Storage System

Product Release Notes for Version 9.16

EC1667-1-E2, Second Edition

Proprietary Rights Notice

This document contains proprietary information of LSI Logic Corporation and Sun Microsystems, Inc. The information contained herein is not to be used by or disclosed to third parties without the express written permission of an officer of LSI Logic or Sun. Any products described herein are licensed products of LSI Logic and Sun.

Document Description

Document EC1667-1-E2, Second Edition. April 2006.
This document will remain the official reference source for all revisions and releases of this product untill rescinded by an update.

Disclaimer

It is the policy of LSI Logic and Sun. to improve products as new technology, components, software, and firmware become available. We reserve the right to make changes to any products herein at any time without notice. All features, functions, and operations described herein may not be marketed in all parts of the world. In some instances, photographs and figures are of equipment prototypes. Therefore, before using this document, consult your sales representative or account team for information that is applicable and current. WE DO NOT ASSUME ANY RESPONSIBILITY OR LIABILITY FOR THE USE OF ANY PRODUCTS DESCRIBED HEREIN EXCEPT AS EXPRESSLY AGREED TO IN WRITING BY LSI LOGIC.

License Restriction

The purchase or use of an LSI Logic Corporation/Sun Microsystems, Inc. solution does not convey a license under any patent, copyright, trademark, or other intellectual property right of LSI Logic, Sun, or its third parties.

Copyright Notice

Trademark Acknowledgments

Engenio, the Engenio design, Sun, the Sun logo, SANtricity, HotScale, and SANshare are trademarks or registered trademarks of LSI Logic Corporation or Sun Microsystems, Inc. All other brand and product names may be trademarks of their respective companies.

Regulatory Compliance Statements

FCC Radio Frequency Interference Statement

This equipment has been tested and found to comply with the limits for a Class A digital device, pursuant to Part 15 of the Federal Communications Commission (FCC) Rules. These limits are designed to provide reasonable protection against harmful interference in a commercial installation. This equipment generates, uses, and can radiate radio frequency energy and, if not installed and used in accordance with the instructions, may cause harmful interference to radio communications. Operation of this equipment in a residential area is likely to cause harmful interference, in which case the user will be required to correct the interference at his/her own expense.

LSI Logic Corporation, and Sun Microsystems, Inc. are not responsible for any radio or television interference caused by unauthorized modification of this equipment or the substitution or attachment of connecting cables and equipment other than those specified by LSI Logic and Sun. It is the user's responsibility to correct interference caused by such unauthorized modification, substitution, or attachment.

Laser Products Statement

This equipment uses Small Form-factor Pluggable (SFP) optical transceivers, which are unmodified Class 1 laser products pursuant to 21 CFR, Subchapter J, Section 1040.10. All optical transceivers used with this product are required to be 21 CFR certified Class 1 laser products. For outside the USA, this equipment has been tested and found compliant with Class 1 laser product requirements contained in European Normalization standard EN 60825-1 1994+A11. Class 1 levels of laser radiation are not considered to be hazardous and are considered safe based upon current medical knowledge. This class includes all lasers or laser systems which cannot emit levels of optical radiation above the exposure limits for the eye under any exposure conditions inherent in the design of the laser products.

LSI Logic and Sun are not responsible for any damage or injury caused by unauthorized modification of this equipment or the substitution or attachment of connecting cables and equipment other than those specified by LSI Logic and Sun. It is the user's responsibility to correct interference caused by such unauthorized modification, substitution, or attachment.

This Class A digital apparatus meets all requirements of the Canadian Interference-Causing Equipment Regulations.

Cet appareil numérique de la classé A respecte toutes les exigences du Règlement sure le matèriel brouilleur du Canada.

この装置は、第一種情報装置（商工業地域において使用されるべき情報装置）で商工業地域での電波障害防止を目的とした情報処理装置等電波障害自主規制協議会（ＶＣＣＩ）基準に適合しております。

従って、住宅地域またはその隣接した地域で使用すると、ラジオ、テレビジョン受信機等に受信障害を与えることがあります。

取扱説明書に従って正しい取り扱いをして下さい。

警告使用者：　這是甲類的資訊產品，在居住的環境中使用時，可能會造成射頻干擾，在這種情況下，使用者會被要求採取某些適當的對策。

Contents

List of Tables

Chapter 1: ABOUT THIS RELEASE

Document Conventions

Definitions of Safety Notices

DANGER Indicates an imminently hazardous situation that will result in death or severe personal injury.

WARNING Indicates a potentially hazardous situation that could result in death or severe personal injury.

CAUTION Indicates a potentially hazardous situation that could result in moderate or minor personal injury.

Definitions of Informational Notices

CAUTION Indicates a potentially hazardous situation that could result in data loss (or other interruption) or equipment damage.

IMPORTANT Indicates information or criteria that is necessary to perform a procedure correctly.

NOTE Indicates a clarification of a concept or presents a maintenance tip.

Typographic Conventions

Italic indicates the title of documents, variables and placeholders in text, emphasized words, and new terms.

Bold indicates choices in procedures and other emphasized text.

`Monospace` indicates arguments, code examples, command-line text, command options, commands, directories, error messages, file names, folders, on-screen text, and user input.

`Monospace italic` indicates command variables in code, parameters, and placeholders and variables in code.

`Monospace bold` indicates keywords and values.

Revision Record

Edition or Revision	Date	Affected Pages or Remarks
First edition	December 2005	New document
Second edition	April 2006	Added information and restrictions for the next release of version 9.16 software and new hardware, including the 6140 storage system

Part Number: EC1667-1-E2

About This Document

The *Storage System Product Release Notes for Version 9.16* provide the following information:

- A list of any product restrictions that might be visible to the end user. If there are any workarounds for the restrictions, they are explained with the restriction itself.
- A list of usage notes that might be useful for an end user's environment.
- Specifications for third-party Fibre Channel hardware, including links to current firmware and drivers.
- Firmware configuration and support information.
- Last-minute update descriptions about the SANtricity® Storage Manager version 9.16 software (version 6.16 firmware) and associated hardware.
- Document set errata and corrections.

The information in this document covers all the operating systems that the software supports.

Intended Readers

This document is intended for end users, system operators, system administrators, and service personnel who are responsible for preparing for or maintaining an installation of one or more storage systems that will be managed using the SANtricity Storage Manager software. Readers must be familiar with computer system operation, maintenance, and repair. In addition, they should understand disk storage system, Redundant Array of Independent Disks (RAID), network, and Fibre Channel technologies.

Chapter 1

About This Release

This chapter provides release-specific information about this release of the SANtricity® Storage Manager version 9.16 (hereinafter referred to as the storage management software). The information includes a description of any software or firmware enhancements, and the features, operating systems, and hardware supported by this release of the storage management software.

What's New?

This section contains information about the new features and enhancements of the storage management software and the hardware that it supports.

NOTE The Enterprise Management Window and the Array Management Window each contain online help systems that provide context-sensitive procedural and conceptual information. For additional information about the features of the software, refer to the online help.

The following enhancements have been made in version 9.16 of the software:

- **Support for 6140 storage system –** A new storage system is now supported in this release.
- **Support for CSM200 drive module –** The control modules or storage systems can connect to either the CSM200 drive module or FLA300 drive module with this version of the software. See Table 1-9 on page 1-9 for more information about the drive modules.
- **Auto firmware synchronization for ESM** – Download firmware to one ESM and the storage management software will automatically synchronize the firmware to the other ESM.
- **Automatic tray ID conflict resolution** – The drive modules are now capable of resolving tray ID conflicts.
- **Improved module alarm support** – The software and hardware has been enhanced to allow alarms to be muted or disabled from SANtricity Storage Manager.
- **Mismatched ESM alert notification** – The Recovery Guru will now report when mismatched ESM cards are present in the controllers.

- **Improved diagnostics –** The storage management software now reports more Fibre Channel and loop diagnostic data.

Software Information

This section contains information about the operating systems supported by and the features included in this release of the storage management software.

Operating Systems

Table 1-1 lists the operating systems that have been tested for compatibility with all functions of the storage management software.

Table 1-1 Supported Operating Systems

Operating System	Version or Edition	Service Pack or Patch
Windows 2000[1]	Server Edition, Advanced Server Edition	Service Pack 4 (SP4)
Windows Server 2003	Standard Edition, Standard Server and Enterprise Server Editions (32-bit and 64-bit)	Service Pack 1 (SP1)
Solaris	8, 9, and 10	
HP-UX	11.11	PA-RISC
	11.23	PA-RISC and IA64
IRIX	6.5.27	
AIX	5.2 and 5.3	
NetWare	6.0	Service Pack 5 (SP5)
	6.5	Service Pack 4a (SP4a)
Red Hat Linux	Enterprise 3.*x* – AS, ES, WS (32- and 64-bit, 2.4 kernel)	Update 6
	Enterprise 4.*x* – AS, ES, WS (32- and 64-bit, 2.6 kernel)	Updates 2 and 3
SuSE Linux	Linux Enterprise Server 8.0 (32- and 64-bit, 2.4 kernel)	Service Pack 4+
	Linux Enterprise Server 9.0 (32- and 64-bit, 2.6 kernel)	Service Pack 3

[1] Windows XP is supported for the storage management station only.

Table 1-2 lists the operating systems that have been tested to be used with storage management stations only.

Table 1-2 Supported Operating Systems (Storage Management Station Only)

Operating System	Version or Edition	Service Pack or Patch
Windows XP	Professional	Service Pack 2 (SP2)
SuSE Linux	Professional 9.0 (32-bit only)	

Restrictions for 256 Volumes per Partition Support

Many hosts can support 256 volumes (0 to 255) per storage partition. However, the hosts listed in Table 1-3 have restrictions that result in less than 256 volumes supported per partition. For additional information about these restrictions and defining host ports, refer to the Array Management Window online help.

The maximum number of mappings differs because of any number of environmental factors. Operating system variables and failover driver issues influence the number of mappings available.

IMPORTANT If you try to map a volume to a logical unit number (LUN) that exceeds the restriction on these operating systems, the host is unable to access the volume.

Table 1-3 Maximum Volumes per Partition by Operating System (OS)

Operating System	Maximum Volumes per Partition	Highest LUN That the OS Looks For
Windows 2000	64[1]	255
Windows Server 2003	255	255
IRIX	255	~1000[2]
Solaris	32 or 255[3]	255[4]
AIX	N/A	31/1000+[5]
NetWare 6.0 and 6.5	255	255

Table 1-3 Maximum Volumes per Partition by Operating System (OS) (continued)

Operating System	Maximum Volumes per Partition	Highest LUN That the OS Looks For
HP-UX 11.11 and 11.23	127	127
Linux	255	255

[1] This is an OS registry problem and is restricted to 64 volumes. This is not enforced with any firmware mechanism.

[2] Operating system installation option.

[3] Maximum is 32 when using the Redundant Dual Active Controller (RDAC) driver and 255 when using VERITAS Dynamic Multi-Pathing (DMP) driver or MPXIO.

[4] The number of volumes supported is driver dependent.

[5] The maximum number of volumes is 32 when using the RDAC driver. The actual operating system limit is much higher.

Features

This section describes the features and functionality supported in this version of the storage management software. For more specific information about any of the supported features, refer to the *SANtricity Storage Manager Concepts Guide* or the appropriate feature guide.

Premium Features

This release of the storage management software supports the following premium features:

- **Volume Copy** – The Volume Copy premium feature copies data from one volume (source volume) to another volume (target volume) in a single storage system. For additional information, refer to the *SANtricity Storage Manager Volume Copy Feature Guide* or the Array Management Window online help.
- **Remote Volume Mirroring (RVM)** – The RVM premium feature is used for online, real-time data replication between storage systems over a remote distance. For additional information, refer to the *SANtricity Storage Manager Remote Volume Mirroring Feature Guide* or the Array Management Window online help.
- **SANshare® Storage Partitioning** – The SANshare Storage Partitioning premium feature allows hosts with different operating systems (heterogeneous hosts) access to the same storage system. For additional information, refer to the *SANtricity Storage Manager Concepts Guide* or the Array Management Window online help.
- **Snapshot Copy** – The Snapshot Copy premium feature creates a point-in-time image of a volume for backup and data recovery. For additional information, refer to the *SANtricity Storage Manager Concepts Guide* or the Array Management Window online help.

For additional information about hardware support of the premium features, see "Supported Premium Features by Control Module or Storage System" on page 1-8.

Standard Features

Dynamic Volume Expansion (DVE) is a modification operation used to increase the capacity of standard or snapshot repository volumes. The increase in capacity can be achieved by using any free capacity that is available on the volume group of the standard or snapshot repository volume.

Increasing the capacity of a standard volume is only supported on the following operating systems:

- Linux
- NetWare
- Windows 2000 Dynamic Disks
- Windows Server 2003 Dynamic Disks (depending on disk partition format used)

If you increase the volume capacity on a host operating system that is unsupported, the expanded capacity is unusable, and you cannot restore the original volume capacity.

Firmware and NVSRAM

Information about controller firmware, environmental services monitor (ESM) firmware, and nonvolatile static random access memory (NVSRAM) is available in Chapter 2, "Firmware Information."

Failover Protection

The storage management software supports several types of failover protection using multi-path drivers. For an overview of the default failover drivers and settings by operating system, see Table 1-4 on page 1-6.

IMPORTANT If you are using SANtricity Storage Manager RDAC version 09.10.00.00 on Solaris, ensure you have installed the patch. Please see the readme file in the rdac directory for installation instructions. See "Failover Software Usage Notes" in Chapter 6 for details about this problem.

Table 1-4 Default Failover Settings by Operating System

Operating System	Multi-Path Driver	Default Failover Setting	
		AVT[1] Disabled	AVT Enabled
AIX	VERITAS DMP	Default	
HP-UX	HP Logical Volume Manager (LVM)		Default
	VERITAS DMP		Default
IRIX	Failover support not available	Default	
Linux	SANtricity Storage Manager RDAC	Default	
NetWare	Netware native MPE failover driver		Default
Solaris 8 and 9	SANtricity Storage Manager RDAC	Default	
Solaris 8, 9, and 10	VERITAS DMP (Volume Manager 4.1)		Default
	MPxIO[2]	Default	
Windows 2000	SANtricity Storage Manager RDAC	Default	
Windows Server 2003	SANtricity Storage Manager RDAC	Default	
	VERITAS DMP (Volume Manager 4.1)		Default
	Microsoft MPIO (Storport)	Default	

[1] Auto Volume Transfer

[2] See Chapter 5, “Restrictions” for restriction information about using this failover solution.

Hardware Information

This section contains information about the supported hardware and any functionality or configurations specific to this release of the storage management software.

Control Modules and Storage Systems

This section describes the features and functionality of the supported hardware that is control module dependent. For more specific information about any of the supported control modules or storage systems, refer to the appropriate user guide for the control module or the storage system.

Supported Control Modules and Storage Systems

Table 1-5 lists the control modules and storage systems that are supported for this release and the type of controller supported for each model.

