

Sun StorEdge™ Fast Write Cache 2.0 System Administrator's Guide



THE NETWORK IS THE COMPUTER™

Sun Microsystems, Inc.
901 San Antonio Road
Palo Alto, CA 94303-4900 USA
650 960-1300 Fax 650 969-9131

Part No. 806-2064-10
February 2000, Revision A

Send comments about this document to: docfeedback@sun.com

Copyright 2000 Sun Microsystems, Inc., 901 San Antonio Road • Palo Alto, CA 94303-4900 USA. All rights reserved.

This product or document is protected by copyright and distributed under licenses restricting its use, copying, distribution, and decompilation. No part of this product or document may be reproduced in any form by any means without prior written authorization of Sun and its licensors, if any. Third-party software, including font technology, is copyrighted and licensed from Sun suppliers.

Parts of the product may be derived from Berkeley BSD systems, licensed from the University of California. UNIX is a registered trademark in the U.S. and other countries, exclusively licensed through X/Open Company, Ltd. For Netscape Communicator™, the following notice applies: Copyright 1995 Netscape Communications Corporation. All rights reserved.

Sun, Sun Microsystems, the Sun logo, AnswerBook2, docs.sun.com, Sun StorEdge, Fast WriteCace, Instant Image, and Solaris [] are trademarks, registered trademarks, or service marks of Sun Microsystems, Inc. in the U.S. and other countries. All SPARC trademarks are used under license and are trademarks or registered trademarks of SPARC International, Inc. in the U.S. and other countries. Products bearing SPARC trademarks are based upon an architecture developed by Sun Microsystems, Inc.

The OPEN LOOK and Sun™ Graphical User Interface was developed by Sun Microsystems, Inc. for its users and licensees. Sun acknowledges the pioneering efforts of Xerox in researching and developing the concept of visual or graphical user interfaces for the computer industry. Sun holds a non-exclusive license from Xerox to the Xerox Graphical User Interface, which license also covers Sun's licensees who implement OPEN LOOK GUIs and otherwise comply with Sun's written license agreements.

RESTRICTED RIGHTS: Use, duplication, or disclosure by the U.S. Government is subject to restrictions of FAR 52.227-14(g)(2)(6/87) and FAR 52.227-19(6/87), or DFAR 252.227-7015(b)(6/95) and DFAR 227.7202-3(a).

DOCUMENTATION IS PROVIDED "AS IS" AND ALL EXPRESS OR IMPLIED CONDITIONS, REPRESENTATIONS AND WARRANTIES, INCLUDING ANY IMPLIED WARRANTY OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE OR NON-INFRINGEMENT, ARE DISCLAIMED, EXCEPT TO THE EXTENT THAT SUCH DISCLAIMERS ARE HELD TO BE LEGALLY INVALID.

Copyright 2000 Sun Microsystems, Inc., 901 San Antonio Road • Palo Alto, CA 94303-4900 Etats-Unis. Tous droits réservés.

Ce produit ou document est protégé par un copyright et distribué avec des licences qui en restreignent l'utilisation, la copie, la distribution, et la décompilation. Aucune partie de ce produit ou document ne peut être reproduite sous aucune forme, par quelque moyen que ce soit, sans l'autorisation préalable et écrite de Sun et de ses bailleurs de licence, s'il y en a. Le logiciel détenu par des tiers, et qui comprend la technologie relative aux polices de caractères, est protégé par un copyright et licencié par des fournisseurs de Sun.

Des parties de ce produit pourront être dérivées des systèmes Berkeley BSD licenciés par l'Université de Californie. UNIX est une marque déposée aux Etats-Unis et dans d'autres pays et licenciée exclusivement par X/Open Company, Ltd. La notice suivante est applicable à Netscape Communicator™ : Copyright 1995 Netscape Communications Corporation. All rights reserved.

Sun, Sun Microsystems, the Sun logo, AnswerBook2, docs.sun.com, Sun StorEdge, Fast WriteCace, Instant Image, et Solaris sont des marques de fabrique ou des marques déposées, ou marques de service, de Sun Microsystems, Inc. aux Etats-Unis et dans d'autres pays. Toutes les marques SPARC sont utilisées sous licence et sont des marques de fabrique ou des marques déposées de SPARC International, Inc. aux Etats-Unis et dans d'autres pays. Les produits portant les marques SPARC sont basés sur une architecture développée par Sun Microsystems, Inc.

L'interface d'utilisation graphique OPEN LOOK et Sun™ a été développée par Sun Microsystems, Inc. pour ses utilisateurs et licenciés. Sun reconnaît les efforts de pionniers de Xerox pour la recherche et le développement du concept des interfaces d'utilisation visuelle ou graphique pour l'industrie de l'informatique. Sun détient une licence non exclusive de Xerox sur l'interface d'utilisation graphique Xerox, cette licence couvrant également les licenciés de Sun qui mettent en place l'interface d'utilisation graphique OPEN LOOK et qui en outre se conforment aux licences écrites de Sun.

CETTE PUBLICATION EST FOURNIE "EN L'ETAT" ET AUCUNE GARANTIE, EXPRESSE OU IMPLICITE, N'EST ACCORDEE, Y COMPRIS DES GARANTIES CONCERNANT LA VALEUR MARCHANDE, L'APTITUDE DE LA PUBLICATION A REPOUDRE A UNE UTILISATION PARTICULIERE, OU LE FAIT QU'ELLE NE SOIT PAS CONTREFAISANTE DE PRODUIT DE TIERS. CE DENI DE GARANTIE NE S'APPLIQUERAIT PAS, DANS LA MESURE OU IL SERAIT TENU JURIDIQUEMENT NUL ET NON AVENU.



Regulatory Compliance Statements

Your Sun product is marked to indicate its compliance class:

- Federal Communications Commission (FCC) — USA
- Industry Canada Equipment Standard for Digital Equipment (ICES-003) - Canada
- Voluntary Control Council for Interference (VCCI) — Japan
- Bureau of Standards Metrology and Inspection (BSMI) — Taiwan

Please read the appropriate section that corresponds to the marking on your Sun product before attempting to install the product.

FCC Class A Notice

This device complies with Part 15 of the FCC Rules. Operation is subject to the following two conditions:

1. This device may not cause harmful interference.
2. This device must accept any interference received, including interference that may cause undesired operation.

Note: This equipment has been tested and found to comply with the limits for a Class A digital device, pursuant to Part 15 of the FCC Rules. These limits are designed to provide reasonable protection against harmful interference when the equipment is operated in a commercial environment. This equipment generates, uses, and can radiate radio frequency energy, and if it is not installed and used in accordance with the instruction manual, it 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 own expense.

Shielded Cables: Connections between the workstation and peripherals must be made using shielded cables to comply with FCC radio frequency emission limits. Networking connections can be made using unshielded twisted-pair (UTP) cables.

Modifications: Any modifications made to this device that are not approved by Sun Microsystems, Inc. may void the authority granted to the user by the FCC to operate this equipment.

FCC Class B Notice

This device complies with Part 15 of the FCC Rules. Operation is subject to the following two conditions:

1. This device may not cause harmful interference.
2. This device must accept any interference received, including interference that may cause undesired operation.

Note: This equipment has been tested and found to comply with the limits for a Class B digital device, pursuant to Part 15 of the FCC Rules. These limits are designed to provide reasonable protection against harmful interference in a residential 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. However, there is no guarantee that interference will not occur in a particular installation. If this equipment does cause harmful interference to radio or television reception, which can be determined by turning the equipment off and on, the user is encouraged to try to correct the interference by one or more of the following measures:

- Reorient or relocate the receiving antenna.
- Increase the separation between the equipment and receiver.
- Connect the equipment into an outlet on a circuit different from that to which the receiver is connected.
- Consult the dealer or an experienced radio/television technician for help.

Shielded Cables: Connections between the workstation and peripherals must be made using shielded cables in order to maintain compliance with FCC radio frequency emission limits. Networking connections can be made using unshielded twisted pair (UTP) cables.

Modifications: Any modifications made to this device that are not approved by Sun Microsystems, Inc. may void the authority granted to the user by the FCC to operate this equipment.

ICES-003 Class A Notice - Avis NMB-003, Classe A

This Class A digital apparatus complies with Canadian ICES-003.

Cet appareil numérique de la classe A est conforme à la norme NMB-003 du Canada.

ICES-003 Class B Notice - Avis NMB-003, Classe B

This Class B digital apparatus complies with Canadian ICES-003.

Cet appareil numérique de la classe B est conforme à la norme NMB-003 du Canada.

VCCI 基準について

クラス A VCCI 基準について

クラス A VCCI の表示があるワークステーションおよびオプション製品は、クラス A 情報技術装置です。これらの製品には、下記の項目が該当します。

この装置は、情報処理装置等電波障害自主規制協議会 (VCCI) の基準に基づくクラス A 情報技術装置です。この装置を家庭環境で使用すると電波妨害を引き起こすことがあります。この場合には使用者が適切な対策を講ずるよう要求されることがあります。

クラス B VCCI 基準について

クラス B VCCI の表示  があるワークステーションおよびオプション製品は、クラス B 情報技術装置です。これらの製品には、下記の項目が該当します。

この装置は、情報処理装置等電波障害自主規制協議会 (VCCI) の基準に基づくクラス B 情報技術装置です。この装置は、家庭環境で使用することを目的としていますが、この装置がラジオやテレビジョン受信機に近接して使用されると、受信障害を引き起こすことがあります。取扱説明書に従って正しい取り扱いをしてください。

BSMI Class A Notice

The following statement is applicable to products shipped to Taiwan and marked as Class A on the product compliance label.

