

- b. If the ROM BIOS is v2.20.0, continue to *Setting the BIOS Variables* on page 20-9.
 - c. If the ROM BIOS is not v2.20.0, upgrade the firmware as described in *Upgrading the HBA BIOS* on page 20-8, then go to *Setting the BIOS Variables* on page 20-9.
- ◆ To determine the firmware version from the Adaptec AHA-2944UW SCSISelect menu:
 - a. Boot the host. As soon as the Adaptec AHA-2944UW banner appears, press <Ctrl> <A> to enter the SCSISelect menu.
 - b. In the header, note the BIOS version:


```
Adaptec <SCSISelect> Utility v2.20.0
```
 - ◆ If the version is v2.20.0, continue to *Setting the BIOS Variables* on page 20-9.
 - ◆ If the version is not v2.20.0, upgrade the firmware as described in *Upgrading the HBA BIOS* which follows, then go to *Setting the BIOS Variables* on page 20-9.

Upgrading the HBA BIOS

The HBA BIOS can be obtained from the Adaptec website at the following url: <http://www.adaptec.com>.

1. From the main page, click on Downloads.
2. Scroll to the High Voltage Differential PCI SCSI listing.
3. Click on AHA-2944W/AHA-2944UW.
4. Scroll to the EROM/BIOS section and click on the 2944V220.exe.
5. Save the BIOS v2.20.0 file (2944v220.exe) to a local directory.
6. The file is a self-extracting file. Run the 2944v220.exe.
7. Copy the expanded files to a bootable DOS floppy diskette.
8. Use an IBM formatted diskette on any Microsoft Windows machine to create a DOS bootable diskette:
 - a. Open a DOS window.
 - b. Type `format /s a:` to format the floppy diskette.
9. When the formatting is complete, change to the directory in which you saved the BIOS v2.20.0 file (2944v220.exe). Copy the expanded BIOS files to the formatted floppy diskette.

To upgrade the BIOS, complete the following steps:

1. Insert the installation diskette in the host's diskette drive.
2. Reboot the host.
3. After the host has rebooted, the DOS prompt appears. To upgrade the BIOS, type `UPDATE` at the `a:\` DOS prompt.
4. After the upgrade is complete, remove the diskette and reboot the host. During the boot process, the Adaptec AHA-2944UW banner (Figure 20-1) should display BIOS version v2.20.0.

The Adaptec utility only updates one controller at a time. If you have multiple controllers, you will need to remove all controllers except the one you intend to update.

Setting the BIOS Variables

1. As the host boots, press Ctrl A when prompted Press <Ctrl> <A> for SCSISelect™ Utility.
2. From the SCSISelect menu, select Configure/View Host Adapter Settings.
3. Under the Additional Operations heading, select Advanced Configuration Options.
4. Select Host Adapter BIOS and change to Disabled: NOT scan.
5. Press ESC to exit the SCSISelect menu and save the changes.
6. Repeat steps 1 - 5 for each additional AHA-2944UW adapter.
7. Reboot the host.

Kernel Configuration & Compilation with Adaptec HBA

Introduction

This section applies to kernel configuration and compilation with the Adaptec AHA-2944UW HBA.

The host must be configured with Linux first. Install Linux from the RedHat v6.x or v7.0 CD following the procedure provided in the RedHat installation guide. Partition the boot drive and select the packages and services necessary for the host.

1. In order to load the modular `aic7xxx` driver automatically at boot time, the driver should be listed in `/etc/conf.modules`. If it is not added by default, `/etc/conf.modules` must be edited.

```
vi /etc/conf.modules
```

Add the appropriate entry:

```
alias scsi_hostadapter aic7xxx
```

2. Proceed to the *Required Kernel Modifications* section below to modify the Linux kernel.

Required Kernel Modifications

Edit `hosts.h`

In order to see more than four SCSI hard drives, edit `hosts.h` as follows:

```
cd /usr/src/linux/drivers/scsi
vi hosts.h
```

The text of `hosts.h` differs depending upon the RedHat distribution:

- ◆ For v6.x

```
#define SD_EXTRA_DEVS 40
#else
#define SD_EXTRA_DEVS 4      (Change to 128.)
#endif
```

- ◆ For v7.0

```
#define SD_EXTRA_DEVS 40      (Change to 128.)
#define ST_EXTRA_DEVS 3
#define SR_EXTRA_DEVS 8
```

Edit sg.h

The value for the SCSI generic buffer must be increased. Change to the appropriate directory and increase the value for `SG_SCATTER_SZ` to the maximum of 131072.

```
cd /usr/src/linux/include/scsi
vi sg.h
```

Around line 195, look for the following code block:

```
#define SG_SCATTER_SZ (8 * 4096) /* PAGE_SIZE not available to user */
/* Largest size (in bytes) a single scatter-gather list element can have.
The value must be a power of 2 and <= (PAGE_SIZE * 32) [131072 bytes on
i386]. The minimum value is PAGE_SIZE. If scatter-gather not supported
by adapter then this value is the largest data block that can be
read/written by a single scsi command. The user can find the value of
PAGE_SIZE by calling getpagesize() defined in unistd.h . */
```

Change the value for `SG_SCATTER_SZ` (`8 * 4096`) to (`32 * 4096`).

After the kernel compilation and system reboot, using the following command can validate this change:

```
cat /proc/sys/kernel/sg-big-buff
```

Proceed to the Kernel Configuration and Compilation section to compile the Linux kernel.

Additional Notes

Dynamic Device Reconfiguration

Linux currently lacks a command built into the kernel that allows for a dynamic SCSI channel reconfiguration (such as the commands `drvconfig` or `ioscan`). Depending upon the hardware configuration and on whether the SCSI driver has been loaded into the kernel as a module, the SCSI driver can be unloaded and then reloaded. If the SCSI driver has been loaded as a module, this unloading and reloading will cause a scan of the SCSI bus. The driver should detect the newly added device(s). If the SCSI driver is loaded into the kernel statically, the system will need to be rebooted in order for the devices to be detected reliably. An example of a system requiring a reboot is one that has an internal hard drive booting from an AIC7xxx on-board SCSI chip plus an Adaptec AHA-2944UW for external device connectivity. In this case, even if the AIC7xxx driver were loaded automatically as a module at boot, it would not be possible to unload and reload the SCSI driver since the internal hard drive is dependent upon it.

If the host has an on-board SCSI chip other than the AIC7xxx, it is possible to use the AIC7xxx driver for the AHA-2944UW as a module so it can be unloaded and reloaded to cause a rescan of the SCSI bus.

The `modprobe` command is a wrapper for or an extension to `insmod` and `rmmod`. It uses and maintains a set of files that describes all the modules that are available for the current kernel in `/lib/modules`. You can use `modprobe` to load and unload an entire set of modules in addition to just a single module. A set of modules can be a single module or a stack of dependent modules. When `modprobe` is used, it automatically loads all of the modules needed in a module stack based on the information in `modules.dep`.

The unloading of the module can be done with either the `modprobe` or `rmmod` commands. When unloading modules, `modprobe` needs to be used with the `-r` switch. These commands are used to unload the loadable modules from the running kernel if they are not in use and if other modules are not dependent upon them.

The loading of the module is accomplished with either the `modprobe` or `insmod` commands. These commands are used to install loadable kernel modules in a running kernel. They will attempt to link a module into the running kernel by resolving the symbols from the kernel's exported symbol table.

In the standard Linux v2.2.x kernel, the SCSI addresses are not used in the device names as they are in other types of Unix such as Sun, SGI, HP-UX, BSD, etc. Block device filenames take the form `/dev/sdln` where `l` is the letter denoting the physical drive and `n` is the number denoting the partition on that physical drive. Device names and minor numbers are assigned dynamically at boot time or device loading time in the order of discovery.

Depending upon the hardware configuration, if a device is added and the system rebooted, the device numbering can and will change. This change makes the host's mount table inaccurate. For the most consistent results and to reduce the possibility of mount table inaccuracies, new devices should be appended to the list of already attached devices. For example, if the host contains multiple HBAs, it would be best to append the new device to the last HBA and to the end of the device list attached to that HBA. This would eliminate the need to alter the pre-existing entries in the mount table since the new device could be appended to that as well. If a new device were added to the first out of two HBAs and the system rebooted, the devices would all shift by one number and the mount table entries would also need to be shifted by one device. If there is only one HBA, the

new device can more easily be appended to the list of the regularly attached devices and the mount table altered accordingly.

tasks.h

There is a limit to the number of processes that can run simultaneously on a host. The default is set to 512. If the environment requires it, this value can be increased up to 4092.

To change the value,

```
cd /usr/src/linux/include/linux
vi tasks.h
```

```
line 14      #define NR_TASKS 512      /*On x86 Max 4092, or 4090 w/ APM configured.*/
```

Change the value of 512 to whatever is required up to a maximum of 4092, then proceed to *Configuring and Compiling the Kernel* on page 20-16.

BLIST_NOLUN

Hosts with a GEM SCSI chip on the motherboard may incorrectly report multiple processors when booting. This incorrect reporting can be fixed by adding an entry to BLIST_NOLUN in the file `scsi.c`.

For example, in the NEC ExpressServer 5800 120Ra2, the GEM chip incorrectly reports back 8 LUNs instead of 1 LUN as it should. This problem can be corrected by editing `scsi.c` as follows:

```
cd /usr/src/linux/drivers/scsi
vi scsi.c
```

Then search for

```
static struct dev_info device_list [] =
{"Aashima", "IMAGERY 2400sp", "1.03", BLIST_NOLUN},
                                /*Locks up if polled for lun != 0 */
{"CHINON", "CD-ROM CDS-431", "H42", BLIST_NOLUN},
                                /*Locks up if polled for lun != 0*/
```

where

```
Aashima = Make
IMAGERY2 2400sp = Model
1.03 = firmware
BLIST_NOLUN = flag
```

Go to the end of the `BLIST_NOLUN` list and add the device.

In the case of the NEC ExpressServer 5800 120Ra2, the entry would be similar to the following:

```
{ "NEC", "GEM312R2-G7CGQ", "4.1a", BLIST_NOLUN}, /*Reports back incorrect LUNs*/
```

After this addition is made, proceed to *Configuring and Compiling the Kernel* on page 20-16.

Reverse Scan

If you are using a host that is a bridged PCI system, you may experience problems with regard to the order of detection of the SCSI controllers. Occasionally, the SCSI controller used to boot the boot driver is detected last, detected after the SCSI controllers connected to the Symmetrix. This causes the system to panic and prevents it from booting. You can fix this problem by setting the `reverse_scan` parameter to a non-zero number so that the sort order of the SCSI controllers will be reversed.

To change the value to a non-zero number:

```
cd /usr/src/linux/drivers/scsi
#vi aic7xxx.c

static int aic7xxx_reverse_scan = 0;
```

Change 0 to a non-zero number (1 will work) and save the file.

Proceed to *Configuring and Compiling the Kernel* on page 20-16.

Modular vs. Static SCSI Driver Discussion

At boot time, the Linux kernel loads statically linked drivers first and modular drivers second, in alphabetical order. When using the Adaptec AHA-2944UW, hosts with on-board SCSI chips (other than the `aic7xxx` chipset) may panic and fail to boot if the SCSI drivers are selected as all modular or as all static.

For example, the Compaq Proliant 1850 has a Symbios chip on the motherboard. The Symbios `sym53c8xx` driver must be loaded before the Adaptec `aic7xxx` driver so that the internal boot drive will be

detected before the Symmetrix devices attached to the AHA-2944UW. If both drivers are selected as either modular or static, the Adaptec `aic7xxx` driver will load prior to the Symbios `sym53c8xx` driver.

The failure to boot can be resolved by selecting the `sym53c8xx` driver as static and the `aic7xxx` driver as modular. This will ensure that the on-board SCSI driver will always load first.

Follow the procedures listed in *Configuring and Compiling the Kernel* on page 20-16.

```
cd /usr/src/linux
make menuconfig
```

Under the SCSI low-level drivers menu, select Adaptec AIC7xxx support as modular (M or m):

```
[M] Adaptec AIC7xxx support
    [*] Enable Tagged Command Queuing (TCQ) by default
    (32) Maximum number of TCQ commands per device
    [*] Collect statistics to report in /proc
    (5) Delay in seconds after SCSI bus reset
```

Select SYM53C8XX SCSI support as static (* or y):

```
<*> SYM53C8XX SCSI support
    (4) Default tagged command queue depth
    (32) maximum number of queued commands
    (20) synchronous transfers frequency in MHz
```

Tagged Command queuing should be enabled for both drivers.

Configuring and Compiling the Kernel

If this is a new installation, proceed to the *New Installation* section below.

If this is an upgrade to a newer kernel revision, proceed to *Kernel Upgrade* on page 20-19.

New Installation

Install Linux from the RedHat v6.x or v7.0 CD following the procedure provided in the RedHat installation guide. Partition the boot drive and select the packages and services necessary for the host.

1. Prior to configuring the host, change to the appropriate directory:

```
cd /usr/src/linux
```

2. Configure the kernel

```
# make config/menuconfig/xconfig
```

3. From the Main Menu, select the Processor type and features menu:

a. Select the appropriate Processor family for the processors in the system.

```
( ) 386
( ) 486/Cx486
( ) 586/K5/5x86/6x86
( ) Pentium/K6/TSC
( ) PPro/6x86MX/PII
( ) PIII/Xeon/Deschutes
```

b. Enable Symmetric multi-processing support as static *only* if the system being used has multiple processors.

c. Select EXIT to return to the Main Menu.

4. From the Main Menu, select the SCSI support menu. The menu shown in Figure 20-2 should appear.




Figure 20-2 SCSI Support Menu

In the SCSI support menu shown in Figure 20-2, make the following selections:

Select SCSI support as static (* or y).

Under

— SCSI support type (disk, tape, CD-ROM)
select the options SCSI disk support and SCSI generic support as static (* or y).

Under

— Some SCSI devices (e.g. CD jukebox) support multiple LUNs
select Probe all LUNs on each SCSI device as static (* or y).

5. Scroll down the SCSI support menu to the SCSI low-level drivers submenu, as shown in Figure 20-3.




Figure 20-3 SCSI Low-Level Drivers Submenu

As shown in Figure 20-3, under SCSI low-level drivers —> select Adaptec AIC7xxx support as modular or static depending upon the host system. (See *Modular vs. Static SCSI Driver Discussion* on page 20-14.)

- ```

[M] Adaptec AIC7xxx support
 [*] Enable Tagged Command Queuing (TCQ) by default
 (33) Maximum number of TCQ commands per device
 [*] Collect statistics to report in /proc
 (5) Delay in seconds after SCSI bus reset

```

Select Exit to return to the Main Menu.

In the Main Menu, select Exit and select Yes to save your configuration in response to the window Do you wish to save your new kernel configuration?

By selecting Yes in response to this query, the configuration will be saved to the file /usr/src/linux/.config.

## 6. Compile the kernel:

```
make dep
make clean
make install
make modules
make modules_install
cd /boot
```

## 7. Build a ramdisk image:

```
mkinitrd /boot/initrd-$1.img $1
```

where \$1 = kernel version.

For example,

```
mkinitrd /boot/initrd-2.2.16.img 2.2.16
```

---

The `mkinitrd` command is used to create a new ramdisk image and is dependent upon the loopback device being available. If the loopback driver is selected in the kernel as a module, ensure that it is loaded with `#lsmod`. If the loopback driver was not loaded, insert it with `#insmod loop`.

---

## 8. Run `lilo` to update the loading map and reboot:

```
lilo -v
reboot
```

---

## Kernel Upgrade

Install Linux from the RedHat v6.x or v7.0 CD following the procedure in the installation guide. Partition the boot drive and select the packages and services necessary for the host. The host can be upgraded from RedHat v6.x to up to v2.2.16 of the Linux kernel. To upgrade the host to a newer version of the kernel, the kernel must be downloaded. The kernel can be downloaded from the following url: <http://www.kernel.org/mirrors/>.

### 1. Prepare the host for the new version of the kernel:

```
cd /usr/src
rm -f linux
```

2. Untar the kernel and create a directory for it:

```
tar zxvf linux-$1.tar.gz
mv linux linux-$1
ln -s linux-$1 linux
cd linux
```

where \$1 = kernel version.

3. Delete stale .o files and dependencies:

```
make mrproper
```

4. Configure the kernel:

```
make config/menuconfig/xconfig
```

5. From the Main Menu, select the SCSI Support menu. The menu shown in Figure 20-4 and in Figure 20-2 should appear.




Figure 20-4 SCSI Support Menu

In the SCSI support menu shown in Figure 20-4, make the following selections:

Select SCSI support as static (\* or y).

Under

— SCSI support type (disk, tape, CD-ROM)  
select the options SCSI disk support and SCSI generic support as static (\* or y).

Under

— Some SCSI devices (e.g. CD jukebox) support multiple LUNs  
select Probe all LUNs on each SCSI device as static (\* or y).

6. Scroll down the SCSI support menu to the SCSI low-level drivers submenu, as shown in Figure 20-5.




Figure 20-5 SCSI Low-Level Drivers Submenu

As shown in Figure 20-5, under

SCSI low-level drivers —>

select Adaptec AIC7xxx support as modular or static depending upon the host system. (See *Modular vs. Static SCSI Driver Discussion* on page 20-14.)

[M] Adaptec AIC7xxx support

- [\*] Enable Tagged Command Queuing (TCQ) by default
- (34) Maximum number of TCQ commands per device
- [\*] Collect statistics to report in /proc
- (5) Delay in seconds after SCSI bus reset

Select `Exit` to return to the Main Menu.

In the Main Menu, select `Exit`. Then select `Yes` to save your configuration in response to the window, `Do you wish to save your new kernel configuration?`

By selecting `Yes` in response to this query, the configuration will be saved to the file `/usr/src/linux/.config`.

## 7. Compile the kernel:

```
make dep
make clean
make install
make modules
make modules_install
cd /boot
```

## 8. Add the new kernel version to `lilo.conf` so that it will be bootable:

```
vi /etc/lilo.conf
```

For example:

```
image=/boot/vmlinuz-2.2.16
label=linux-2.2.16
root=/dev/sda1
initrd=/boot/initrd-2.2.16.img
read-only
```

## 9. Build a ramdisk image:

```
mkinitrd /boot/initrd-$1.img $1
```

where `$1` = kernel version.

For example,

```
mkinitrd /boot/initrd-2.2.16.img 2.2.16
```

The `mkinitrd` command is used to create a new ramdisk image and is dependent upon the loopback device being available. If the loopback driver is selected in the kernel as a module, ensure that it is loaded with `#lsmod`. If the loopback driver was not loaded, insert it with `#insmod loop`.

10. Run `lilo` to update the loading map and reboot:

```
lilo -v
reboot
```

### **Incorporating Symmetrix SCSI**

Once the Symmetrix has devices assigned with device addresses and the directors have been switched on line, the host can see all the target Symmetrix devices assigned to that host interface.

### **Partitioning and Labeling Devices**

Refer to *Partitioning and Labeling a New Device* on page 20-24.

### **Creating and Mounting a File System**

Refer to *Creating and Mounting a File System* on page 20-29.

---

## Partitioning and Labeling a New Device

Block device filenames take the following form:

`/dev/sdln`

where `l` is the letter denoting the physical drive and `n` is the number denoting the partition on that physical drive. Usually, the partition number is not included when the device corresponds to the entire drive.

Following this format, the filenames would appear as follows:

`/dev/sd[a-d][a-z][1-15]`

Linux supports a total of 128 SCSI devices per system. As a result, the host could support disk devices from `/dev/sda` through `/dev/sddx`.

Linux supports from 1-15 partitions per disk device. Partitions 1 through 4 are the primary partitions, while partitions 5 and greater are the logical or extended partitions. (These limitations are specific to the Intel platform.) By default, slices are not used in Linux.

---

At this time, true raw devices have not yet been incorporated into the kernel.