Table 1-5 Supported Control Modules and Storage Systems

Control Module or Storage System	Controller Type
FLX240 Drive Limited storage system	2880
FLX240 storage system	2882
FLX280 control module	5884
6540 control module	6998
6140 storage system	3994

Supported Drives, Volumes, and Volume Groups

Table 1-6 shows the maximum number of physical drives and total volumes that are supported by the control modules and storage systems. For more information on supported drive modules, see “Drive Modules” on page 1-9.

Table 1-6 Maximum Drives and Volumes Supported by the Control Modules and Storage Systems

Control Module or Storage System (Controller Type)	Maximum Drives per Storage System	Maximum Number of Volumes/ Mirrors with RVM Enabled	Maximum Volumes per Storage System[1]	Maximum Snapshots Supported	Maximum Volume Copy Requests
FLX240 Drive Limited (2880)	14	Not supported	1024	Not supported	Not supported
FLX240 (2882)	112	1024/32	1024	512	1023
FLX280 (5884)	224	2048/64	2048	1024	2047
6540 (6998)	224	2048/64	2048	1024	2047
6140 (3994)	112	1024/32	1024	512	1023

[1] Snapshot and Remote Volume Mirroring repository volumes are included in the total number of volumes supported.

Supported Premium Features by Control Module or Storage System

Table 1-7 shows the premium features supported by each control module or storage system and the total supported volumes for each feature.

Table 1-7 Premium Features Supported by the Control Modules and Storage Systems

	RVM	Snapshot		Volume Copy	
Control Module or Storage System (Controller Type)	**Maximum Mirror Pairs**	**Maximum Total Snapshots**	**Maximum per Base RAID Volume**	**Maximum Requests per Source Volume**	**Maximum Copy In-progress Volumes**
FLX240 Drive Limited (2880)	Not supported	512	4	Not supported	Not supported
FLX240 (2882)	32	512	4	1023	8
FLX280 (5884)	64	1024	4	2047	8
6540 (6998)	64	1024	4	2047	8
6140 (3994)	32	512	4	2046	8

Supported SANshare Storage Partitioning by Control Module or Storage System

Table 1-8 lists information about SANshare Storage Partitioning support.

Table 1-8 Supported SANshare Storage Partitioning and Host Ports

Control Module (Controller Type)	Maximum Defined Host Ports (HPs)	Maximum Hosts	Maximum Total Partitions	Maximum Volumes per Partition[2]
FLX240 Drive Limited (2880)	256	Host ports[1]	8	256
FLX240 (2882)	256	Host ports[1]	64	256
FLX280 (5884)	512	Host ports[1]	64	256
6540 (6998)	512	Host ports[1]	64	256
6140 (3994)	256	Host ports[1]	64	256

[1] The maximum number of hosts is the same as the number of host ports that are currently defined, up to the maximum number of defined host ports.

[2] Some operating systems do not support the maximum of 256 volumes per partition. For information on these restrictions, see "Restrictions for 256 Volumes per Partition Support" on page 1-3.

Drive Modules

This section describes the features and functionality of the supported hardware that are drive module dependent. For more specific information about any of the supported drive modules, refer to the appropriate user guide for the drive module.

Supported Drive Modules

IMPORTANT Only the drive modules listed are compatible with this version of the controller firmware. Downloading this version of controller firmware to an older system not listed here will result in the inability to manage the system using the storage management software.

Table 1-9 lists the drive modules that are supported for this release and specifications for each model.

Table 1-9 Supported Drive Modules

Drive Module	Drives per Module	Drive Technology
CSM200 (SBOD[1])	16	2-Gb/s Fibre Channel
FLA300 (SBOD)	14	2-Gb/s Fibre Channel

[1] Switched Bunch of Disks

Supported Mixed Drive Types

Table 1-10 lists the mixed drive types per drive module for this release. For the drive technology in each drive module, see Table 1-9.

IMPORTANT If your storage system includes model FLA300 drive modules with other models of drive modules, group the FLA300 drive modules together within a drive loop. Avoid any configuration in which two FLA300 drive modules are separated along the loop by a drive module that is of a different model.

Table 1-10 Supported Mixed-Drive Types

Drive Module	Supported Speeds	Mix Drive Types within a Drive Module?[1]	Mix Drive Types in a Volume Group?	Mix Drive Modules Behind a Controller Pair?[2]
CSM200	2-Gb/s only	No	No	Yes
FLA300	2-Gb/s only	No	No	Yes

[1] Additionally, global hot spares must be of the same drive type as the drives they are protecting.

[2] Drive modules must all be at the same speed behind the controller pair.

Host Adapters and Other Hardware

Information about host adapters and Fibre Channel switches is available in Chapter 3, “Fibre Channel Hardware Specifications.”

Chapter 2

Firmware Information

This chapter describes the firmware packages on the firmware installation CD. These files are needed to properly use the storage management software and all of its features.

NVSRAM Files

CAUTION **Potential loss of data –** Do not download nonvolatile static random access memory (NVSRAM) before contacting Customer and Technical Support. Do not download the NVSRAM packages unless instructed to do so in the SANtricity Storage Manager installation guide for your operating system or by Customer and Technical Support. Inappropriate application of these files could cause serious problems with your storage system.

Located in the `nvsram` directory, the NVSRAM downloadable packages specify default settings for the storage system controllers. Install these files on the host system, and download them to the storage system controllers as necessary. For instructions on when and how to download these files, refer to the SANtricity Storage Manager installation guide for your operating system.

Controller Firmware Files

Located in the `firmware` directory, the controller firmware files are installed onto the host system for use on the storage system controllers. For instructions on when and how to download these files to the storage system controllers, refer to the SANtricity Storage Manager installation guide for your operating system. A description of the firmware packages has been included under each specific firmware file name.

Environmental Services Monitor Firmware Files

Located in the `esm` directory, the environmental services monitor (ESM) firmware files are copied to the host system for use on the ESMs, which are located on the back of the drive modules. For instructions on when and how to download these files to the ESMs, refer to the Array Management Window online help. A description of the firmware packages has been included under each specific firmware file name.

Chapter 3

Fibre Channel Hardware Specifications

This chapter provides information about the available Fibre Channel equipment that is part of the tested hardware solution for this version of the storage management software. Included in this information are Internet links for the supported host adapters, switches, tested driver and firmware levels, and specific hardware restrictions.

More information on host adapters and switches is available through the following sources:

- Any documentation that came with a particular host adapter or switch
- Vendor and manufacturer support web sites
- Customer and Technical Support

Fibre Channel HBAs

NOTE The SANtricity Storage Manager Installation CD does not contain any configuration files or drivers for the host bus adapters (HBAs) tested with this version of the storage management software. Current Internet download sites have been provided to assist in obtaining these files.

Information about supported HBAs for this release of the storage management software can be obtained from the Certified Compatibility Matrix, which can be found at the following Internet location:

http://www.engenio.com/html/partners/compatible_matrix/interop.asp

Driver Information

The Certified Compatibility Matrix contains information about the files needed to support the Fibre Channel HBAs. Use the following Internet locations to obtain the drivers listed in the Certified Compatibility Matrix.

QLogic Drivers

Drivers for QLogic HBAs can be obtained from the following Internet location:

http://www.qlogic.com/support/oem_detail_all.asp?oemid=164

NOTE Internet Uniform Resource Locators (URLs) on the QLogic Web site are case sensitive.

To configure a QLogic HBA, see one of the following sections:

- "Configuring QLogic HBA BIOS Settings" on page 3-2
- "Configuring QLogic HBAs for the NetWare Operating System" on page 3-3
- "Configuring QLogic HBAs for the Windows Operating System" on page 3-6

LSI Logic Drivers

Drivers for LSI Logic HBAs can be obtained from the following Internet location:

http://www.lsilogic.com/downloads/selectDownload.do

To configure an LSI Logic HBA, see one of the following sections:

- "Configuring LSI Logic HBAs for the Solaris Operating System" on page 3-8
- "Configuring LSI Logic HBAs for the Windows Operating System" on page 3-12
- "Configuring LSI Logic HBA BIOS Settings" on page 3-13

QLogic HBA Configuration Information

Configure the HBA according to the information provided in these procedures.

Configuring QLogic HBA BIOS Settings

The HBA BIOS settings allow for changes to the HBA configuration. You can change the values through the BIOS interface at boot time. Use the following procedure to ensure proper operation of the QLogic HBA with different operating systems.

1. Press Alt-Q upon booting to gain access to QLogic BIOS utility.
2. Select the HBA to change.
3. Select Configuration Settings.

4 Select Host Bus Adapter Settings, and verify the following settings:

Frame Size = 2048
Loop Reset Delay = 8
Adapter Hard Loop ID = Enabled
Hard Loop ID = * (assign a unique value)

IMPORTANT In an Arbitrated Loop environment, it is always recommended to assign unique Arbitrated Loop-Physical Addresses (AL_PAs) to the devices attached. The QLogic HBA is automatically assigned an AL_PA
of 125 decimals during installation.

---

5 Select Advanced Adapter Settings, and verify the following settings:

Execution Throttle = 255 (cannot exceed I/O Control Block [IOCB] allocation)
LUNs per Target = 0 (“0” activates maximum logical unit number [LUN] support)
Enable LIP Reset = NO
Enable LIP Full Login = YES
Enable Target Reset = YES
Login Retry Count = 30
Port Down Retry Count = 70
IOCB Allocation = 256

Exceptions based on operating system include the following:

- In the NetWare environment, set LUNs per Target = 32.
- In the Linux environment, set Port Down Retry Count = 12.

6 Repeat steps 2 through step 5 on page 3-3 for other QLogic HBAs.

7 Reboot the system.

End Of Procedure

Configuring QLogic HBAs for the NetWare Operating System

Use the following procedure to configure your QLogic HBAs for use with the NetWare operating system. Perform this procedure after installing the HBA.

1 Set the BIOS tunable parameters according to the procedures in “Configuring QLogic HBA BIOS Settings” on page 3-2.

2 Load the NetWare server.

3 Ensure that the HBA initialization line in the `sys:\system\STARTUP.NCF` file contains all of the following settings:

```
LOAD driver-name SLOT=x /LUNS /XRETRY=600
```

where the *driver-name* (`ql2200.ham` or `ql2300.ham`) and the SLOT=*x* settings are correct for your HBA. If the host system contains additional HBAs, enter a copy of this line into the `STARTUP.NCF` file for each HBA.

4 Using the SANtricity Storage Manager, set the default host type to NWRFO to enable the NetWare failover NVSRAM settings.

5 Reboot the controllers.

6 Connect the controller to the host.

7 Reboot the NetWare host.

NOTE The QLogic configuration utility for the NetWare operating system can be obtained from the following Internet location:

http://www.qlogic.com/lsi/2200/NetWare/cfg.zip

---

8 Set up load balancing by using the executable file `CFG.NLM`.

a Copy `CFG.NLM` to `sys:\system` on the server.

IMPORTANT The load-balancing settings stored in the `QL2X00.CFG` file are valid until a configuration change occurs that involves adding or replacing controllers or storage systems. If the old `QL2X00.CFG` file is present when you are attempting to add new hardware, the devices are not detected.

When you are connecting to a new storage system, it is suggested that the `QL2X00.CFG` file be deleted, the devices added using the `hot_add` utility, and the following steps be used to re-create the `QL2X00.CFG` file and set up load balancing. Also note that if the storage system has never been connected to the server, a *restart server* might be necessary for the devices to be detected, rather than using the `hot_add` utility.

b If the `c:\nwserver\QL2X00.CFG` file exists, delete it.

c Type the following command, and press Enter:

```
cfg /I
```

d Write down the node name for each HBA.

e Type the following command, and press Enter:

```
cfg /s
```

f Verify that each HBA is listed, followed by the controllers to which the HBA is attached.

NOTE You can obtain controller node and port names by selecting Storage System >> Profile in the storage management software.

g To create a new `c:\nwserver\QL2X00.CFG` file, type the following command, and press Enter:

```
cfg /fs
```

You can also use this command to edit the `QL2X00.CFG` file.

h Edit the `QL2X00.CFG` file similar to the following:

IMPORTANT Two typical path masks are provided in the following example. The path mask that starts with the numeral 5 designates access to the odd volumes through the listed controller and HBA. The path mask that starts with the letter A designates access to the even volumes through the listed controller and HBA. For more information, refer to the `cfg1.txt` and `cfg2.txt` readme files that accompany the `cfg.nlm` executable.

```
2
2
210000E08B01613A (1st host bus adapter world wide name)
3  (number of devices)
200400A0B809368E (controller world wide node name)
200500A0B809368F (controller world wide port name)
5555555500000000 0000000000000000 0000000000000000
   0000000000000000 (path mask)
0000000000000000 0000000000000000 0000000000000000
   0000000000000000 (lun mask)
0
200400A0B807B16E
200500A0B807B16F
5555555500000000 0000000000000000 0000000000000000
   0000000000000000
0000000000000000 0000000000000000 0000000000000000
   0000000000000000
0
200400A0B80B1DE1
200500A0B80B1DE2
5555555500000000 0000000000000000 0000000000000000
   0000000000000000
0000000000000000 0000000000000000 0000000000000000
   0000000000000000
0
210000E08B00C8BA (2nd host bus adapter wwn)
3
200400A0B807B16E
200400A0B807B16F
```

```
AAAAAAAA00000000 0000000000000000 0000000000000000
   0000000000000000
0000000000000000 0000000000000000 0000000000000000
   0000000000000000
0
200400A0B809368E
200400A0B809368F
AAAAAAAA00000000 0000000000000000 0000000000000000
   0000000000000000
0000000000000000 0000000000000000 0000000000000000
   0000000000000000
0
200400A0B80B1DE1
200400A0B80B1DE2
AAAAAAAA00000000 0000000000000000 0000000000000000
   0000000000000000
0000000000000000 0000000000000000 0000000000000000
   0000000000000000
0
```

NOTE A volume does not need to be present to define the path mask for that volume. Therefore, if the path masks for all possible volumes (such as 32 or 128) are defined in step h on page 3-5, reconfiguration of the `QL2X00.CFG` file is not necessary after any volume deletion or creation changes (provided, of course, that the preferred path specified is still the desired preferred path for the new volume).

i To load the new settings, type the following command, and press Enter:

```
cfg /fl
```

End Of Procedure

Configuring QLogic HBAs for the Windows Operating System

The HBA BIOS and Windows Registry settings allow for minor adjustments to the HBA configuration. The BIOS values can be changed through the BIOS interface at boot time. The Windows Registry settings are automatically modified when the registry configuration file is executed. Use the following procedure to ensure proper operation of the QLogic HBA with the Windows 2000 or Windows Server 2003 operating system.

1 Press Alt-Q upon booting to gain access to the QLogic BIOS utility.

2 Select the HBA you want to modify.

3 Select Configuration Settings.

4 Select Host Bus Adapter Settings, and verify the following settings:

Frame Size = 2048
Loop Reset Delay = 8
Adapter Hard Loop ID = Enabled
Hard Loop ID = * (assign a unique value)

IMPORTANT In an Arbitrated Loop environment, it is always recommended to assign unique Arbitrated Loop-Physical Addresses (AL_PAs) to the devices attached. The QLogic HBA is automatically assigned an AL_PA
of 125 decimals during installation.