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

- 1. Regulatory Compliance Statements iii**
 - FCC Class A Notice iii
 - FCC Class B Notice iii
 - ICES-003 Class A Notice - Avis NMB-003, Classe A iv
 - ICES-003 Class B Notice - Avis NMB-003, Classe B iv
 - BSMI Class A Notice v

- Preface xvii**
 - How This Book Is Organized xvii
 - Using UNIX Commands xviii
 - Shell Prompts xviii
 - Documentation Conventions xix
 - Related Documentation xx
 - Sun Documentation on the Web xx
 - Sun Welcomes Your Comments xxi

- 2. Overview 1-1**
 - Capabilities 1-1

- 3. Fast Write Cache Hardware 2-1**
 - Features 2-2

EMC Compliance	2-2
Error Detection and Correction Capability	2-2
Electrostatic Discharge Precautions	2-3
Qualified Platforms	2-3
Enabling the NVRAM Card Batteries	2-4
SBus NVRAM Card	2-4
▼ To Enable the SBus NVRAM Card Batteries	2-5
PCI NVRAM Card	2-7
▼ To Enable the PCI NVRAM Card Batteries	2-7
Installing the NVRAM Cards	2-8
▼ To Install the SBus NVRAM Cards	2-8
▼ To Install the PCI NVRAM Cards	2-10
Configuration Data	2-11

4. Using the Fast Write Cache Management Console 3-1

Overview	3-1
Installation	3-1
Configuration	3-3
Control	3-6
Health	3-8
Volumes	3-12
Alarm Viewer	3-16
Log Viewer	3-17

5. Using the Command Line Interface 4-1

Reading Fast Write Cache Man Pages	4-1
CLI Compatibility	4-2
Using the CLI	4-3
▼ To Start the Command Line Interface Shell	4-3

Command Summary	4-4
Using the cache Commands	4-6
Command Examples	4-7
▼ To start the cache	4-7
▼ To stop the cache	4-7
▼ To change the number of threads	4-7
▼ To change the cache memory size	4-7
▼ Checking the overall health	4-8
▼ Checking the volume health	4-8
▼ Changing the cache read/write options on all volumes	4-8
▼ Changing the cache read/write options on a single volume	4-9
▼ List current statistics of the overall health	4-9
▼ Changing the data acquisition interval	4-9
Data Service Command <code>fwcadm</code>	4-10
<code>fwcadm cache</code> Command	4-10
Syntax	4-10
Parameters	4-11
<code>fwcadm cache -s</code> Display	4-11
<code>fwcadm nvram</code> Command	4-16
Syntax	4-16
Options	4-16
Example	4-16
<code>fwcadm volume</code> Command	4-21
Syntax	4-21
Options	4-21
<code>sd.cf</code> Configuration File	4-22
<code>sv.cf</code> Configuration File	4-23

6. Normal Operations 5-1

Removing Fast Write Cache Hardware 5-2

▼ To Remove an NVRAM Card 5-2

Disabling the Batteries on the NVRAM Card 5-3

▼ To Disable the Batteries on the SBus NVRAM Cards 5-3

▼ To Disable the Batteries on the PCI Bus NVRAM Cards 5-4

Selecting the File Systems and Volumes to Cache 5-4

▼ To Enable Caching for a File System 5-6

▼ To Enable Caching for a Volume 5-6

▼ To Display Cache Statistics 5-7

Removable Disk Packs 5-7

▼ To Remove Disk Packs 5-7

7. Abnormal Operations 6-1

Unclean and Clean Shutdowns 6-2

 Booting a Different Kernel 6-2

 Cache `redevi`d 6-2

Moving an NVRAM Card Containing Data 6-3

Handling Disk Failures 6-4

 Recovering Trapped Data After a Disk Failure 6-5

 Clearing the Offline Device State 6-5

A. Theory of Operation A-1

Introduction A-2

Buffer Management A-2

Fast Write Cache Modes A-3

Recovery A-4

Hints for Safe Operation A-4

B. Error Messages B-1

C. SunVTS Diagnostic Software C-1

SunVTS Overview C-2

Preparing to Run SunVTS C-3

▼ To Prepare to Run SunVTS C-4

Running SunVTS C-4

▼ To Invoke SunVTS Locally (OpenWindows™) C-4

▼ To Invoke SunVTS Locally (CDE) C-5

▼ To Invoke SunVTS Remotely C-5

Glossary Glossary-1

Index Index1

Figures

- FIGURE 2-1 SBus NVRAM Card 2-5
- FIGURE 2-2 NVRAM Card Showing Location of Jumpers and Jumper Pins 2-6
- FIGURE 2-3 Detail of Jumpers and Jumper Pins With Batteries Disabled 2-6
- FIGURE 2-4 Detail of Jumpers and Jumper Pins With Batteries Enabled 2-6
- FIGURE 2-5 PCI NVRAM Card 2-7
- FIGURE 3-1 Sun StorEdge Management Splash Screen 3-2
- FIGURE 3-2 Fast Write Cache Management Console 3-2
- FIGURE 3-3 Configuration Window 3-3
- FIGURE 3-4 Cache Disable and Re-enable Dialog 3-5
- FIGURE 3-5 Control Panel 3-7
- FIGURE 3-6 Stop Cache Confirmation 3-8
- FIGURE 3-7 Health Panel 3-9
- FIGURE 3-8 Cache Statistics Graph 3-11
- FIGURE 3-9 Volume Configuration Window 3-13
- FIGURE 3-10 Applying Volume Configuration 3-14
- FIGURE 3-11 Volume Health 3-15
- FIGURE 3-12 Alarm Viewer 3-17
- FIGURE 3-13 Log viewer 3-18
- FIGURE 5-1 Jumpers and Jumper Pins With Batteries Disabled 5-3

Tables

TABLE P-1	Shell Prompts	xviii
TABLE P-2	Documentation Conventions	xix
TABLE P-3	Related Documentation	xx
TABLE 4-1	Command Parameters	4-5
TABLE 4-2	<code>fwcadm cache -s</code> Display Options	4-11
TABLE 4-3	Device State Description	4-17
TABLE 4-4	SBus NVRAMBit Description	4-18
TABLE 4-5	Dirty LED States	4-18
TABLE 4-6	Battery Status Description	4-19
TABLE B-1	Known Error Conditions	B-1
TABLE C-1	SunVTS Documentation	C-2

Preface

The *Sun StorEdge Fast Write Cache 2.0 System Administrator's Guide* describes how to install the Sun StorEdge™ Fast Write Cache hardware and operate Fast Write Cache 2.0 for Sun workstations and servers.

All of the information that you need to install and operate the NVRAM hardware and Fast Write Cache software is included in this guide, installation guide, the online man pages, the Release Notes, and the Service Manual that came with your system. Have the *SunVTS User's Guide* ready for reference when running diagnostics at the end of your installation.

The procedures in this manual are for system or network administrators experienced in installing similar hardware and software in systems running the Solaris™ operating environment.

How This Book Is Organized

Chapter 1 contains an overview of the Fast Write Cache product.

Chapter 2 describes the SBus NVRAM card, provides compliance and safety information, and describes how to enable the batteries and install the NVRAM card.

Chapter 3 describes how to use the Sun StorEdge Management Console

Chapter 4 describes how to use the `cache` and Data Services commands.

Chapter 5 describes how to manage the Fast Write Cache under normal operating conditions.

Chapter 6 describes how to manage the Fast Write Cache under abnormal operating conditions.

Appendix A describes the operational theory of the Fast Write Cache.

Appendix B describes some of the error messages, notices, and solutions.

Appendix C provides an overview of the SunVTS™ diagnostic tool and describes how to use it with NVRAM cards.

Glossary defines terms used throughout this manual.

Using UNIX Commands

This document may not contain information on basic UNIX® commands and procedures such as shutting down the system, booting the system, and configuring devices.

Refer to one or more of the following for this information:

- *Solaris Handbook for Sun Peripherals*
- AnswerBook™ online documentation for the Solaris operating environment
- Other software documentation that you received with your system

Shell Prompts

TABLE P-1 Shell Prompts

Shell	Prompt
C shell	<i>machine_name%</i>
C shell superuser	<i>machine_name#</i>
Bourne shell and Korn shell	\$
Bourne shell and Korn shell superuser	#

Documentation Conventions

TABLE P-2 Documentation Conventions

Typeface	Meaning	Examples
AaBbCc123	The names of commands, files, and directories; on-screen computer output	Edit your <code>.login</code> file. Use <code>ls -a</code> to list all files. % You have mail.
AaBbCc123	What you type, when contrasted with on-screen computer output	% su Password:
<i>AaBbCc123</i>	Book titles, new words or terms, words to be emphasized	Read Chapter 6 in the <i>User's Guide</i> . These are called <i>class</i> options. You <i>must</i> be superuser to do this.
	Command-line variable; replace with a real name or value	To delete a file, type <code>rm filename</code> .
{arg arg}	In syntax, braces and pipes indicate that one of the arguments must be specified.	In this example, one of the options must be specified with the <code>fwcadm cache</code> command. <code>fwcadm cache {-d -e}</code>
command(<i>n</i>)	The form <code>command(number)</code> , where the number in parentheses ranges from 1 through 6 and is followed by letters, indicates the presence of an online reference man page.	<code>fwcadm(1FWC)</code>

Related Documentation

TABLE P-3 Related Documentation

Application	Title	Part Number
Release	<i>Sun StorEdge Fast Write 2.0 Cache Release Notes</i>	806-3165
Release	<i>Sun StorEdge Fast Write 2.0 Cache Installation Guide</i>	806-
Configuration	<i>Sun StorEdge Fast Write 2.0 Cache Configuration Guide</i>	806-xxxx
Installation	<i>fwcache man page</i>	N/A
Installation	<i>Solaris 2.6 Handbook for SMCC Peripherals</i>	802-7295
Installation	<i>Solaris Handbook for SMCC Peripherals</i>	805-7404
Installation	<i>Solaris 2.x Handbook for SMCC Peripherals</i>	801-5488
Options	<i>OpenBoot 3.x Command Reference Manual</i>	802-3242
Diagnostics (Solaris 2.6)	<i>SunVTS 2.1 User's Guide</i>	802-7299
Diagnostics (Solaris 2.6)	<i>SunVTS 2.1.3 Test Reference Manual</i>	805-4163
Diagnostics (Solaris 2.6)	<i>SunVTS 2.1 Quick Reference</i>	802-7301
Diagnostics (Solaris 7)	<i>SunVTS 3.0 User's Guide</i>	805-4442
Diagnostics (Solaris 7)	<i>SunVTS 3.0 Test Reference Manual</i>	805-4443
Diagnostics (Solaris 7)	<i>SunVTS 3.0 Quick Reference</i>	805-4444
Diagnostics (Solaris 7)	<i>SunVTS 3.1 User's Guide</i>	805-7406
Diagnostics (Solaris 7)	<i>SunVTS 3.1 Test Reference Manual</i>	805-7407
Diagnostics (Solaris 7)	<i>SunVTS 3.1 Quick Reference</i>	805-7408

Sun Documentation on the Web

The `docs.sun.comsm` web site enables you to access Sun technical documentation on the Web. You can browse the `docs.sun.com` archive or search for a specific book title or subject at:

<http://docs.sun.com>

Sun Welcomes Your Comments

We are interested in improving our documentation and welcome your comments and suggestions. You can email your comments to us at:

docfeedback@sun.com

Please include the part number of your document in the subject line of your email.