---

The corresponding character device filenames take the following form:

`/dev/sg[a-d][a-z]`

`sg` is one of four high-level device drivers in the SCSI subsystem.

The four high-level device drivers are:

|                 |                           |
|-----------------|---------------------------|
| <code>sg</code> | SCSI generic interface    |
| <code>sd</code> | direct access, i.e. disks |
| <code>st</code> | tapes                     |
| <code>sr</code> | data CD-ROMs              |

The `sg` driver is a character-based device while the other three drivers are block-based devices. `sg` is used primarily for scanners, CD writers, printers, and for raw access, although `sg` is not a true raw device. The `sg` device files are dynamically mapped to SCSI IDs/LUNs on the SCSI bus starting with the first SCSI controller.



---

The Linux kernel assigns minor SCSI device numbers dynamically by assigning them only to devices that are actually connected to the host in the order of their SCSI IDs. This means that connecting an external SCSI device can change the minor numbers of all the internal SCSI devices with a higher SCSI ID.

---



---

## Partitioning Symmetrix Devices

Once the Adaptec driver is loaded, the Symmetrix devices will be represented in `/proc/scsi/scsi`. They can be viewed by using `more/proc/scsi/scsi`. Only `/dev/sda` through `/dev/sdp` and `/dev/sga` through **`/dev/sgh`** will be created by default upon kernel compilation. In order to create subsequent device files for the Symmetrix devices attached to the host, the root user will need to create them via the `mknod` command.

Using the `fdisk` command allows multiple partitions to reside on a single disk device.

To partition a Symmetrix device for the use of Linux, log in as root and follow the procedure below for each new device.

To partition and label new devices:

1. `fdisk /dev/sdb`
2. Type `m` (for help) to display the following menu of options:

```

Command action
a toggle a bootable flag
b edit bsd disklabel
c toggle the dos compatibility flag
d delete a partition
l list known partition types
m print this menu
n add a new partition
o create a new empty DOS partition table
p print the partition table
q quit without saving changes
s create a new empty Sun disklabel
t change a partition's system id
u change display/entry units
v verify the partition table
w write table to disk and exit
x extra functionality (experts only)

```

Type `x` (for expert functionality), then `m` to see the following menu of options:

```

Command action
b move beginning of data in a partition
c change number of cylinders
d print the raw data in the partition table
e list extended partitions
g create an IRIX partition table
h change number of heads
m print this menu
p print the partition table
q quit without saving changes
r return to main menu
s change number of sectors
v verify the partition table
w write table to disk and exit

```

3. From the main `fdisk` menu, type `n` to add a new partition
4. Type `t` to change a partition's system ID
5. Type `L` for a list of Hex codes for system ID's. (Linux is type 83. Linux swap is type 82.)
6. Type `p` to print the partition table and to verify that the selections look correct.
7. Type `w` to write the table to disk and exit `fdisk`.
8. Repeat steps 3 through 7 for each new Symmetrix device.

You can use `cfdisk` instead of `fdisk`. (`cfdisk` is a curses-based disk partition table manipulator for Linux.) If used with the `-z` option, `cfdisk` allows the device's partition table to be zeroed out entirely. This is useful if you are converting a device previously used by another host.

Example:

```
cfdisk -z /dev/sdb
```



### CAUTION

**Caution: Changing the partition table results in loss of any previous data.**

1. To show the partition table, type `cfdisk -z /dev/sdb`. You should see a display like the one below:

```

cfdisk 0.8n

 Disk Drive: /dev/sdb
Heads: 255 Sectors per Track: 63 Cylinders: 550

Name Flags Part Type FS Type [Label] Size (MB)

 Pri/Log Free Space 4314.34

[Help] [New] [Print] [Quit] [Units]
[Write]

```

2. Select `[ New ]`. The following options then appear at the bottom of the window:

```
[Primary] [Logical] [Cancel]
```

3. Select `[Primary]` to create a primary Linux partition and select the size of the partition that is required.

In this case, to use the entire device, select:

```
Size (in MB): 4314.34
```

4. After Step 3 above is completed, the options listed at the bottom change as shown:

```

cfdisk 0.8n

 Disk Drive: /dev/sdb
Heads: 255 Sectors per Track: 63 Cylinders: 550

Name Flags Part Type FS Type [Label] Size (MB)

sdb1 Primary Linux 4314.34

[Bootable] [Delete] [Help] [Maximize] [Print]
[Quit] [Type] [Units] [Write]

```

`cfdisk` automatically selects Linux as the file system type. If a different type of file system or a swap partition is required, select `[Type]` and choose the necessary file system type.

5. Select `[Write]` to write the partition table to the disk. `cfdisk` prompts the user with the following warning:

```
Are you sure you want to write the partition table to disk (yes or no)?
Warning!! This may destroy data on your disk!
```

6. Type `yes` to write the partition table. Once the partition table has been written to the disk, select `[Quit]` to exit out of `cfdisk`.

---

## Creating and Mounting a File System

To create standard UNIX file systems under Linux, log in as root and proceed as follows for each previously labeled device.

### Create a New File System

Once each Symmetrix disk device has been formatted, partitioned, and labeled, a file system needs to be created in order to make the partition usable. Use the `mkfs` command to create a file system.

For a single partition, use

```
mkfs /dev/sdb1
```

You will be prompted to confirm the construction of a new file system. Type `Y` and press `ENTER`.

### Create a Mount Directory

Once the file system for each Symmetrix device has been created, a mount directory for each device must be created.

To create a mount directory, use

```
mkdir /mnt/sdb
```

where `/mnt/sdb` is the complete path for the new file system.

---

By default, `/mnt` is created upon installation, but a separate directory can be created in which to mount the devices' file systems.

---

### Mount the File System

To mount each file system, enter a statement similar to the following:

```
mount -t ext2 /dev/sdb1 /mnt/sdb
```

This mounts the raw device `/dev/sdb1` onto the mount point `/mnt/sdb`.

---

## System and Error Messages

Linux logs system and error messages to a file called `/var/log/messages`. This file includes SCSI error messages. Linux also logs boot messages in a file called `/var/log/boot.log`. The host will also display these messages at the system console.

## Symmetrix Device Definition Files

This appendix describes how to obtain special drive definition files needed for use with open systems host.

- ◆ Overview .....A-2
- ◆ Using the World Wide Web .....A-3
- ◆ Connecting to EMC's FTP Server.....A-4
- ◆ Transferring the Device Definition Files .....A-5
- ◆ The inq Command .....A-6

---

## Overview

The chapters in this manual describe how to configure Symmetrix to host environments it currently supports. At the beginning of each chapter is a table that lists key operating system features when using Symmetrix with that host. The last item in this table specifies if special drive definition files are needed for use with that host. If a file name or “yes” is listed, you must obtain these files from EMC and copy them to your host before configuring Symmetrix in your host environment.

You or your system administrator can FTP these files from EMC’s anonymous ftp server, **ftp.emc.com**, to your host. Depending on your particular host, you may need to copy one or more files. Please refer to your host section in this chapter for the specific files needed. If you cannot FTP these files, each host section provides a “manual” device definition alternative.

You may access EMC’s ftp server through EMC’s home page on the world wide web or through an FTP software package. This section describes both access methods.



---

## Using the World Wide Web

You can connect to EMC's home page through any web browser using the address **www.emc.com**. You can also go directly to EMC's anonymous ftp server by entering the URL **ftp://ftp.emc.com**.

Perform the following steps to access EMC's ftp server through the EMC home page:

1. Launch your web browser and type `ftp://ftp.emc.com` at the prompt.
2. Choose the **pub/** directory.
3. Choose **symm3000/** from the listing.
4. Choose the appropriate directory for your environment and ftp the device definition files to your host.

## Connecting to EMC's FTP Server

Perform the following steps to connect to EMC's anonymous FTP server:

1. At the host, log in as **root** and create the directory **/usr/ftp\_emc** by typing

```
mkdir /usr/ftp_emc
```

2. Go to the directory **/usr/ftp\_emc** by typing:

```
cd /usr/ftp_emc
```

3. Connect to EMC's FTP server by typing

```
ftp ftp.emc.com
```

4. At the FTP server login prompt, log in as **anonymous**

5. At the password prompt, type your e-mail address.

You are now connected to EMC's FTP server. To display a listing of FTP commands available to you, type **help** at the prompt.

---

## Transferring the Device Definition Files

The device definition files for the various hosts are contained in directories within **/pub/symm3000** on the server.

---

EMC recommends that you copy these files to a **/usr/ftp\_emc** directory.

---

1. Go to the directory containing the device definition files for your host. If you are unsure of the directory name, please refer to the appropriate host section in this chapter.

For example, if you are using a Pyramid host, the device definition files are contained in the **pyramid** directory. Go to this directory by typing `cd /pub/symm3000/pyramid`

2. Confirm you are in the correct directory and note the names and number of files present by typing:

```
pwd
ls
```

3. Confirm that your host's current directory is **ftp\_emc** by typing  
`lcd /usr/ftp_emc`
4. Disable the interactive mode by typing `prompt`. This allows you to copy several files without intervention.
5. Copy all files in the directory to your host by typing `mget *`
6. Confirm that all files copied to the directory on your host by typing `!ls` at the host prompt.
7. Exit the FTP session by typing `quit`

## The inq Command

The **inq** command displays several fields that you may use to determine which Symmetrix volume is associated with a particular device as seen by the host. The figure below illustrates a sample output of the **inq** command when run from the host console.

```

DEVICE:VEND:PROD:REV:SER NUM:CAP (kb)

/dev/hdisk0:IBMRISC:DFHSS2F:96C3:54843215:2202244
/dev/hdisk1:EMC:SYMMETRIX:5062:38000040:2080800
/dev/hdisk2:EMC:SYMMETRIX:5062:38001040:2080800
/dev/hdisk3:EMC:SYMMETRIX:5062:38002040:2080800
/dev/hdisk4:EMC:SYMMETIRX:5062:38003040:2080800
```

**Figure A-1 Sample Output of inq Command**

The output fields are as follows:

|          |                                                                                                                                                                                                                |
|----------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| DEVICE   | UNIX device name (full pathname) for the SCSI device.                                                                                                                                                          |
| VEND     | Vendor Information. This field will be "EMC" for a Symmetrix device.                                                                                                                                           |
| PROD     | Product Name. This field will be "SYMMETRIX" for a Symmetrix device.                                                                                                                                           |
| REV      | Revision number. For a Symmetrix device, this will be the microcode version.                                                                                                                                   |
| SER NUM  | Serial Number. The third, fourth, and fifth digits from the left in the SER NUM field are the Symmetrix volume number. For example, in the bottom row in the figure above, 003 is the Symmetrix volume number. |
| CAP (kb) | Size of the device in kilobytes.                                                                                                                                                                               |

## Supported Hosts

The **inq** command is supported by several host systems compatible with Symmetrix. These host systems include:

- ◆ IBM RS/6000 (AIX)
- ◆ Compaq/DEC (Tru64 UNIX)
- ◆ Data General (DG/UX)
- ◆ Fujitsu Siemens (Solaris)
- ◆ NCR (SVR4)
- ◆ Pyramid
- ◆ SGI (IRIX)

- ◆ Sun (Solaris)
- ◆ Siemens (Reliant Unix)
- ◆ Windows NT
- ◆ Sequent (DYNIX)
- ◆ HP (HP-UX)

EMC has placed an executable copy of the `inq` command for each of these host systems on its anonymous FTP server, **ftp.emc.com**.

These commands are located in the **/pub/sym3000/inquiry** directory on this server and will have the suffix of its respective host operating system. For example:

```
inq.AIX41
inq.AIX43
inq.bcslinux
inq.bcsuxw
inq.celerra
inq.celerra1
inq.dynix42
inq.dynix44
inq.dynix45
inq.hpux
inq.irix
inq.linux
inq.mpras
inq.osf1
inq.sinix
inq.solaris
inq.tru64
inq.unixsv
inq.unixware
inq.wnt
inq.xsolaris
```

Refer to the sections at the beginning of this appendix for instructions on how to access and transfer files from the anonymous FTP server.



# SCSI Directors and Cabling

---

---

|                                                                                                                                              |     |
|----------------------------------------------------------------------------------------------------------------------------------------------|-----|
| This appendix lists SCSI directors for models of Symmetrix and describes several SCSI cabling configurations in the open system environment. |     |
| ◆ SCSI Directors .....                                                                                                                       | B-2 |
| ◆ SCSI Cabling .....                                                                                                                         | B-3 |

# SCSI Directors

The following table lists SCSI directors (also referred to as SCSI adapters) for 3xxx, 5xxx and 8xxx models of the Symmetrix.

**Table B-1      Symmetrix SCSI Directors**

| Symmetrix Models | Director                                        | Low/High Voltage Director | EMC Part Number |
|------------------|-------------------------------------------------|---------------------------|-----------------|
| 3xxx<br>5xxx     | Fast-Wide Differential SCSI                     | High                      | 200-881-903     |
| 3xxx<br>5xxx     | Fast-Wide Differential SCSI (60 MHz processor)  | High                      | 201-207-917     |
|                  | Fast-Wide Differential SCSI (66 MHz processor)  | High                      | 201-207-927     |
|                  | Ultra-Wide Differential SCSI (60 MHz processor) | High                      | 201-277-917     |
|                  | Ultra-Wide Differential SCSI (66 MHz processor) | High                      | 201-277-927     |
| 8xxx             | Ultra-Wide Differential SCSI (40 MB/s)          | High                      | 201-406-907     |
|                  | Ultra2-Wide Differential SCSI (80 MB/s)         | Low                       | 201-340-907     |



## SCSI Cabling

### Multiple SCSI Hosts to Symmetrix

The following cabling configurations illustrate how to connect multiple SCSI hosts to a Symmetrix unit. Both hosts are configured for the same Symmetrix volumes. One configuration shows a split cable configuration and the other configuration is a “Y-cable” configuration. For more information on connecting multiple hosts to Symmetrix, please refer to the *Symmetrix High Availability Product Guide*, EMC P/N 200-999-566.

#### Split Cable Configuration

The cabling configuration in Figure B-1 uses two SA ports on Symmetrix. The SCSI bus is terminated at the SA ports on Symmetrix and at the Fast-Wide SCSI adapter in each host. Symmetrix is configured such that both hosts address the same volumes for redundancy even though the hosts connect to different ports. This is a split SCSI bus implementation.



#### CAUTION

**The SA port and SCSI controller must be terminated in this configuration.**




Figure B-1 Split Cable Configuration (2 SA Ports)

## Y-cable Configuration

The cabling configuration in Figure B-2 uses one SA port on Symmetrix. This is a shared SCSI bus implementation.

The SCSI bus is terminated at the end of the Y-cables that connect to the hosts. This allows you to service the SCSI controllers without affecting system availability.

The SA port and SCSI controller must be non-terminated in this configuration. Use SCSI Adapter P/N 201-215-900, and set port termination switch to OFF (LED out).




Figure B-2 “Y” Cable Configuration (1 SA Port)

## SCSI Interface Tests

These tests are run while Symmetrix is offline from all hosts to which it is connected. Connect each end of the Fast-Wide or Ultra-Wide SCSI cable(s) used to attach Symmetrix to the host to the specified pair of ports.

This section tests one pair of SCSI ports at a time. The SCSI ports are referred to as port A, B, C, and D.

1. Attach each end of a Fast-Wide SCSI cable to SCSI ports A and B on each SCSI adapter.
2. Choose **INLINES** from the Symmetrix Main Menu and type **N** <CR> at the “LOG SESSION IN A FILE?” prompt.

3. Type `$RS <CR>` to reset communication.
4. Type `$Rx` (where “x” = director number of the SCSI channel director being tested) `<CR>` to establish communication with that channel director.

---

The “\$Rx” command allows you to communicate with each director. It is important to reset communications (via \$RS) before selecting the next director.

---

5. Type `F4, 'OTST', 'DD'` to switch to the diagnostic mode to run the SCSI interface tests. Enter `Y` at the prompt, *Are you sure you want to proceed?*
6. Type `BD <CR>` to run the tests on all directors simultaneously.
7. Type `41 <CR>` to execute the SCSI loopback tests for ports A and B. Enter `Y` at the prompt, *Are you sure you want to proceed?*
8. Move the Fast-Wide SCSI cables to SCSI ports C and D on each SCSI adapter.
9. Type `BD <CR>` to run the tests on all directors simultaneously.
10. Type `42 <CR>` to execute the SCSI loopback tests for ports C and D.
11. Symmetrix reports “Passed” or “Failed” for the tests.
12. Type `$RS <CR>` to reset communication.
13. Type `$FF <CR>` to exit the diagnostic mode.

---

## Cable Tests

To test the cables that are to be used to connect Symmetrix with the host, use the cables as the loopbacks for the SCSI interface tests in the previous section.



## Host Bus Adapter Configuration

This appendix lists the host bus adapters (also called controller cards) used with the host systems in Volume I.

For other host bus adapters used with MicroChannel, EISA, and PCI buses, see Appendix C in Volume II.

- ◆ Host Bus Adapters Tables..... C-2

## Host Bus Adapters Tables

The following tables list the host bus adapters used on the hosts in Volume I. The host systems are listed in the same order as the chapters in this Volume. For host bus adapters used on other systems, see Appendix C in Volume II.

For an updated listing in the *EMC Symmetrix Open Systems Support Matrix* of all qualified hosts and for more detailed information including cabling and Fibre Channel hub and switch support guidelines, you must consult your EMC sales representative.

**Table C-1 NCR**

| Host Models or Series          | Earliest OS Qualified | Latest OS Qualified  | Bus | Adapter Type | Host Bus Adapter                                         | Comments                                                                              |
|--------------------------------|-----------------------|----------------------|-----|--------------|----------------------------------------------------------|---------------------------------------------------------------------------------------|
| 34xx, 35xx, 41xx, 45xx, 3600AP | UNIX SVR4 2.03.01     | UNIX SVR4 3.02       | MCA | FND          | 53C700-Q720                                              | Limited support by NCR - requires RPQ <sup>a</sup>                                    |
| 34xx, 35xx, 3600               |                       |                      |     | FWD          | NCR 53C720-Q720                                          |                                                                                       |
| 4300, 4380                     | NT 4.0 SP3            | NT 4.0 SP5           | PCI | UWD          | PQS2.0, PQS2.1                                           |                                                                                       |
|                                | UNIX SVR4 MPRAS 3.01  | UNIX SVR4 MPRAS 3.02 |     |              | PQS2.1                                                   | Dual-port SCSI is not offered on the 4700                                             |
| 4300, 4400, S50                | NT 4.0 SP3            | NT 4.0 SP5           |     |              | 4400-F282<br>4400-F280                                   |                                                                                       |
|                                | UNIX SVR4 MPRAS 3.02  | UNIX SVR4 MPRAS 3.02 |     |              | Qualified by NCR, no RPQ necessary.                      |                                                                                       |
|                                | Solaris 7 x86         | Solaris 7 x86        |     | UWD          | NCR pqs2.1 Quad                                          | Firmware BIOS 3.02                                                                    |
|                                |                       |                      |     | FC-AL        | QLogic QLA2100F-EMC <sup>b</sup>                         | Supports 128 LUNs per HBA; requires driver version 2.20 and HBA firmware version 1.37 |
|                                |                       | QLogic QLA2200F-EMC  |     |              | BIOS requires driver version 1.08 and BIOS version 1.44. |                                                                                       |
| 4300, 4400                     | Solaris 7 x86         | Solaris 7 x86        |     | UWD          | NCR pqs2.1 Quad                                          | Firmware BIOS 2.13FPGA                                                                |

a. NCR UNIX SVR4 Release 2.xx requires special installation procedures and support for Symmetrix.

b. In NCR, with the current SVR4 MPRAS 3.02 driver, only 32 LUNs are supported. Requires maintenance version 03.02.00.39 or higher.