---

5 Select Advanced Adapter Settings, and verify the following:

Execution Throttle = 255 (cannot exceed IOCB allocation)
LUNs per Target = 0 (‘0’ activates maximum LUN support)
Enable LIP Reset = NO
Enable LIP Full Login = YES
Enable Target Reset = YES
Login Retry Count = 30
Port Down Retry Count = 70
IOCB Allocation = 256

6 Repeat steps 2 through step 5 on page 3-7 for all remaining QLogic HBAs.

7 Save any changes, and exit the QLogic BIOS utility.

8 Reboot the system.

9 When the host is finished booting, insert the SANtricity Storage Manager Installation CD.

10 Using Windows Explorer, navigate to the `\hostadapter` directory.

11 Double-click one of the following files based on the operating system of the host:

- **Windows NT** – `ql23xx_nt.reg`
- **Windows 2000 or Windows Server 2003** – `ql23xx_w2k-w2k3.reg`

The Windows Registry is automatically updated with the proper settings for the QLogic HBA.

12 For the changes to take effect, reboot the host.

End Of Procedure

LSI Logic HBA Configuration Information

Configure the HBA according to the information provided in these procedures.

Configuring LSI Logic HBAs for the Solaris Operating System

After installation of the LSI Logic HBA driver, the following message is displayed:

```
Entries added. For support of more than 15 targets or nonzero
LUNs it may be necessary to edit /kernel/drv/ssd.conf to add
additional entries.
```

To successfully connect the LSI Logic HBA to the storage system, complete these configuration steps after the message is displayed.

1. Assign target identification numbers to the controllers on the storage system. Go to "Binding at the Software Level" on page 3-8.
2. Configure the driver to probe for multiple LUNs on the assigned targets. Go to "Editing the `ssd.conf` File" on page 3-10.
3. If you are configuring the Solaris operating system to boot from a storage system, go to "Root Boot Support" on page 3-11, and set the persistent bindings at the firmware level.

End Of Procedure

Binding at the Software Level

Specific instructions about this procedure are detailed in the `/kernel/drv/itmpt.conf` file. Use the following reference material to ensure proper configuration.

Persistent Bindings of WWN to Target

To persistently map a Fibre Channel World Wide Name (WWN) to a given target, use the following syntax:

`target-`*`X`*`-wwn="`*`portWWN`*`"`

For example:

```
target-4-wwn="2200002037102d0f"
```

This persistently maps the WWN *2200002037102d0f* to target *4* on all *itmpt* buses.

Restrict Mapping to a Single Bus

To restrict the mapping to a single bus, use the following syntax:

`hba-`*`X`*`-target-`*`Y`*`-wwn="`*`portWWN`*`"`

For example:

```
hba-1-target-4-wwn="2200002037102d0f"
```

This persistently maps the WWN *2200002037102d0f* to target *4* on bus *itmpt1* only.

General Information on WWNs

It is important to use the World Wide Port (WWP) for the device you want to map, rather than the WWN, or only part of the WWN. The name displayed in the boot log or displayed by the `probe-scsi-all` command at the OpenBoot programmable read-only memory (PROM) `ok` prompt is the correct name to use (such as the WWP).

A node WWN for devices generally looks something like the following example.

```
2000002037102d0f
```

If the device is connected to the Fibre Channel bus using port A, the WWP for that device would be similar to the following example.

```
2100002037102d0f
```

If the device is connected to the Fibre Channel bus using port B, the WWP for that device would be similar to the following example.

```
2103002037102d0f
```

Many times, the WWN printed on the physical device is only part of the full WWN. For example, for the WWN used in the preceding example, the WWN listed on the disk case itself is missing the first four characters.

```
002037102d0f
```

Editing the `ssd.conf` File

When installing the LSI Logic HBA driver, the installation program puts entries into the `ssd.conf` file for logical unit number (LUN) 0 only.

To successfully detect all available devices on the storage system, add entries for LUNs 0 through 31 by using the following procedure.

1 Open the `ssd.conf` file in an editor, and find the following section:

```
name="ssd" parent="itmpt" target=0;
name="ssd" parent="itmpt" target=1;
name="ssd" parent="itmpt" target=2;
name="ssd" parent="itmpt" target=3;
name="ssd" parent="itmpt" target=4;
      ******
name="ssd" parent="itmpt" target=15;
```

2 Locate the entry that references the target identification number that you used to bind to one of the storage system controllers. Edit that line to contain LUN=0, for example:

```
name="ssd" parent="itmpt" target=7 lun=0;
```

3 Following that line, add entries for LUNs 1 through 31 so that the driver probes each possible LUN on that target.

NOTE This example list is truncated. When editing the `ssd.conf` file, enter each line from LUN=1 to LUN=31.

```
name="ssd" parent="itmpt" target=7 lun=1;
name="ssd" parent="itmpt" target=7 lun=2;
name="ssd" parent="itmpt" target=7 lun=3;
name="ssd" parent="itmpt" target=7 lun=4;
name="ssd" parent="itmpt" target=7 lun=5;
name="ssd" parent="itmpt" target=7 lun=6;
name="ssd" parent="itmpt" target=7 lun=7;
      ******
name="ssd" parent="itmpt" target=7 lun=29;
name="ssd" parent="itmpt" target=7 lun=30;
name="ssd" parent="itmpt" target=7 lun=31;
```

4 Repeat this procedure for each target identification number that is associated with a storage system controller.

End Of Procedure

Root Boot Support

The LSI Logic Fcode BIOS that resides on the HBA enables the system administrator to configure a target device that will become the root boot device. The HBA keeps this information persistently in nonvolatile memory. When the root boot capability is required by the system, you might need to configure the LSI Logic Fibre Channel HBA to allow the target device to be locked to a persistent unit number.

The number of devices attached to the HBA determines the need to configure a persistent unit number as follows:

- If only one target device is attached to the LSI Logic HBA, do *not* configure a persistent unit number with the LSI Logic HBA.
- If more than one target device is attached to the LSI Logic HBA, configure *only* one of the target devices as a persistent unit number with the LSI Logic HBA.

NOTE Follow these guidelines to ensure that the root boot feature functions as expected.

The following is an example of the user interaction with the Fcode BIOS to select target devices to be mapped in the persistent device table.

All configuration command requests are specified from the `ok` prompt in the Fcode BIOS. The command to be entered is shown in **`bold`** type.

EXAMPLE

1 Select the Fibre Channel HBA to configure by typing the bold entries at the prompts specified.

```
ok show-disks
   a) /pci@1f,0/pci@1/IntraServer,fc@2/disk
   b) /pci@1f,0/pci@1/IntraServer,Ultra2-scsi@1/disk
   c) /pci@1f,0/pci@1,1/ide@3/cdrom
   d) /pci@1f,0/pci@1,1/ide@3/disk
   e) /pci@1f,0/pci@1,1/ebus@1/fdthree@14,3203f0
   q) NO SELECTION

   Enter Selection, q to quit: a

   /pci@1f,0/pci@1/IntraServer,fc@2/disk has been selected.
```

2 Type ^Y (Ctrl-Y) to insert the information in the command line. For example, to create devalias mydev for `/pci@1f,0/pci@1/IntraServer,fc@2/disk`, type the following commands.

```
ok nvalias mydev ^Y
```

```
ok select /pci@1f,0/pci@1/IntraServer,fc@2
ok show-children

   MPT Firmware Version 1.01
      Target 0
       Unit 0   Disk     SEAGATE ST39173FC       6615
       WWN 200000203710c4e8  PortID a3

ok set-persistent
```

3 Enter the commands without parameters to print the following help information:

usage is *current-target-ID persistent-target-ID* set-persistent

```
ok 0 0 set-persistent
ok show-persistent

   Entry 1  WWN 200000203710c4e8  Target 0

ok
```

4 To clear an entry in the persistent device map, use the `clear-persistent` command.

```
ok 1 clear-persistent
ok show-persistent
ok
```

Entry 1 has been deleted from the table, and the table is now empty, so there is no response.

End Of Procedure

Configuring LSI Logic HBAs for the Windows Operating System

The HBA BIOS and Windows Registry settings allow for minor adjustments to the HBA configuration. You can change the BIOS values through the BIOS interface at boot time. The Windows Registry settings are automatically modified when the registry configuration file is executed. Use the following procedure to ensure proper operation of the LSI Logic HBA with the Windows NT, Windows 2000, or Windows Server 2003 operating systems.

1 Press Ctrl-C during boot to gain access to the LSI Logic BIOS utility.

2 Select the HBA port you want to modify.

3 Change Multi-pathing to YES.

4 Press Esc to return to the list of ports.

5 Repeat steps 2 through 4 for all remaining HBA ports.

6 Select Save Changes Then Exit This Menu.

The host automatically reboots upon exiting the LSI Logic BIOS utility.

7 When the host is finished booting, insert the SANtricity Storage Manager Installation CD into the CD-ROM drive.

8 Using Windows Explorer, navigate to the `/hostadapter` directory.

9 Based on the host's operating system, double-click one of the following files:

- **Windows NT** – `lsi_nt.reg`
- **Windows 2000 or Windows Server 2003** – `lsi_w2k-w2k3.reg`

The Windows Registry is automatically updated with the proper settings for the LSI Logic HBA.

10 For the changes to take effect, reboot the host.

End Of Procedure

Configuring LSI Logic HBA BIOS Settings

The HBA BIOS settings allow for minor adjustments to the HBA configuration. You can change these values through the BIOS interface at boot time. Use the following procedure to ensure proper operation of the LSI Logic HBA with different operating systems.

1 Press Ctrl-C during boot to gain access to the LSI Logic BIOS utility.

2 Select the HBA port you want to modify.

3 Change Multi-pathing to YES.

4 Press Esc to return to the list of ports.

5 Repeat steps 2 through 4 for all remaining HBA ports.

6 Select Save Changes Then Exit This Menu.

The host automatically reboots upon exiting the LSI Logic BIOS utility.

End Of Procedure

Fabric Switches

You can obtain information about supported fabric switches from the Certified Compatibility Matrix, which can be found at the following Internet location:

http://www.engenio.com/partners/certified_compatible.html

Chapter 4

User Documentation Updates

This chapter provides corrections and updates to the user documentation. Information in this chapter is presented as it would be displayed or printed in its appropriate location; therefore, references to other material in the original document might seem out of context within this Product Release Notes document.

IMPORTANT Always refer to this chapter before installing or modifying a storage system. This chapter contains documentation corrections and additions, and important last-minute information that was not available when the documentation set was released.

Document Set Corrections

The following sections provide corrections and updates to the user documentation.

FLA300 Drive Module User Guide

The following correction pertains to the *FLA300 Drive Module User Guide*.

Correction: Environmental Services Monitors (ESMs)

The statement in the opening paragraph on page 1-6 should read “The ESMs support both fiber-optic and copper-fiber interfaces for the drive channels.”

Chapter 5

Restrictions

IMPORTANT Always refer to this chapter before installing or modifying a storage system. This chapter contains important, late-breaking information that was not available when the documentation set was released.

This chapter provides two types of information:

- Restrictions that apply to using hardware associated with this release, including storage systems and host bus adapters. See the following sections for hardware restrictions:
  - “Storage System Restrictions” on page 5-2
  - “Host Adapter Restrictions” on page 5-9
  - “General Hardware Restrictions” on page 5-19
- Restrictions that apply to using the software associated with this release, including the storage management software, failover software, and utility software. See the following sections for software restrictions:
  - “SANtricity Storage Manager Restrictions” on page 5-21
  - “Failover Software Restrictions” on page 5-33
  - “Utilities Software Restrictions” on page 5-47
  - “General Software Restrictions” on page 5-52

Information about which operating systems might be affected is kept with each applicable restriction. After installing the software, see Chapter 6, “Usage Notes” for usage information that might be useful for your hardware and software setup.

Storage System Restrictions

CAUTION **Risk of permanent damage to the control module –** The controllers in the control modules are not interchangeable. Installing the wrong controller permanently damages the control module and renders the controller inoperable.

The restrictions in this section apply specifically to the hardware components in storage systems, including modules, CRUs, hard drives, and so on.

Host Cannot Access the 6540 Control Module after the Cables are Disconnected and Moved

Operating System
AIX

Problem or Restriction
If cables are disconnected from a host port and then reconnected to a different host port, the AIX host will not be able to access the 6540 control module. The AIX operating system requires that the controller retain World Wide Name (WWN) information after a logout—which includes a disconnected cable. This condition can cause two ports to present the same port ID information to the host if the cables are pulled and moved to a different host port.

Workaround
Reboot the controllers after the cables are disconnected and moved to different host ports.

Control Module Powered up Without Drive Modules Might Not Recognize Drives or Modules When They Are Attached

Problem or Restriction
A control module is powered up without any drive modules attached. Then one or more drive modules are connected to the control module. The control module might not recognize the drives or drive modules correctly. This condition results in error messages or the inability to do operations, such as creating volumes or initializing volume groups from the storage management software.

Workaround
Make sure that the drive modules are connected to the control module when you turn on the control module’s power.

Drives Fail to Spin up if Inserted While the Storage System Reboots

Problem or Restriction

Removing the drives while a storage system is online and then waiting to reinsert the drives until the storage system is starting after a reboot might cause the drives to be marked as failed after the storage system comes back online.

Workaround

Wait until the storage system is back online before reinserting the drives. If the storage system still does not recognize the drives, reconstruct the drives by using the storage management software.

Drive Firmware Download Fails on Seagate SATA Drives

Problem or Restriction

If the firmware file is larger than 256 KB, the drive firmware fails to download to Seagate SATA drives.

Workaround

None. Firmware files must be smaller than 256 KB for the drive firmware download to complete.

Volumes Transfer During Host Reboot or LUNs Formatting

Operating System
Solaris with VERITAS Dynamic Multi-Pathing (DMP)

Problem or Restriction
After creating a maximum number of volumes that include remote volume mirroring (RVM) and Snapshot, volume transfer occurs when two Solaris hosts, running with VERITAS Volume Manager 4.0, are rebooted.

Workaround
Because SMagent is responsible for the volume transfer, perform one of the following actions:

- Unmap the Access Volume Logical Unit Number (LUN) in the mapping view (the Access Volume is usually LUN 7), and then reboot the hosts.
- (Recommended method): After each host reboots, type the following command to stop the SMagent:

```
# SMagent stop
```

Parallel Drive Firmware Download to Mixed Drive Types Causes the Drives to Fail

Problem or Restriction
Downloading drive firmware to mixed drive types causes drives to fail.

Workaround
Download to only one type of drive at a time. For example, if you have Fibre Channel and SATA drives in your storage system, download firmware to only Fibre Channel drives or SATA drives during the parallel drive firmware download.

Two In-Band Storage Devices are Displayed in the Enterprise Management Window as a Single Node

Operating System
Linux Red Hat

Problem or Restriction
When two hosts are connected to the same storage system, both of the attached hosts display in the Enterprise Management Window as a single node, showing the IP addresses of both hosts rather than displaying them as separate nodes. As a single node, it is not possible to remove just one of the hosts.

Workaround

If the host has a static IP address, then modify the `/etc/hosts` file to map the host name to the host address. For more information, refer to http://bugs.sun.com/bugdatabase/view_bug.do?bug_id=4665037.