Overview

This chapter contains an overview of Fast Write Cache.

Fast Write Cache reduces the frequency of disk I/O access by caching the written data blocks in nonvolatile memory and then destaging the cached data to disk asynchronously.

Note – Fast Write Cache is implemented using a pair of NVRAM cards installed on your system.

Capabilities

Fast Write Cache provides the following capabilities:

- Accelerated NFS™ server performance
 - Write cache for synchronous I/O operations
 - Synchronous writes at memory speed
 - Synchronous write requests to disk are intercepted, and the data is stored in nonvolatile memory
- Accelerated mail server performance
 - Significantly improves performance of small file writes on UFS
- Reliable, high-speed performance
 - Reliable write caching
 - Accelerate write performance of UFS
 - Accelerated database logs
 - Accelerate ISP logging performance
 - Write coalescing improves small sequential write performance and allows full stripe writes for RAID 5

- Accelerate write performance of UFS
- Accelerate ISP logging performance
- Caching is definable at the volume level
 - File systems and raw volumes can be cached
- Support of Sun Dynamic Reconfiguration (DR)

Sun Dynamic Reconfiguration allows NVRAM boards to be removed or added to a system without having to power down or reboot. For information on using the DR commands, refer to `cfgadm(1M)` for the Exx0 servers and `dr(1M)` for the E10000 server.

Before attempting any dynamic reconfiguration operations, ensure that dynamic reconfiguration is enabled on your system. Refer to your system-specific dynamic reconfiguration documentation. See “Related Documentation” in the *Fast Write Cache Release Notes* for the part numbers of these documents or go to the dynamic reconfiguration web site at <http://sunsolve5.sun.com/sunsolve/Enterprise-dr>. This web site also contains patches and updates.

The following two lines must be added to the `/etc/system` file to enable dynamic reconfiguration of the NVRAM cards on a 10000:

```
set dr:detach_safe_list1="nvmem"
set hswp:suspend_safe_list1="nvmem"
```

For Fast Write Cache, the cache controller software must be disabled with `fwcadm(1M)` before performing dynamic reconfiguration operations. After successfully disabling the cache, you can perform dynamic reconfiguration operations as documented in the appropriate Dynamic Reconfiguration User’s Guide for your system. The following is an outline of a typical sequence of commands to replace a board:

```
# fwcadm cache -d
  Type DR commands to detach NVRAM boards
  (replace the board)
  Type DR command to attach NVRAM boards
# fwcadm cache -e
```

Fast Write Cache Hardware

The Fast Write Cache software is implemented using SBus or PCI nonvolatile random access memory (NVRAM) cards. The SBus NVRAM is a single-wide SBus memory card with battery backed-up error detection and correction (EDC) static memory. The PCI NVRAM is a full-length PCI memory card with battery backed-up error detection and correction (EDC) static memory. A maximum of one pair of cards is supported per system.

Nonvolatile memory remains intact even if the system loses power. Therefore, no data captured by the NVRAM is lost if the system goes down unexpectedly (during a power outage, for example).

The data stored on one NVRAM card is mirrored on the second NVRAM card preventing the failure of one NVRAM card from causing the loss of data.

This chapter contains the following topics:

- Features
- EMC Compliance
- Error Detection and Correction Capability
- Electrostatic Discharge Precautions
- Enabling the NVRAM Card Batteries
- Installing the NVRAM Cards
- Configuration Data

Features

Following is a list of the SBus NVRAM card features:

- Maximum power dissipation less than 12 watts
- Error detection and correction
- Battery backed-up NVRAM
- Battery redundancy for increased data integrity during power failure
- Data transfer rates of up to 108 Mbytes per second

Following is a list of the PCI NVRAM card features:

- Maximum power dissipation less than 12 watts
- Error detection and correction
- Battery backed-up NVRAM
- Battery redundancy for increased data integrity during power failure
- 64bit, 33 MHz PCI bus compliant
- Data transfer rates of up to 90 Mbytes per second

EMC Compliance

The NVRAM card has been tested for electromagnetic compatibility (EMC) compliance (USA FCC, Canada DOC, Japan VCCI, and EU CE Mark). When installed and operated in accordance with this guide, the EMC class marked on your Sun™ server label remains the same.

Error Detection and Correction Capability

Data is stored in 64-bit double long words, with 8 check bits for each double long word. This gives the NVRAM card the capability to detect all double-bit errors and correct all single-bit errors in a single 64-bit double long word. The NVRAM card cannot detect errors of 3 or more bits within the same nibble.

Because the check bits are generated and checked for the entire 64-bit data word, all writes to memory that are less than 8 bytes wide require a read-modify-write cycle to memory. The following occurs in the read-modify-write cycle:

1. Existing 64 data bits and 8 check bits are read.
2. The bytes to be written are merged into the data.
3. New check bits are generated.
4. The resulting 64 data and 8 check bits are written back to memory.

Electrostatic Discharge Precautions

Circuit board components are vulnerable to damage by electrostatic discharge (ESD). An electrostatic charge can build up on the human body and discharge when you touch a board. Such discharge can be produced by walking across a carpet and touching a board, or by other similar causes. Touch a conductive surface of the chassis or other element connected to common earth ground to discharge the static electricity present in your body.

To minimize the risk of ESD damage:

- Handle the card by the edges only
- Store the card in the antistatic bag provided
- Use a grounding strap whenever you work on a circuit board

Qualified Platforms

The following platforms are qualified for the Sbus and PCI NVRAM cards.

SBus NVRAM	PCI NVRAM
Sun Enterprise™ 3000 and 3500	Sun Enterprise 250
Sun Enterprise 4000 and 4500	Sun Enterprise 450
Sun Enterprise 5000 and 5500	
Sun Enterprise 6000 and 6500	
Sun Enterprise 10000	

Enabling the NVRAM Card Batteries

The NVRAM card is shipped with the batteries disabled to prevent battery drainage during shipment.



Caution – Handle the NVRAM card with care to avoid damage. Wear an antistatic wrist strap when handling the card. Hold the card by its edges only, and do not touch any of the components on the card.

Do not place the NVRAM card on a metal surface, as it may discharge the batteries regardless of the battery jumper settings. If you must put the card down, place it in its antistatic bag.

Note – You must enable the batteries using the following procedure before you install the hardware.

SBus NVRAM Card

The SBus NVRAM card uses three permanently mounted lithium batteries that protect data when the system power is off. These batteries last at least ten years in a running system or in storage with the batteries disabled. The batteries last approximately 18 months in a powered down system or in storage with the batteries enabled.



Caution – There are Eagle Picher Model # LTC16M lithium batteries on the SBus NVRAM card. Batteries are not customer replaceable parts. They may explode if mishandled. Do not dispose of the battery in fire. Do not disassemble it or attempt to recharge it. If a battery fails, call your authorized Sun service provider.

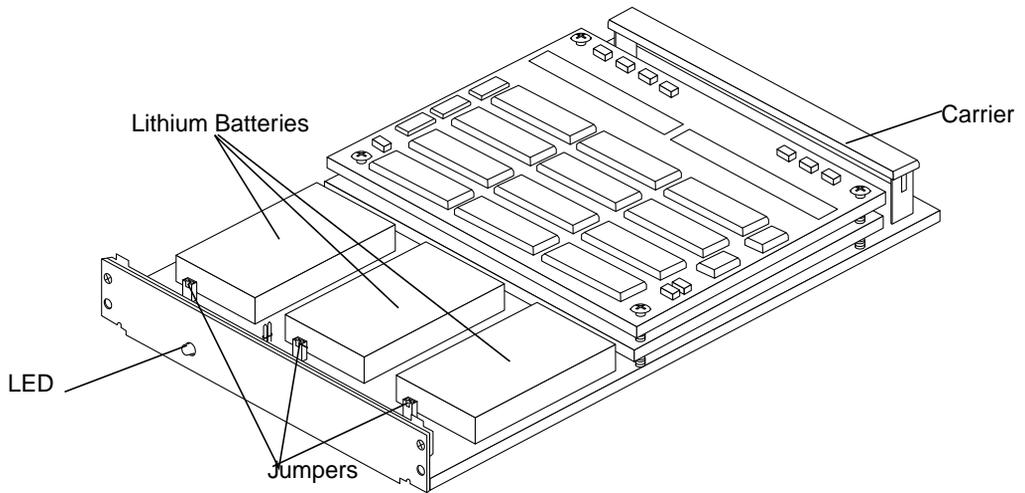


FIGURE 2-1 SBus NVRAM Card

▼ To Enable the SBus NVRAM Card Batteries

- 1. Attach the antistatic wrist strap's adhesive copper strip to the metal casing of the power supply. Wrap the other end twice around your wrist, with the adhesive side against your skin.**
- 2. Remove the card from the antistatic bag. Hold the card by its edges and place it, component side up, on the bag.**
Use antistatic bags as a protective cushion for the SBus NVRAM card.
- 3. Look for the battery jumpers.**
The SBus NVRAM holds three batteries. Each has its own jumper, and all three batteries must be activated.

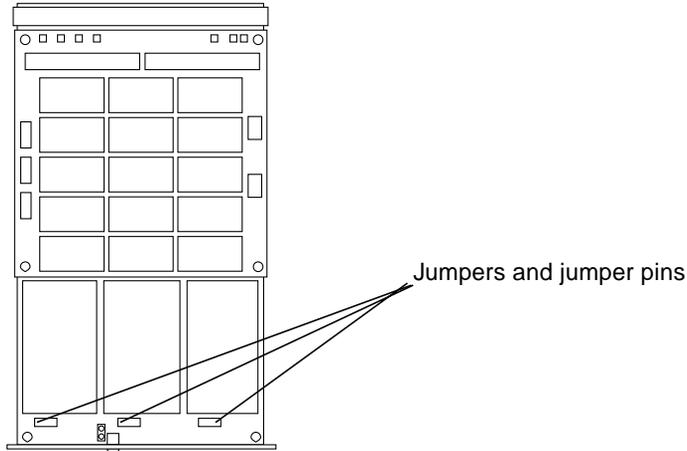


FIGURE 2-2 NVRAM Card Showing Location of Jumpers and Jumper Pins

4. Put the jumpers on the jumper pins.

When you unpack the card, the jumpers are inserted on only one pin of each pair, leaving the battery disabled.