**Table C-2 IBM NUMA-Q (Sequent)**

| Host Models or Series          | Earliest OS Qualified | Latest OS Qualified                               | Bus | Adapter Type   | Host Bus Adapter                                      | Comments                                                                                                                                                                           |
|--------------------------------|-----------------------|---------------------------------------------------|-----|----------------|-------------------------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Symmetry 2000<br>Symmetry 5000 | DYNIX/ptx 2.1.5       | DYNIX/ptx 4.2.4,<br>4.5.1 <sup>a</sup> , 4.4.8    |     | FWD            | QCIC-W-CTLR-01<br>QCIC-E                              |                                                                                                                                                                                    |
| NUMA-Q                         | DYNIX/ptx 4.4.2       | DYNIX/ptx<br>4.5.1 <sup>a</sup> , 4.4.8,<br>4.6.0 |     | FC-BR          | FC to SCSI Bridge                                     | Attachment through an integrated Fibre Channel to SCSI Bridge<br>C10V cable requires switched termination.<br>ptx 4.4.2 requires a minimum Symmetrix microcode level of 5264.25.23 |
| NUMA-Q 2000                    | DYNIX/ptx 4.4.4       |                                                   |     |                |                                                       |                                                                                                                                                                                    |
| NUMA-Q,<br>NUMA-Q2000          | DYNIX/ptx 4.4.4       | DYNIX/ptx<br>4.5.1 <sup>a</sup> , 4.4.8,<br>4.6.0 | PCI | FC-SW<br>FC-AL | IOC-210-52<br>(LP6500),<br>IOC-210-54<br>(LP7000E-N1) | FC-AL support only available in DYNIX/ptx 4.5.0 and above.                                                                                                                         |

a. Requires a minimum Symmetrix microcode level of 5265.47.29.

**Table C-3 HP 9000**

| Host Models or Series                                                    | Earliest OS Qualified | Latest OS Qualified <sup>a</sup> | Bus          | Adapter Type | Host Bus Adapter       | Comments                                                                                                                              |
|--------------------------------------------------------------------------|-----------------------|----------------------------------|--------------|--------------|------------------------|---------------------------------------------------------------------------------------------------------------------------------------|
| 700 Series:<br>725/100                                                   | HP-UX 10.01           | HP-UX 10.20                      | HSC          | FWD          | HP A4107A              | External EISA bus not supported on 700 Series; 714, 725, 735, 755, etc. all use Single-ended Differential interfaces on the EISA bus. |
| 735/99, 735/125<br>755/99, 755/125                                       |                       |                                  | Internal     |              | Internal Bus           |                                                                                                                                       |
| 715/100XC,<br>B-Series (all).<br>C-Series: J-200,<br>J-210, 210CX, J-250 |                       |                                  | HSC          |              | HP A4107A              |                                                                                                                                       |
| 800 Series:<br>E, F, G, H, I, T500                                       |                       |                                  | HP-PB        |              | HP 28696A              | Not qualified with Symmetrix 5 code 66                                                                                                |
| T520                                                                     |                       |                                  |              |              |                        |                                                                                                                                       |
| K-Class (SCSI)                                                           |                       | HP UX 11.ACE                     | HSC<br>HP-PB |              | HP A2969A<br>HP 28696A | PA8000 RISC architecture supported on K series models                                                                                 |
| T600 (SCSI)                                                              | HP-UX 10.20<br>TFC    |                                  | HSC<br>HP-PB |              | HP A3644A<br>HP 28696A |                                                                                                                                       |
| D-Class, R-Class                                                         | HP-UX 10.01           |                                  | HSC          |              | HP A4107A              |                                                                                                                                       |

Table C-3 HP 9000 (continued)

| Host Models or Series                                          | Earliest OS Qualified        | Latest OS Qualified <sup>a</sup> | Bus | Adapter Type           | Host Bus Adapter                    | Comments                                       |
|----------------------------------------------------------------|------------------------------|----------------------------------|-----|------------------------|-------------------------------------|------------------------------------------------|
| V2500, SCA, V2600 SCA                                          | HP-UX 11.10                  | HP-UX 11.10                      | PCI | UWD                    | HP A4800A                           | See <sup>b</sup>                               |
| V-Class, V2200, V2250, V2500 <sup>c</sup> , V2600 <sup>d</sup> | HP-UX 11.0                   | HP-UX 11.ACE                     | PCI | UWD                    | HP A4800A                           |                                                |
| N-Class, L-Class, L1000, L2000                                 | HP-UX 11.0 990P              | HP-UX 11.ACE                     | PCI | UWD                    | HP A4800A                           |                                                |
|                                                                |                              |                                  |     |                        | HP A5159A                           |                                                |
| A-Class (Staccato A180, A180C)                                 | HP-UX 10.20                  | HP-UX 11.ACE                     | HSC | FWD                    | HP A4107A                           | Requires minimum microcode release 5265.39.24. |
| A-Class A500/440MHz, A400/440MHz                               | HP-UX 11.ACE                 | HP-UX 11.ACE                     | PCI | UWD                    | HP A4800A                           | See <sup>b</sup>                               |
|                                                                |                              |                                  |     |                        | HP A5159A                           | See <sup>e</sup>                               |
| D-Class, R-Class                                               | HP-UX 10.20 TFC              | HP-UX 11.ACE                     | HSC | FC-AL                  | HP A3591A<br>HP A3591B <sup>f</sup> |                                                |
| K-Class (Fibre Channel)                                        | HP-UX 10.20 TFC              | HP-UX 11.ACE                     | HSC |                        | HP A3404A                           |                                                |
| T600 (Fibre Channel)                                           | HP-UX 10.20 TFC              | HP-UX 11.ACE                     | HSC |                        | HP A3636A                           |                                                |
| V-Class, V2200, V2250, V2500 <sup>c</sup> , V2600 <sup>d</sup> | HP-UX 11.0                   | HP-UX 11.ACE                     | PCI |                        | FC-AL, FC-SW                        |                                                |
|                                                                | HP-UX 11.0 990P              | HP-UX 11.ACE                     |     | HP A5158A <sup>g</sup> |                                     |                                                |
| N-Class, L-Class, L1000, L2000                                 | HP-UX 11.0 990P              | HP-UX 11.ACE                     | PCI | FC-AL                  | HP A3740A                           |                                                |
|                                                                |                              |                                  |     | FC-AL, FC-SW           | HP A5158A <sup>h, i</sup>           |                                                |
| A-Class (Staccato A180, A180C)                                 | HP-UX 11.0 9906              | HP-UX 11.ACE                     | PCI | FC-AL                  | HP A3740A                           | Requires minimum microcode release 5265.39.24. |
|                                                                | HP-UX 10.20 <sup>j</sup> TFC | HP-UX 11.ACE                     | HSC |                        | HP A3591B                           |                                                |
|                                                                |                              | HP-UX 10.20 TFC                  |     |                        |                                     |                                                |
| V2500 SCA, V2600 SCA                                           | HP-UX 11.10                  | HP-UX 11.10                      | PCI | FC-AL                  | HP A5158A <sup>h</sup>              | Non-HA only                                    |
| A-Class A500/440MHz, A400/440MHz                               | HP-UX 11.ACE                 | HP-UX 11.ACE                     | PCI | FC-AL, FC-SW           | HP A5158A <sup>h, i</sup>           |                                                |



- a. Install Bad Block Reallocation Patch Pair for all UP/UX 10.xx versions.
- b. The A4800A and A5159A HBAs are both capable of Ultra 1 speeds, but ship from HP with a default setting of FWD. Symmetrix is supported at Ultra 1 speeds with both these HBAs. Ultra 1 can be enabled in the V-Class, N-Class, and L-Class PDC firmware. V-Class supports the A4800A only.
- c. Minimum OS revision is HP-UX 11.0 Rev 9812.
- d. Minimum OS revision is HP-UX 11.0 990P.
- e. The cables are the same as the Cxx mini 68S ("VHDCI" 8mm to SCSI3). Cables require termination at host end for N-Class and L-Class. Can be terminated either on cable, or on HBA, but not both.
- f. Qualified and supported by HP.
- g. HP A5158A is enabled in the March 2000 HWCR bundle XSWHWCR1100.48. Additional patches may be required for support.
- h. HP A5158A is enabled in the March 2000 HWCR bundle XSWHWCR1100.48 and requires the FC-AL driver released in the March 2000 applications CD. Additional patches may be required for support.
- i. Support for fabric is enabled with: fabric device driver version B.11.00.03; minimum operating system level HP/UX 11.0 990P with March 2000 HW-CR bundle and dependency patch PHKL\_21381; Symmetrix microcode versions 5265.49.31, 5266.20, 5566.22.
- j. Symmetrix attached to FC-AL loops with servers running HP-UX 10.20 require minimum Symmetrix microcode rev 5265.44.xx.

Table C-4 NEC

| Host Models or Series                                                                                                                                                        | Earliest OS Qualified                             | Latest OS Qualified  | Bus | Adapter Type | Host Bus Adapter                    | Comments                                                                                                                                 |
|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|---------------------------------------------------|----------------------|-----|--------------|-------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------|
| HP9000 D-Class:<br>NX7000/22,<br>NX7000/23,<br>NX7000/27,<br>NX7000/D280,<br>NX7000/32,<br>NX7000/33,<br>NX7000/37,<br>NX7000/D380,<br>NX7000/D390,<br>TX7/D390,<br>TX7/D280 | HP UX 11.0                                        | HP UX 11.0           | HSC | FWD          | HP A4107A                           |                                                                                                                                          |
|                                                                                                                                                                              |                                                   |                      | HSC | FC-AL        | HP A3591A<br>HP A3591B <sup>a</sup> |                                                                                                                                          |
|                                                                                                                                                                              |                                                   |                      | HSC | FC-AL        | HP A3404A                           |                                                                                                                                          |
|                                                                                                                                                                              |                                                   |                      | HSC | FWD          | HP A3644A                           |                                                                                                                                          |
| HP9000 K-Class:<br>NX7000/220,<br>NX7000/260,<br>NX7000/K370,<br>NX7000/K380,<br>NX7000/420,<br>NX7000/460,<br>NX7000/K570,<br>NX7000/K580,<br>TX7/K370, TX7/K380            | HP UX 11.0                                        | HP UX 11.0           | HSC | FC-AL        | HP A3636A                           | Adapter is Ultra-capable, but disabled by HP; HP “Commercial Server” negotiates bus as FWD; HP “Technical Server” negotiates bus as UWD. |
| NX7000/T600                                                                                                                                                                  |                                                   |                      |     | FC-AL        | HP A3636A                           |                                                                                                                                          |
| NX7000/V2250,<br>TX7/V2250<br>NX7000/V2500 <sup>b</sup> ,<br>TX7/V2500                                                                                                       |                                                   |                      | PCI | FC-AL        | HP A3740A                           |                                                                                                                                          |
|                                                                                                                                                                              |                                                   |                      |     | UWD          | HP A4800A                           |                                                                                                                                          |
| UP4800/760EX, 860, 860AD, 760REX, 860R, 860RAD, 780, 790, 880, 890, 880AD, 890AD                                                                                             | UX4800 R13.5 <sup>c</sup> ,<br>R13.6 <sup>c</sup> | UX4800 R13.7 (Rev A) | PCI | UWD          | N4209-54                            |                                                                                                                                          |

a. Qualified and supported by HP.

b. Minimum OS revision is HP-UX 11.0 Rev 9812.

c. Limited support is available for R13.5 and R13.6; requires an RPQ and a patch from NEC.

Table C-5 IBM <sup>a</sup>

| Host Models or Series                                                                                                     | Earliest OS Qualified | Latest OS Qualified | Bus | Adapter Type | Host Bus Adapter                  | Comments                                                                                                                                                                        |
|---------------------------------------------------------------------------------------------------------------------------|-----------------------|---------------------|-----|--------------|-----------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| MCA-based RS/6000 (7015-R30, 7015-R40, 7015-R50) & SP2                                                                    | AIX 4.2.1             | AIX 4.3.2, 4.3.3    | MCA | FWD          | IBM 2416, 2412                    |                                                                                                                                                                                 |
| PCI-based RS/6000 (7025-F50, 7026-H50) & SP2                                                                              | AIX 4.2.1             |                     | PCI | FWD          | IBM 2409 <sup>b</sup><br>IBM 6209 |                                                                                                                                                                                 |
|                                                                                                                           |                       |                     |     | UWD          | IBM 6207                          |                                                                                                                                                                                 |
| PCI-based SP2:<br>9076-2xx, 9076-3xx,<br>9076-4xx, 9076-500,<br>9076-550, 9076-T70                                        | AIX 4.3.3             | AIX 4.3.3           |     | UWD          | IBM 6204                          | Feature codes supported:<br>2054: power3 high node<br>2055: SP expansion I/O<br>2056: 375MHz power3 thin node<br>2057: 375MHz power3 wide node<br>2058: 375MHz power3 high node |
| 7013-S70, 7015-S70,<br>7025-H70, 7026-H70,<br>7017-S70, 7017-S70 as SP2 node                                              | AIX 4.3.1             | AIX 4.3.2, 4.3.3    |     | FWD          | IBM 6209                          |                                                                                                                                                                                 |
|                                                                                                                           |                       | AIX 4.3.2, 4.3.3    |     | UWD          | IBM 6207                          |                                                                                                                                                                                 |
| 7013-S7A, 7015-S7A,<br>7017-S7A, 7017-S7A as SP2 node                                                                     | AIX 4.3.2             | AIX 4.3.2, 4.3.3    |     | FWD          | IBM 6209                          |                                                                                                                                                                                 |
|                                                                                                                           |                       | AIX 4.3.2, 4.3.3    |     | UWD          | IBM 6207                          |                                                                                                                                                                                 |
| 7044-170, 7044-270,<br>7025-F80, 7026-H80,<br>7026-M80, 7017-S80,<br>7017-S80 as SP2 node<br>pSeries 640 (B80), 680 (S85) | AIX 4.3.3             | AIX 4.3.3           |     | FWD          | IBM 6209                          |                                                                                                                                                                                 |
|                                                                                                                           |                       |                     |     | UWD          | IBM 6207, 6204                    |                                                                                                                                                                                 |

Table C-5 IBM <sup>a</sup> (continued)

| Host Models or Series                                                                                                                         | Earliest OS Qualified | Latest OS Qualified | Bus | Adapter Type   | Host Bus Adapter    | Comments                                                                                                                                                                  |
|-----------------------------------------------------------------------------------------------------------------------------------------------|-----------------------|---------------------|-----|----------------|---------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| PCI-based RS/6000:<br>7025-F50, 7026-H50<br>PCI-based SP2 nodes:<br>9076-2xx,<br>9076-3xx,<br>9076-4xx,<br>9076-500,<br>9076-550,<br>9076-T70 | AIX 4.3.0             | AIX 4.3.3           | PCI | FC-AL<br>FC-SW | EMC<br>CKIT-E70-AIX | HBA and driver available exclusively as EMC Fibre Channel Interface V2.0 for AIX Platforms.                                                                               |
| 7013-S70, 7015-S70,<br>7025-H70, 7026-H70,<br>7017-S70, 7017-S70 as<br>SP2 node                                                               | AIX 4.3.1             |                     |     |                |                     | Latest PTF Package can be downloaded from the EMC FTP server<br><a href="ftp://ftp.emc.com/pub/symm3000/aix/EMC-FC-Kit">ftp://ftp.emc.com/pub/symm3000/aix/EMC-FC-Kit</a> |
| 7013-S7A, 7015-S7A,<br>7017-S7A, 7017-S7A as<br>SP2 node                                                                                      | AIX 4.3.2             |                     |     |                |                     |                                                                                                                                                                           |
| 7044-170, 7044-270,<br>7025-F80, 7026-H80,<br>7026-M80, 7017-S80,<br>7017-S80 as SP2 node<br>pSeries 640 (B80), 680 (S85)                     | AIX 4.3.3             |                     |     |                |                     |                                                                                                                                                                           |

Table C-5 IBM <sup>a</sup> (continued)

| Host Models or Series                                                                                                                                                                                                                                                                                                                                                  | Earliest OS Qualified | Latest OS Qualified | Bus | Adapter Type | Host Bus Adapter | Comments                                                                                                                                                                                                                                                                            |
|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-----------------------|---------------------|-----|--------------|------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| PCI-based RS/6000:<br>7044-170, 7044-270,<br>7025-F50, 7025-F80,<br>7025-H70, 7026-H50,<br>7026-H70, 7026-H80,<br>7026-M80, 7017-S70,<br>7017-S70 as SP2 node,<br>7017-S7A, 7017-S7A as<br>SP2 node, 7017-S80,<br>7017-S80 as SP2 node,<br>pSeries 640 (B80), 680<br>(S85)<br>PCI-based SP2 nodes:<br>9076-2xx, 9076-3xx,<br>9076-4xx, 9076-500,<br>9076-550, 9076-T70 | AIX 4.3.3             | AIX 4.3.3           | PCI | FC-AL        | IBM 6227         | Requires minimum Symmetrix microcode level 5265.39.25.<br>Requires minimum AIX APAR IY05369. Supports adapter firmware revisions 2.22x1, 2.23, 3.20A5, 3.20A9.<br>Requires minimum EMC ODM fileset support (4.3.3.0 IBM Fibre Channel interface support for Symmetrix).             |
|                                                                                                                                                                                                                                                                                                                                                                        |                       |                     |     | FC-SW        |                  | Requires minimum Symmetrix microcode level 5265.48.30.<br>Requires minimum AIX APAR IY08960, IY03872, IY06844. Supports adapter firmware revision 2.23 and 3.20A5, 3.20A9.<br>Requires minimum EMC ODM fileset support (4.3.3.1 IBM Fibre Channel interface support for Symmetrix). |

a. For AS/400, see Volume II of this Product Guide.

b. Support is for legacy RS/6000 systems.

**Table C-6 Silicon Graphics**

| Host Models or Series                                     | Earliest OS Qualified   | Latest OS Qualified | Bus | Adapter Type   | Host Bus Adapter         | Comments                                                                     |
|-----------------------------------------------------------|-------------------------|---------------------|-----|----------------|--------------------------|------------------------------------------------------------------------------|
| Challenge S <sup>a b</sup> ,<br>L <sup>a b</sup> , DM, XL | IRIX 5.3                | IRIX 6.5.9          | HIO | FWD            | P-S-HIO SCSI             | Challenge not supported by SGI as part of SGI-ware program. See <sup>c</sup> |
| Origin 200 <sup>a b</sup>                                 | IRIX 6.4.1 <sup>d</sup> | IRIX 6.5.9          | PCI | FWD            | QLogic QLA1041B          |                                                                              |
| Origin 2000 <sup>a b</sup>                                |                         |                     | XIO | Ultra          | XT-SCSIB-4P              |                                                                              |
| Origin 2000 <sup>a</sup>                                  | IRIX 6.4.1 <sup>d</sup> | IRIX 6.5.9          | XIO | FC-AL          | XT-FC-2P <sup>e f</sup>  |                                                                              |
| Origin 200 <sup>a</sup>                                   | IRIX 6.4.1 <sup>d</sup> | IRIX 6.5.9          | PCI |                | PCI-FC-1P <sup>e f</sup> |                                                                              |
| Origin 200, 2000                                          | IRIX 6.5.9              | IRIX 6.5.9          | XIO | FC-AL<br>FC-SW | XT-FC-IP-OPT-A           |                                                                              |
|                                                           |                         |                     | PCI |                | PCI-FC-IP-OPT-A          |                                                                              |

- a. I/Os will immediately fail if a device becomes not ready. This is an SGI bug.
- b. I/Os will fail if a SCSI cable is temporarily disconnected.
- c. Platform and current OS (as well as previous, and anticipated for next) have the ability to add devices online, add a physical volume to a logical volume online, and extend a RAID or stripeset while online. Extending logical volumes requires a separate plexing license.
- d. Requires SGI patch set 10/1/98.
- e. Requires copper to optical MIA. Contact Methode Electronics at 800-323-6858 (P/N MDB-9-6-1) or Gadzoox Microsystems at 410-838-2108 (P/N MDM1063TC-G).
- f. SGI has replaced this adapter with the QLogic 2200 card. New installs should order part number PCI-1P-OPT-A for PCI, or XT-FC-1P-OPT-A for XIO from SGI. This adapter does not require a copper to optical MIA.