Additional Drive Modules Cannot be Added during an ESM Firmware Download

Problem or Restriction

If you add an additional drive module by using the loop topology option 2 during ESM firmware download, the ESM firmware download process might fail due to a disconnected loop. The drive module would come up correctly after being added to the loop.

Workaround

When adding the additional drive module, do not follow the loop topology option 2. If you add the additional drive module by connecting the ports to the end of the storage system without disconnecting the loop, the ESM firmware download is successful.

Drives Cannot be Removed during a Drive Firmware Download

Problem or Restriction

Removing and reinserting drives during the drive firmware download process might cause the drive to be shown as unavailable, failed, or missing.

Workaround

Remove the drive, and either reinsert it or reboot the controllers to recover the drive.

Controller Replacement in a Single Controller Environment Causes Unauthorized Premium Features

Problem or Restriction

After you replace a controller with a new controller in an storage system containing a single controller, the premium features become unauthorized and are out of compliance. Any premium features not supported by the firmware of the new controller are unavailable.

Workaround

Re-enable the premium features by using the keys you were provided when the premium features were originally enabled. Re-enabling the premium feature is only possible if the new level of firmware supports the premium feature. If the keys do not work, contact Customer and Technical Support.

Host Fails to Boot in a Root Boot Environment

Operating System
Windows Server 2003

Problem or Restriction
You are unable to boot your Windows Server 2003 host in a root boot environment. This problem occurs because the controller ownership of the boot volume has been changed to the alternate controller.

Workaround
From a different host or storage management station, you must manually change the Ownership/Preferred Path of the boot volume to the preferred path. For additional information about changing controller ownership of a volume group or volume, refer to the Array Management Window online help.

Too Many Devices Can Result in Registry Corruption

Operating System
Windows 2000

Problem or Restriction
Due to limitations in the operating system, the system registry file might get too large for the system to handle if the system is attached to a storage system that is configured with large number of volumes. This problem prevents the system from booting the next time it is started. The maximum number of volumes that can be attached to a host varies depending on what other resources are affecting the registry.

Workaround
Periodically back up the following system registry file:

```
C:\winnt\system32\config\system
```

NOTE If the system cannot reboot, restore the backed-up version of the registry file to its
original location. You can also back up the registry file by pressing F8 in the early stage of a reboot and selecting “Last Known Configuration.” The system then uses a backup copy of the registry.

Low Voltage Current of Full Rack Might Exceed Breaker Capacity

Problem or Restriction

If the available AC power supply voltage drops to 180V, the current requirements of a 72-inch cabinet that is fully-populated with FLA200 drive modules will exceed the current capacity of the 20-Amp circuit breaker by nearly 2 Amps. At nominal voltages, the circuit breaker is within limits.

Workaround

Both of the qualified part suppliers have a 25-Amp circuit breaker in the same form factor. Migrating the power supplies to the 25-Amp circuit breaker ensures continued operation in situations where the voltage might fluctuate.

I/O Errors Appear when Changing Host Cache Settings

Operating System

AIX

Problem or Restriction

Changing the AIX host cache settings while the storage system is accepting I/O requests can cause I/O errors or loss of access to data.

Workaround

Stop all I/O activity to the storage system, and then change the host cache setting. After the host cache settings have been changed, start I/O activity.

ESM Firmware Does Not Download through a Serial Connection

Operating System

Windows 2000

Problem or Restriction

When using a serial connection to download ESM firmware, the following error message is displayed:

```
Cannot communicate with the controller to complete this
request
```

Workaround
Confirm that no one is making configuration changes to the storage system, such as adding volumes or modifying volume groups, and then restart the ESM firmware download.

Tray ID of 0 (Zero) is Restricted

Problem or Restriction
Because of the potential conflict between a drive module intentionally set to 0 (zero) and a tray ID switch error that causes a tray ID to be accidentally set to 0, do *not* set your tray ID to 0.

Workaround
None.

Hibernate Does Not Work in a Root Boot Environment

Operating System
Windows Server 2003

Problem or Restriction
When you configure a storage system as a boot device, the system displays a blue screen and does not respond when it is manually or automatically set to hibernate.

Workaround
If you use a storage system as a boot device for Windows Server 2003, you cannot use the hibernation feature.

To ensure that the hibernation feature is disabled, complete the following steps:

1. Select Start >> Settings >> Control Panel >> Power Options.
2. Select the Hibernate tab.

IMPORTANT If the Hibernate tab does not appear in this dialog, then your system does not support hibernate. No additional action is required.

---

3. Is the Enable Hibernate Support check box selected?
   - **Yes** – Deselect the check box, select Apply, and then go to step 4.
   - **No** – Go to step 4.
4. Select the Power Schemes tab.
5. From the Turn Off Hard Disks drop-down menu, select Never.

6 Select OK.

End Of Procedure

Moving Drives in a Storage System Might Cause Data Loss

Problem or Restriction
When you pull a drive from a storage system to move it to a different slot, wait 60 seconds before inserting the drive into the new slot. Inserting the drive sooner might cause lost data or data corruption.

IMPORTANT The Recovery Guru recommends waiting only 30 seconds before replacing a drive. This recommendation refers to situations where the drive is being replaced within the same slot.

Workaround
None.

Host Adapter Restrictions

The restrictions in this section apply specifically to host adapters.

IBM xSeries 346 Server Cannot Connect to an Emulex HBA's BIOS While Booting

Problem or Restriction
When you attempt to connect to Emulex host bus adapters (HBAs) through an IBM xSeries 346 server, the option to enter the HBA BIOS is not available (using Alt-E) at boot-time.

Workaround
None. The HBA BIOS (using Alt-E) is needed at boot-time to set up the HBA or for root boot. If you need to configure the HBA, use Emulex HBAnywhere software instead of the boot-time BIOS.

Emulex HBA Driver Causes a Panic on the Loop

Operating System
Red Hat Enterprise Linux (2.6 kernel)

Problem or Restriction
Emulex HBA driver version 8.0.16.6_x2 running on Red Hat Enterprise Linux (2.6 kernel) causes a panic on the Fibre Channel loop.

Workaround
Upgrade the Emulex HBA driver to version 8.0.16.17.

Emulex HBA Connected to a Brocade Switch Initializes Incorrectly

Problem or Restriction
An Emulex HBA connected to a Brocade Silkworm 4100 Fabric Switch initializes incorrectly after the switch is re-enabled. The host attached to the HBA can no longer see devices on the fabric.

Workaround
Open the settings on the switch, and set the switch port to G port only. Then reset the HBA.

Cannot Load QLogic HBA Drivers on an HP rx2620 Server Running SuSE Linux

Operating System
SuSE Linux Enterprise Server 8

Problem or Restriction
You can successfully compile the QLogic HBA driver version 7.07.00b6 and install the QLogic object files in the `/lib/modules` tree on a Hewlett-Packard (HP) Integrity rx2620 server. However, when the server is rebooted and the driver attempts to load, a Machine Check Architecture (MCA) error message is displayed and a system crash occurs.

Workaround
None. HP does not support running SuSE Linux Enterprise Server 8 on the HP Integrity rx2620 server.

Root Boot Fails with QLogic HBA Port 1 Connected to a Target

Operating System
Linux

Problem or Restriction
QLogic HBA port 1, instead of port 0, is connected to a target during installation of the operating system. With this hardware setup, the operating system gives you the option to load the Master Boot Record to `/dev/sdb`. The QLogic HBA driver skips the scan of `/dev/sda/`.

If you accept this option and install the operating system on `/dev/sdb`, root boot will fail after the operating system is installed. This action occurs because the host will try to scan for the operating system at `/dev/sda` when it is actually loaded on `/dev/sdb` instead.

Workaround
If port 1, instead of port 0, on the QLogic HBA is connected to a target while installing the operating system, make sure that the "Configure advanced boot loader options" box is checked on the Boot Loader Configuration window. Edit the settings in the "Change Drive Order" window by prioritizing `/dev/sdb` as the boot device instead of `/dev/sda`.

HBA Host Port Fails to Reinitialize after a Switch Port Failure

Operating System
Solaris

Problem or Restriction
A host port fails to reinitialize after a switch port failure. This condition will cause I/O errors if the second port fails.

Workaround
Reinitialize the host port manually by resetting the HBA host port from the host.

QLogic HBA with Driver Version 4.17 Causes System Panic

Operating System
Solaris

Problem or Restriction
If you use QLogic HBAs with driver version 4.17 on a Sun Fire V440 server, a system panic could occur.

Workaround
Do not use QLogic HBAs in a Sun Fire V440 server.

HBA Hot Swap Causes System Panic

Operating System
Linux

Problem or Restriction
If you hot swap an HBA while I/O activity is present on dual paths, a system panic could occur.

Workaround
Do not hot swap an HBA while I/O activity is present on both data paths. If the data path on the HBA is active, place the controller attached to the HBA into service mode to transfer volume ownership to the other controller. Then perform the hot swap on the HBA after the data path is inactive.

HBA Driver Cannot Load After the Kernel is Upgraded

Operating System
Linux 2.6 kernel

Problem or Restriction
The Red Hat Enterprise Linux 4 operating system cannot load the HBA driver after the kernel is upgraded to a later version. When the Red Hat Enterprise Linux 4 operating system is initially installed, the installation creates a symbolic link to the kernel source tree. When a kernel is upgraded to a later version, the operating system does not create the symbolic link, which prevents the HBA driver from loading.

Workaround
Create symbolic links similar to the following:

```
ln -s /usr/src/kernels/2.6.9-5.EL-ia64/ build
ln -s build/ source
```

Qlogic HBA is Causing Volume Thrashing

Operating System
Linux 2.6 kernel

Problem or Restriction
Controllers connected to hosts through Qlogic HBAs can undergo volume thrashing. (Volume thrashing is the condition where volume ownership moves from one controller to a second controller, then back to the first controller.) The reason for volume thrashing is

that the Qlogic HBAs are not using the timeout value `portDownRetryCount` before initiating failover. Failover is then initiated before a controller can respond. Failover and failback occur without data loss.

Workaround

None. This is an issue with the Qlogic HBAs.

FLX280 Loses Connection When 6540 is on the Same Fibre Channel Loop Using Emulex HBA

Problem or Restriction

When an FLX280 control module and an 6540 control module are on the same Fibre Channel loop with an Emulex HBA, the Emulex HBA might not log back into the FLX280 if the FLX280 is reset. FLX280 control modules and 6540 control modules cannot be used with an Emulex HBA on the same Fibre Channel loop.

Workaround

Connect the Emulex HBA to the FLX280 control module and the 6540 control module using a Fibre Channel switch.

Sun Cluster 3.*x* is Not Functional in Solaris 8 Operating Environment

Operating System

Solaris 8

Problem or Restriction

Sun Cluster 3.*x* is not functional in a Solaris 8 operating environment when VERITAS DMP is the failover driver for HBAs that use the “sd” boot driver.

Workaround

None. Changing the “sd” driver to work in the Sun Cluster 3.*x* configuration would require modifications to the devid generation algorithm, and Sun does not plan to fix it.

Solaris Hosts Configured with LSI HBAs Fail to Boot Using the VERITAS `devalias` Entry

Operating System
Solaris 7 and 8

Problem or Restriction

When mirroring the encapsulated boot disk to a volume on a storage system, VERITAS has a utility that creates a devalias entry, which enables users to boot the system (provided the
`use-nvramrc?` switch is set to `true`). With Solaris 7 and 8 hosts that are configured with LSI HBAs, the hosts fail to boot (the error message states that the "device is not found").

Workaround

For the volume on the storage system to successfully boot, manually change the `use-nvramrc?` switch to `true` using the devalias entry that the VERITAS utility created.

JBODs Connected to HBA Causes Kernel Panic at Boot

Problem or Restriction

A kernel panic is a message displayed by an operating system upon detecting an internal system error from which it cannot recover. When three JBOD devices are connected to an HBA, a kernel panic condition occurs at boot; however, this occurs only when persistent data bindings are loaded.

Workaround

To avoid kernel panic, do not load the persistent bindings data.

NOTE The purpose of persistent bindings for Fibre Channel (FC) HBAs is to assign specific target IDs to specific FC SCSI devices. This target ID association is retained from reboot to reboot, unless you change it.

QLogic HBA Has Memory Allocation Failure

Linux Red Hat Operating System

Problem or Restriction

When Linux hosts that are configured with QLogic HBA drivers attempt to allocate DMA-capable memory by way of the kernel's `pci_alloc_consistent()` function, the allocation sometimes fails.

Workaround
None.

Linux Red Hat Locks Due to No Memory Error

Linux Red Hat Operating System

Problem or Restriction
Linux hosts that are configured with SteelEye® LifeKeeper® cluster version 4.4.3 lock up due to Error No Memory (ENOMEM) problems. The error messages, located in the `/var/log/messages` file are displayed as follows:

```
ENOMEM in journal_alloc_journal_head, retrying
```

Workaround
Use only one resource (two LUNs).

Root Boot on 64-bit Servers is Not Supported with Emulex HBAs on Linux

Operating System
Red Hat Enterprise Linux or SuSE Linux

Problem or Restriction
64-bit AMD or Itanium servers with Emulex HBAs cannot be set up for root boot. At this time, the correct driver set cannot be loaded onto this type of system.

Workaround
None.

Root Boot on a 64-bit Itanium Server is Not Supported with Emulex HBAs

Operating System
Windows Server 2003

Problem or Restriction
A 64-bit Itanium server with Emulex HBAs cannot be set up for root boot. At this time, the correct driver set cannot be loaded onto this type of system.

Workaround
None.

Root Boot on a 64-bit Itanium Server is Not Supported with LSI Logic HBAs

Operating System
Windows Server 2003

Problem or Restriction
A 64-bit Itanium server with LSI Logic HBAs cannot be set up for root boot. At this time, the correct driver set cannot be loaded onto this type of system.

Workaround
None.

Sun HBA Driver Fails to Communicate with Controllers after Link Initialization

Problem or Restriction
After link initialization, a driver associated with the Sun SG-XPCI1FC-QF2 (formerly X6767A) and SG-XPCI2FC-QF2 (formerly X6768A) HBAs fails to communicate with the controllers, causing path errors in a Fibre Channel loop configuration.

Workaround
Only use Sun SG-XPCI1FC-QF2 (formerly X6767A) and SG-XPCI2FC-QF2 (formerly X6768A) HBAs in a fabric environment.

HBA Conflict Occurs on a Storage System Network

Operating System
Solaris

Problem or Restriction
There is a known issue when JNI and LSI Logic HBAs are being used in the same storage system network. At system boot, the JNI HBA detects the LSI Logic HBA as a volume (target) rather than as a host (initiator). The genfa script of the storage management software locates and binds to the correct volumes, but these actions add approximately one minute per target to the boot time. This issue has been sent back to both vendors to find a solution.

Workaround
To ensure proper performance, use only one manufacturer’s HBA in a storage system network.

Path Errors Occur when Running I/O Activity and sysReboot in a Fibre Channel Loop Configuration with Sun HBAs

Problem or Restriction

Running I/O activity and sysReboot in a Fibre Channel loop configuration can cause the Sun SG-XPCI1FC-QF2 (formerly X6767A) and SG-XPCI2FC-QF2 (formerly X6768A) HBAs to stop recognizing the controllers, leading to path failures. If the problem continues and the remaining path is lost, I/O activity errors might occur as well.