FIGURE 2-3 Detail of Jumpers and Jumper Pins With Batteries Disabled

a. Use tweezers or needle-nosed pliers to lift each jumper off the single pin and align it with both pins.

b. Press the jumper in place on top of the pair of jumper pins.



FIGURE 2-4 Detail of Jumpers and Jumper Pins With Batteries Enabled

Note – Be sure to *disable* the batteries if you remove the NVRAM card from the system and store it. Failure to do so drains the batteries and severely reduces their backup time.

PCI NVRAM Card

The PCI NVRAM card uses four permanently mounted lithium batteries that protect data when the system power is off. These batteries last at least ten years in a running system or in storage with the batteries disabled. The batteries last approximately 34 months in a powered down system or in storage with the batteries enabled.



Caution – There are Eagle Picher Model # LTC16M lithium batteries on the PCI NVRAM card. Batteries are not customer replaceable parts. They may explode if mishandled. Do not dispose of the battery in fire. Do not disassemble it or attempt to recharge it. If a battery fails, call your authorized Sun service provider.

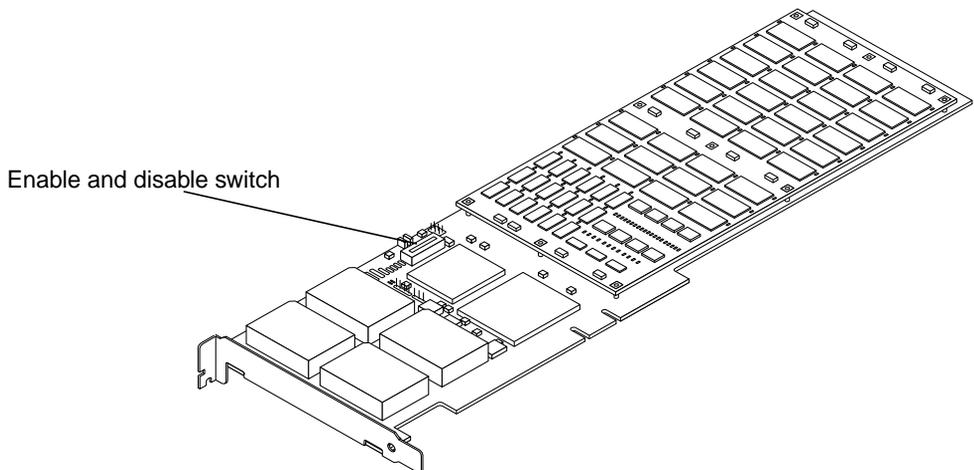


FIGURE 2-5 PCI NVRAM Card

▼ To Enable the PCI NVRAM Card Batteries

1. Attach the antistatic wrist strap's adhesive copper strip to the metal casing of the power supply. Wrap the other end twice around your wrist, with the adhesive side against your skin.
2. Remove the card from the antistatic bag. Hold the card by its edges and place it, component side up, on the bag.
Use antistatic bags as a protective cushion for the PCI NVRAM card.
3. Look for the **Enable and Disable** switch.

The PCI NVRAM holds four batteries. There is one switch to enable and disable all the batteries.

4. **Move the switch from Disable to Enable. Make sure that you pull or push the switch from the middle of the switch bar to ensure that the switch enables or disables all batteries.**

When you unpack the card, the switch is set to Disable.

Note – Be sure to *disable* the batteries if you remove the PCI NVRAM card from the system and store it. Failure to do so drains the batteries and severely reduces their backup time.

Installing the NVRAM Cards

You must enable the NVRAM card batteries before installing the cards. Make sure you have performed the procedure in the previous section before continuing.

Note – To use Fast Write Cache, a pair of SBus or PCI NVRAM cards must be installed on your system.

If your server contains more than one I/O card, install each NVRAM card on a different I/O card or bus to prevent the failure of a single I/O card from rendering the cache unavailable during a system recovery.

▼ To Install the SBus NVRAM Cards

Note – Refer to your system installation or service manual for detailed instructions for the following tasks.

1. **Power off your system.**
2. **Open the system unit.**
3. **Attach the antistatic wrist strap's adhesive copper strip to the metal casing of the power supply. Wrap the other end twice around your wrist, with the adhesive side against your skin.**
4. **Identify the slot in which you want to insert the NVRAM card.**

5. Disconnect the SBus filler panel from the slot that you selected.
6. Position the NVRAM on the system board over the SBus socket.
7. Carefully insert the NVRAM into the SBus socket.
The carrier is keyed so the NVRAM can be installed only one way.
8. Push the NVRAM into the socket until it is properly seated.
9. Detach the wrist strap, replace the unit cover, and power on the system.
10. Install the Fast Write Cache software. Refer to the *Sun StorEdge Fast Write Cache 2.0 Installation Guide* for installation instructions.

▼ To Verify the Installation

Note – The Fast Write Cache software must be installed before verifying the installation of the NVRAM cards. See the *Sun StorEdge Fast Write Cache 2.0 Installation Guide* for the Fast Write Cache installation instructions.

- As superuser, type the following command to verify that the NVRAM cards are properly installed.

```
# /usr/opt/SUNWesm/sbin/fwcdm nvram -s
Device      Size  Inst  Parent  Status  Dirty  Battery  UCE  CE  Mirror
-----
nvram0      32MB  #0   sbus #0  Active   0      Good     0    0  nvram1
nvram1      32MB  #1   sbus #0  Active   0      Good     0    0  nvram0
```

The system returns the status of all cards installed including the total amount of memory in Mbytes, the dirty bit status, the battery status, the ECC errors, and the mirror configuration.

Once the cards are installed it is transparent to you unless problems arise or the batteries become low.

See Chapter 4 for a detailed description of the `fwcdm nvram -s` display.

▼ To Install the PCI NVRAM Cards

Note – Refer to your system installation or service manual for detailed instructions for the following tasks.

1. **Power off your system.**
2. **Open the system unit.**
3. **Attach the antistatic wrist strap's adhesive copper strip to the metal casing of the power supply. Wrap the other end twice around your wrist, with the adhesive side against your skin.**
4. **Holding the PCI card by the edges, unpack it and place it on an antistatic surface.**
5. **Identify the slot number in which you want to insert the PCI NVRAM card.**
6. **Remove the PCI filler panel from the slot that you selected.**
7. **Holding the PCI card by the edges, align the card edge connector with the PCI slot. Slide the card face plate into the small slot at the end of the PCI opening.**
8. **Applying even pressure at both corners of the card, push the PCI card until it is firmly seated in the slot.**
9. **Screw the card in place.**
10. **Detach the wrist strap, replace the unit cover, and power on the system.**
11. **Install the Fast Write Cache software. Refer to the *Sun StorEdge Fast Write Cache 2.0 Installation Guide* for installation instructions.**

▼ To Verify the Installation

Note – The Fast Write Cache software must be installed before verifying the installation of the NVRAM cards. See the *Sun StorEdge Fast Write Cache 2.0 Installation Guide* for the Fast Write Cache installation instructions.

- As superuser, type the following command to verify that the NVRAM cards are properly installed.

```
# fwcadm nvram -s
Device      Size  Inst  Parent  Status  Dirty  Battery  UCE  CE  Mirror
-----
nvram0     64MB  #0   pci #0  Active   0      Good     0    0  nvram1
nvram1     64MB  #1   pci #0  Active   0      Good     0    0  nvram0
```

The system returns the status of all cards installed including the total amount of memory in Mbytes, the dirty bit status, the battery status, the ECC errors, and the mirror configuration.

Once the cards are installed it is transparent to you unless problems arise or the batteries become low.

See Chapter 4 for a detailed description of the `fwcadm nvram -s` display.

Configuration Data

The operating system keeps configuration data on the NVRAM card to keep track of information such as the host ID, mirror pairs (primary and secondary), and the number of NVRAM cards in the machine. During boot time, the operating system uses this information to verify the NVRAM configuration.

Using the Fast Write Cache Management Console

This chapter describes the operation of the Fast Write Cache Management Console graphical user interface.

The Fast Write Cache Management Console has tab panes for cache management.

- Configuration
- Control
- Health

Overview

Installation

After the Fast Write Cache Management Console has been installed, as superuser, bring up the GUI with the following command:

```
/usr/opt/SUNWesm/bin/esm_gui
```

An installation window appears for a few seconds before the Management Console.