Table C-7 Sun

| Host Models or Series                                                                                                                                                                                                  | Earliest OS Qualified      | Latest OS Qualified                                                          | Bus  | Adapter Type   | Host Bus Adapter                                                         | Comments                                                                                                                              |
|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|----------------------------|------------------------------------------------------------------------------|------|----------------|--------------------------------------------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------|
| SPARCServer & SPARCcenter 10, 20, 1000, 2000                                                                                                                                                                           | Solaris 2.3                | Solaris 2.5.1 <sup>a</sup>                                                   | SBUS | FWD            | X1062A/A5                                                                | X1062A/A5 requires firmware version 1.24 or later.                                                                                    |
| SBUS-based Servers:<br>Ultra Workstation (Ultra 2)<br>Ultra Enterprise (Ultra 3000, 3500, 4000, 4500, 5000, 5500, 6000, 6500)                                                                                          | Solaris 2.5.1 <sup>a</sup> | Solaris 2.6 <sup>b</sup><br>Solaris 7 <sup>c</sup><br>Solaris 8 <sup>d</sup> | SBUS | FWD            | X1062A/A5                                                                | scsi_options=0x3F8                                                                                                                    |
|                                                                                                                                                                                                                        |                            |                                                                              |      | UWD            | X1065A                                                                   | This adapter is compatible with all FWD directors.<br>scsi_options=0x7F8                                                              |
|                                                                                                                                                                                                                        |                            | Solaris 2.6 <sup>b</sup><br>Solaris 7 <sup>c</sup><br>Solaris 8 <sup>d</sup> | SBUS | FC-AL<br>FC-SW | Jaycor<br>FC-1063-EMC <sup>e f</sup><br>FC64-1063-EMC <sup>g f</sup>     | Contact EMC for Fibre Channel hub & switch support guidelines.                                                                        |
| PCI-based Workstations & Servers:<br>Netra 1120, 1125, 1400, 1405, T1;<br>Ultra 5, 10, 30, 60, 80, 220(R), 250, 420(R), 450<br>Ultra Enterprise w/PCI IO Boards (Ultra 3000, 3500, 4000, 4500, 5000, 5500, 6000, 6500) | Solaris 2.5.1 <sup>h</sup> | Solaris 2.6 <sup>i</sup><br>Solaris 7 <sup>j</sup><br>Solaris 8 <sup>k</sup> | PCI  | UWD            | X6541A <sup>l m</sup>                                                    | For "SunFire" servers with PCI I/O board, only the left slot can be utilized due to physical constraints of the Sun PCI/SCSI adapter. |
| Ultra Enterprise 10000                                                                                                                                                                                                 | Solaris 2.5.1 <sup>a</sup> | Solaris 2.6 <sup>b</sup><br>Solaris 7 <sup>c</sup><br>Solaris 8 <sup>d</sup> | SBUS | FWD            | DWIS (X1062A/A5)                                                         |                                                                                                                                       |
|                                                                                                                                                                                                                        |                            |                                                                              |      | UWD            | Ultra DWIS (X1065A)                                                      |                                                                                                                                       |
|                                                                                                                                                                                                                        |                            |                                                                              |      | FC-AL<br>FC-SW | Jaycor<br>FC-1063-EMC <sup>e f n</sup><br>FC64-1063-EMC <sup>f g n</sup> | Contact EMC for Fibre Channel hub & switch support guidelines.                                                                        |

Table C-7 Sun (continued)

| Host Models or Series                                                                                                                                                                                                  | Earliest OS Qualified    | Latest OS Qualified                              | Bus | Adapter Type   | Host Bus Adapter                                            | Comments                                                        |
|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|--------------------------|--------------------------------------------------|-----|----------------|-------------------------------------------------------------|-----------------------------------------------------------------|
| PCI-based Workstations & Servers:<br>Netra 1120, 1125, 1400, 1405, T1;<br>Ultra 5, 10, 30, 60, 80, 220(R), 250, 420(R), 450<br>Ultra Enterprise w/PCI IO Boards (Ultra 3000, 3500, 4000, 4500, 5000, 5500, 6000, 6500) | Solaris 2.6 <sup>i</sup> | Solaris 2.6 <sup>i</sup>                         | PCI | FC-AL          | Emulex LP6502E-EMC <sup>f o n</sup>                         | LP6502E-EMC supports 32 LUNs/HBA and FC-AL point-to-point only. |
|                                                                                                                                                                                                                        |                          |                                                  |     |                | QLogic QLA2100F-EMC <sup>f p</sup>                          |                                                                 |
|                                                                                                                                                                                                                        | Solaris 7, 8             | Solaris 7 <sup>i</sup><br>Solaris 8 <sup>k</sup> | PCI | FC-AL<br>FC-SW | LP7000E-EMC <sup>f n q</sup><br>LP8000-EMC <sup>f n q</sup> | Contact EMC for Fibre Channel hub & switch support guidelines.  |
|                                                                                                                                                                                                                        |                          |                                                  |     |                | QLA2200F-EMC <sup>r</sup>                                   |                                                                 |
|                                                                                                                                                                                                                        |                          |                                                  |     |                | Jaycor FCIN-1063-EMC <sup>f, s</sup>                        |                                                                 |

- a. EMC required Sun patches for SBUS at Solaris 2.5.1:
  - 103640-26 SunOS 5.5.1: kernel patch
  - 103622-14 SunOS 5.5.1: sd driver patch
  - 103934-08 SunOS 5.5.1: kernel/drv/isp patch
  - 104708-16 SunOS 5.5.1: ssd, pln, soc, ssaadm and ssa firmware patch
- b. EMC required Sun patches for SBUS at Solaris 2.6:
  - 105181-20 SunOS 5.6: kernel update patch
  - 105356-13 SunOS 5.6: /kernel/drv/ssd and /kernel/drv/sd patch
  - 105600-17 SunOS 5.6: /kernel/drv/isp patch
- c. EMC required Sun patches for SBUS at Solaris 7:
  - 106541-12 SunOS 5.7: kernel update patch
  - 107458-10 SunOS 5.7: dad, sd, ssd, uata drivers patch
  - 106924-06 SunOS 5.7: /kernel/drv/isp and /kernel/drv/sparcvg/isp patch
  - 107469-08 SunOS 5.7: sf and socal driver patch
- d. EMC required Sun patches for SBUS at Solaris 8:
  - 108528-03 SunOS 5.8: kernel update patch
  - 108974-06 SunOS 5.8: data, uata, dad, sd, and scsi patch
  - 109524-02 SunOS 5.8: ssd driver patch.
  - 109657-01 SunOS 5.8: isp driver patch
  - 109460-03 SunOS 5.8: sf and socal driver patch
- e. 32 bit HBA: requires Driver revision 2.5.9 or later, Fcode 13.3.7 and Rev H. (Driver revisions below 2.2.1.EMC will only support FC-AL, not switches.)
- f. Jaycor adapter must be purchased from Polaris Service, Inc. (877-574-7600) or McData Corporation (800-545-5773) in the US or from ACAL Electronics (44-1344-723237) in Europe. Polaris also carries QLogic and Emulex adapters.
- g. 64 bit HBA: requires Driver revision 2.4.1.EMC or later and Fcode 13.3.7 and Rev B, C, D, and E. (Driver revisions below 2.2.1.EMC will only support FC-AL, not switches.)



- h. EMC required Sun patches for PCI at Solaris 2.5.1:
  - 103640-33 SunOS 5.5.1: kernel patch
  - 105004-11 SunOS 5.5.1: pci\_pci, ebus, pci and rootnex driver patch
  - 103622-15 SunOS 5.5.1: /kernel/drv/sd driver patch
- i. EMC required Sun patches for PCI at Solaris 2.6:
  - 105181-23 SunOS 5.6: kernel update patch.
  - 105356-16 SunOS 5.6: /kernel/drv/ssd and /kernel/drv/sd patch
  - 105604-09 SunOS 5.6: ebus/pci/rootnex driver patch
  - 105580-15 SunOS 5.6: /kernel/drv/glm patch
- j. EMC required Sun patches for PCI at Solaris 7:
  - 106541-12 SunOS 5.7: kernel update patch
  - 107458-10 SunOS 5.7: dad, sd, ssd, uata drivers patch
  - 107147-08 SunOS 5.7: pci driver patch
  - 106925-04 SunOS 5.7: GLM driver patch
  - 107469-08 SunOS 5.7: sf and socal driver patch
- k. EMC required Sun patches for PCI at Solaris 8:
  - 108528-03 SunOS 5.8: kernel update patch
  - 108974-06 SunOS 5.8: data, uata, dad, sd, and scsi patch
  - 109524-02 SunOS 5.8: ssd driver patch
  - 109460-03 SunOS 5.8: sf and socal driver patch
- l. PowerPath support for this controller is only with Solaris 2.6.
- m. Disconnecting and reconnecting SCSI cables on an active system may induce the system to crash.
- n. Refer to EMC Fibre Channel Product Guide (<http://symmdoc.lss.emc.com:6336/docs/pdf/20999642.pdf>) for required /etc/system and other OS settings.
- o. Requires driver lpfc-SPARC-2.5.1 version 2.03 for Solaris 2.5.1 and lpfc-SPARC-2.6 version 2.03 for Solaris 2.6.
- p. Requires driver QLA2100F-EMC version 2.20 or later which can be downloaded from <http://www.qlc.com> under the "EMC approved drivers directory."
- q. Requires driver lpfc-SPARC-V4.10e for Solaris 2.6, Solaris 7, and Solaris 8 only. Requires firmware 3.20x3.
- r. Requires driver QLA2200 v1.10 or later which can be downloaded from <http://www.qlc.com> under the EMC approved drivers directory.
- s. Jaycor PCI adapter requires driver version 2.5.9 or later.

Table C-8 Sun Solaris x86

| Host Models or Series                                                                                                                                                                                                           | Earliest OS Qualified      | Latest OS Qualified | Bus | Adapter Type   | Host Bus Adapter                 | Comments                                                                                                                                                 |
|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|----------------------------|---------------------|-----|----------------|----------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------|
| Siemens Primergy 460, 670, 700, and 870 series <sup>a</sup> , H400, K400, N400, N800<br><br>NCR Worldmark 4300, 4380, 4400 <sup>a</sup> , 47XX, 48XX, 52XX<br><br>Dell Poweredge 2300, 2450, 4400, 6100, 6300, 6350, 6450, 8450 | Solaris 7 x86 <sup>b</sup> | Solaris 7 x86       | PCI | FC-AL          | QLogic QLA2100F-EMC <sup>c</sup> | QLogic supports 128 LUNs/HBA                                                                                                                             |
|                                                                                                                                                                                                                                 |                            |                     |     | FC-SW<br>FC-AL | QLogic QLA2200F-EMC <sup>d</sup> |                                                                                                                                                          |
|                                                                                                                                                                                                                                 |                            |                     |     | UWD            | Adaptec AHA-2944UW               | Driver released with 2.7 x86 firmware version 2.11; no LUN skipping and no LUN out-of-sequence                                                           |
|                                                                                                                                                                                                                                 |                            |                     |     |                | NCRpqs2.0 Quad                   | Driver version V1.1, Rev 1998.09.09.11.23; Firmware BIOS v2.13FPGA                                                                                       |
| NCR 4300 <sup>a</sup> , 47XX                                                                                                                                                                                                    |                            |                     |     |                |                                  |                                                                                                                                                          |
| NCR 4400 <sup>a</sup> , 48XX, 52XX, S50                                                                                                                                                                                         |                            |                     |     |                | NCRpqs2.1 Quad                   | Driver version V1.1, Rev 1998.09.09.11.23; Firmware BIOS V3.02. NCR requires using only bus 1, PCI slot 6-10 in NCR 4400 with NCR pqs2.1 with BIOS 3.02. |

a. Refer to "Host BIOS Chart," on page 15.

b. EMC has qualified the following patches to not interfere with Symmetrix operation: 107545-03, 109105-03, 106794-05, 106961-01, 107039-01, 108375-03, 107023-06, 107457-01, 107588-01, 107637-05, 107888-09, 108344-03, 107201-12, 106945-03, 106953-01, 106979-10, 107116-05, 107260-01, 107455-05, 107685-01, 107793-01, 107973-01, 108302-02, 106737-03, 107339-01, 107894-06, 106935-03, 107181-22, 108220-01, 108222-01, 107886-06, 108483-02, 108663-01, 106328-08, 108722-01, 108485-01, 108377-10, 108332-01, 108328-01, 106939-04, 107710-06, 107019-02, 109254-01, 109252-01, 109405-01, 108799-01, 109950-01.

c. Requires driver QLA2100F-EMC version 2.20 or later which can be downloaded from <http://www.qlc.com> under the "EMC approved drivers directory" and firmware version 1.37.

d. Requires driver QLA2200F-EMC version 3.01 or later which can be downloaded from <http://www.qlc.com> under the "EMC approved drivers directory" and firmware version 1.59.

**Table C-9 Host BIOS Chart**

|     | Host                 | BIOS                                                              |
|-----|----------------------|-------------------------------------------------------------------|
| (1) | HP LH (LH-PRO)       | 4.05.14 PF                                                        |
|     | HP LX (LX-PRO)       | 1.00.11 CDOC                                                      |
|     | HP LXR (LXR-PRO)     | 1.00.11 CDOC                                                      |
|     | HP LXR-PRO8          | 2.00 (PB 4.0, rel 6.20)                                           |
|     | HP LH-II             | 4.05.14 PF                                                        |
|     | HP LH-3              | 4.06.25 PL                                                        |
|     | HP LH-4              | 4.06.12 PS                                                        |
|     | HP LXR-8000          | 4.0 Release 6.0 Navigator L.18.02.L3.2.C                          |
|     | HP LXR-8500          | 4.0 Release 6.0 Navigator L.18.02.L3.2.C                          |
|     | HP LT6000r           | Phoenix BIOS 4.06.26 PW Navigator L.18.02.L3.2.C                  |
|     | HP LH6000            | Phoenix BIOS 4.06.21 PU Navigator L.18.02.L3.2.C                  |
|     | HP LH3000            | Phoenix BIOS 4.06.26 PT L.18.02.L3.2.C                            |
| (2) | Compaq 1600          | P08 (8/17/98)                                                     |
|     | Compaq 1850          | 4.11A                                                             |
|     | Compaq 2500          | 4.12A                                                             |
|     | Compaq 3000          | 4.07A                                                             |
|     | Compaq 5000          | E16 (1/26/98)                                                     |
|     | Compaq 5500          | P12 (9/18/98)                                                     |
|     | Compaq 6000          | P43 (12/07/99)                                                    |
|     | Compaq 6400R         | P11 (12/10/99)                                                    |
|     | Compaq 6500          | P11 (12/30/98)                                                    |
|     | Compaq 7000          | P43 (12/07/99)                                                    |
|     | Compaq 8000          | P41 (07/06/2000)                                                  |
|     | Compaq 850R          | P04 (9/25/97)                                                     |
|     | Compaq 8500          | P42 (6/16/2000)                                                   |
|     |                      |                                                                   |
|     |                      |                                                                   |
| (3) | SNI 460              | 4.05 Release 1.03.992                                             |
|     | SNI 760              | 1.08.887                                                          |
|     | SNI 870              | 4.0 Release 6.1.35                                                |
|     | FSC Primergy N800    | Phoenix BIOS 4.0 Rel. 60 SABR 1.86B.0009.P09.9909301531 Release 9 |
|     | FSC N400, K400, H400 | 4.06R1.07, CFW is 2.07.119                                        |

**Table C-9 Host BIOS Chart (continued)**

|      | Host                   | BIOS                                                               |
|------|------------------------|--------------------------------------------------------------------|
| (4)  | IBM 5500               | KIE165AUS                                                          |
|      | IBM 7000               | 1.00.14.CD0                                                        |
|      | IBM 5500-M10           | KIE166AUS                                                          |
|      | IBM 5500-M20           | KQE117AUS                                                          |
|      | IBM 5000               | MOKT20AUS                                                          |
|      | IBM 6000R              | ARE110AUS                                                          |
|      | IBM 7000-M10           | OSE126AUS                                                          |
|      | IBM 7600               | MWE119AUS                                                          |
|      | IBM 8500               | MMKT18AUS                                                          |
| (5)  | NCR 4300               | 1.00.14.CD0                                                        |
|      | NCR 4380               | 2.01.01.003                                                        |
|      | NCR 4400               | 4.0 rel 6.0                                                        |
| (6)  | Dell 2300              | 1.10A9                                                             |
|      | Dell 2450              | 1.10A3                                                             |
|      | Dell 4400              | 1.10A5                                                             |
|      | Dell 6100              | 1.00.10.CD0L                                                       |
|      | Dell 6300              | 1.10.A03                                                           |
|      | Dell 6350              | 1.10.A09                                                           |
|      | Dell 6450              | 1.10A5                                                             |
|      | Dell 8450              | 1.10A00                                                            |
| (7)  | DG AV1400              | Phoenix BIOS 4.0 Release 6.0 L440GX Production Release 11.1        |
|      | DG AV3704R             | Phoenix BIOS 4.0 Release 6.0 S450NX BIOS Release 11                |
|      | DG AV8600              | 1.04                                                               |
|      | DG AV8900/8950         | Phoenix BIOS 4.0 Release. 6.0 SABR1.86b.0027.A.0004141435 ALPHA 27 |
| (8)  | Unisys AQUANTA QS/V2   | Phoenix BIOS 4.0 Release 6.0 SNX450 BIOS Release 3.3 Build 31      |
| (9)  | NEC Express 5800/140Ma | Phoenix BIOS 4.0 Rel. 6.0 SNX450 BIOS Release 1 Build 17           |
|      | NEC Express 5800/140Ha | Phoenix BIOS 4.0 Rel. 6.0 0325                                     |
|      | NEC Express 5800/180Ha | Phoenix BIOS 4.0 Release 6.0 1339                                  |
| (10) | ICL Trimetra P2000     | Phoenix BIOS 4.0 Rel. 6.0                                          |
| (11) | Network Engines P7000  | BIOS V0.99 (Build 03)                                              |

**Table C-10 Compaq (DEC)**

| Host Models or Series                                                                                                                                                                             | Earliest OS Qualified <sup>a</sup> | Latest OS Qualified                                                             | Bus           | Adapter Type | Host Bus Adapter                               | Comments                                                                 |
|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|------------------------------------|---------------------------------------------------------------------------------|---------------|--------------|------------------------------------------------|--------------------------------------------------------------------------|
| Alpha Station and Alpha Server Series 3000 AXP Models 300-900, 1000, 2000, 2100, 4000/4100, 8200, 8400 DS10, DS20, ES40, GS60, GS80 <sup>b</sup> , GS140, GS160 <sup>b</sup> , GS320 <sup>b</sup> | Digital UNIX V3.2A (DEC OSF/1)     | Digital UNIX V4.0B, V4.0D<br>Tru64 UNIX V4.0E, V4.0F <sup>c</sup> , V5.0A, V5.1 | Turbo-Channel | FWD          | KZTSA-AA                                       |                                                                          |
|                                                                                                                                                                                                   |                                    |                                                                                 | PCI           | FWD          | KZPSA-BB <sup>d</sup>                          |                                                                          |
|                                                                                                                                                                                                   |                                    |                                                                                 |               | UWD          | KZPBA-CB <sup>e</sup>                          |                                                                          |
|                                                                                                                                                                                                   |                                    | Tru64 UNIX V4.0F <sup>f</sup> , V5.0A <sup>g</sup> , V5.1 <sup>g</sup>          |               | FC-SW        | KGPSA-BC <sup>h</sup><br>KGPSA-CA <sup>i</sup> | Requires Fibre Channel Connectrix switch or Compaq Fibre Channel switch. |

a. Compaq (Digital) & Tru64 UNIX only. For Open VMS, see Volume II of this Product Guide

b. Tru64 UNIX V5.1 only.

c. V4.0x BL13 Patch Kits may introduce problems with KZPBA adapters. Install the ISP.mod fix (kzpba\_v40\_bl13.tar) if running 4.0D Patch Kit 5, 4.0E Patch Kit 3, or 4.0F Patch Kit 2.

d. Minimum firmware revision is A12.

e. Minimum firmware revision is 5.57.

f. V4.0F: 8 LUNs/Symmetrix Fibre director port.

g. V5.0A, V5.1: 128 LUNs/Symmetrix fibre director port; requires a minimum Symmetrix microcode level of 5265.48.30 or 5566.26.19 and OVMS director flag.

h. Minimum firmware revision is 2.20x2.

i. Minimum firmware revision is 3.03x2 (interim firmware release, check for updates).