Workaround

Only use Sun SG-XPCI1FC-QF2 (formerly X6767A) and SG-XPCI2FC-QF2 (formerly X6768A) HBAs in a fabric environment.

JNI FCE-1473 is Restricted to a Fabric Environment Only

Operating System

Solaris 8 or 9

Problem or Restriction

If a Solaris host is equipped with JNI FCE-1473 HBAs, the host must be in a fabric environment. Direct connection is not supported for this HBA.

Workaround

None.

Controller Reboots Cause Volume Transfers

Operating System

Solaris

Problem or Restriction

If a Solaris host is equipped with JNI HBAs, controller reboots cause a transfer of ownership to the alternate controller.

Workaround

None. This is the way JNI HBAs are designed to function. You can change back ownership by using the storage management software.

Using QLogic HBAs with Stratos Transceivers Might Cause Host I/O Timeouts and Degradation

Problem or Restriction

If you are using a direct connect configuration with QLogic QLA2310F or QLA2342F HBAs with a *Stratos* or *Infineon* transceiver, host I/O timeouts and degradation in I/O processing performance might occur. Stratos transceivers can also cause additional conflicts in non-direct connect (fabric) configurations.

Workaround

Use Qlogic HBAs with *Finisar* transceivers.

Communication to the Storage System Might be Lost

Operating System

Solaris

Problem or Restriction

When the host has a single Fibre Channel connection from a JNI HBA to the host-side minihub of the storage system, a situation might occur after controller reboot where the HBA locks up and fails to initialize the connection. Communication to that storage system controller is lost.

Workaround

Use a fabric switch topology instead of a direct Fibre Channel connection.

I/O Errors Might Occur if a Controller is Reset

Operating System

AIX

Problem or Restriction

In a non-failover environment, the storage system becomes unresponsive after a controller is rebooted. This is a problem with hosts that are configured with IBM FC6228 HBAs. The host does not allow enough time for the controller to reset before sending I/O requests.

Workaround

Before rebooting the controller, stop all I/O activity to the controller. When the reboot is complete, start I/O activity.

Hosts Using JNI HBAs Cannot Use the `hot_add` Utility

Operating System
Solaris

Problem or Restriction
If a Solaris host contains some JNI HBAs, you cannot use the `hot_add` utility to detect newly created volumes.

Workaround
After creating new volumes, you must reboot the operating system.

General Hardware Restrictions

The restrictions in this section apply to hardware components in a SAN setup that are not specific to storage systems or host adapters.

Removed SFP Hardware is Not Recognized When Reinserted

Problem or Restriction
Small Form-factor Pluggable (SFP) hardware that is removed and then reinserted within 10 seconds might not be recognized by the controller. Device polling is done at 10-second intervals. Any device that is removed and replaced within the 10-second window will not be recognized.

Workaround
After removing SFP hardware, wait at least 10 seconds before reinserting it.

Cannot Connect to a Cisco MDS 9020 Fabric Switch

Problem or Restriction
An “Unknown User” error message appears while the host is trying to connect to a Cisco MDS 9020 fabric switch using Fabric Manager Version 2.1 with SNMPv3 enabled.

Workaround
None. This is a known issue with the Cisco MDS 9020 fabric switch. The fabric switch does not support SNMPv3.

NFS Time-Out after Node Failover

Operating System
Solaris

Problem or Restriction
A stale Network File System (NFS) file handle is seen on a Sun Ultra 80 server after a node failover. This problem is most likely caused by an NFS time-out because of the slower speed of the server. This condition could possibly cause an I/O error.

Workaround
Retry the I/O.

64-Bit SuSE Linux Host Reboots after `cmd_timeout`

Operating System
SuSE Linux

Problem or Restriction
Disconnecting and reinserting a cable in an out-of-band (directly connected), 64-bit host running SuSE Linux could cause it to lock up or reboot. This is a problem in the SuSE kernel and is due to a `cmd_timeout`.

Workaround
None. Ensure I/O is stopped prior to re-cabling so that the server does not lock up or reboot.

Hosts With Multiple EM64T Processors Will Not Boot Up On SMP 2.4 Kernel

Operating System
Linux 2.4 kernel

Problem or Restriction
Configurations with a host running multiple EM64T (Xeon 64) processors cannot run the Red Hat Enterprise Linux Advanced Server 3.0 (RHELAS) operating system. On system boot up, the SMP kernel does not install. If the SMP kernel is manually installed when the host boots up, the host generates a kernel panic with the following results:

```
Kernel panic: Fatal exception
In interrupt handler - not syncing4)
```

On an AMD64 bit architecture type, the host boots up and installs the SMP kernel without any errors.

Workaround
None.

Switch Port Identifier is Misconfigured

Problem or Restriction
If the switch port identifier is substituted as a host port identifier, the host system loses access to the storage system.

Workaround
Type the World Wide Names (WWNs) manually rather than relying on those presented by the software interface.

SANtricity Storage Manager Restrictions

The restrictions in this section apply specifically to SANtricity Storage Manager.

Snapshot Creation Fails When Using an 6540 Control Module in a Cluster Environment

Problem or Restriction
Snapshot creation fails on an 6540 control module because a suitable logical unit number (LUN) could not be found for the target host or host group. This usually occurs if all LUNs are already in use for the host or host group.

Workaround
Map the access volume to the host or host group.

Error Messages Are Not Written to the System Log on UNIX Platforms

Operating System
All UNIX platforms.

Problem or Restriction
Error messages do not get written to the system log because UDP port 514 is disabled.

Workaround
None.

“Error 33” When Redistributing Volumes after Controller Replacement

Problem or Restriction

After warm swapping both controllers in a control module or an storage system, the following error message appears when trying to redistribute the volumes back to their preferred owners:

```
Error 33 - The operation cannot complete because the volume
group specified in the request is not valid (unknown volume
group reference). The volume group may have been deleted or
modified by a user on another storage management station
accessing this storage system.
```

Workaround

Reboot both controllers to synchronize data in the controller firmware, and then redistribute the volumes.

Uninstalling the Software Leaves the JRE Directory Behind

Problem or Restriction

This problem might occur if you uninstall a copy of the SANtricity Storage Manager software that has been upgraded from a previous release. The newer installation package cannot delete the JRE directory created with an older installation package.

Workaround

Manually delete any directories left behind after uninstalling the software.

“Error 173” When Creating Remote Volume Mirrors through the Script Engine

Problem or Restriction

When creating remote volume mirrors through the Script Engine in SANtricity Storage Manager, the following error might display during execution, stopping the remote volume mirror creation process:

```
Error 173 - The operation cannot complete because the
configuration number on the remote storage system is invalid.
```

Workaround
Some remote volume mirrors might have been created. To finish creating the rest, re-run the script to complete the operation.

Dialogs Are Blank on Solaris Host Running the Agent Software if Network Cable is Removed

Operating System
Solaris

Problem or Restriction
Removing the Ethernet network cable from a Solaris host while the SANtricity Storage Manager Agent software is running causes networking problems that subsequently cause screen refresh problems in SANtricity Storage Manager. Dialogs will go blank for any functions that need access to files on the host, including downloading firmware files, saving profiles, saving error information, and so on.

Workaround
Do not remove the network cable from the Solaris host while the agent software is running.

"Error 62" When Mapping Volumes to an Unused Logical Unit Number through the Agent Software

Problem or Restriction
When connected through the SANtricity Storage Manager Agent and mapping volumes using SANshare Storage Partitioning on an in-band storage system, the following error message might appear after a successful volume-to-LUN mapping operation:

```
Error 62 - The operation cannot complete because the logical
unit number (LUN) is already in use. Please enter another LUN.
```

Workaround
Verify that the volume-to-LUN mapping was created correctly. If it was, this error is erroneous and can be ignored.

Disable Premium Feature Dialog is Empty

Problem or Restriction
After you disable a premium feature, the Disable Premium Feature dialog might go blank and no longer show any other enabled premium features to select.

Workaround
Close and then re-open the dialog.

File-Selection Causes Firmware Download Dialog to Go Blank

Operating System
Linux, Solaris

Problem or Restriction
After you select a file on the firmware download or NVSRAM download dialog, double-clicking on ".." in the dialog no longer moves the directory back one level.

Workaround
Do not select a firmware file until you are finished navigating through your directory structure.

Cut-and-Paste from Online Help Causes Commands to Fail

Problem or Restriction
Cutting CLI or Script Engine examples from the online help and then pasting them into any application (Script Editor, Command Prompt, and so on) causes extra spaces to be added to the script where they are not needed. The pasted commands fail to execute properly because of the extra spaces.

Workaround
None.

JavaScript Bug Causes Glossary Pop-Up to Stay Open

Problem or Restriction
A JavaScript bug causes the Glossary pop-up window to stay open when the help contents on the storage management software GUI are closed and reopened.

Workaround
None. This is a vendor defect.

HP-UX GUI Does Not Properly Show NVSRAM Files

Operating System
HP-UX

Problem or Restriction
The storage management software GUI, when installed on an HP-UX machine, does not properly display All Files when the user selects the All Files option from the Files of Type drop-down list.

Workaround
Open a different directory using the "*" filter, and then switch back to the directory that contains the files that you were trying to display using the GUI.

Remote Volume Mirroring Activation Fails

Problem or Restriction
Remote Volume Mirroring activation fails if there is too little free storage available after Remote Volume Mirroring has been deactivated.

Workaround
Free some storage space (add drives, delete unused volumes, and so on) before attempting to reactivate Remote Volume Mirroring.

Diagnostics Do Not Run in a Single-Controller Storage System

Problem or Restriction
Controller diagnostics do not run in a single-controller storage system.

Workaround
To run controller diagnostics, you must have two controllers in the storage system.

Some SMcli Commands Do Not Work on NetWare 6.5

Problem or Restriction
The NetWare 6.5 operating system (Service Pack 1) server returns an error message when you run SMcli commands that include parameters surrounded by double quotation marks. In addition, NetWare 6.5 does not support escape characters and ignores single ticks. For example:

```
SYS:\storagemanager\client\SMcli 123.45.678.99 -c "show
storagearray profile;"
```

produced the following error message:
`Unknown or unresponsive address (hostname or IP address): storagearray`

Workaround
There is no workaround for NetWare CLI commands; however, you can use the SANtricity Storage Manager GUI instead of the NetWare server command line to execute commands.

Error Occurs When Opening a PDF File after Installing Adobe Reader from the Installation CD

Operating System
Windows Server 2003

Problem or Restriction
Adobe® Portable Document Format (PDF) files on the SANtricity Storage Manager Installation CD do not open after installing Adobe Reader. When the first PDF file is opened, a License Agreement dialog displays behind the application window. Until the License Agreement is accepted, you cannot open the PDF files.

Workaround
Select the License Agreement dialog in the background, and accept the License Agreement.

Storage System Unmanageable through the Agent after a Configuration Reset

Problem or Restriction
A loss of communication dialog appears after you reset the configuration using a host-agent connection.

Workaround
Re-establish an access volume mapping using a direct connection.

Host Reports Errors When a Controller is Placed Offline with Remote Volume Mirroring Enabled

Operating System
Linux

Problem or Restriction

I/O errors might occur when a volume group is in a Degraded state, and a controller is placed offline. Because of the Remote Volume Mirroring operations being performed on the degraded volume group, queued write commands cause a timeout on the host.

Small Computer System Interface (SCSI) aborts are reported in the `/var/log/message` file, but the kernel tries the I/O command again. If the I/O activity still does not complete in the allotted time, an I/O error is reported. The result is temporary loss of data accessibility.

Workaround
Reduce the number of volumes involved in Remote Volume Mirroring relationships to 16 per storage system (eight primary and eight secondary), or add drives to the volume group that contains the volume for which the SCSI abort was reported.

Target Volume Might Not Contain Exact Data as Source Volume after Volume Copy Creation

Operating System
NetWare

Problem or Restriction
If the source volume in a volume copy is mounted to a NetWare host, the target volume might not contain an exact copy of the source volume.

Workaround
Before creating a volume copy, unmount the source volume. After the volume copy is complete, remount the source volume.

Truncated Titles are Displayed in the Volume Creation Wizard Dialogs

Operating System
Solaris

Problem or Restriction
The title bar text might be truncated in some of the Volume Creation Wizard dialogs.

Workaround
Resize the window by using the mouse or similar input device.

Storage Management Software Crashes during Volume Copy Creation

Operating System
Solaris

Problem or Restriction
During a volume copy, the storage management station software crashes, and the system displays the following information:

```
Unexpected Signal : 10 occurred at PC=0xff0340c0
Function name=JVM_NewArray
Library=/opt/SM8/jre/lib/sparc/client/libjvm.so
...
```

Workaround
This is a known issue with the Java Runtime Environment (JRE). For assistance with fixing this issue, contact Customer and Technical Support.

Host Does Not See Changes to a Base Volume When the Snapshot Volume is Modified

Operating System
Windows Server 2003

Problem or Restriction
If you disable and then re-create a snapshot volume that is mapped to a Windows Server 2003 host, the operating system does not detect any changes to the base volume that occurred between the time you disabled and re-created the snapshot volume. This occurs on both the 32-bit and 64-bit versions of the operating system.

Workaround
When you finish modifying the snapshot volume, reboot the Windows Server 2003 host.

Windows 2000 Hosts are Restricted to 64 LUNs per Partition

Operating System
Windows 2000

Problem or Restriction
Due to a possibility of corrupting the Windows Registry, the Windows 2000 operating system is restricted to 64 LUNs per partition for this release of the storage management software.

Workaround
Do not use more than 64 LUNs per partition.

Repeated Opening and Closing of the Array Management Window Can Cause Excessive Memory Usage

Operating System
Windows 2000

Problem or Restriction
Leaving the Enterprise Management Window open for long periods of time, combined with opening and closing the Array Management Window repeatedly, can cause excessive system memory usage.

Workaround
If the Enterprise Management Window is left open for long periods of time, close the Enterprise Management Window occasionally, and then reopen it.

Long Boot Time Occurs after Mapping Volumes to a Host

Operating System
Windows Server 2003

Problem or Restriction
Occasionally, after mapping volumes to a Windows Server 2003 host, the host can take a longer time to boot than usual. The amount of time is based on how many volumes are mapped to the host.

Workaround
None. This is a known issue with the Plug-and-Play (PnP) Manager in the Windows Server 2003 operating system. Regardless of the long boot time, this cycle must complete.

Replaced Drive Appears Offline

Problem or Restriction

When you replace an optimal drive with a spare drive, the new drive displays an Offline status in the Array Management Window for approximately 10 seconds.

Workaround

Wait at least 30 seconds between removing a drive and inserting a replacement. The storage management software can take up to 10 seconds to detect changes in the hardware. The software must find the removed drive before the replacement is discovered.

Dynamic Segment Sizing (DSS) Modification Operation Error Occurs

Problem or Restriction

When attempting to increase the capacity of a volume group using Dynamic Segment Sizing (DSS), the following error message might be received:

```
Error 40 - The operation cannot complete because either (1)
the segment size requested is not valid, or (2) the segment
size you specified is not allowed because this volume has an
odd number of segments. Therefore, you can only decrease the
segment size for this volume to a smaller number.
```

Workaround

Use the Dynamic Capacity Expansion (DCE) modification operation to increase the capacity of the volume group, and then retry the DSS operation. If you continue to receive this error, contact Customer and Technical Support.