FIGURE 3-1 Sun StorEdge Management Splash Screen

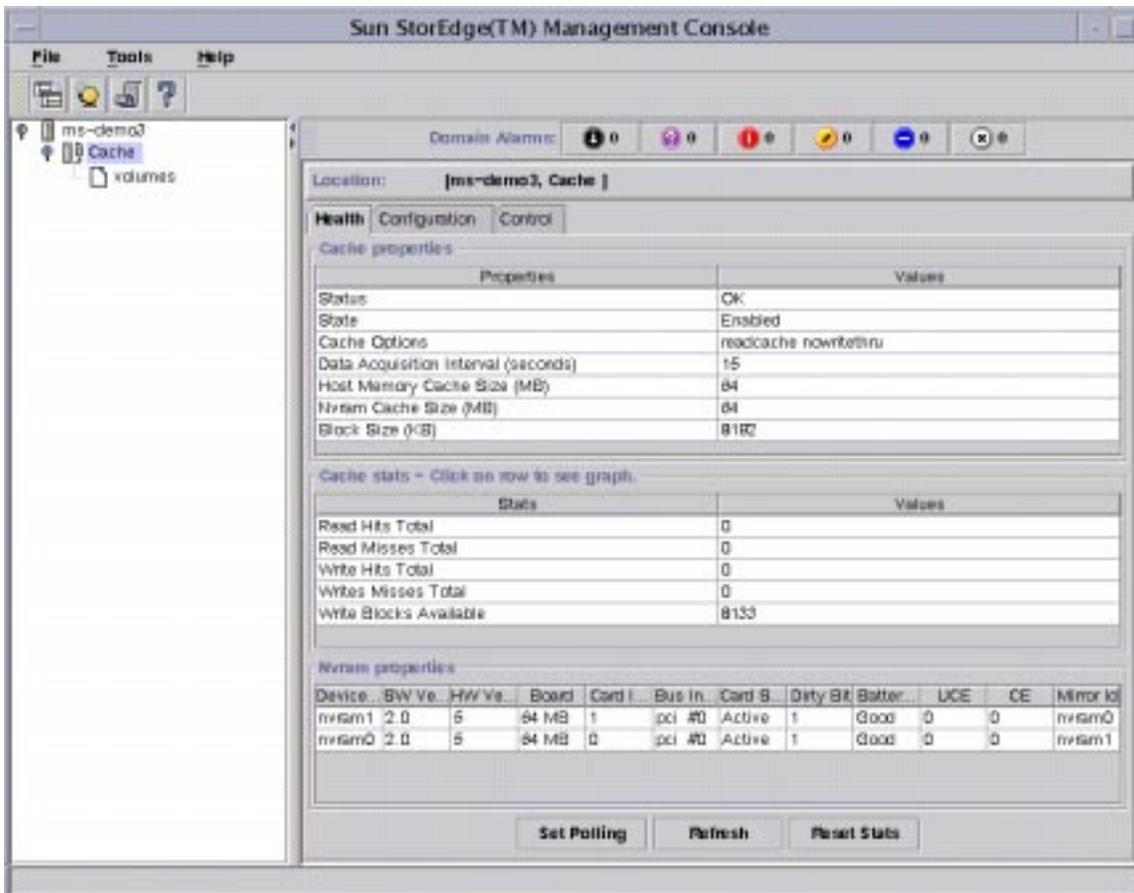


FIGURE 3-2 Fast Write Cache Management Console

When the Console appears, the Cache option in the menu bar to the left is selected. You have three primary tabs to work through:

- Health
- Configuration
- Control

The first time you bring up the Console you may want to make changes to the configuration.

Configuration

You can change the cache size, the number of threads, or the read and write options.

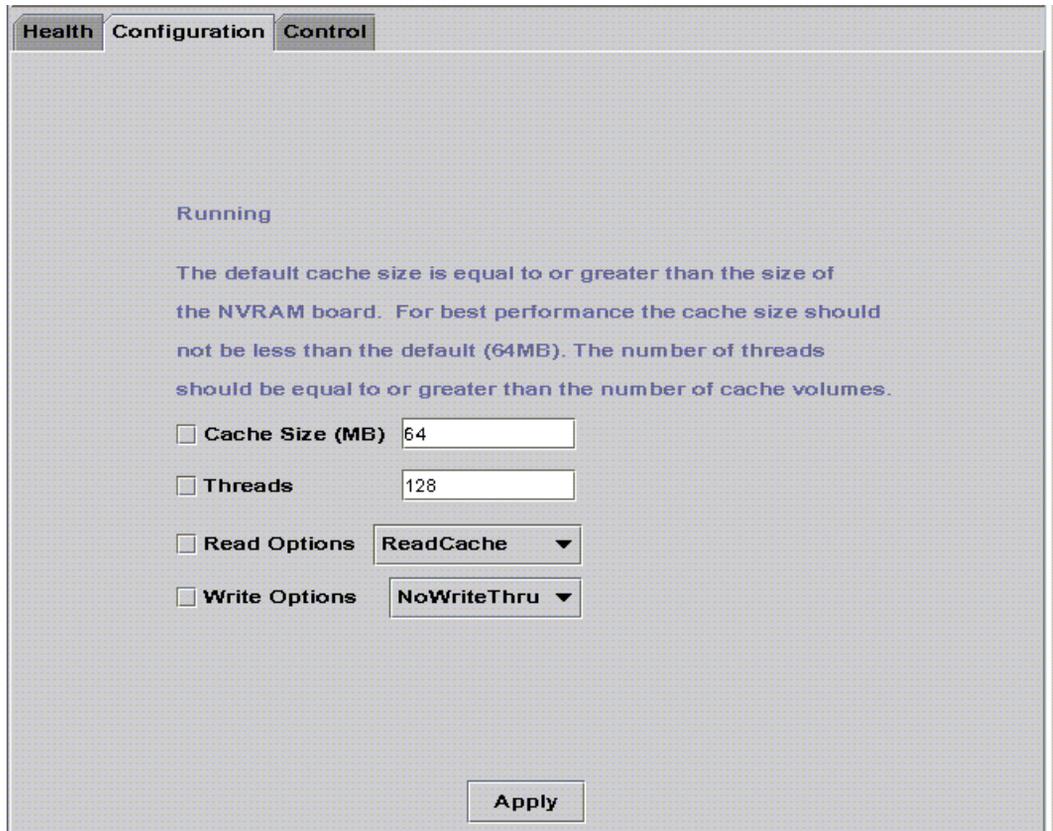


FIGURE 3-3 Configuration Window

Cache Size

Note – The Cache Size denotes the amount of host memory that is used to back the NVRAM and to satisfy read hits. If Cache Size is less than the nominal size of an NVRAM card (32 Mbytes or 64 Mbytes), write performance will be degraded. If Cache Size is greater than the NVRAM card size, read hits may be improved, but write performance will not be affected.

- **In the Cache Size window, change the cache size by typing in the number of Mbytes you want between 8 and 64.**

The number must be in multiples of 8. This number denotes the amount of host memory used to back the NVRAM and to satisfy read hits. If Cache Size is less than the nominal size of a single NVRAM card (32Mb or 64Mb), write performance will be degraded. If Cache Size is greater than the size of the NVRAM card, read hits may be improved, but write performance will not be affected.

Note – If you change the cache size, a dialog box informs you that you must re-enable the cache. See FIGURE 3-4.

To disable and re-enable the cache, you must go to the Control tab. See “Control” later in this chapter.

Threads

Threads are separate execution paths. In FWC, each cached volume gets its own thread, which allows multiple volumes to be cached at the same time.

If you add more than 128 cache volumes, you must change the number of threads used by FWC. Use the Configuration menu to change the thread value and the Control menu to disable and re-enable the cache.

- **In the Threads window, type in the number of threads you want between 128 and 1024.**

Specify a number of threads that is equal to or greater than the number of cached volumes. Threads in excess of the number of cached volumes are not used.

Note – If you change the number of threads, a dialog box informs you that you must re-enable the cache. See FIGURE 3-4.

To disable and re-enable the cache, you must go to the Control tab. See “Control” later in this chapter.

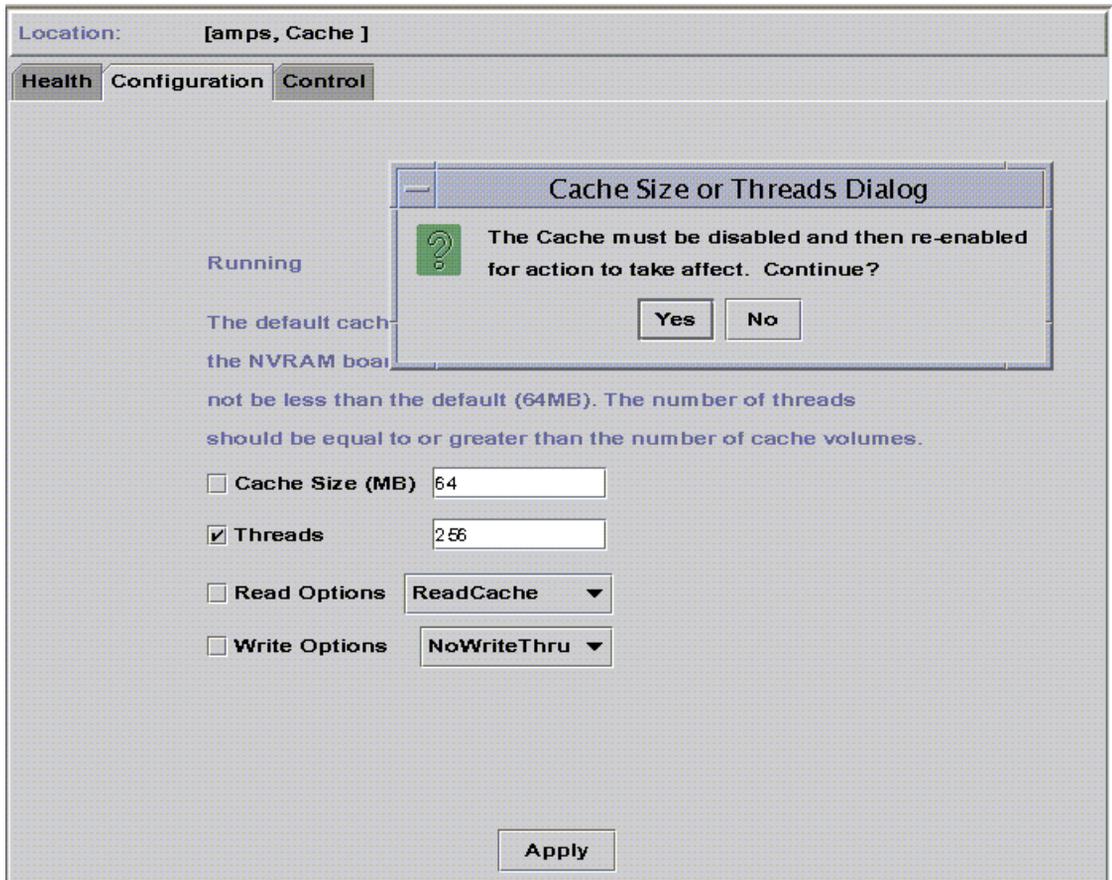


FIGURE 3-4 Cache Disable and Re-enable Dialog

Read/Write Options

The default overall cache options are ReadCache/NoWriteThru. Cache options for each volume are also available (see “Volumes”). You might want to turn off the caching on a volume (NoReadCache/WriteThru) to compare the volume to a one using caching (ReadCache/NoWriteThru).

Read Options

You have two read options. ReadCache is the default.

- Use the menu box to select NoReadCache or ReadCache.

Write Options

You have two write options. `NoWriteThru` is the default. When you select `WriteThru`, data is written directly to disk, bypassing cache.

In normal operation, it is not necessary to set the mode to `WriteThru`. To test applications without cacheing, disable the `sv` layer. By doing this, you bypass the Fast Write Cache software layer. This option could be set to `WriteThru` if there is an NVRAM board failure. In this case, the Fast Write Cache software destages cached blocks from the remaining good board (remember, the data is mirrored) , and puts the cache into `WriteThru` mode.