Table C-11 Fujitsu

| Host Models or Series                                            | Earliest OS Qualified            | Latest OS Qualified                                                | Bus | Adapter Type | Host Bus Adapter                              | Comments                                                                                                                                |
|------------------------------------------------------------------|----------------------------------|--------------------------------------------------------------------|-----|--------------|-----------------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------|
| DS90 (Japan only)                                                | SVR4 V4.20                       | SVR4 V4.20                                                         |     | FWD          | F7958HS1                                      | Requires Fujitsu patch V20L20                                                                                                           |
| GP7000D                                                          | UXP/DS V20L10                    | UXP/DS V20L40                                                      |     | FWD          | F7958HS1                                      | Requires Fujitsu patch V20L30 or higher                                                                                                 |
| PRIMEPOWER (GP7000F) models M200, M400, M600, M800, M1000, M2000 | Fujitsu Solaris 2.6              | Fujitsu Solaris 7 Aug 99 version & Nov 99 version<br><br>Solaris 8 | PCI | UWD          | Sun X6541A-X <sup>a b</sup><br><br>GP70F-CS02 | No PowerPath support on these SCSI controllers. Disconnecting and reconnecting SCSI cables on an active system may induce system crash. |
| PRIMEPOWER (GP7000F) model M200                                  | Fujitsu Solaris 7 Aug 99 version | Solaris 7 Aug 99 version                                           | PCI | FC-AL        | Fujitsu GP7B8FC1                              |                                                                                                                                         |

- a. EMC required Solaris patches for Fujitsu servers running Solaris 7 (must be obtained from Fujitsu):

106541-10 SunOS 5.7: kernel update patch

106925-04 SunOS 5.7: GLM driver patch

107147-05 SunOS 5.7: pci driver patch

107458-08 SunOS 5.7: dad, sd, ssd, uata drivers patch

- b. EMC required Solaris patches for Fujitsu servers running Solaris 2.6 (must be obtained from Fujitsu):

105181-19: SunOS 5.6: kernel update patch

105356-12: SunOS 5.6: /kernel/drv/ssd and /kernel/drv/sd patch (this patch corresponds to Sun Solaris patch 105356-13)

105580-14: SunOS 5.6: /kernel/drv/glm patch

105604-09 SunOS 5.6: ebus/pci/rootnex driver patch

**Table C-12 Fujitsu Siemens/FSC**

| Host Models or Series <sup>a</sup>                                  | Earliest OS Qualified    | Latest OS Qualified (preferred)          | Bus | Adapter Type   | Host Bus Adapter                                                        | Comments                                                                                                                                                                                                                                |
|---------------------------------------------------------------------|--------------------------|------------------------------------------|-----|----------------|-------------------------------------------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| PRIMEPOWER (GP7000F)<br>Models M200, M400, M600, M800, M1000, M2000 | Solaris 2.6 <sup>b</sup> | Solaris 7 Nov99 version<br><br>Solaris 8 | PCI | UWD            | Sun X6541A-X <sup>c</sup><br>GP70F-CS02                                 | No dynamic partition support.<br>Disconnecting and reconnecting SCSI cables on an active system may induce the system to crash.                                                                                                         |
| PRIMEPOWER (GP7000F)<br>Models M200, M400, M600, M800, M1000, M2000 | Solaris 7 Aug 99 Version | Solaris 7 Aug 99 Version                 | PCI | FC-AL<br>FC-SW | Emulex LP8000-EMC <sup>d e f g</sup><br>(GP70F-CF10) <sup>d e f g</sup> | Volume Logix V2.2 (Solaris) is supported on Fujitsu PRIMEPOWER (GP7000F) hosts. Fujitsu-Siemens Reliant UNIX hosts also support the use of VolumeLogix V2.2 through the use of Fujitsu PRIMEPOWER or NT hosts acting as a control host. |
|                                                                     |                          | Solaris 7 Nov 99 Version                 |     |                |                                                                         |                                                                                                                                                                                                                                         |
|                                                                     |                          | Solaris 8                                |     |                |                                                                         |                                                                                                                                                                                                                                         |

a. For RM series, see Volume II of this Product Guide.

b. EMC required Sun patches for PCI at Solaris 2.6:

105181-20 SunOS 5.6: kernel update patch.

105356-13 SunOS 5.6: /kernel/drv/ssd and /kernel/drv/sd patch

105604-09 SunOS 5.6: ebus/pci/rootnex driver patch

105580-14: SunOS 5.6: /kernel/drv/glm patch

c. PowerPath support for this controller is only with Solaris 2.6.

d. Requires Emulex driver 4.10e and firmware 3.20x3.

e. EMC required Solaris patches for Fujitsu servers running Solaris 7 (must be obtained from Fujitsu):

106541-10 SunOS 5.7: kernel update patch

107458-10 SunOS 5.7: dad, sd, ssd, uata drivers patch

107147-05 SunOS 5.7: pci driver patch

106925-04 SunOS 5.7: GLM driver patch

f. PowerPath support requires version 1.5.0.f1c2 with patch 1.5.0.3.Patch.B09 and patch 1.5.0.5.Patch.B13. Also supports PowerPath v2.0.

g. Additional patches required for PowerPath to run on a Fujitsu host (Solaris 7 Fujitsu August 99 version):

106300-06: shared library patch for 64-bit C++

106327-05: patch requires patches for 106300-06

107058-01: assembler patch

**Table C-13 Stratus**

| Host Models or Series                                             | Earliest OS Qualified | Latest OS Qualified            | Bus | Adapter Type | Host Bus Adapter                               | Comments                                                                                                                                                                                                                          |
|-------------------------------------------------------------------|-----------------------|--------------------------------|-----|--------------|------------------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Continuum 400 Series hosts (Models AC POLO 7100 and AC POLO 8000) | HP-UX B.10.20.1       | HP-UX 11.00.01 with PHKL_19773 | PCI | FWD          | STRATUS SCSI Controller AA-U50300 <sup>a</sup> | Targets supported: 0-D (E, F are reserved for the controller)<br>LUNs supported: 0-7 (no LUN skipping).<br>Maximum of 32 devices supported.<br>Symm flags: T, L, Y, W, B, C, D, and "Enable Propagate Reset?" must be set to "N". |

a. Requires Firmware Version S06 (u5030fw0S06raw).

**Table C-14 Lucent**

| Host Models or Series | Earliest OS Qualified | Latest OS Qualified            | Bus | Adapter Type | Host Bus Adapter                               | Comments                                                                                                                                                                                                                          |
|-----------------------|-----------------------|--------------------------------|-----|--------------|------------------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Lucent DNCP-400       | HP-UX B.10.20.1       | HP-UX 11.00.01 with PHKL_19773 | PCI | FWD          | STRATUS SCSI Controller AA-U50300 <sup>a</sup> | Targets supported: 0-D (E, F are reserved for the controller)<br>LUNs supported: 0-7 (no LUN skipping).<br>Maximum of 32 devices supported.<br>Symm flags: T, L, Y, W, B, C, D, and "Enable Propagate Reset?" must be set to "N". |

a. Requires Firmware Version S06 (u5030fw0S06raw).



Table C-15 SCO UnixWare 7

| Host Models or Series                                                                       | Earliest OS Qualified    | Latest OS Qualified      | Bus | Adapter Type | Host Adapter Card  | Comments                                                                                                                   |
|---------------------------------------------------------------------------------------------|--------------------------|--------------------------|-----|--------------|--------------------|----------------------------------------------------------------------------------------------------------------------------|
| FSC Primergy 460, 470, 670, 870 series, H400, K400, N400, N800                              | 7.0.1 <sup>a b c d</sup> | 7.0.1 <sup>a b c d</sup> | PCI | UWD          | Adaptec AHA-2944UW | Driver version 7.0.1a release supports LUNs 0-7 (no LUN skipping, no LUNs out of order); LUN series must start with LUN 0  |
| Dell Poweredge 2300, 2450, 4400, 6100, 6300, 6350, 6450, 8450                               |                          |                          |     |              | Adaptec AHA-2944UW |                                                                                                                            |
| HP Netserver LH (LH-PRO), LHII, LHIII, LH-4, LX (LX-PRO), LXR (LXR-PRO), LXR-PRO8, LXR-8000 |                          |                          |     |              | QLogic QLA1041     | Driver version 2.07: LUN series must start with LUN 0; always select auto configure option in HBA BIOS utility.            |
| NCR Worldmark Server 4300, 4380, 4400                                                       |                          |                          |     |              | QLogic QLA1240D    | Driver version 2.04: LUN series must start with LUN 0; always select auto configure option in HBA BIOS utility.            |
| ICL XtraSERVER P1000                                                                        | 7.1.0 <sup>e f</sup>     | 7.1.0 <sup>b e f</sup>   |     |              | QLogic QLA1041     | Driver version 2.07: LUN series must start with LUN 0; always select auto configure option in HBA BIOS utility.            |
|                                                                                             |                          |                          |     |              |                    | Driver version 7.1.0. LUN series must start with LUN 0. Always select auto configure option in a QLA1041 HBA BIOS utility. |

Table C-15 SCO UnixWare 7 (continued)

| Host Models or Series                                                                                                                                                                                                                                                   | Earliest OS Qualified | Latest OS Qualified    | Bus | Adapter Type   | Host Adapter Card                 | Comments                                                                                   |
|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-----------------------|------------------------|-----|----------------|-----------------------------------|--------------------------------------------------------------------------------------------|
| FSC Primergy 460, 470, 670, 870 series, H400, K400, N400, N800<br>Dell Poweredge 2300, 2450, 4400, 6100, 6300, 6350, 6450, 8450<br>HP Netserver LH (LH-PRO), LHII, LHIII, LH-4, LX (LX-PRO), LXR (LXR-PRO), LXR-PRO8, LXR-8000<br>NCR Worldmark Server 4300, 4380, 4400 | 7.1.0 <sup>e f</sup>  | 7.1.0 <sup>b e f</sup> | PCI | FC-AL          | QLogic QLA2200F-EMC, QLA2200A-EMC | Requires driver version 3.11, BIOS version 1.59. EMC supports FC-AL (direct-connect only). |
|                                                                                                                                                                                                                                                                         | 7.1.1 <sup>g h</sup>  | 7.1.1 <sup>g h</sup>   |     | FC-AL<br>FC-SW |                                   |                                                                                            |

- a. Required patches for UnixWare 7.0.1: ptf7034a, ptf7036e, ptf7045c, ptf7051d, ptf5052b, ptf7054b, ptf7068c, ptf7069e, ptf7072a, ptf7074d, ptf7079b, ptf7124a. Patches are available at <http://www.sco.com>.
- b. EMC's API requires for 7.0 and 7.1 a restamp driver 'rstmp'; the patch is available at <http://iqweb>.
- c. The following patches for UnixWare 7.0.1 have been qualified not to interfere with Symmetrix operation: ptf7053, ptf7058, ptf7059, ptf7061, ptf7064, ptf7070, ptf7073, ptf7075, ptf7076, ptf7080, ptf7082, ptf7090, ptf70104, ptf7407. Patches are available at <http://www.sco.com>.
- d. MPIO feature of SCO UnixWare 7 must be either disabled or not installed.
- e. SCO patch ptf7418 is not supported by EMC.
- f. The following are required patches for SCO UnixWare 7.1.0: ptf7045, ptf7402a, ptf7408e, ptf7410f, ptf7417a, ptf7418c, ptf7419d, ptf7421c, ptf7428d, ptf7430f, ptf7439e, ptf7446c, ptf7624b, ptf7635a, ptf7636a, ptf7637c, ptf7639a, ptf7640a.
- g. EMC requires the following patches for proper host behavior with the Symmetrix for SCO UnixWare 7.1.1: ptf7045f, ptf7602d, ptf7603l, ptf7616b, ptf7624b.
- h. EMC has qualified the following patches not to interfere with Symmetrix operation for SCO UnixWare 7.1.1: ptf7045f, ptf7058d, ptf7080d, ptf7083c, ptf7130a, ptf7139a, ptf7140a, ptf7141a, ptf7407c, ptf7410c, ptf7426b, ptf7438b, ptf7440a, ptf7441d, ptf7602d, ptf7603l, ptf7604c, ptf7607d, ptf7609c, ptf7610b, ptf7611b, ptf7612b, ptf7613b, ptf7614b, ptf7616b, ptf7617b, ptf7619b, ptf7620a, ptf7624b, ptf7625b, ptf7627b, ptf7631c, ptf7632c, ptf7633b, ptf7634a, ptf7635a, ptf7636a, ptf7639a, ptf7640a, ptf7641a, ptf7642a, ptf7643a, ptf7644a, ptf7646b, ptf7647a, ptf7650a, ptf7651a, ptf7654b, ptf7656a, ptf7659a, ptf7660a.

Table C-16 Linux

| Host Models or Series                                          | Earliest OS Qualified <sup>a</sup> | Latest OS Qualified <sup>a</sup> | Bus | Adapter Type | Host Bus Adapter                | Comments                                                                                                                                                                                   |  |
|----------------------------------------------------------------|------------------------------------|----------------------------------|-----|--------------|---------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|--|
| Compaq Intel-based platforms                                   | RedHat v2.2.9 <sup>b</sup>         | RedHat v2.2.16 <sup>c</sup>      | PCI | UWD          | Adaptec AHA-2944UW <sup>i</sup> | Requires microcode 5265.29.19 or later.<br>Targets 0-F supported with 7 reserved for the controller. LUNS 0-7 supported.<br>Maximum of 128 devices supported (per system not per adapter). |  |
| NEC ExpressServer 5800 120Ra2                                  | RedHat v2.2.10 <sup>d</sup>        | RedHaxt v2.2.16 <sup>c</sup>     |     |              |                                 |                                                                                                                                                                                            |  |
| HP Netserver LH Series                                         | RedHat v2.2.11 <sup>e f g</sup>    | RedHat v2.2.16 <sup>c</sup>      |     |              |                                 | LH-Pro and Siemens 460 require microcode 5265.38.24 or later.                                                                                                                              |  |
| FSC Primergy 460, 470, 670, 870 Series, H400, K400, N400, N800 |                                    |                                  |     |              |                                 |                                                                                                                                                                                            |  |
| Dell PowerEdge 2300                                            |                                    |                                  |     |              |                                 |                                                                                                                                                                                            |  |
| HP Netserver LX Series                                         | RedHat v2.2.15 <sup>h</sup>        | RedHat v2.2.16 <sup>c</sup>      |     |              | Adaptec AHA-2944UW <sup>i</sup> | Maximum of 128 devices supported.<br>Requires microcode 5265.50.32 or later or 5566.19.16s or later.<br>Usable card slots are P5 and P6.                                                   |  |
| VALinux FullOn 2200 Series, 3500                               | RedHat v6.2 <sup>j</sup>           | RedHat v2.2.16 <sup>c</sup>      |     |              | Adaptec AHA-2944UW <sup>i</sup> | Maximum of 128 devices supported.<br>Requires microcode 5265.50.32 or later or 5566.22.17s or later.                                                                                       |  |

Table C-16 Linux (continued)

| Host Models or Series                                                                                                                                                                                                                                                                                                       | Earliest OS Qualified <sup>a</sup> | Latest OS Qualified <sup>a</sup> | Bus | Adapter Type                | Host Bus Adapter                    | Comments                                                                                                                                          |
|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|------------------------------------|----------------------------------|-----|-----------------------------|-------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------|
| NEC<br>ExpressServer<br>5800 120Ra2<br>Compaq<br>AP200/Proliant<br>800, 850, 1600,<br>1850, 8500<br>Dell Poweredge<br>2300<br>FSC Primergy<br>460, 470, 670,<br>870 Series,<br>H400, K400,<br>N400, N800<br>Data General<br>AViiON 1400<br>HP Netserver<br>LH, LX Series<br>VA Linux FullOn<br>2200 Series<br>VA Linux 3500 | RedHat v6.2<br>(v2.2.14)           | RedHat v2.2.17 <sup>k</sup>      | PCI | FC-SW <sup>l</sup><br>FC-AL | QLogic<br>QLA2200F <sup>m n o</sup> | Maximum of 128<br>devices supported<br>per system, not per<br>adapter.<br>Requires microcode<br>5265.50.32 or later<br>or 5566.22.17 or<br>later. |

- a. Please note that this lists the specific kernel version, not just the RedHat distribution. If there is a customer requirement for a different distribution, please submit an RPQ.
- b. Requires Adaptec Patch v5.1.17 available from <http://people.redhat.com/dledford/aic7xxx.html>
- c. OS is RedHat V6.x upgraded to v2.2.16.
- d. OS is RedHat V6.0 upgraded to v2.2.10 (per request by NEC) plus aic7xxx patch v5.1.19 patch (from <http://people.redhat.com/dledford/aic7xxx.html>)
- e. OS is RedHat v6.0 upgraded to v2.2.11.
- f. When running Inquiry and I/O simultaneously on an SMP SCSI system, the host hangs and requires a power cycle. A new SCSI generic driver (v2.1.36) that fixes this problem via a locking mechanism is available at <http://www.torque.net/sg/> under subdirectory test with a file name of sg2136.tgz.
- g. RedHat v2.2.11 should be patched to fix a memory leak. The patch is available from <http://www.linux.org.uk/VERSION/relnotes.2211.html> and listed as 'TCP Memory Leak'. Patch name is tcp.diff.
- h. OS is RedHat v6.x upgraded to v2.2.15.
- i. Requires BIOS v2.20.0 and driver version 5.1.28.
- j. VAlinux release v2.2.14-5.0.14 (smp).
- k. OS is RedHat v6.x upgraded to v2.2.17.
- l. Single HBA zoning is required regardless of the switch being utilized.
- m. Requires QLogic driver version 2.19.15 and BIOS v1.61.
- n. QLogic driver v2.19.15 supports persistent binding and only supports Class 3.
- o. Driver can be obtained from [http://www.qlogic.com/bbs-html/csg\\_web/adapter\\_pages/driver\\_pages/22xx/22emc.html](http://www.qlogic.com/bbs-html/csg_web/adapter_pages/driver_pages/22xx/22emc.html).

## VI Commands

This appendix describes some of the most commonly used “vi” editor commands.

- ◆ Editing Tasks.....D-2

# Editing Tasks

The following tables list various UNIX tasks and the most commonly used vi editor commands to accomplish these tasks.

Table D-1 Starting vi

| Command        | Keyword  | Description                         |
|----------------|----------|-------------------------------------|
| vi <file_name> | start vi | Edit <file_name> starting at line 1 |

Table D-2 Moving the Cursor<sup>a</sup>

| Command | Keyword           | Description                                 |
|---------|-------------------|---------------------------------------------|
| h or →  | Move right        | Move one space to the right                 |
| j or ↑  | Move up           | Move up one line                            |
| k or ↓  | Move down         | Move down one line                          |
| l or ←  | Move left         | Move one space to the left                  |
| \$      | end of line       | End of the current line                     |
| 0       | beginning of line | Beginning of the current line               |
| nG      | to line n         | Move to line n (without n to the last line) |

a. You must be in the “command mode” to move the cursor.