Downgrading Storage Management Software to Previous Versions

Problem or Restriction

To prevent accidental data loss because of architectural differences, the version 8.*x* or version 9.*x* storage management software prevents downgrades to previous versions of the software.

Workaround

If you need to downgrade your storage management software from version 8.*x* or 9.*x*, contact Customer and Technical Support for assistance.

Controllers in the Same Storage System Appear as Separate Devices

Problem or Restriction

After cabling a storage system, the controllers appear as separate storage systems after they are added to the Enterprise Management Window.

Workaround

The cabling in the storage system showing the error is incorrect. Check the drive-side cabling connections, and ensure that the cabling is a supported configuration.

Drives Show Failed Status after a Power Cycle or an Interruption

Problem or Restriction

When power is disconnected or interrupted from a drive module, drives might appear Offline or show a Failed status in the Array Management Window. This problem typically happens when the drive module is shut down or loses power, even though the control module might not lose power or be shut down.

Workaround

After power has been completely restored to the control module and drive modules, select the failed drives in the Array Management Window, and then select Advanced >> Reconstruct Drive. Repeat this action for every drive with Failed status.

Tray IDs are Listed Incorrectly in the Array Management Window

Problem or Restriction

Tray IDs might be listed incorrectly in the Array Management Window, even if the tray ID switch settings appear to be correct.

Workaround

Because of the physical design and movement of the tray ID switch, it is possible to leave the switch in a “dead zone” between ID numbers, which return an incorrect tray ID to the storage management software. The most commonly returned tray ID is 0 (zero).

When setting the tray ID, ensure that the switch has actuated completely and has settled so that the value is clearly visible in the viewing window.

Storage System Disappears on a Path Failure

Operating System
HP-UX

Problem or Restriction
The automatic discovery of the access volume is only performed at startup, and the SMagent does not remove or activate the access volume dynamically. When a failure occurs on the path to the access volume, the SMagent can lock up during a write command to the access volume, and the storage system disappears from the SMclient window.

Workaround
Stop the SMagent, and then restart it. The host software does not detect the path after the restart (because of the failure) and does not attempt to send any I/O through it. The SMclient allows the storage system to be added underneath the host icon again, listing it as a partially managed device.

Some Keyboard Shortcuts Do Not Function as Expected

Problem or Restriction
Some typical keyboard shortcuts do not function as expected. This is a known issue in the Java Runtime Environment (JRE). Many keyboard accessibility issues will be corrected in a future release of the JRE. For information about which keyboard shortcuts are available for this release, refer to the Array Management Window online help.

Workaround
None. Use the mouse or similar input device for options that are not keyboard accessible.

Update Button in the Download Dialogs Does Not Function Correctly

Problem or Restriction
The Update button in the NVSRAM Download dialog and the Firmware Download dialog does not function as expected on some UNIX platforms.

Workaround
In the file selection area, leave the selected directory, and reselect it again.

Failover Software Restrictions

The restrictions in this section apply specifically to failover software.

Errors Occur after a Controller Comes Back Online While Using MPxIO with an Emulex HBA

Operating System
Solaris 8

Problem or Restriction
I/O errors and “busy too long” error messages occur on a host after a controller goes offline and comes back online while using the MPxIO failover driver with an Emulex HBA.

Workaround
Use VERITAS DMP as the failover solution instead of MPxIO.

Host Running MPxIO on Solaris 10 Panics after Controller Exceptions Occur

Operating System
Solaris 10

Problem or Restriction
After four or five controller exceptions, MPxIO might cause the Solaris 10 host system to panic.

Workaround
Use VERITAS DMP as the failover solution instead of MPxIO.

Persistent Reservation Not Found on a Node That Has Been Rebooted

Operating System
Solaris 9

Problem or Restriction
Persistent reservation type WE_RO is not found on volumes in a Sun Cluster Server node after any node in the cluster has been rebooted.

Workaround
Contact Customer and Technical Support for a fix to this problem.

Application Errors Occur Before MPxIO Can Fail Over to the Other Controller

Operating System
Solaris

Problem or Restriction
Link instability caused by intermittent hardware failures has caused an I/O error. Because the instability might not affect the controller directly, MPxIO does not fail over correctly, and an application error occurs.

Workaround
Find the component causing the instability, and remove it or replace it.

Large Number of Volumes Takes a Long Time to Fail Back to Preferred Controller Using MPxIO

Operating System
Solaris

Problem or Restriction
After a controller failover, a large number of volumes might take a long time to fail back to the preferred controller when the host is running the MPxIO driver.

Workaround
None. The MPxIO driver must process each volume individually. On a large configuration, this can take some time to complete.

Failovers Occur Immediately in a Fabric Environment

Problem or Restriction
After removing a cable, failovers might occur immediately when running in a fabric environment. This problem is more likely to occur in a topology that contains 4-Gb switches and has been seen with a variety of failover drivers, including MPIO and RDAC.

Workaround
None. The failover driver will move the volumes back to the preferred path after the cable is reinserted.

Failover Delayed Due to RSCN Processing

Problem or Restriction
If a failover occurs while a device attached to the SAN is causing excessive registered state change notifications (RSCNs), failover processing of I/O is delayed. As a result, I/O will be queued for a longer than normal amount of time.

Workaround
Determine the cause of the excessive RSCNs, and adjust the device or remove it from the SAN.

NOTE To find out which device is causing excessive RSCNs to be generated, either look at the switch log or look in the host log (`/var/log/messages`) for entries similar to the following:
`RSCN database changed -0x1,0xf00`
In this example, the port causing excessive RSCNs is `0x1,0xf00`.

.Failover Events Cause I/O Errors Using MPxIO Driver

Operating System
Solaris

Problem or Restriction
Multiple external problems that cause failover events could lead to timing issues where volumes are temporarily unavailable when using the MPxIO failover driver. If this happens, I/O errors could occur.

Workaround
None. This is a limitation of the MPxIO driver.

“Cluster Lost Operational Quorum” Error Message

Operating System
Solaris

Problem or Restriction
All nodes in a cluster running Sun Cluster 3.1 will panic with a “Cluster lost operational quorum” error message if the node owning the MPxIO-enabled quorum is taken offline. This is a problem if the `ssd` driver patch is at level 113277-28 or later.

Workaround
Do not install a patch to the `ssd` driver to level 113277-28 or later on nodes in an MPxIO-enabled quorum.

Controller Cannot Fail Back after Replacement

Operating System
Irix

Problem or Restriction
A controller that has recently been replaced will fail volumes over to the second controller. However, the volumes will not fail back to the primary controller because the server is not aware of the device name.

Workaround
After replacing a controller, manually scan for the device name from each server that is connected to the storage system.

I/O Errors Occur after a Host Receives Fatal Errors from the `ssd` Driver

Operating System
Solaris

Problem or Restriction
A controller taking a long time to complete the start-of-day (SOD) boot process or placed in Service mode could cause the `ssd` driver to time-out and send fatal errors to the host. This problem could lead to I/O errors.

Workaround
None.

Device Node Name Entries in the VERITAS Volume Manager are Incorrect

Operating System
Solaris

Problem or Restriction
The device node name entries in VERITAS Volume Manager (after using the command `vxdisk list`, for example) are incorrect after a change has been made to the devices in the fabric topology. This problem could cause inconsistent disk usage and possible I/O errors.

Workaround
Run the `vxdiskadm` utility and choose option 20 to force VERITAS Volume Manager to update its device node name entries. Run the utility a second time to restore the disk-naming conventions.

INQ-83 Errors on the Host Cause I/O Failures

Operating System
Solaris

Problem or Restriction
A host using the MPxIO failover driver reports INQ-83 errors, which means that the storage system did not return the correct volume World Wide Name (WWN) to the host in the allocated time.

Workaround
To correct this problem, you must force the host to rescan the affected path for the devices. Use the following procedure:

1. Have VERITAS Volume Manager rescan for disks by entering the following command:

   ```
   vxdisk scandisks #
   ```

2. Make sure that all paths are optimal and available by using the following command:

   ```
   vxdisk list <disk accessname> #
   ```

3. Find the affected path by sending the following command:

   ```
   luxadm fcode_download -p
   ```

4. After the path address is known, force a LIP to cause a link initialization on the affected port using this command:

   ```
   luxadm -e forcelip <physical path to port>
   ```

NOTE If running this procedure does not correlate devices and WWNs, reboot the host.

Erroneous Message Appears When Using the MPxIO Failover Driver

Operating System
Solaris

Problem or Restriction
After a host reboot, the following error message might appear during `vxdctl enable` or `vxconfigd` startup.

```
get_geometry_info_common: /dev/vx/rdmp//fabric_331
fmt_page_code failed. ret 0x5
get_geometry_info_common: /dev/vx/rdmp//fabric_331
fmt_page_code failed. ret 0x5
get_geometry_info_common: /dev/vx/rdmp//fabric_330
fmt_page_code failed. ret 0x5
get_geometry_info_common: /dev/vx/rdmp//fabric_330
fmt_page_code failed. ret 0x5
get_geometry_info_common: /dev/vx/rdmp//fabric_332
fmt_page_code failed. ret 0x5
get_geometry_info_common: /dev/vx/rdmp//fabric_332
fmt_page_code failed. ret 0x5
get_geometry_info_common: /dev/vx/rdmp//fabric_328
fmt_page_code failed. ret 0x5
get_geometry_info_common: /dev/vx/rdmp//fabric_328
fmt_page_code failed. ret 0x5
```

Workaround
This error message is harmless and can be ignored.

Access Volume Not Using MPxIO Is Offline after a Controller Reboot

Operating System
Solaris

Problem or Restriction
Several volumes on a storage system are configured to use MPxIO, but the access volume on the storage system is not. If an event triggers a reboot of a controller, it might cause all of the volumes, except the access volume, to come back online.

Workaround
Use the `luxadm` utility to bring the access volume back online, or force the host to rescan for devices. As an example, you could use this command:

```
luxadm -e online <physical path to access volume>
```

See the `luxadm` man page for more information about using this command.

Failure Event Does Not Trigger VERITAS DMP 4.1 to Fail Over

Operating System
Solaris 8 and 9

Problem or Restriction
A link failure event (such as a controller failure) does not trigger VERITAS DMP 4.1 to fail over to the alternate path on the Solaris 8 or 9 operating system. After this event occurs, it is possible that I/O might not continue to be written to the volumes on the storage system, and the host will not report any I/O errors.

Workaround
Trigger another failover event to move the volumes back to their preferred owners. If this does not work, reboot the host to clear the error.

RDAC Compile Warning Message

Operating System
Red Hat Enterprise Linux 4 (2.6 kernel)

Problem or Restriction
During installation of the RDAC driver, the following warning message might appear:

```
Your Kernel version is 2.6.9-5.0.5.ELsmp
Preparing to install MPP driver against this kernel version

pcilib: Resource 5 in /sys/bus/pci/devices/0000:00:1f.1/
resource has a 64-bit address, ignoring
pcilib: Resource 5 in /sys/bus/pci/devices/0000:00:1f.1/
resource has a 64-bit address, ignoring
pcilib: Resource 5 in /sys/bus/pci/devices/0000:00:1f.1/
resource has a 64-bit address, ignoring
```

Workaround
None. The message is erroneous and can be ignored. The kernel believes that the device has been assigned to 64-bit addressing but it has not.

VERITAS Cluster Server Switchover Error Results in Diskgroup Failures

Operating System
Solaris

Problem or Restriction
After setting up a VERITAS Cluster Server (VCS) service resource group with Storage Foundation-High Availability (SF-HA) and Fencing, a resource switch to the other node fails and possibly causes data corruption. Future import operations also fail, and the `vxdisk list` command outputs the disk with the “altused” tag.

Workaround
Do not use VCS with SF-HA and Fencing. This configuration is not currently supported.

NOTE Other possible problems you might see while attempting to use VCS with SH-FA and Fencing include cluster nodes that have stopped operating after controller failovers and I/O errors.

VERITAS DMP Does Not Fail Over with Controller in Service Mode

Problem or Restriction
VERITAS DMP will not work while the controllers are in service mode because DMP relies on Auto Volume Transfer (AVT) to move the volumes to the controller that receives the I/O request. When a controller is in service mode, the controller is not capable of successfully completing the I/O request. Since the path to that controller is not "broken," DMP has no knowledge to fail over to the alternate controller. The I/O request gets tried repeatedly until the I/O timeout period has been reached and the request times out.

NOTE Any failover driver that relies on AVT has this limitation.

Workaround
None.

Persistent Reservations Cannot Be Used with VERITAS DMP

Operating System
Solaris

Problem or Restriction
If persistent reservations are enabled, the VERITAS DMP failover driver fails to detect the path to the controller after a start-of-day operation or a controller reboot. This results in I/O failures.

Workaround
If you are using VERITAS DMP, disable persistent reservations.

VERITAS DMP Fails to Disable Offline Path

Operating System
Solaris

Problem or Restriction
If a cable is disconnected from a storage system, the HBA driver on the host fails to detect the offline path correctly because VERITAS DMP does not mark the offline path as disabled. This results in I/O activity not getting through to the correct volumes, and the VERITAS DMP path scanning daemon `vxdctl enable` stops responding.

Workaround
None.

Split-Brain Condition Occurs When Hosts Fail to Register with VERITAS Disks Correctly

Operating Systems
Solaris 8

Problem or Restriction
A split-brain condition (when one node survives and the other must be shut down, which can cause data corruption), occurs when a Solaris host fails to correctly register with VERITAS vxfendg disks.

Workaround
When you encounter a split-brain condition, clear the keys manually using the `vxfenclearpre` command, which removes Small Computer System Interface 3 (SCSI-3) registrations and reservations on the coordinator disks, as well as on the data disks in all shared disk groups.

NOTE Clearing the keys manually does not always work; sometimes it is necessary to bring all nodes down.

To clear keys manually, perform the following steps:

1. Shut down all other systems in the cluster that have access to the shared storage. This prevents data corruption.
2. Start the script:

```
# cd /opt/VRTSvcs/vxfen/bin
# ./vxfenclearpre
```

3. Read the script's introduction and warning. Then, you can choose to let the script run.
   Do you still want to continue: [y/n] (default : n)

   ```
   Y
   ```

   Note Informational messages resembling the following might appear on the console of one of the nodes in the cluster when a node is ejected from a disk or LUN:

   ```
   <date> <system name> scsi: WARNING: /sbus@3,0/lpfs@0,0/sd@0,1(sd91):
   <date> <system name> Error for Command: <undecoded cmd 0x5f> Error
   Level:Informational
   <date> <system name> scsi: Requested Block: 0 Error Block 0
   <date> <system name> scsi: Vendor: <vendor> Serial Number:
   0400759B006E
   <date> <system name> scsi: Sense Key: Unit Attention
   <date> <system name> scsi: ASC: 0x2a (<vendor unique code 0x2a>), ASCQ:
   0x4,FRU: 0x0
   ```

   You may ignore these informational messages.

```
Cleaning up the coordinator disks...

Cleaning up the data disks for all shared disk groups...

Successfully removed SCSI-3 persistent registration and reservations from
the coordinator disks as well as the shared data disks.

Reboot the server to proceed with normal cluster startup...
#
```

4. Restart all systems in the cluster.

End Of Procedure

Controller Reboots Cause I/O Failures

Operating System
Solaris

Problem or Restriction
If a host using the MPxIO failover driver cannot access all volume information during start-of-day, the storage system could go into a “silent” degraded mode. The problem causes the “volume not on preferred path” error to be logged if a forced failover has occurred. In this mode of operation there is no failover protection on the host and any event that causes disruption on the remaining path would cause I/O failures.