- **Use the menu box to select `WriteThru` or `NoWriteThru`.**

After you reconfigure the Fast Write Cache options, you may want to go to the Control panel.

Control

In the Control function, you may enable or disable the cache.

- **Select `Control` at the top of the window.**

The `Control` panel appears.

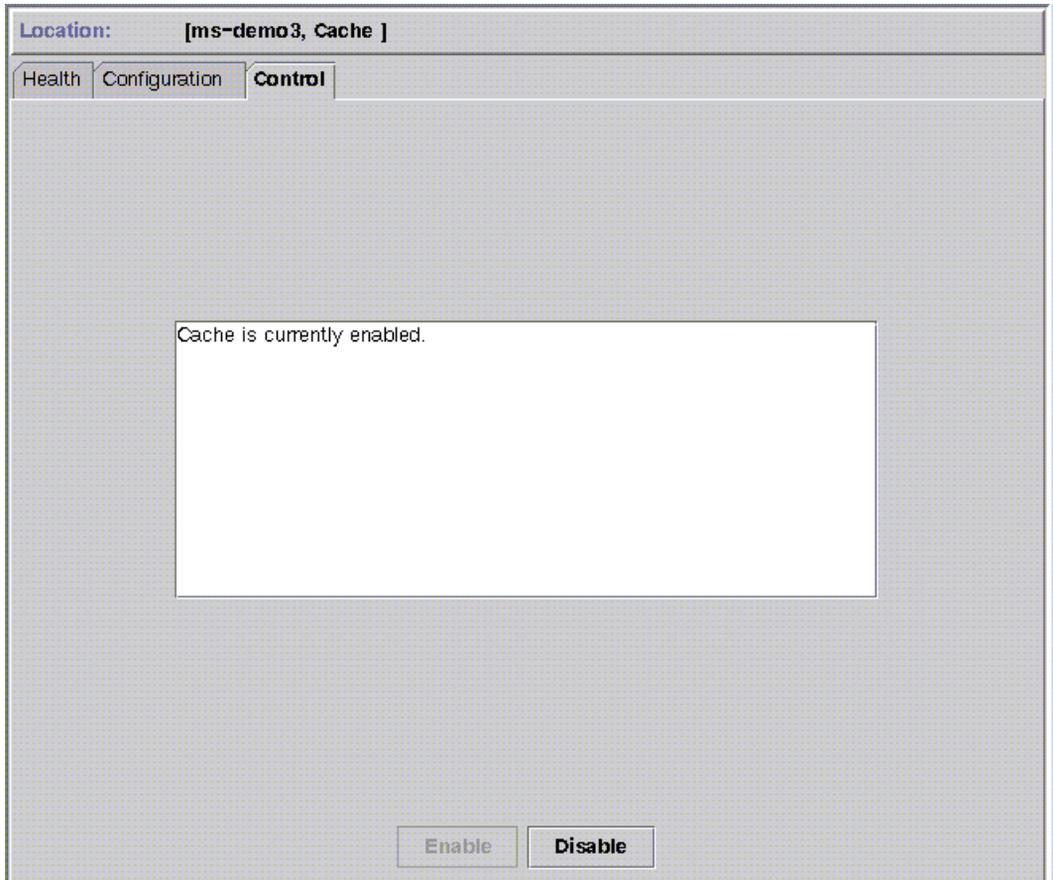


FIGURE 3-5 Control Panel

When you elect to disable the cache, a confirmation window appears.

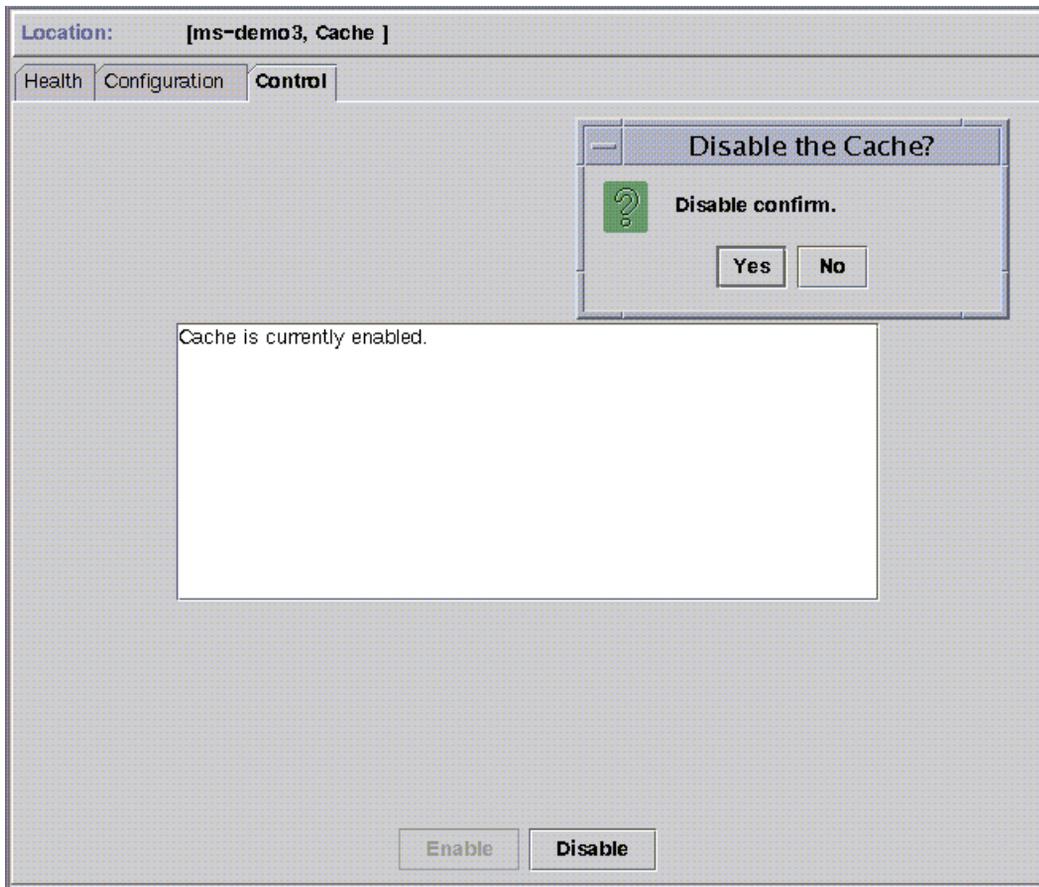


FIGURE 3-6 Stop Cache Confirmation

Health

After you configure the software, you are likely to use the `Health` portion of the Fast Write Cache Management Console most often. The `Health` panel reports the properties and statistics of the cache.

- **Select the `Health` tab at the top of the window.**

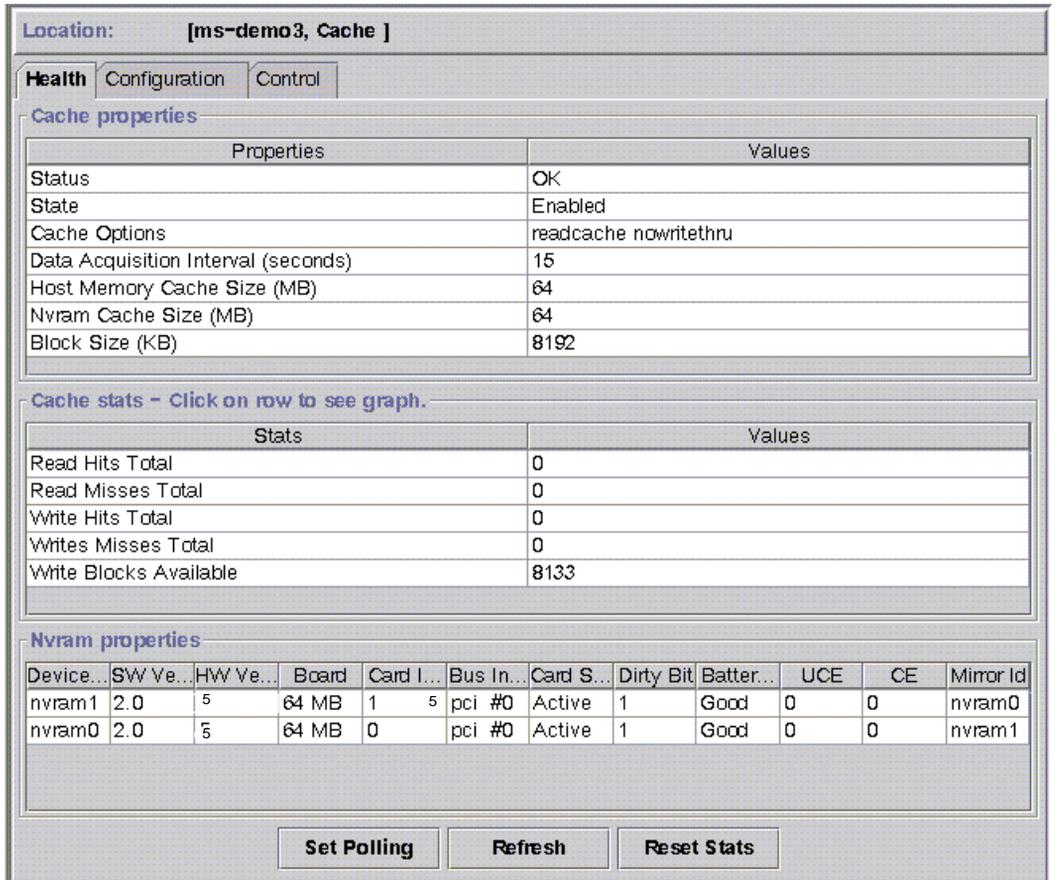


FIGURE 3-7 Health Panel

Cache Properties

- Status—The Status is OK or NotOK
 - Cache Status becomes NotOK when
 - A volume fails
 - An NVRAM card fails
- State—Fast Write Cache is Enabled or Disabled

- **Cache Options**—Configurations are read/write options selected in the Configuration folder. See “Configuration” earlier in this chapter. The default configuration is `ReadCache/NoWriteThru`. The configurations are defined as follows:
 - `ReadCache`: Data blocks are likely to be referenced again and will remain in cache
 - `NoReadCache`: Data blocks are unlikely to be referenced again and will be treated as least recently used so that other blocks can remain in cache longer
 - `WriteThru`: Writes go directly to the volume
 - `NoWriteThru`: Writes are stored in NVRAM and later destaged to the volume
- **Data Acquisition Interval (seconds)**—The default cache statistics polling and display period is every 15 seconds. The polling frequency can be changed by using the `Set Polling` button. The data acquisition interval can be from 5 to 120 seconds.
- **Host Memory Cache Size (Mbytes)**—May be 8-64 Mbytes. See “Configuration” earlier in this chapter.
- **NVRAM Cache Size (Mbytes)**—Not configurable, it always matches the NVRAM board size, 32 or 64 Mbytes
- **Block Size (Kbytes)**—8192 bytes. The block size is not changeable.