Table D-3 Inserting Text<sup>a</sup>

| Command | Keyword | Description                                 |
|---------|---------|---------------------------------------------|
| i       | insert  | Insert text before cursor                   |
| a       | append  | Insert text after the cursor                |
| o       | open    | Insert text in open line below current line |

a. These commands leave “vi” in the input mode. Press <ESC> to return to the Command mode.

**Table D-4 Deleting Text**

| Command | Keyword | Description                                                             |
|---------|---------|-------------------------------------------------------------------------|
| r       | replace | Replace current character; no <ESC> needed                              |
| R       | overlay | Replace each character (start overlay mode, <ESC> to exit overlay mode) |
| ndd     | delete  | Delete the following n lines (if n not specified, delete 1 line)        |

**Table D-5 Searching**

| Command | Keyword | Description                |
|---------|---------|----------------------------|
| /<text> | search  | Search forward for <text>  |
| ?<text> | search  | Search backward for <text> |
| n       | repeat  | Repeat last search         |

**Table D-6 Saving and Exiting vi**

| Command | Keyword      | Description                                                  |
|---------|--------------|--------------------------------------------------------------|
| :q!     | quit vi      | Quit this vi session without saving changes                  |
| :w      | save file    | Write (save) changes to current file                         |
| :wq     | write & quit | Write (save) changes to current file and quit the vi session |

**Table D-7 Miscellaneous vi Commands**

| Command | Keyword          | Description                                                  |
|---------|------------------|--------------------------------------------------------------|
| nyy     | copy to buffer   | copy n lines into a buffer (if n not specified, copy 1 line) |
| p       | copy from buffer | put text in buffer after current line                        |
| .       | repeat           | Repeat last command                                          |

**Table D-7     Miscellaneous vi Commands**

| Command           | Keyword     | Description                                                                  |
|-------------------|-------------|------------------------------------------------------------------------------|
| u                 | undo        | Undo last change                                                             |
| mx                | mark        | Set marker x, where x is a later (a-z)                                       |
| 'x                | go to mark  | Move cursor to marker x, where x is a letter (a-z)                           |
| :r <file_name>    | insert file | Insert the file <file_name> after the cursor                                 |
| :w <file_name>    | save        | Save changes to file <file_name>                                             |
| vi +n <file_name> | start vi    | Edit <file_name> starting at line n (if n not specified, start at last line) |
| CTRL, z           | suspend     | UNIX command to suspend the current vi session                               |
| fg                | resume      | UNIX command to resume the suspended vi session                              |



---

This glossary contains terms related to the Symmetrix open systems and mainframe operation as well as disk storage subsystems. Many of these terms are used in this manual.

### A

|                                  |                                                                                                                                                                                                                                                                            |
|----------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Actuator</b>                  | A set of access arms and their attached read/write heads, that move as an independent component within a head and disk assembly (HDA).                                                                                                                                     |
| <b>Adapter</b>                   | A circuit board that provides the physical interface between the director and disk devices (Dual-initiator adapter), director and SCSI host channels (SCSI adapter), director and parallel channels (Bus & Tag adapter), or director and serial channels (Serial adapter). |
| <b>AdvFS</b>                     | Advanced File System. A journaled filesystem provided with Digital UNIX.                                                                                                                                                                                                   |
| <b>ANSI</b>                      | American National Standards Institute. A standards-setting, non-government organization which develops and publishes standards for voluntary use in the USA.                                                                                                               |
| <b>Arbitration</b>               | The process of selecting from multiple devices competing for use of the SCSI bus concurrently.                                                                                                                                                                             |
| <b>Asynchronous transmission</b> | A handshaking protocol that requires each byte be requested and acknowledged before the transmission of the next byte begins.                                                                                                                                              |

|                                       |                                                                                                                                                                                                                                                                                                                                     |
|---------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Auxiliary Storage Pool (ASP)</b>   | An ASP is a logical grouping of physical disks in the IBM AS/400 environment. An ASP is similar to a Volume Group in a UNIX environment. There can be up to 16 ASPs on any one AS/400 unit. ASP #1 is the System ASP.                                                                                                               |
| <b>Availability</b>                   | A measure of a system's ability to provide client access to applications and data. Expressed as a percentage of available time to total time. For example, an availability of 99.9 represents an unplanned application down time of 9 hours per year (available 8,751 hours or 8,760 hours per year).                               |
| <b>B</b>                              |                                                                                                                                                                                                                                                                                                                                     |
| <b>Backplane</b>                      | Card that accommodates the director, cache, and adapter cards.                                                                                                                                                                                                                                                                      |
| <b>Bus free phase</b>                 | A SCSI command phase that indicates that no SCSI device is actively using the SCSI bus.                                                                                                                                                                                                                                             |
| <b>Byte</b>                           | A group of 8 logical bits.                                                                                                                                                                                                                                                                                                          |
| <b>C</b>                              |                                                                                                                                                                                                                                                                                                                                     |
| <b>Cache</b>                          | Random access electronic storage used to retain frequently used data for faster access by the channel.                                                                                                                                                                                                                              |
| <b>Cache Slot</b>                     | Unit of cache equivalent to one disk track.                                                                                                                                                                                                                                                                                         |
| <b>Campus Solution</b>                | An SRDF solution that uses combinations of multimode LED fiber, single-mode laser fiber, ESCON directors, and ESCON repeaters to provide a link distance of up to 60 km. This solution typically uses a real-time (synchronous) operational mode for its remote mirroring technique. An asynchronous (journaling) mode is optional. |
| <b>Channel Director</b>               | The component in the Symmetrix subsystem that interfaces between the host channels and data storage. It transfers data between the channel and cache.                                                                                                                                                                               |
| <b>Channel Failover</b>               | The process of automatically changing the device path from a failed channel to a backup channel within the same host.                                                                                                                                                                                                               |
| <b>Command Descriptor Block (CDB)</b> | The structure used to communicate commands from an initiator to a target. This structure may be 6 bytes, 10 bytes, or 12 bytes in size.                                                                                                                                                                                             |

**Command queue** The queue in which I/O processes are queued for execution in the target. The head of the queue refers to the next queued I/O process to be executed. The tail of the queue refers to the last process to be executed in the queue.

**Connect** The function that occurs when an initiator selects a target to start an operation, or a target reselects an initiator to continue an operation.

**Contingent allegiance** A condition during which a target preserves sense data for a corresponding initiator. This condition is typically indicated by a Check Condition status.

**Control signals** A set of 9 signals lines that place the SCSI bus into its different phases. The combination of asserted and negated signal lines define the SCSI phases.

**Controller ID** Controller identification number of the director the disks are channeled to for EREP usage. There is only one controller ID for Symmetrix.

**Command Descriptor Block (CDB)** The structure used to communicate commands from an initiator to a target. This structure may be 6 bytes, 10 bytes, or 12 bytes in size.

## D

**DASD** Direct access storage device.

**Data availability** Access to any and all user data by the user.

**Delayed Write** There is no room in cache for the data presented by the write operation.

**Destage** The asynchronous write of new or updated data from cache to disk device.

**Device** A uniquely addressable part of the Symmetrix subsystem that consists of a set of access arms, the associated disk surfaces, and the electronic circuitry required to locate, read, and write data. See also *Volume*

**Device Controller** The hardware that interface devices (disks) to the SCSI bus. Symmetrix SCSI Adapter cards operate as device controllers on the host-to-Symmetrix SCSI interface.

|                                       |                                                                                                                                                                                                                                                                                                                                           |
|---------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Device-level parity protection</b> | The main feature of the Symmetrix RAID-S data protection option in which XOR parity calculations are performed on the disk device rather than the disk director, allowing the disk director controller to continue to handle I/O requests.                                                                                                |
| <b>Diagnostics</b>                    | System level tests or firmware designed to inspect, detect, and correct failing components. These tests are comprehensive and self-invoking.                                                                                                                                                                                              |
| <b>Differential interface</b>         | An electrical signal configuration that uses a pair of lines for transfer. This configuration (as compared to a single-ended interface) has a higher tolerance for common-mode noise and little crosstalk when used with twisted pair cables. It supports cables lengths to 20 meters (82 feet). See also <i>Single-ended interface</i> . |
| <b>Director</b>                       | See <i>Channel Director</i> or <i>Disk Director</i> .                                                                                                                                                                                                                                                                                     |
| <b>Disk Director</b>                  | The component in the Symmetrix subsystem that interfaces between cache and the disk devices.                                                                                                                                                                                                                                              |
| <b>Distributed Parity</b>             | The RAID-S feature achieved with Hyper Volume Extension in which the logical volumes (rather than data) in a RAID-S group are spread across the physical disk.                                                                                                                                                                            |
| <b>DMA</b>                            | Direct Memory Access.                                                                                                                                                                                                                                                                                                                     |
| <b>Dynamic Sparing</b>                | A Symmetrix feature that automatically transfers data from a failing disk device to an available spare disk device without affecting data availability. This feature supports all non-mirrored devices in the Symmetrix subsystem.                                                                                                        |
| <b>F</b>                              |                                                                                                                                                                                                                                                                                                                                           |
| <b>Failover</b>                       | The process of moving a system resource (application, volume, filesystem) that has failed from one node to another.                                                                                                                                                                                                                       |
| <b>Fast SCSI</b>                      | A synchronous transmission rate defined to be between 5 MHz and 10 MHz.                                                                                                                                                                                                                                                                   |
| <b>Fast-Wide Differential</b>         | A SCSI interface using two wires for each signal (differential), and capable of 16-bit data transfers (wide) at 20 MB/sec (fast).                                                                                                                                                                                                         |
| <b>Fast Write</b>                     | In Symmetrix, a write operation at cache speed that does not require immediate transfer of data to disk. The data is written directly to cache and is available for later destaging.                                                                                                                                                      |

|                                       |                                                                                                                                                                                                                                                                                       |
|---------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Filesystems</b>                    | Filesystems are the organizing structure for data. The filesystem is perhaps the most important part of the UNIX operating system. The filesystem goes beyond being a simple repository for data by providing the means of organizing the layout of the data storage in complex ways. |
| <b>Field Replaceable Unit (FRU)</b>   | A component that is replaced or added by service personnel as a single entity.                                                                                                                                                                                                        |
| <b>Fixed Block Architecture (FBA)</b> | A data recording format that employs data blocks of fixed size. The blocks are addressed by block number relative to the beginning of the particular file.                                                                                                                            |

## G

|                      |                                                                                                            |
|----------------------|------------------------------------------------------------------------------------------------------------|
| <b>Gigabyte (GB)</b> | 1024 x 1024 x 1024 bytes ( $2^{30}$ ).                                                                     |
| <b>Group</b>         | The physical disks related to each other for common parity protection when implementing the RAID-S option. |

## H

|                                     |                                                                                                                                                                                                                                                                                                                     |
|-------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>HACMP</b>                        | IBM's clustering product for AIX-RS/6000.                                                                                                                                                                                                                                                                           |
| <b>Head and Disk Assembly (HDA)</b> | A field replaceable unit in the Symmetrix subsystem containing the disk and actuator.                                                                                                                                                                                                                               |
| <b>Host</b>                         | A computer connected to a network or cluster system.                                                                                                                                                                                                                                                                |
| <b>Host Adapter</b>                 | A device (circuit board) that provides an interface between the SCSI bus and the computer I/O bus (e.g., PCI, EISA, Microchannel).                                                                                                                                                                                  |
| <b>Hostname</b>                     | A unique name by which a computer is known on a network or cluster system.                                                                                                                                                                                                                                          |
| <b>Hyper-Volume Extension</b>       | The ability to define more than one logical volume on a single physical disk device making use of its full formatted capacity. These logical volumes are user-selectable in size. The minimum volume size is one cylinder and the maximum size depends on the disk device capacity and the emulation mode selected. |

## I

|             |                               |
|-------------|-------------------------------|
| <b>ICDA</b> | Integrated Cached Disk Array. |
|-------------|-------------------------------|

|                                            |                                                                                                                                                                                                                                                                                                                                         |
|--------------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Identifier (ID)</b>                     | A sequence of bits or characters that identifies a program, device, controller, or system.                                                                                                                                                                                                                                              |
| <b>IML</b>                                 | Initial microcode program loading.                                                                                                                                                                                                                                                                                                      |
| <b>Initiator</b>                           | A SCSI device (usually a host system) that requests an operation to be performed by another SCSI device.                                                                                                                                                                                                                                |
| <b>Inodes</b>                              | Inodes are disk blocks that contain information relating to the file/directory. Some of the functions of the inodes are: the file's size, access permissions and ownership, and most importantly pointers to the data blocks associated with the file/directory. Inodes are very similar to the FAT (File Allocation Tables) in MS/DOS. |
| <b>Intermediate Status</b>                 | A status code sent from a target to an initiator upon completion of each command, except the last command, in a set of linked commands.                                                                                                                                                                                                 |
| <b>inq</b>                                 | Inquiry utility. An EMC utility program used to query SCSI devices for device configuration information.                                                                                                                                                                                                                                |
| <b>I/O Device</b>                          | An addressable input/output unit, such as a disk device.                                                                                                                                                                                                                                                                                |
| <b>IP Address</b>                          | Internet Protocol Address or Internet Address. The 32-bit host address defined by the Internet Protocol. Usually represented in dotted decimal notation (e.g., 123.456.7.8).                                                                                                                                                            |
| <b>K</b>                                   |                                                                                                                                                                                                                                                                                                                                         |
| <b>Kernel</b>                              | The kernel is the part of the system that manages the resources of the computer system it lives on. Its role is to keep track of the disks, tapes, printers, terminals, communication links, and any other devices.                                                                                                                     |
| <b>Kilobyte (K)</b>                        | 1024 bytes ( $2^{10}$ ).                                                                                                                                                                                                                                                                                                                |
| <b>L</b>                                   |                                                                                                                                                                                                                                                                                                                                         |
| <b>Least Recently Used Algorithm (LRU)</b> | The algorithm used to identify and make available the cache space by removing the least recently used data.                                                                                                                                                                                                                             |
| <b>Logical Partition</b>                   | In the AIX environment, it is a collection of physical partitions. A logical partition consists of 1 to 3 physical partitions.                                                                                                                                                                                                          |

|                                        |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
|----------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Logical Unit</b>                    | A physical or virtual device addressable through a target. A physical device can have more than one logical unit.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| <b>Logical Unit Number (LUN)</b>       | An encoded three-bit identifier for the logical unit.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| <b>Logical Volume (LV)</b>             | <p>A user-defined storage device. In the Symmetrix subsystem, the user can define a physical disk device as one, two, or three logical volumes or one, two, or three physical devices as one logical volume.</p> <p>In the AIX environment, logical volumes are groups of physical partitions. A typical LV is a collection of disk space. The physical partitions that make up an LV can reside on separate physical volumes and need not be contiguous when they are on the same physical volumes. All physical partitions in a logical volume must be on physical volumes in the same volume group.</p> |
| <b>Logical Volume Manager (LVM)</b>    | The part of the RS/6000 AIX system responsible for disk subsystems. Interaction with the LBM is done via SMIT menus or via stand-alone commands.                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
| <b>Long Miss</b>                       | Requested data is not in cache and is not in the process of being fetched.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| <b>Longitude Redundancy Code (LRC)</b> | Exclusive OR (XOR) of the accumulated bytes in the data record.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| <b>M</b>                               |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| <b>Media</b>                           | The disk surface on which data is stored.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| <b>Megabyte (MB)</b>                   | 1024 x 1024 bytes (2 <sup>20</sup> ).                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| <b>MII</b>                             | Machine Initiated Interrupt.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| <b>Mirroring</b>                       | The Symmetrix maintains two identical copies of a designated volume on separate disks. Each volume automatically updates during a write operation. If one disk device fails, Symmetrix automatically uses the other disk device.                                                                                                                                                                                                                                                                                                                                                                           |
| <b>Mirrored Pair</b>                   | A logical volume with all data recorded twice, once on each of two different physical devices.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |

**Model Unique  
Licensed Internal  
Code (MULIC)**

The MULIC is an integral part of the OS/400. It needs to be loaded upon initial installation of the unit only!

**Mount Point**

A directory file that identifies a mounted filesystem.

**Multi-initiator SCSI**

A SCSI bus with more than one SCSI initiator. SCSI bus connections are made using Y-cables to "daisy chain" multiple hosts to shared disks.

**N****Nexus**

The connection between the initiator and target. This connection begins with the initial connection and ends with the completion of the I/O process.

**O****Object Database  
Manager (ODM)**

The ODM is a data manager that controls the storage of AIX and HP-UX system data. All devices, volume groups, logical volumes, physical volumes, physical partitions, and logical partitions have an entry in the ODM. ODM also contains information on commands that are initiated via SMIT (System Manager Interface Tool) and SAM (System Administrator Manager).

**P****Page**

Several commands use regular parameter structures that are referred to as pages. These pages are identified with a value known as a page code.

**Parallel channel**

Block multiplexer channel (mainframe hosts).

**PCI**

Peripheral Component Interconnect. A self-configuring local bus used on most Pentium-based PCs and multi-processor machines. Provides a synchronous, asymmetrical I/O channel. The PCI bus is processor independent, and supports Power PC and Motorola 680x0 processors.

**Physical Partition (PP)**

A physical partition is the smallest unit of disk space that can be allocated in a volume group in an AIX environment. Any disk space allocated is an integral number of physical partitions. By default, a PP is 4 MB in size.



**Physical Volume** In LVM terminology (AIX environment), a physical volume is each physical disk drive connected to the system. A physical volume is an addressable disk on the SCSI bus. By default, AIX refers to the physical volumes as *hdisk0*, *hdisk1*, *hdisk2*, etc.

**Physical Volume Identifier (PVID)** The PVID is information that the RS/6000 writes to block 0 of each disk. Every disk that is part of a volume group must have a PVID. The PVID is part of the VGDA. Volume groups use the PVID, therefore only one volume group can reside on a physical volume (disk). You can have multiple physical volumes in a volume group. The PVID and all the volume group information are kept in two places: the ODM and on the physical volume in block 0 of the volume group descriptor area.

**Private Bus** A SCSI bus that connects a device used exclusively by one node in a cluster.

**Promotion** The process of moving data from a track on the disk device to cache. See also *Stage*.

## Q

**QCIC** Quad Channel Integrated Controller. Sequent Symmetrix SCSI host adapter with four separate SCSI ports.

**Queue** This term refers to the command queue used in tagged queuing. It is a first-in/first-out (FIFO) structure.

**Queue Tag** The value associated with an I/O process that uniquely identifies it from other queued I/O processes on the logical unit for the same initiator.

## R

**RAID-S** A Symmetrix feature that provides parity data protection on the disk device level using physical parity volumes. A RAID-S group normally consists of 3 data drives and 1 parity drive. The RAID-S option can be also used with Hyper Volume Extension to establish distributed parity.