Workaround
In order to correct this problem, you must force the host to rescan the affected path for the devices. The following options can be done:

- Remove the cable to the affected HBA port and reinsert it after 30 seconds.
- Find the affected path by sending the following command:

  ```
  luxadm fcode_download -p
  ```

  After the path address is known, force a LIP to cause a link initialization on the affected port using this command:

  ```
  luxadm -e forcelip <physical path to port>
  ```

“Busy Too Long” Error Using MPxIO Failover Driver

Operating System
Solaris

Problem or Restriction
During failover, the operating system might encounter a busy status from the storage system. If the busy status is received after another storage system status, the host might inadvertently fail I/O. To check if “busy too long” has been logged on the host, check the `/var/adm/messages` file by using this command:

```
grep “busy too long” /var/adm/messages
```

Workaround
Contact Customer and Technical Support for a fix to this problem.

NetWare Does Not Support a Controller in Service Mode

Operating Systems
NetWare

Problem or Restriction
NetWare does not support the In Service mode; therefore, the host port driver does not recognize failover attempts. Due to the lost path to the controller that was placed into Service mode, the I/O failed.

Workaround
Restrict Service mode in NetWare.

The `modprobe` Utility Causes the System to Freeze or the Kernel to Panic

Operating System
Linux

Problem or Restriction
Using the `modprobe` or `modprobe -r` command on the MPP driver module's stack on a Linux host can cause the operating system to freeze or can cause kernel panic.

Workaround
Do not use the `modprobe` or `modprobe -r` command on the MPP driver module's stack. Instead, use the `insmod` and `rmmod` commands on each of the modules in the MPP driver stack in the proper order.

Path Loss in the Cluster Environment Causes Repeated Volume Failover

Operating System
Windows 2000, Windows Server 2003, and Linux SteelEye environments

Problem or Restriction
If one physical path to a storage system is lost in a cluster environment, volumes repeatedly fail over on one node and then another until the path is restored.

Workaround
For Windows, change the failover driver's registry settings to keep the volumes from failing over repeatedly. Run the file called `DisableLunRebalance.reg`, located on the SANtricity Storage Manager Installation CD, on each node in the cluster that has the MPP driver installed. Then reboot each node for the changes to take effect.

For Linux SteelEye environments, perform the following steps:

1. Open the `/etc/mpp.conf` file, and change the `DisableLunRebalance` parameter to 3. Save the changes.
2. From a shell prompt, type the following command, and press Enter:

   `mppUpdate`
3. Reboot the computer for the changes to take effect.
4. Repeat this procedure on any system in the cluster that has the MPP driver installed.

End Of Procedure

I/Os Do Not Use All Available Paths after Releasing a Legacy Reservation

Operating System

Windows 2000 and Windows Server 2003

Problem or Restriction

If a host places a legacy reservation on a volume, all future I/O requests to the volume will use the reserved path. For example, if a reservation command is sent to LUN 0 where there are two paths to Controller A, only one path can receive the request (reservation path).

If the same host places a reservation on another volume (LUN 1), RDAC uses a different path to Controller A, if one is available. If an alternative path is not available, RDAC uses the path for LUN 1 that is also used for LUN 0. This process applies to any additional volume reservations and applies as well to Controller B.

If the reservation for LUN 0 is cleared, all controller paths to LUN 0 should become available. If a host is running Windows 2000 or Windows Server 2003 with Microsoft's Server Cluster software, and the user discontinues use of Microsoft's Server Cluster software without rebooting the host, a bug makes all controller paths to LUN 0 unavailable.

Workaround

To recover from this situation, reboot the host machine after you finish using Microsoft's Server Cluster software. This has been a known issue since version 8.0 of the storage management software.

Cluster “is Alive” Timeout Causes Volumes to Move Between the Controllers

Operating System
Windows 2000

Problem or Restriction
Cluster hosts and resources might fail over to another host system when you are downloading firmware to a storage system with I/O running.

During a firmware upgrade, the volumes are being moved back and forth between the controllers. If the cluster service is unable to access the volumes within the “is alive” timeout value, this problem might result.

Workaround
Change the “is alive” timer value to 120,000 milliseconds for all resources. The default value is 60,000. For the procedure to change this value, refer to the operating system documentation.

Files Needed Dialog Appears after Removing the Redundant Dual Active Controller (RDAC)

Operating System
Windows Server 2003

Problem or Restriction
The Files Needed dialog appears, requesting the missing `rdacdisk.sys` file. This condition occurs when a removal of the RDAC software package is unsuccessful and files still remain that must be removed.

You receive a Files Needed dialog for every volume mapped to the host that contains a file system.

Workaround
For assistance in removing the remaining files, contact Customer and Technical Support.

Utilities Software Restrictions

The restrictions in this section apply specifically to the software utilities, such as SMdevices, SMutil, `hot_add`, and so on, that configure disk drives, operating systems, or other hardware or software components.

SMmonitor Displays an Error When Starting or Stopping

Operating System
Windows 2000

Problem or Restriction
When the SMmonitor agent starts or stops, an error similar to the following example is displayed in the Event Log:

```
The description for Event ID ( 2 ) in Source ( SMMONI~1 )
cannot be found. The local computer may not have the necessary
registry information or message DLL files to display messages
from a remote computer. The following information is part of
the event: SMMONI~1:, Event Monitor stopped..
```

The problem occurs because the installation wizard sets the path to the SMmonitor executable in the Windows registry to a condensed 8.3 path name (`C:\PROGRA~1\STORAG~1\client\monitor\SMMONI~1.EXE`).

Workaround
None. This message is erroneous and can be ignored. The SMmonitor service will still start and stop correctly.

Non-Contiguous LUNs Cause Problems with SMagent and SMdevices

Operating System
Linux (2.4 Kernel)

Problem or Restriction
In Linux, the SCSI device discovery operation scans all contiguous LUNs and stops scanning after a break in the numbers. For example, if LUNs 0, 1, 2, 3, and 7 are being used, 0 through 3 will be scanned and 7 will be skipped. This scenario will cause several problems:

- Hosts will not have access to any volumes associated with skipped LUNs.

- SMagent will be stopped at boot time if it does not detect an Access Volume in the block of contiguous LUNs.
- SMdevices will not be able to report on volumes associated with any skipped LUNs.

Workaround
Run `hot_add` to force the operating system to discover the skipped LUNs.

NOTE The `hot_add` utility is included with the SANtricity Storage Manager RDAC package for Linux.

SMdevices Does Not Work Because of Missing SCSI Generic Device Node

Operating System
Linux (Kernel 2.6)

Problem or Restriction
SMdevices requires a SCSI generic device node to function properly. In the Linux 2.6 release, the kernel generates device creation and device deletion events when devices are discovered and a device driver is unloaded from the kernel. The user space hotplug scripts handle the device creation and device deletion events. Device nodes such as `/dev/sg*` are dynamically created by the hotplug scripts.

There are some timing problems between the kernel and user space scripts. Sometimes the device node creation script gets timed out and results in a sg device node not being created.

Workaround
Run the following script to correct this problem:

```
#!/bin/bash
############################################################
#This is sh function that checks whether or not all sg devices

#reported by sysfs have their /dev/sgX node. If it doesn't

#exist, make the device node.

#

#This function returns immediately if the system is not

#an Linux 2.6 system
############################################################
checkAndMkSgNode()
```

```
{
   SG_CLASS_DIR="/sys/class/scsi_generic"
   if [ ! -d ${SG_CLASS_DIR} ]
   then
      #not an linux 2.6 system or sg driver is not loaded
      #leaving...
      return
   fi

   /bin/ls -1 ${SG_CLASS_DIR}/sg*/dev |/usr/bin/awk -F"/" \
   '{
      #the format of the input line looks like
      #/sys/class/scsi_generic/sg89/dev
      sg_dev_node="/dev/" $5;
      #check whether or not /dev/sgX exist
      # NoOp if the /dev/sgX exist
      # Otherwise, make a node
      if( system("test -c " sg_dev_node) != 0)
      {
         #create the sg node
         #the content in the /sys/class/scsi_generic/sg89/dev
looks like
         #21:89
         getline < $0;
            split($0,major_minor,":")
            major=major_minor[1];
            minor=major_minor[2];
       mknod_cmnd="/bin/mknod " sg_dev_node " c " major " " minor
         system(mknod_cmnd);
      }
   }'
}


checkAndMkSgNode
```

Asynchronous Remote Volume Utility on the NetWare Operating System Might Cause Volume Failures

Operating System
NetWare

Problem or Restriction
The Asynchronous Remote Volume Mirroring Utility runs poorly on NetWare and could lead to volume failures.

Workaround
None. Do not use the Asynchronous Remote Volume Mirroring Utility on the NetWare operating system.

I/O Errors Occur on Linux Blast

Operating System
64-bit Linux 3.0 kernel

Problem or Restriction
The Read Ahead portion of the Linux 3.0 kernel code triggers I/O errors, where writes are sent to one of the controllers, and that controller is not able to accept I/Os. If the Linux MPP cannot communicate with all controllers in the storage system, no I/O is read or written, resulting in read/write failures.

Workaround
Turn off the Linux 3.0 Kernel Read Ahead feature.

CAUTION **Degraded performance** – Turning off the Linux Read Ahead feature results in performance degradation, because each read must now be routed to the disk.

SMdevices Does Not Return any Information

Operating System
AIX

Problem or Restriction
SMdevices does not return any information if the access volume has not been mapped to the AIX host on the storage system.

Workaround
Map the access volume to the AIX host on the storage system. For instructions on mapping the access volume, refer to the Array Management Window online help.

The `hot_add` Utility Does Not Run

Operating System
NetWare

Problem or Restriction
On newly installed versions of NetWare, the `hot_add` utility does not run if it is not in the search path in the `autoexec.ncf` file.

Workaround
Edit the `autoexec.ncf` file, and add the following line to the search path:

```
search add sys:\
```

After adding the information to the search path, be sure to save the file, and restart the computer.

IMPORTANT If you have installed the `hot_add` utility to a directory other than the default, add the appropriate directory path to the `autoexec.ncf` file instead.

Installation Failure Occurs if the Array Management Window is Open

Problem or Restriction
If the Array Management Window is open while you install the SMutil package, the installation fails, and errors are displayed.

Workaround
Close all Array Management Windows before installing the SMutil package.

Snapshot Volume is Not Assigned a Drive Letter

Operating System
Windows Server 2003

Problem or Restriction
After enabling a snapshot volume and running the `hot_add` utility, the Windows Server 2003 operating system does not assign a drive letter to the snapshot volume.

Workaround
You must manually assign a drive letter to the snapshot volume using the Disk Administrator. For additional information about assigning drive letters using the Disk Administrator, refer to the operating system online help.

SMdevices Displays Disabled Snapshot Volumes

Operating System
Windows Server 2003

Problem or Restriction
The SMdevices utility displays snapshot volumes in its output, even if a snapshot volume is disabled.

Workaround
None. This is normal operation of the utility.

General Software Restrictions

The restrictions in this section apply to general software components. Many of these apply to operating system-specific restrictions that might affect your storage environment or setup.

Snapshot Requests Fail from VDS/VSS-Aware Applications

Operating System
Windows Server 2003

Problem or Restriction
If you try to create more than three snapshot copies by using Virtual Disk Service (VDS) or Virtual Shadow Copy Service (VSS) -aware applications, the snapshot request will fail.

Workaround
Do not request more than three snapshot copies at one time from VDS/VSS-aware applications.

Mismatched Patch Version Error Message When Using SAN Software Version 4.4.6 or 4.4.7

Operating System
Solaris

Problem or Restriction
After patching SAN Software to version 4.4.6 or 4.4.7, the following error message appears on all hosts during the boot process:

```
"Driver/Binary usr/lib/libsun_fc.so.1 does not match version
installed with Patch 113767-08 on ..."
```

Workaround
Use the following procedure to correct this error:

1. Make sure that you have the SAN 4.4.7 Software package downloaded to a location you can access.
2. As the host is booting, note the version number of the patch that is causing the error message.
3. Remove the patch with the `patchrm` command.
4. Manually install the patch matching that number from the downloaded SAN package.

File System Corruption Using LifeKeeper Cluster Software

Operating System
Linux

Problem or Restriction
File system corruption occurs when connecting two variants of Linux hosts (such as a Red Hat Enterprise Linux 3 host and a Red Hat Enterprise Linux 4 host) in a cluster by using SteelEye® LifeKeeper® cluster software.

Workaround
None. The hosts in the LifeKeeper cluster must use the exact same version of the operating system software (homogeneous).

LifeKeeper Might Try to Recover Resources during Controller Failover or Reboot

Problem or Restriction
A device undergoing a failover or reboot might not respond to a poll from the SteelEye LifeKeeper cluster software within its timeout threshold. If this happens, logical resource failovers could occur and LifeKeeper will mark the resource with an unresponsive status. LifeKeeper will then attempt recovery of the resources through another node until the resource or device responds.

Support Information
None.

Sun Cluster Node Fails to Completely Reboot

Operating System
Solaris 9

Problem or Restriction
A rebooting node running Sun Cluster 3.1 on the Solaris 9 operating system fails to complete the reboot process after several hours. The node continues trying to contact the alternate node to get permission to join the cluster, but the process is not successful. I/O operations continue to run on the alternate node, but cluster commands might start failing because of the failed node.

Workaround
Reboot both nodes. This is a known issue with Sun Cluster 3.1.

Cluster Resources Fail When a Host Comes Back Online after a Reboot

Operating System
Windows Server 2003

Problem or Restriction
Cluster resources fail on a host after rebooting if the “Failback immediately” option is checked in the Microsoft Cluster Server software. With this option checked, you might see failed resources or devices not associated with their preferred node.

Workaround
Make sure the “Prevent failback” option is checked under Group Properties. This is the default option. Do not use the “Failback immediately” option.

Cluster Node Configured with VERITAS DMP Takes a Long Time to Boot

Operating System
Windows 2000

Problem or Restriction
In a Microsoft Cluster Server (MSCS) environment configured with VERITAS DMP, rebooting a server node owning the resources takes a long time (10 to 15 minutes). It is possible that the long server boot time might cause some Windows delayed write failures.

Workaround
Before rebooting, move all the resources from the server node you are planning to reboot to another node.