Cache Statistics

The Cache Statistics will update on each data acquisition interval. The monitored cache statistics are in FIGURE 3-7. You can display these in graph format by clicking the mouse over a cache statistic row. For example, usage of the cache write blocks can be displayed by clicking the mouse over the Write Blocks Available row.

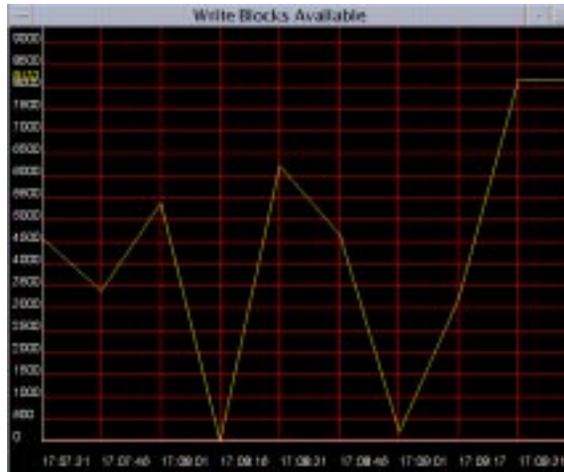


FIGURE 3-8 Cache Statistics Graph

Cache Statistics change with the cache activity.

- Read Hits Total—The number of times that all data for a read is in cache
- Read Misses Total—The number of times that some data for a read is not in cache
- Write Hits Total—The number of times a cache block was allocated in FastWrite mode
- Write Misses Total—The number of times a write did not go through the cache. Writes will not go through the cache if the cache is full and the block being written is not already in the cache.
- Write Blocks Available—The number of free 8Kbyte blocks in NVRAM.

NVRAM Properties

The NVRAM Properties give you information about your NVRAM boards. In the example, nvram1 and nvram0 are mirrors.

- SW Ver — Software Version
- HW Ver — Hardware Version
- Board—Board Size. PCI boards are 64 Mbytes; SBus boards are 32 Mbytes
- Card Instance—Instance number assigned by system (0 or 1)
- Bus Instance—PCI or SBus Id number
- Card Status—Either Active or Disabled

- **Battery**—Charge status. If the battery charge is low, Active status goes to Disabled.
- **UCE**—Uncorrectable Errors. Any value other than 0 causes the cache state to go to `Not Running` and makes it necessary to replace the card.
- **CE**—Errors corrected automatically by the NVRAM card (CE errors do not affect NVRAM card performance)

Note – The number of CE errors is not important, as the NVRAM cards are mirrored. If a CE error becomes a UCE error, the cache will be brought down and the card will have to be replaced. At mission-critical sites, if a card is generating CE errors, you may want to replace it at the next scheduled downtime. This will prevent cache loss during normal operation.

Buttons

- **Set Polling**—Updates the frequency of cache data acquisition
- **Refresh**—Instantly updates the statistics without waiting 15 seconds
- **Reset Stats**—Resets the cumulative statistics back to zero

Volumes

Volumes may be of two types:

- **Raw disk partitions**
- **Logical volumes presented by a volume manager**

Examples:

raw disk: `/dev/rdisk/c2t65d0s0`

Veritas Volume: `/dev/vx/rdisk/vol01`

Solstice DiskSuite™ volume: `/dev/md/rdisk/vol01`

The volumes are created and designated cache volumes by entering them in the `sv.cf` file. See “`sv.cf` Configuration File”.

Volume Configuration

You can change options for individual volumes.

- Choose volumes in the menu bar at the left side of the Management Console. See “Fast Write Cache Management Console”.

The volume configuration window appears.

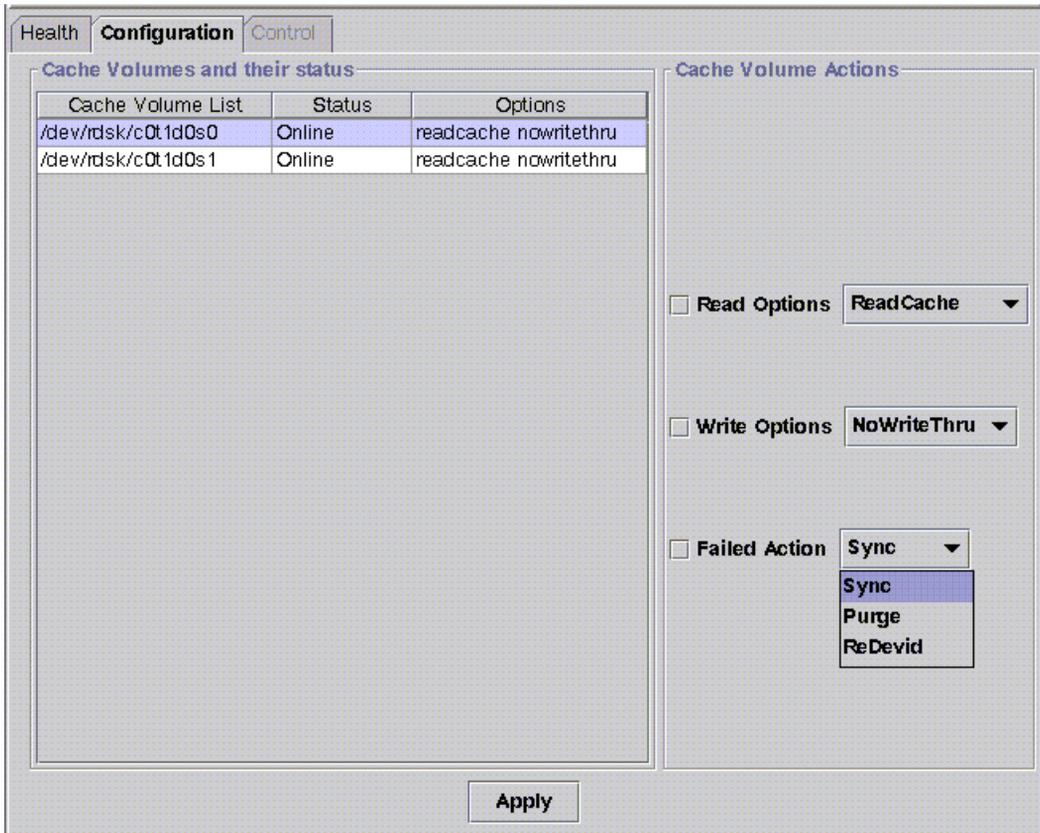


FIGURE 3-9 Volume Configuration Window

The volume configuration panel tells you what volumes you have, their status, and the configurable options for each volume.

- Status—Online or Offline
- Option—readcache or noreadcache
Option—writethru or nowritethru

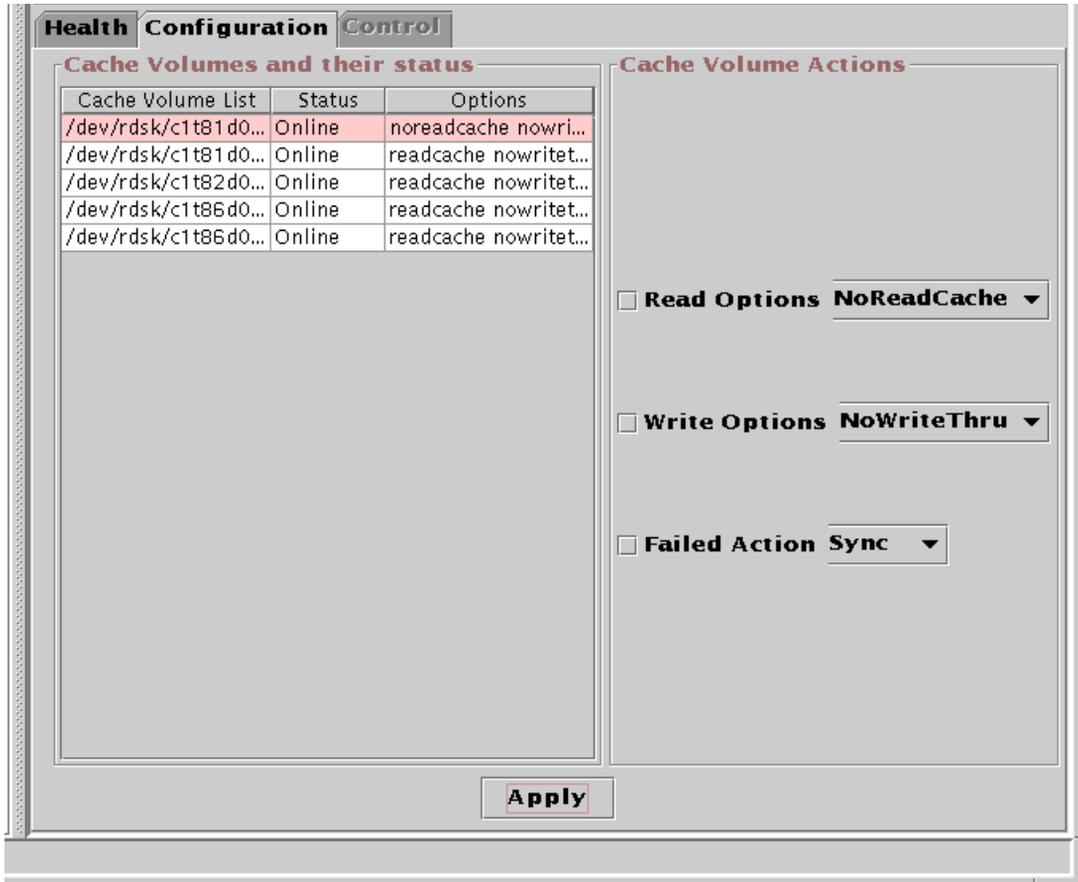


FIGURE 3-10 Applying Volume Configuration

1. Highlight the volume you want to configure.
2. Select the options you want for that volume at the right.
3. Press the `Apply` button at the bottom of the screen.