**Rank** The logical volumes on a disk which are related to each other for common parity protection. A single RAID-S group can contain 1 to 4 ranks.

|                     |                                                                                                                                                                                                                                                                                                                                                  |
|---------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Read Hit</b>     | Data requested by the read operation is in cache.                                                                                                                                                                                                                                                                                                |
| <b>Read Miss</b>    | Data requested by the read operation is not in cache.                                                                                                                                                                                                                                                                                            |
| <b>Reconnect</b>    | The function that occurs when a target selects an initiator to continue an operation after a disconnect.                                                                                                                                                                                                                                         |
| <b>Reconnection</b> | A reconnection exists from the assertion of the BSY signal in a RESELECTION phase until the next BUS FREE phase occurs. A reconnection can only occur between a target and an initiator.                                                                                                                                                         |
| <b>Reselect</b>     | The function that occurs when a target disconnects from an initiator to perform a time-consuming function and, then after performing that function, reestablishes the connection.                                                                                                                                                                |
| <b>Reserved</b>     | The term used for bits, bytes, fields, and code values that are set aside for future standardization.                                                                                                                                                                                                                                            |
| <b>Resource</b>     | Hardware or software that provides a service, e.g., an application, a filesystem, or a disk device.                                                                                                                                                                                                                                              |
| <b>Root</b>         | Root is the highest level of authority on UNIX platforms. Most higher level functions require you to have root privileges.                                                                                                                                                                                                                       |
| <b>S</b>            |                                                                                                                                                                                                                                                                                                                                                  |
| <b>SAM</b>          | System management tool for HP-UX.                                                                                                                                                                                                                                                                                                                |
| <b>SCSI</b>         | Small Computer System Interface. A standard interface used to connect disks and other devices to a computer. SCSI devices can be configured in series to allow multiple devices on the same bus.                                                                                                                                                 |
| <b>SCSI Adapter</b> | A circuit board used to interface the SCSI bus with an internal bus. Used in this manual to identify the Symmetrix SCSI Adapter that interfaces the SCSI bus from the host(s) to the Symmetrix SCSI Channel Director. Symmetrix SCSI Adapters, located in the rear card cages, have four SCSI connectors for attaching cables from the hosts(s). |
| <b>SCSI Address</b> | The octal representation of the unique address (0-7) assigned to a SCSI device. This address would normally be assigned and set in the SCSI device during system installation.                                                                                                                                                                   |

|                                |                                                                                                                                                                                                                                                                                                                                    |
|--------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>SCSI ID</b>                 | The bit significant representation of the SCSI address referring to one of the signal lines DB(7-0).                                                                                                                                                                                                                               |
| <b>SCSI device</b>             | A host computer adapter or peripheral controller or intelligent peripheral that can be attached to the SCSI bus.                                                                                                                                                                                                                   |
| <b>SCSI Termination</b>        | Resistor network connected to SCSI signal lines for impedance matching to prevent signal reflection and degradation. A SCSI bus requires termination at both ends, and only at both ends, of the bus.                                                                                                                              |
| <b>Service Guard (HP)</b>      | HP clustering product. This product is available with HP-UX 10.                                                                                                                                                                                                                                                                    |
| <b>Shared Device</b>           | Device that may be accessed by more than one host node.                                                                                                                                                                                                                                                                            |
| <b>Shell</b>                   | The shell is the command interpreter. The shell is just a utility program and is not properly a part of the system. The shell “listens” to your terminal and translates your requests into actions on the part of the kernel and the many utility programs. Two of the more common shells are the Korn shell and the Bourne shell. |
| <b>Single Point of Failure</b> | A hardware or software item that will cause an application or a whole system to become unavailable if it fails. High availability design attempts to eliminate SPOF through redundancy, recovery, or failover.                                                                                                                     |
| <b>Short Miss</b>              | Requested data is not in cache, but is in the process of being fetched.                                                                                                                                                                                                                                                            |
| <b>Signal Assertion</b>        | The act of driving a signal to the true state.                                                                                                                                                                                                                                                                                     |
| <b>Signal Negation</b>         | The act of driving a signal to the false state or allowing the cable terminators to bias the signal to the false state (by placing the driver in a high impedance condition).                                                                                                                                                      |
| <b>Signal release</b>          | The act of allowing the cable terminators to bias the signal to the false state (by placing the driver in the high impedance condition).                                                                                                                                                                                           |
| <b>Single-ended interface</b>  | An electrical signal configuration that uses a single line for each signal and references a ground path common to the other signal lines. The single-ended configuration is susceptible to common mode noise and has a maximum cable length of 6 meters. See also <i>Differential interface</i> .                                  |
| <b>SMIT</b>                    | System management tool for AIX.                                                                                                                                                                                                                                                                                                    |

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
|---------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>SRDF</b>                     | Symmetrix Remote Data Facility. SRDF consists of the microcode and hardware required to support Symmetrix remote mirroring.                                                                                                                                                                                                                                                                                                                                                                                                      |
| <b>Stage</b>                    | The process of writing data from a disk device to cache. See also <i>Promotion</i> , <i>Destage</i> .                                                                                                                                                                                                                                                                                                                                                                                                                            |
| <b>Striping</b>                 | The process of segmenting logically sequential data and writing the segments to multiple physical disk devices.                                                                                                                                                                                                                                                                                                                                                                                                                  |
| <b>Superblock</b>               | Each filesystem begins with a special disk block called a "superblock". This superblock contains global information regarding the filesystem, such as block and fragment size. The superblock is checked when the system boots and when a filesystem check (fsck) is run from a command line. The superblock is in a known location. Redundant copies of the superblock are kept in every cylinder group. If the main superblock (primary copy) is corrupted, the system can use the other copies to reconstruct the filesystem. |
| <b>Switchover</b>               | Transferring the control of the application from one host to another when a failure occurs.                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| <b>Synchronous transmission</b> | A timing protocol that uses a master clock with a clock period and allowable offset that holds the sending and receiving devices in the desired phase relationship.                                                                                                                                                                                                                                                                                                                                                              |
| <b>System-Level Mirroring</b>   | RAID-1 implemented by the host using volume manager software. The term "system-level" is used to differentiate this type of mirroring from Symmetrix mirroring, which is implemented within Symmetrix and is transparent to the host.                                                                                                                                                                                                                                                                                            |
| <b>T</b>                        |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| <b>Target</b>                   | A SCSI device that performs an operation requested by an initiator.                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
| <b>Target ID</b>                | See <i>SCSI ID</i> .                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
| <b>Terabyte</b>                 | 1,024 gigabytes (2 <sup>40</sup> ).                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
| <b>Terminator</b>               | A resistor network connected to SCSI signal lines for impedance matching to prevent signal reflection and degradation. Termination must be provided at both ends, and only at both ends, of a SCSI bus. Refers to the module that connects to a SCSI port or Y-cable to provide external termination.                                                                                                                                                                                                                            |

**Terminator Power** A +5VDC level provided to the TERMPWR contact of the SCSI bus to power the terminator circuits (see *Terminator*). Terminator power is normally provided by the host initiator. Note that a non-terminated host adapter (terminator resistors removed) can continue to supply terminator power to terminators at the end of the bus.

## U

**UNIX** UNIX is an interactive, multi-tasking, multi-user operating system. UNIX is written in “C” language.

There are three types of UNIX files: directories, data files, and special files. A directory is a file containing certain information about another file. A directory contained within another directory is a sub-directory.

The two most common types of UNIX are BSD (Berkeley Software Distribution) and System VR4 (developed by AT&T). Most UNIX systems are a “mix” of both types.

## V

**Vertical Licensed Internal Code (VLIC)** The VLIC is the part of the O/S 400 that is loaded first during a system load/restore. The VLIC resides on the load source disk (unit 1) that is the first addressable disk in the AS/400 residing in ASP #1.

**Volume** A general term referring to a storage device.

**Volume Group** A group of physical volumes in the AIX environment. The name of the volume group is chosen by the person who defines the group.

**Volume Group Descriptor Area (VGDA)** Contains the control information for a volume group (AIX environment). At least one VGDA resides on each physical volume in the volume group.

## W

**Write Hit** There is room in cache for the data presented by the write operation.

**Write Miss** There is no room in cache for the data presented by the write operation.

## X

**XFS (SGI)** Logical volume manager (Silicon Graphics Inc.)

## Y

**Y-Cable** A SCSI cable with three connectors used to split the bus in a multi-initiator bus configuration.

## Z

**Zero** A logical false signal value (negation).

## Sales and Service Locations

### **Corporate Headquarters**

EMC Corporation  
Hopkinton, Massachusetts  
01748-9103 U.S.A.  
508-435-1000  
1-800-424-EMC2

### **Asia Pacific Headquarters**

EMC Japan K.K.  
P.O. Box 323  
Shinjuku-Mitsui Bldg., 55  
1-1, Nishi-Shinjuku  
2-chome Shinjuku-ku  
Tokyo 163-0466 JAPAN  
Tel: 81-3-3345-3211  
Fax: 81-3-3345-3221

### **North American Sales and Service Locations**

Atlanta, GA 404-705-4750  
Baltimore, MD 410-850-4324  
Bellevue, WA 425-201-6400  
Bethesda, MD 301-530-0091  
Boston, MA 617-449-8100  
Charlotte, NC 704-521-9773  
Chicago, IL 708-390-8800  
Cincinnati, OH 513-745-0300  
Cleveland, OH 216-573-1162  
Columbus, OH 614-436-3900  
Dallas, TX 972-233-5676  
Denver, CO 303-294-9966  
Farmington Hills, MI 313-553-4810  
Ft. Lauderdale, FL 305-776-3622  
Greensboro, NC 910-665-1077

Houston, TX 713-621-9800  
 Indianapolis, IN 317-577-9766  
 Inglewood, CA 310-364-1222  
 Irvine, CA 714-833-1442  
 Lenexa, KS 913-469-9292  
 Long Island, NY 516-393-5045  
 Minneapolis, MN 612-835-1994  
 Montreal, Quebec 514-856-6166  
 Nashville, TN 615-781-4394  
 New York City, NY 212-564-6866  
 Ottawa, Ontario 613-233-0111  
 Orlando, FL 407-855-4087  
 Phoenix, AZ 602-955-0702  
 Philadelphia, PA 610-834-7740  
 Pittsburgh, PA 412-922-5222  
 Portland, OR 503-293-8450  
 Raleigh, NC 919-420-0405  
 Rochester, NY 716-387-0970  
 Salt Lake City, UT 801-532-1454  
 San Diego, CA 619-576-1880  
 San Francisco, CA 415-871-1970  
 Southington, CT 860-620-6730  
 Springfield, NJ 201-467-7979  
 St. Louis, MO 314-469-9005  
 Tampa, FL 813-282-0274  
 Toronto, Ontario 905-206-1580  
 Tulsa, OK 918-663-2255  
 Vancouver, BC 604-270-1657  
 Vienna, VA 703-893-8400

### International Sales Locations

Berlin, Germany (49) 30 25 49 31 86  
 Bologna, Italy (39) 51 522579  
 Brescia, Italy (39) 30 2421791  
 Brussels, Belgium (32) 2 725 74 25  
 Cape Town, South Africa (27) 21 686 1430  
 Dublin, Ireland (353) 1 475 4172  
 Durban, South Africa (27) 31 83 6611  
 Dusseldorf, Germany (49) 2131 9191 0  
 Frankfurt, Germany (49) 6196 4 72 80  
 Hamburg, Germany (49) 4069 69 610  
 Hong Kong, Taikooshing (852) 2839 9600  
 Johannesburg, South Africa (27) 11 807 5300  
 Leatherhead, Surrey, UK (44) 1372 36 0000



Lyon, France (33) 7 860 1330  
Marseille, France (33) 4292 2549  
Melbourne, Australia (61) 3 653 9519  
Milan, Italy (39) 02 409081  
Munich, Germany (49) 89 14 31 320  
Nagoya, Japan (81) 52 223 1900  
Newton Le Willows, Merseyside, UK (44) 1942 275 511  
Nieuwegein, Netherlands (31) 03 6055777  
Nuernberg, Germany (49) 911 2379260  
Osaka, Japan (81) 6 373 8300  
Padua, Italy (39) 49 8235853  
Paris, France (33) 1 3082 5100  
Pretoria, South Africa (27) 12 663 6635  
Rome, Italy (39) 06 8552116  
Schlieren, Switzerland (411) 743 4949  
Seoul, Korea (82) 2 539 4455  
Singapore (65) 338 9265  
Stuttgart, Germany (49) 7152 979340  
Sydney, Australia (61) 2 922 7888  
Toulouse, France (33) 6131 6262  
Turin, Italy (39) 11 746527  
Vienna, Austria (43) 1 599 99 627



## Symbols

/etc/.osm file 4-19  
/usr/adm/msgbuf 5-40  
/var/adm/messages 12-48, 13-19  
/var/adm/ras/errlog 9-24  
/var/adm/SYSLOG 10-25, 11-23  
/var/adm/syslog.cont 15-50

## A

Adaptec AHA2944UW 19-6  
Adding devices online  
    AIX 9-23  
    Fujitsu 15-16  
    HP-UX 6-27, 7-20, 17-18  
    NCR 4-18  
    SGI/Challenge 10-24, 11-22  
    Solaris 12-14, 16-12  
    Solaris x86 13-18  
Adding logical volumes 9-20  
AIX  
    adding devices online 9-23  
    configuring devices 9-15 to 9-19  
    creating a file system 9-18  
    creating a mount point 9-18  
    creating logical volumes 9-18  
    incorporating Symmetrix 9-12 to 9-14  
    Logical Volumes (LV) 9-3  
    LVM 9-2  
    operating system features 9-4  
    overview 8-2 to ??, 9-2 to 9-3, 18-2 to ??  
    Physical Partitions (PP) 9-3  
    physical volume (PV) 9-2  
    Symmetrix configuration 9-11

system and error messages 9-24 to 9-29  
terminology 9-2  
utilities and functions 9-5 to 9-9, 9-10  
Volume Group Descriptor Area (VGDA) 9-3  
Volume Groups 9-3  
AIX 4.3 9-12  
Alternate links  
    setting up 6-22  
Ascend 18-1

## B

bdf  
    HP-UX 6-7, 7-7, 17-6

## C

Cable connections  
    split cable B-3  
    Y-cable B-4  
cfgmgr  
    AIX 9-5  
Channel failover  
    current links 6-24  
    HP-UX 6-22  
    primary link failure 6-23  
    setting up alternate links 6-22  
Chapter Contents 5-1, 6-1, 9-1, 10-1, 11-1, 13-1,  
    14-1, 15-1, 18-1  
Command Tag Queueing 12-2  
Configuring a file system  
    using the command line 4-16  
Configuring devices  
    adding a volume group 9-20  
    adding logical volumes 9-20

- creating a filesystem 6-18, 7-18, 9-18, 10-12, 11-10, 15-13, 17-16
- creating a mount directory 6-18, 7-18, 12-41, 12-42, 12-43, 12-44, 12-45, 12-46, 15-14, 15-43, 15-44, 15-45, 15-46, 15-47, 15-48, 16-39, 16-40, 16-41, 16-42, 16-43, 16-44, 17-16
- creating a mount point 9-18
- creating a UFS filesystem 12-41, 12-43, 12-44, 15-43, 15-45, 15-47, 16-39, 16-41, 16-42
- creating a VxFS filesystem 12-42, 12-43, 12-45, 15-44, 15-46, 15-48, 16-40, 16-41, 16-43
- creating filesystems 9-21
- creating physical volumes 6-17, 7-17, 17-15
- making a volume group 9-15
- making a volume group available 9-17
- mounting a directory 6-18, 7-18, 17-16
- partitioning and labeling 10-8, 11-7
- using dskconfig.ATT 4-14 to 4-15
- using SMIT 9-20 to 9-22
- using sysadm 4-11 to 4-14
- using the command line 4-15 to 4-17, 9-15 to 9-19
- Configuring new devices
  - Fujitsu 15-12 to 15-14
  - HP-UX 6-17 to 6-19, 7-17 to 7-18, 17-15 to 17-16
  - NCR 4-11 to 4-17
  - SGI/Challenge 10-7 to 10-9, 11-6 to 11-9
  - Solaris 12-12 to 12-41, 15-43, 16-11 to 16-39
  - Solaris x86 13-13
- Configuring the kernel 15-8 to 15-9
- create 4-14
- Creating a bootable disk
  - HP-UX 6-9 to 6-12, 7-9 to 7-11
- Creating a concatenated volume
  - Solstice DiskSuite 12-15, 15-17, 16-13
- Creating a concatenated volume for UFS logging 12-16, 15-18, 16-14
- Creating a filesystem 9-18, 10-12 to 10-13, 11-10 to 11-11
  - Fujitsu 15-13
  - HP-UX 6-18, 7-18, 17-16
  - Sequent 5-24
  - SGI/Challenge 10-12, 11-10
  - Solstice DiskSuite 12-40 to 12-47, 15-42 to 15-49, 16-38 to 16-45
  - Veritas Volume Manager 12-40 to 12-47, 15-42 to 15-49, 16-38 to 16-45
- Creating a mirrored striped volume
  - Solstice DiskSuite 12-20, 15-22, 16-17
- Creating a mirrored striped volume for UFS logging 12-21, 15-23, 16-19
- Creating a mirrored volume
  - using vxassist 12-37, 15-40, 16-36
- Creating a mount directory
  - Fujitsu 15-14
  - HP-UX 6-18, 7-18, 17-16
  - Sequent 5-24
  - Solaris 12-41, 12-42, 12-43, 12-44, 12-45, 12-46, 15-43, 15-44, 15-45, 15-46, 15-47, 15-48, 16-39, 16-40, 16-41, 16-42, 16-43, 16-44
- Creating a mount point 9-18
- Creating a stripe volume set 10-14 to 10-19, 11-12 to 11-14
- Creating a striped mirrored volume
  - using vxmake 12-34, 15-36, 16-32
  - Veritas Volume Manager 12-33, 15-36, 16-32
- Creating a striped volume
  - Solstice DiskSuite 12-17, 15-19, 16-15
  - using vxassist 12-33, 15-35, 16-31
  - using vxmake 12-30, 15-32, 16-28
  - Veritas Volume Manager 12-30, 15-32, 16-28
- Creating a striped volume for UFS logging 12-18, 15-20, 16-16
- Creating a UFS filesystem
  - Solaris 12-41, 12-43, 12-44, 15-43, 15-45, 15-47, 16-39, 16-41, 16-42
- Creating a VxFS filesystem
  - Solaris 12-42, 12-43, 12-45, 15-44, 15-46, 15-48, 16-40, 16-41, 16-43
- Creating an advanced filesystem 14-16
- Creating an xfs filesystem 10-12, 11-10
- Creating concatenated volumes
  - using vxassist 12-29, 15-31, 16-27
  - using vxmake 12-26, 15-28, 16-24
  - Veritas Volume Manager 12-26, 15-28, 16-24
- Creating device files
  - NCR 4-15
- Creating filesystems
  - NCR 4-13

- using SMIT 9-21
  - using sysadm 4-13
- Creating logical volumes 6-17, 7-17, 9-18, 17-15
- Creating physical volumes
  - creating a filesystem 5-24
  - creating a mount directory 5-24
  - HP-UX 6-17, 7-17, 17-15
  - mounting a directory 5-24
  - Sequent 5-23 to 5-25
- Creating volume groups 6-17, 7-17, 17-15
- Creating volumes
  - using Sequent Volume Manager 5-27 to 5-29
- crfs
  - AIX 9-8
- Cursor movement commands D-2
- Cylinders per Logical Volume Physical Devices 2-13

## D

- Data and command formats
  - FBA devices 2-3
- Deleting a stripe volume set 10-19, 11-14
- Device characteristics 9-14
- Device definition files
  - accessing through FTP server A-4
  - accessing through world wide web A-3
  - AIX 9-12
  - NCR 4-7
  - overview A-2
  - Sequent 5-19
  - Solaris 12-9, 15-10, 16-8
  - Solaris x86 13-11
  - transferring A-5
- df
  - IRIX 10-3, 11-3
- Diagnostic patches
  - HP-UX 10.x 6-28, 7-22, 17-19
  - HP-UX 9.04 6-28, 17-19
  - HP-UX 9.05 7-22
- Digital (Tru64) UNIX
  - creating an advanced filesystem 14-16
  - overview 14-2
  - Symmetrix configuration 14-6
  - system and error messages 14-22
  - utilities and functions 14-4
- Dirty region logging 12-37, 15-40, 16-36

- Disk devices
  - capacity 2-4
- diskid
  - Sequent 5-5
- disklabel
  - Digital UNIX 14-4
- disksecn
  - HP-UX 6-5, 17-4
- disksetup
  - NCR 4-5
- Displaying new devices
  - Sequent 5-15
- dklayout
  - NCR 4-4
- dmesg
  - Sequent 5-5
- DNCP-400 18-2
- dumpconf
  - Sequent 5-5

## E

- edvtoc
  - NCR 4-5
- EMC FTP server A-2, A-4
- Enterprise Backup Solutions 1-23
  - EMC Data Manager 1-23
  - Fast Dump Restore Safeguard Open Storage 1-23
  - Legato Networker for EMC Symmetrix 1-23
  - Omniback II 1-23
- Environment variables
  - GPLOGGEROPTS 4-19
- Error log file 9-24, 10-25, 11-23, 12-48, 13-19, 15-50
- Error log files
  - NCR 4-19
  - Sequent 5-40
- Error messages
  - AIX 9-24 to 9-29
  - Digital UNIX 14-22
  - Fujitsu 15-50
  - HP-UX 6-28, 7-22, 17-19
  - NCR 4-19
  - Sequent 5-40 to 5-41
  - SGI/Challenge 10-25 to 10-26, 11-23 to 11-24
  - Solaris 12-48, 16-46
  - Solaris x86 13-19

**errpt**

AIX 9-9

viewing messages 9-24

Example 10-18, 11-16

Exiting vi commands D-3

**F**

FBA data format 2-3

**fdisk**

NCR 4-4

Features of SCO UnixWare 7 19-2

Filesystem expansion 12-46, 15-49, 16-44

**format**

Solaris 12-4, 16-4

Solaris x86 13-4

**Fujitsu**

adding devices online 15-16

configuring new devices 15-12 to 15-14

configuring the kernel 15-8 to 15-9

creating a filesystem 15-13

creating a mount directory 15-14

incorporating Symmetrix 15-8 to 15-9

overview 15-2

recognizing LUNs 15-8

Symmetrix configuration 15-7

system and error messages 15-50

utilities and functions 15-5

**Functions**

AIX 9-5 to 9-9, 9-10 to??