Red Hat Enterprise Linux 4 Does Not Recognize 256 or More LUNs

Operating System
Linux 2.6 kernel

Problem or Restriction
Red Hat Enterprise Linux 4 does not correctly load the `scsi_mod` module parameters. Both scsi_mod module parameters must be on one line to enable customers to boot up 256 or more LUNS. If the `scsi_mod` module parameters are not on one line, the host returns:

```
ERROR: max_report_luns=256 exited abnormally!
```

Workaround
Manually remove the following line from the modules.conf file:

```
options scsi_mod max_luns=256
```

SteelEye LifeKeeper Cannot Be Installed On an IA64 Platform

Operating System
Linux 2.6 kernel

Problem or Restriction
The SteelEye Lifekeeper cluster software cannot be installed on a configuration with a host using an Intel Itanium processor (IA64 processor) and running the Red Hat Enterprise Linux 4 operating system. At this time, SteelEye does not plan to support LifeKeeper on the IA64 bit platform for Red Hat Enterprise Linux 4.

Workaround
None.

LifeKeeper Software Marks Devices as Not Responding during a System Reboot

Operating System
Linux

Problem or Restriction
SteelEye LifeKeeper software might show that applications or volumes have failed during a system reboot.

Workaround
Set the following parameters in `/etc/default/LifeKeeper`, and reboot the host:

```
LCMNUMHBEATS=15
LCMHBEATTIME=5
```

UDP Broadcast Address Must be Enabled on the Host for Auto Discovery to Work

Operating System
Linux

Problem or Restriction
The transmission of the User Datagram Protocol (UDP) broadcast message that starts auto discovery might be blocked on a host using the Linux operating system because of a configuration issue. The storage management software uses UDP at the broadcast IP address `255.255.255.255`. You must configure the system to allow broadcast messages for auto discovery to work.

Workaround
Manually add storage systems using their IP addresses or host names, or add the following command to enable the broadcast address on that Ethernet port and allow normal auto discovery to proceed:

```
route add -host 255.255.255.255 eth0
```

Mapped NSS Volume is Inaccessible after a Cluster Resource Migration

Problem or Restriction

A Novell Storage Services (NSS) volume mapped to a Windows client is rendered inaccessible after a “cluster leave” command on one cluster node caused the NSS pool to be migrated to the other cluster node.

Workaround

None.

Cluster Administrator Shows a Failed Disk as Online

Operating System

Windows Server 2003

Problem or Restriction

When in a server cluster environment, a failed disk in the storage system does not show as failed in the Cluster Administrator. The host still receives I/O errors, and the disk is shown as failed in the storage management software.

Workaround

This is a known issue with the operating system.

Choose one of the following workarounds:

- **Workaround 1 –** Reboot the host that owns the disk resource. The Cluster Administrator should now show the proper state of the disk resource.
- **Workaround 2 –** Change the ownership of the disk resource to a different node in the server cluster.

Disk Timeout Value is Changed during a Server Cluster Configuration

Operating System

Windows Server 2003

Problem or Restriction

The disk timeout value in the system registry is changed to an incorrect value after using the Microsoft Cluster Administrator to configure a server cluster.

Workaround

Complete the following procedure to change the disk timeout value to the correct value after the server cluster configuration is complete.

1. Select Start >> Run.
2. Type the following command, and press Enter:

   ```
   regedit
   ```

   The Registry Editor is displayed.
3. From the left pane of the Registry Editor, navigate to the following directory:

   ```
   HKEY_LOCAL_MACHINE\System\CurrentControlSet\Services\Disk
   ```
4. From the right pane of the Registry Editor, double-click TimeOutValue.

   The Edit DWORD Value dialog is displayed.
5. In the Value Data text area, type the following number, and press Enter:

   ```
   78
   ```
6. Under Base, select Hexadecimal.
7. Select OK.

End Of Procedure

Mapping Snapshot Volumes with Dynamic Disks Fails

Operating System
Windows Server 2003

Problem or Restriction
When you use Dynamic Disks on a Windows Server 2003 system, a snapshot volume cannot be mapped to the same host as the base volume.

Workaround
This is a known issue in the Windows Server 2003 operating system. If Dynamic Disks are not required, use Basic Disks.

Display Problems Occur with Windows XP

Operating System
Windows XP

Problem or Restriction
Some aspects of the storage management software appear different than expected when using the default Windows XP (display) Theme.

Workaround

For the expected appearance of the storage management software, switch to the Windows Classic Theme. Ensure that the storage management station software is not running when switching themes, because further display problems can occur until the software is restarted.

Complete the following procedure to switch to the Windows Classic Theme.

1. Close the storage management software.
2. Select Start >> Control Panel >> Appearance and Themes >> Display.
3. Select the Themes tab.
4. From the Theme drop-down menu, select Windows Classic.
5. Select OK.
6. Restart the storage management software.

End Of Procedure

Unable to Boot to a Windows Server 2003 Operating System in a Dual-Boot Environment with Windows XP

Operating System

Windows Server 2003 and Windows XP

Problem or Restriction

After upgrading Windows XP to Service Pack 1 (SP1), you are unable to boot to a Windows Server 2003 partition in a dual-boot environment.

Workaround

When running a dual-boot host, you must install the older operating system before you can install the newer operating system. In this instance, you should install Windows XP first, and then install Windows Server 2003. For additional information about installing Windows operating systems, refer to Microsoft’s web site.

Chapter 6

Usage Notes

Usage information, whenever possible, is included in the installation guides, user guides, or online help for a particular product. However, late-breaking usage information not available when a product's documentation was being written is included here for reference. This information will move to the appropriate document at the next major release of the product to which it refers. This chapter provides information for software and hardware products:

Information about which operating systems might be affected is kept with each applicable usage note.

Storage System Usage Notes

The usage notes in this section apply specifically to the hardware components in storage systems, including modules, CRUs, controllers, hard drives, and so on.

Event Log Shows Errors on a Drive that Have Been Removed

Problem or Restriction

When a drive fails, an event is logged, and the drive is removed from the drive module. Before the drive is replaced, the drive failure event might be repeated several times in the Event Log if other configuration changes are made to the storage system. Items that might trigger this event to be logged again include controller resets, major changes to the configuration, or modification of system settings.

Workaround

Replace the failed drive as soon as possible.

Relocating Individual Drives or a Drive Module is Not Recommended

CAUTION **Potential data loss –** Physically relocating drives and drive modules after they have been installed in a storage system and configured as part of a volume group can cause data loss. To avoid data loss, always consult Customer and Technical Support before you relocate drives or drive modules.

Problem or Restriction

Relocating drives or drive modules after they have been installed and configured is not recommended.

Workaround

If you need to relocate storage system components, call Customer and Technical Support for detailed instructions.

Host Adapter Usage Notes

The usage notes in this section apply specifically to host adapters.

LSI Logic HBAs Initialize as Public Loop

Problem or Restriction

When LSI Logic host bus adapters (HBAs) are connected to a Brocade 4-Gb switch, they might initialize as a public loop instead of a dedicated N/F port.

Workaround

None.

New HBAs and Switches Might Not Scan for LUNs Above Zero

Operating System

Linux

Problem or Restriction

It is necessary to force a new HBA or switch to discover LUNs that are greater than 0 (zero), because a new HBA or switch might not scan for LUNs above 0 (zero).

Workaround

To force the adapter to scan for LUNs greater than 0, add the following setting to the `/etc/modules.conf` file.

```
options scsi_mod max_scsi_luns=128
```

After adding the information, rebuild the initial ramdisk, and reboot to use the new setting. The procedure for rebuilding the initial ramdisk differs depending on the Linux distribution. Check your operating system help for more information.

SANtricity Storage Manager Usage Notes

The usage notes in this section apply specifically to SANtricity Storage Manager.

Snapshot Requests from VSS Aware Applications Fail under Heavy I/O Load

Problem or Restriction

If you try to create multiple snapshots using VSS aware applications while the controllers are under heavy I/O load, the snapshot creation requests might fail if the controller does not respond within five seconds.

Workaround

Quiesce all hosts before creating the snapshots.

“Automatic Controller Firmware Synchronization Failed” Error Message

Problem or Restriction

The Major Event Log (MEL) and SANtricity Storage Manager might show a critical error stating “automatic controller firmware synchronization failed,” even though the firmware levels are the same on both controllers.

Workaround

Verify that the firmware level is the same on both controllers. If it is, this error is erroneous and can be ignored.

Linux Systems Display an Error Message with the InstallAnywhere Installation

Operating System

Linux

Problem or Restriction

On a Linux IA64 Red Hat or SuSE system, when you install the software using InstallAnywhere, an error message is displayed before the SANtricity Storage Manager software launches. The following is the error message that is displayed.

```
Strings: /lib/libc.so.6: No such file or directory
```

Workaround

This is a vendor defect. The error message is annoying, but the installation is successful. The next message that is displayed is “Launching installer.”

Global Settings for SMTP Must Be Set before Sending Email Validation

Problem or Restriction

The storage management software returns an error message if you do not set the global settings for Simple Mail Transfer Protocol (SMTP) properly. You must close and then reopen the Configure Alerts dialog for the email validation to be successful.

Workaround

To set the SMTP Mail Server global settings, perform the following steps:

1. Select the Edit >> Configure Alerts >> Mail Server tab.
2. Type `ra.ks.lsil.com` , and press Enter.
3. Type an email address, and press Enter.
4. Click OK.

   The Configure Alerts dialog closes.
5. Reopen the Configure Alerts dialog.
6. Select the Edit >> Configure Alerts >> Mail Server tab again.
7. Select the email tab.
8. Type an email address to where the validation will be sent, and press Enter.
9. Click Validate.

   A validation successful message is displayed.

JAWS Screen Reader Must be Restarted to Load or Unload the New `.dll` Files

Problem or Restriction

The JAWS Screen Reader installs with a speech synthesizer that reads information aloud and outputs Braille displays for the visually impaired. Accessibility to the JAWS Screen Reader is not supported for versions prior to version 9.14 of the storage management software. In addition, navigation of this screen reader is different than other screen readers. For example, instead of using the up and down arrow keys to move between buttons, JAWS uses the Tab key.

Workaround

After installing the Java Access Bridge (JAB), you must restart the JAWS Screen Reader to load the new `.dll` files. After storage management software the JAB, you must restart the JAWS Screen Reader to unload the `.dll` files.

For more information about JAWS Screen Reader navigation, refer to the following Internet location:
`http://tlt.its.psu.edu/suggestions/accessibility/jaws.html`.

Software Packages Do Not Install Using the Installation Wizard

Operating System

Solaris

Problem or Restriction

The Installation Wizard might not install all the software packages on certain configurations of the Solaris operating environment. If the packages do not install correctly, the installation log file shows “non fatal error: interaction required.”

Workaround

Install the packages manually using the Solaris native packages available on the SANtricity Storage Manager Installation CD.

Volume-to-LUN Mappings are no Longer Preserved in Volume Group Relocations

Problem or Restriction

Volume-to-LUN mappings are lost when relocating volume groups from one storage system to another.

Workaround
Re-create the volume-to-LUN mappings manually after the volume groups have been relocated.

SCSI-3 Persistent Reservations are Not Supported

Operating System
AIX

Problem or Restriction
SCSI-3 Persistent Reservation is not supported on AIX with FAStT Storage Systems.

Workaround
None.

Volume Thrashing Possible in Configurations Using Snapshot Volumes and AVT

Problem or Restriction
Volume thrashing occurs if a base volume and its associated snapshot volume are mapped to different hosts, with a different preferred path assigned on each host.

In Figure 6-1 on page 6-7, the base volume is mapped to Host 1, and the snapshot volume is mapped to Host 2 (this is a typical setup when creating a point-in-time backup). Host 1 is set up with Controller A as the base volume's preferred path. Host 2 is set up with Controller B as the snapshot volume's preferred path.

In this case, each time Host 1 sends I/O to the base volume, the auto volume transfer (AVT) mechanism on the controller moves both the base volume and snapshot volume to Controller A. Then, for each I/O to the snapshot volume, the controller moves both the base volume and snapshot volume to Controller B. So, for every I/O, both the base volume and snapshot volume move. The result of this is thrashing that, depending on load, can lead to host hang, stale data, and data corruption.

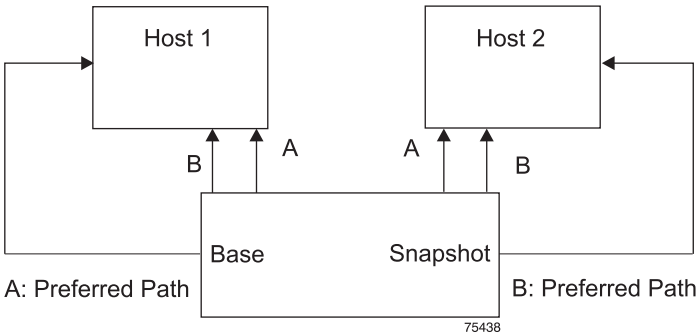

Figure 6-1 Invalid Preferred Path Configuration

Workaround

The base volume and snapshot volume must have the same preferred path when mapped to one or more hosts. The preferred path can be either Controller A or Controller B, as long as the base volume, snapshot volume, and snapshot repository volume are owned by the same controller, and the controller owner is configured as the preferred path from both hosts.

Failover Software Usage Notes

The usage notes in this section apply specifically to failover software.

Restoring I/O Path Failover after a Controller Replacement

Operating System

Windows 2000 and Windows Server 2003

Problem or Restriction

After controller failure occurs in an active/active controller pair, and the volumes are automatically transferred to the alternate controller, you can replace the failed controller and bring it online.

Until the volume group ownership is transferred back to the preferred controller owner, you lose I/O path failover protection for the second controller.

Workaround

It might take one minute after the failed controller is replaced and brought online for a re-scan to detect and return ownership to the preferred controller.

To ensure I/O path failover protection after replacing the controller, *immediately* select the Storage System >> Redistribute Volume Groups option in the Array Management Window to transfer volume groups back to preferred controller owners.

For more information, refer to the Array Management Window online help.

Increased Bandwidth Capabilities (Trunking) are Limited to the Type of Failover Driver

Operating Systems
Linux, Windows, and Solaris

Problem or Restriction
The capability to increase bandwidth (trunking) is possible only for failover drivers that have either trunking capabilities or that have the ability to use the Persistent Reservation Management method. Current failover drivers that are capable of trunking include MPP on Linux, MPP on Windows, Sun MPxIO, and VERITAS DMP 4.0. The Persistent Reservation Management method is only supported using MPP on Windows or Linux, and with the VERITAS Cluster Server with the VERITAS DMP solution on a Solaris machine. The following operating systems do *not* have the capability to increase bandwidth: HP-UX, NetWare, and IRIX.

Workaround
None (information only).

General Software Usage Notes

The usage notes in this section apply to general software problems. These notes apply to operating system-specific restrictions that might affect your storage environment or setup.

Operating System Restrictions for Frameworks Applications

Operating System
Windows 2000

Problem or Restriction
The IBM Director frameworks application software can only be used with the Windows 2000 operating system.

Workaround
None.

Volume Thrashing Possible if Using Snapshot Volumes in a Server Cluster Environment

Problem or Restriction

Volume thrashing occurs if a base volume and its associated snapshot volume are mapped to different hosts, with a different preferred path assigned on each host.

Workaround

For more information and a workaround, see “Volume Thrashing Possible in Configurations Using Snapshot Volumes and AVT” on page 6-6.

Index

T

U

V

W

Numerics