Digital UNIX 14-4 to??

Fujitsu 15-5

HP/9000 Series 700 7-5

HP-UX 6-5 to 6-7, 7-5 to 7-7, 17-4 to 17-6

NCR 4-4 to 4-5

Sequent 5-5, 20-3 to??

SGI/Challenge 10-3 to 10-4,?? to 10-4, 11-3 to 11-4,?? to 11-4

Solaris 12-4, 16-4

Solaris x86 13-4

**fx**

IRIX 10-3, 11-3

**G**

GP7000F 16-1

**H****halt**

Solaris 12-4, 16-4

**Hardware mapping**

HP/9000 Series 700 7-16

HP-UX 10.0 6-13 to 6-14, 7-12 to 7-14, 17-9

Solaris 12-5, 16-5

**hinvt**

IRIX 10-3, 11-3

**Host controller card**

FWD SCSI (SGI S-Series) 11-2

IBM 2412 9-2

IBM 2416 9-2

MDL FAWS 7-2

Model 1062A 12-2, 16-2

Model QCW-CTRL-01 5-2

SGI DF SCSI (IO4) 11-2

**Host controller cards C-2****HP/9000 Series 700**

overview 7-2 to??

recognizing new devices 7-15

Symmetrix configuration 7-8

utilities and functions 7-5 to??

**HP-9000 K-class servers 6-21****HP-UX**

adding devices online 6-27, 7-20, 17-18

channel failover 6-22

configuring new devices 6-17 to 6-19, 7-17 to 7-18, 17-15 to 17-16

creating a filesystem 6-18, 7-18, 17-16

creating a mount directory 6-18, 7-18, 17-16

creating physical volumes 6-17, 7-17, 17-15

Creating volume groups and logical volumes 6-17, 7-17, 17-15

hardware mapping 6-13 to 6-14, 7-12 to 7-14, 17-9

I/O configuration 6-14, 7-13

mounting a directory 6-18, 7-18, 17-16

overview 6-2, 7-3, 17-2

performance issues

Performance issues

**HP-UX 7-21**

recognizing new devices 6-15, 7-15 to 7-16, 17-14

SAM 6-20, 17-19, 17-17

Symmetrix configuration 6-8, 17-7

- system and error messages 6-28, 7-22, 17-19
- utilities and functions 6-5 to 6-7, 7-5 to 7-7, 17-4 to 17-6
- HP-UX 10.x
  - diagnostic patches 6-28, 7-22, 17-19
- HP-UX 9.04
  - diagnostic patches 6-28, 17-19
- HP-UX 9.05
  - diagnostic patches 7-22
- Hyper-volume extension
  - disk capacities 2-16
  - Fast-Wide SCSI devices 2-16 to 2-18
  - general description 2-16
  - multiple logical volumes per physical device 2-16
- Hyper-volumes
  - open systems 2-20

**I**

- I/O configuration
  - HP-UX 6-14, 7-13
- IBM 2412 controller card 9-2
- IBM 2416 controller card 9-2
- Incorporating Symmetrix
  - AIX 9-12 to 9-14
  - Fujitsu 15-8 to 15-9
  - Sequent 5-7 to 5-12
  - Solaris 12-9 to 12-11, 15-11, 16-8 to 16-10
  - Solaris x86 13-11
- Information Management 1-17
  - EMC ControlCenter 1-18, 1-19
  - EMC PowerPath 1-20
- Information Protection 1-14
  - EMC CopyPoint for AS/400 1-16
  - Non-disruptive Symmetrix Data Migration Service 1-15
  - Symmetrix Remote Data Facility 1-14
  - System Data Migration Service 1-15
- Information Sharing 1-17
  - EMC InfoMover 1-17
  - Symmetrix Storage Platform 1-17
- inq command A-6 to A-7
  - output fields A-6
  - supported hosts A-6
- Inserting text commands D-2
- insf

- HP-UX 6-5, 7-5, 17-5
- insight
  - IRIX 10-3, 11-3
- installing 4-7
- Intent logging 12-40, 15-42, 16-38
  - SDS 12-40, 15-42, 16-38
  - VxFS 12-40, 15-42, 16-38
  - VxVM 12-40, 15-42, 16-38
- ioscan
  - HP-UX 6-5, 7-5, 17-4
- iostat 3-3
  - Digital UNIX 14-4

## L

- Linux 20-1
- Logical Storage Manager (LSM)
  - Digital UNIX 14-4
  - examples 14-18 to 14-21
- Logical Volume Manager (LVM) 9-2
- Logical volumes 9-3
- Logical-to-physical volume relationships 2-16
- lsdev
  - AIX 9-5, 9-6
  - HP-UX 6-7, 7-7, 17-6
- LSM examples 14-18 to 14-21
  - creating a 4-way striped volume 14-20
  - creating a mirrored (2-way) volume 14-19
  - setting up 14-18
- lsvg
  - AIX 9-7
- Lucent 18-1
- LUN support
  - Sequent environment 5-10
- lvchange
  - HP-UX 6-7, 7-6, 17-6
- lvcreate
  - HP-UX 6-6, 7-6, 17-6
- lvdisplay
  - HP-UX 6-7, 7-7, 17-6
- lvextend
  - HP-UX 6-7, 7-7, 17-6
- lvreduce
  - HP-UX 6-7, 7-7, 17-6
- lvremove
  - HP-UX 6-7, 7-7, 17-6

**M****MAKEDEV**

- IRIX 10-3, 11-3

- Making a volume group 9-15

- Making a volume group available 9-17

- MDL FAWS 7-2

- menu

- Sequent 5-5

- metaclear

- Solstice DiskSuite 12-24, 15-26, 16-22

- metadb

- Solstice DiskSuite 12-24, 15-26, 16-22

- metahs

- Solstice DiskSuite 12-24, 15-26, 16-22

- metareplace

- Solstice DiskSuite 12-24, 15-26, 16-22

- metastat

- Solstice DiskSuite 12-24, 15-26, 16-22

- Miscellaneous vi commands D-3

- mkdev

- AIX 9-7

- mkdir

- IRIX 10-3, 11-3

- mkfs

- IRIX 10-3, 11-3

- mknod

- HP-UX 6-5, 7-5, 17-5

- mkvg

- AIX 9-7

- Model QCW-CTRL-01 5-2

- monitor

- Sequent 5-5

- mount

- AIX 9-8

- IRIX 10-3, 11-3

- Mounting a directory

- HP-UX 6-18, 7-18, 17-16

- Sequent 5-24

**N****NCR**

- adding devices online 4-18

- configuring new devices 4-11 to 4-17

- creating filesystems 4-13

- defining partitions 4-11

- displaying new devices 4-10

- incorporating Symmetrix 4-7 to 4-9

- overview 4-2 to 4-3

- Symmetrix configuration 4-6

- system and error messages 4-19

- NEC 8-1

- newfs

- Digital UNIX 14-5

- HP-UX 6-7, 7-7, 17-6

- Sequent 5-5

- Solaris 12-4, 16-4

- Solaris x86 13-4

- Number of cylinders

- 23 GB disk devices 2-16

- 4 GB disk devices 2-16

- 9 GB disk devices 2-16

**O**

- odm.add file 9-14

- odmget

- AIX 9-9

- Operating system features

- AIX 9-4

- HP-UX 6-3, 6-4, 7-3, 7-4, 17-2

- NCR 4-2

- Sequent 5-3, 5-4, 20-2, 20-6

- SGI/Challenge 10-2

- Solaris 12-3, 15-4, 16-3

- Solaris x86 13-6, 13-7, 13-8

- Optimizing performance 2-2

**P**

- Partition Table 19-11

- Partitioning devices

- SGI/Challenge environment 10-8, 11-7

- Solaris 2.3 12-12

- using sysadm 4-11 to 4-12

- Partitions

- defining on NCR host 4-11

- Performance

- optimizing 2-2

- Performance tools

- iostat 3-3

- sar 3-2

- UNIX 3-2 to 3-3

- vmstat 3-3

- Phone multiplexer 1-13



Physical Partitions (PP) 9-3

Physical volume (PV) 9-2

Plex 5-28

Primary link failure 6-23

PRIMEPOWER 16-1

probe-scsi-all

SparcServers 12-7, 16-7

SparcStations 12-7, 16-7

UltraSparc Series 12-8

prtconf

NCR 4-4

prtvto

NCR 4-4

PV links 6-22

pvchange

HP-UX 6-5, 7-5, 17-5

pvcreate

HP-UX 6-5, 7-5, 17-5

pvdisplay

HP-UX 6-6, 7-5, 17-5

## Q

qcic\_space.c file 5-10

## R

Recognizing LUNs 15-8

Recognizing new device

HP-UX 10.x 7-16

Recognizing new devices 6-15, 7-15 to 7-16, 17-14

HP/9000 Series 700 7-15

HP-UX 10.x 6-15

HP-UX 9.04 7-15

rmdev

AIX 9-8

Root disk 7-9

## S

SAM 6-20, 7-19, 17-17

Santa Cruz Operation 19-1

sar 3-2, 14-4

Saving vi commands D-3

SCSI protocols

Sequent 5-7, 5-8

scsi\_initiator\_id 12-5, 16-5

scsi\_options 12-11, 15-11, 16-9

scsi\_space.c file 5-10

scsidisk\_space.c file 5-7, 5-8, 5-10

scsimgr

Digital UNIX 14-5

scu

Digital UNIX 14-4

Searching commands D-3

Sequent

creating a filesystem 5-24

creating a mount directory 5-24

creating physical volumes 5-23 to 5-25

displaying new devices 5-15

formatting and partitioning definition files  
5-18, 5-19

incorporating Symmetrix 5-7 to 5-12

LUN support 5-10

mounting a directory 5-24

operating system features 5-3, 5-4, 20-2, 20-6

overview 5-2 to 5-4

Sequent Volume Manager (SVM) 5-26 to 5-39

Symmetrix configuration 5-6

system and error messages 5-40 to 5-41

target ID support 5-7

utilities and functions 5-5, 20-3 to??

Sequent Menu utility 5-40

Sequent Volume Manager 5-26 to 5-39

creating volumes 5-27 to 5-29

display volume manager records 5-39

object control operations 5-36

plex operations 5-34

subdisk operations 5-33

utilities 5-33 to 5-39

voldisk command 5-36

voledit command 5-36

volmend command 5-37

volplex command 5-34

volprint command 5-39

volsd command 5-33

volstat command 5-38

volume command 5-35

volume operations 5-34

volume statistics reporting 5-38

SGI/Challenge

adding devices online 10-24, 11-22

configuring new devices 10-7 to 10-9, 11-6 to  
11-9

- creating a stripe volume set 10-14 to 10-19, 11-12 to 11-14
  - operating system features 10-2
  - overview 10-2, 11-2
  - Symmetrix configuration 10-5, 11-5
  - system and error messages 10-25 to 10-26, 11-23 to 11-24
  - using XLV 10-14 to 10-19, 11-12 to 11-14
  - utilities and functions 10-3 to 10-4, 11-3 to 11-4,
- SGI/CHALLENGE environment 10-1
- Shared volumes 5-27
- showcfg
  - Sequent 5-5
- showprods
  - IRIX 10-3, 11-3
- shutdown
  - IRIX 10-4, 11-4
  - NCR 4-5
- Silicon Graphics 10-2
- Slices
  - defining on NCR host 4-11
- Slicing devices
  - NCR 4-12
  - using sysadm 4-12
- SMIT 9-3, 9-5
  - viewing messages 9-26
- Software options 1-13
- Solaris
  - adding devices online 12-14, 16-12
  - configuring new devices 12-12 to 12-41,?? to 15-43, 16-11 to 16-39
  - creating a mount directory 12-41, 12-42, 12-43, 12-44, 12-45, 12-46, 15-43, 15-44, 15-45, 15-46, 15-47, 15-48, 16-39, 16-40, 16-41, 16-42, 16-43, 16-44
  - creating a UFS filesystem 12-41, 12-43, 12-44, 15-43, 15-45, 15-47, 16-39, 16-41, 16-42
  - creating a VxFS filesystem 12-42, 12-43, 12-45, 15-44, 15-46, 15-48, 16-40, 16-41, 16-43
  - incorporating Symmetrix 12-9 to 15-11, 16-8 to 16-10
  - modifying the system specification file 12-10, 15-10, 16-9
  - operating system features 12-3, 15-4, 16-3
  - overview 12-2 to 12-3 to 15-4, 16-2 to 16-3
  - Solstice DiskSuite 12-15 to 12-39, 15-17 to 15-41, 16-13 to 16-37
  - Symmetrix configuration 12-6, 16-6
  - system and error messages 12-48, 16-46
  - utilities and functions 12-4, 16-4
  - Veritas Volume Manager 12-25 to 12-39, 15-27 to 15-41, 16-23 to 16-37
- Solaris x86
  - adding devices online 13-18
  - configuring new devices 13-13
  - incorporating Symmetrix 13-11
  - operating system features 13-6, 13-7, 13-8
  - overview 13-2
  - Symmetrix configuration 13-6, 13-9
  - system and error messages 13-19
  - utilities and functions 13-4
- Solstice DiskSuite 12-15 to 12-39, 15-17 to 15-41, 16-13 to 16-37
  - creating a concatenated volume 12-15, 15-17, 16-13
  - creating a concatenated volume for UFS logging 12-16, 15-18, 16-14
  - creating a mirrored striped volume 12-20, 15-22, 16-17
  - creating a mirrored striped volume for UFS logging 12-21, 15-23, 16-19
  - creating a striped volume 12-17, 15-19, 16-15
  - creating a striped volume for UFS logging 12-18, 15-20, 16-16
  - filesystem expansion 12-46, 15-49, 16-44
  - state database replicas 12-23, 15-25, 16-21
- Solstice DiskSuite commands 12-24, 15-26, 16-22
- space.c file
  - NCR 4-7
- space.c.emc file
  - NCR 4-7
- space.README file
  - NCR 4-7
- Split cable connections B-3
- Starting vi commands D-2
- State database replicas 12-23, 15-25, 16-21
- Stratus 17-1
- Stratus Continuum 400 18-2
- Striped plex 5-28
- Subdisks 5-27

- Supported hosts 1-3
- Symmetrix
  - options 1-13
- Symmetrix configuration
  - AIX 9-11
  - Digital UNIX 14-6
  - Fujitsu 15-7
  - HP/9000 Series 700 7-8
  - HP-UX 6-8, 17-7
  - NCR 4-6
  - Sequent 5-6
  - SGI/Challenge 10-5, 11-5
  - Solaris 12-6, 16-6
  - Solaris x86 13-6, 13-9
- Symmetrix software options 1-13
- sysadm
  - NCR 4-4, 4-11 to 4-14
- System messages
  - AIX 9-24 to 9-29
  - Digital UNIX 14-22
  - Fujitsu 15-50
  - HP-UX 6-28, 7-22, 17-19
  - NCR 4-19
  - Sequent 5-40 to 5-41
  - SGI/Challenge 10-25 to 10-26, 11-23 to 11-24
  - Solaris 12-48, 16-46
  - Solaris x86 13-19

## T

- Tag Command Queuing Enabled 11-8
- Target IDs
  - restrictions in Sequent environment 5-2, 5-7
  - restrictions in SGI/Challenge environment 11-2
- Test 20 B-5
- the 4-12
- Transferring device definition files A-5
- Tru64 14-1

## U

- uerf command 14-22
- uname
  - IRIX 10-4, 11-4
- UNIX performance tools 3-2 to 3-3
  - iostat 3-3
  - sar 3-2

- vmstat 3-3
- UnixWare 19-1
- use 4-17
- Using XLV 10-14 to 10-19, 11-12 to 11-14
- Utilities
  - AIX 9-5 to 9-9, 9-10
  - Digital UNIX 14-4
  - Fujitsu 15-5
  - HP/9000 Series 700 7-5
  - HP-UX 6-5 to 6-7, 7-5 to 7-7, 17-4 to 17-6
  - NCR 4-4 to 4-5
  - Sequent 5-5, 20-3
  - SGI/Challenge 10-3 to 10-4, ?? to 10-4, 11-3 to 11-4, ?? to 11-4
  - Solaris 12-4, 16-4
  - Solaris x86 13-4
  - SVM 5-33 to 5-39
- Utility B3 B-5

## V

- varyoffvg
  - AIX 9-8
- varyonvg
  - AIX 9-8
- Verifying Symmetrix devices
  - NCR 4-10
- Veritas Volume Manager 12-25 to 12-39, 15-27 to 15-41, 16-23 to 16-37
  - creating a striped mirrored volume 12-33, 15-36, 16-32
  - creating a striped volume 12-30, 15-32, 16-28
  - creating concatenated volumes 12-26, 15-28, 16-24
  - initializing disks 12-25, 15-27, 16-23
- vgchange
  - HP-UX 6-6, 7-6, 17-5
- vgcreate
  - HP-UX 6-6, 7-6, 17-5
- vgdisplay
  - HP-UX 6-6, 7-6, 17-5
- vgexport
  - HP-UX 6-6, 7-6, 17-5
- vgimport
  - HP-UX 6-6, 7-6, 17-6
- vgremove
  - HP-UX 6-6, 7-6, 17-6

vgscan  
     HP-UX 6-6, 7-6, 17-5  
 vi commands D-3  
     editing tasks D-2 to D-3  
     inserting text D-2  
     miscellaneous commands D-3  
     moving the cursor D-2  
     saving and exiting D-3  
     searching commands D-3  
     starting D-2  
 Viewing messages  
     using errpt 9-24  
     using SMIT 9-24, 9-26  
 vmstat 3-3  
 voldisk  
     Sequent Volume Manager 5-36  
 voledit  
     Sequent Volume Manager 5-36  
 volmake command 5-28  
 volmend  
     Sequent Volume Manager 5-37  
 volplex  
     Sequent Volume Manager 5-34  
 volprint  
     Sequent Volume Manager 5-39  
 volsd  
     Sequent Volume Manager 5-33  
 volstat  
     Sequent Volume Manager 5-38  
 volume  
     Sequent Volume Manager 5-35  
 Volume Group Descriptor Area (VGDA) 9-3  
 Volume Groups 9-3  
 vxassist 12-29, 12-33, 12-37, 15-31, 15-35, 15-40,  
     16-27, 16-31, 16-36  
 vxmake 12-26, 12-30, 12-34, 15-28, 15-32, 15-36,  
     16-24, 16-28, 16-32

## W

World wide web A-3

## X

XLV utility 10-14, 11-12

## Y

Y-cable connections B-4