Sync and Purge

The `Sync` and `Purge` options come into play when there are failures, as when a disk drive loses power or fails before cache blocks are written to the disk.

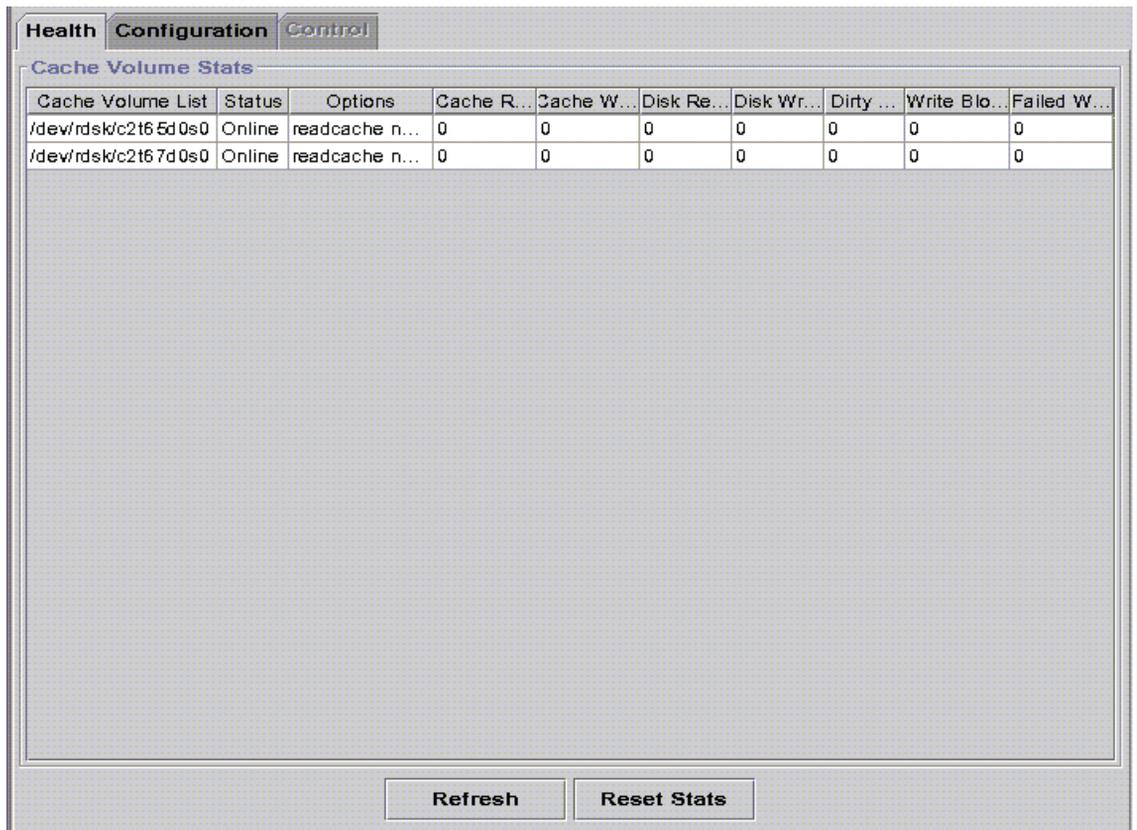
If a power loss or failure occurs, the Status changes to `Offline`. If, at the same time, data was being written to it, the Volume Health screen would show failed write blocks. See FIGURE 3-11.

Sync/Purge Actions

- Sync—Destages the failed blocks and clears the offline status.
- Purge—Discards the failed blocks and clears the offline status. When the disk drive is replaced, for example, purging the failed blocks is usually the correct action.

Volume Health

You can view volume statistics through the Health tab under the Volumes option.



The screenshot shows a window titled 'Volume Health' with three tabs: 'Health', 'Configuration', and 'Control'. The 'Health' tab is selected. Below the tabs is a section titled 'Cache Volume Stats' containing a table with the following data:

Cache Volume List	Status	Options	Cache R...	Cache W...	Disk Re...	Disk Wr...	Dirty ...	Write Blo...	Failed W...
/dev/rdisk/c2t65d0s0	Online	readcache n...	0	0	0	0	0	0	0
/dev/rdisk/c2t67d0s0	Online	readcache n...	0	0	0	0	0	0	0

Below the table are two buttons: 'Refresh' and 'Reset Stats'.

FIGURE 3-11 Volume Health

Cache Read, Cache Write and Disk Read, Disk Write keep track of cumulative statistics, which keep increasing.

- Cache Read—The number of Kbytes read from cache

- **Cache Write**—The number of Kbytes written to cache on a fast write hit
- **Disk Read/Writes**—The number of Kbytes read from disk and written to disk

The Failed Write Block statistics indicate conditions for which you may want to take action, such as `Sync` or `Purge`.

- **Dirty Write Blocks**—The number of 8Kbyte cache blocks having dirty sectors that have not been processed for I/O.
- **Write Blocks to Disk**—The number of 8Kbyte cache blocks being processed to build request to the I/O driver.
- **Failed Write Blocks**—The number of 8Kbyte cache blocks held in cache due to an I/O error.
- **Refresh**—Instantly updates the statistics without waiting until the next data acquisition interval
- **Reset Stats**—Resets the cumulative statistics back to zero

Alarm Viewer

You may see the Alarm Viewer screen.

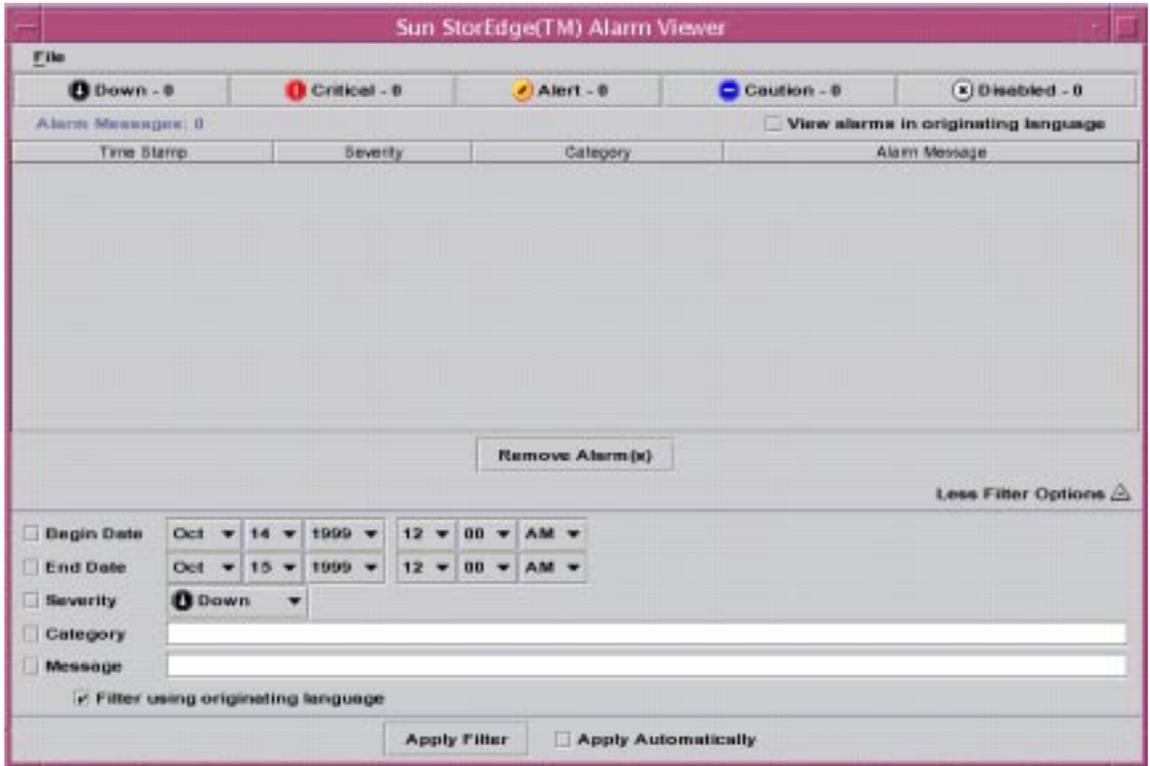


FIGURE 3-12 Alarm Viewer

- **Critical**—The underlying data services software tells the user something in the host is not working, for example, a memory problem, a process problem, or a core dump.
- **Alert**—The cache alerts the user that the status of the cache has changed. For example, the cache (volume?) has gone Offline.

Log Viewer

Any user action is added automatically to the log. The log may be useful in solving problems, as when a status change occurred. It may be useful to know, for example, when an NVRAM went offline.

Sun StorEdge(TM) Log Viewer		
File		
Log Messages: 11		<input type="button" value="Previous"/> <input type="button" value="Next"/>
Time Stamp	Category	Log Message
10/28/1999 05:10PM	SxSH:BOOT_EXCEPTI...	Boot Exception occurred localhost
10/28/1999 05:13PM	`StorageVolume`	SV MO Initial Start
10/28/1999 05:13PM	Instant Image	MO Initial Start
10/28/1999 05:13PM	SYMONDAQ	Server:: Listening on port # 4555
10/28/1999 05:14PM	Instant Image	Config MC Initial Start
10/28/1999 05:14PM	Instant Image	Control MC Initial Start
10/28/1999 05:14PM	Instant Image	Health MC Initial Start
10/28/1999 05:18PM	Cache	Updated the threads used by Cache.
10/28/1999 05:32PM	Cache	Updated the threads used by Cache.
10/28/1999 05:32PM	Cache	Disabled the Cache.
10/28/1999 05:32PM	Cache	Enabled the Cache.

More Filter Options ▾

FIGURE 3-13 Log viewer

Using the Command Line Interface

This chapter describes how to start the Fast Write Cache command line interface (CLI) shell and provides the syntax of the CLI commands.

Through the Fast Write Cache CLI, you can perform the same operations that are available in the Fast Write Cache management service shown in the Sun StorEdge Management Console described in Chapter 3. The Fast Write Cache CLI is useful when you want to script Fast Write Cache operations or when you want to perform Fast Write Cache operations from a terminal.

Note – To use the Fast Write Cache software, you must be the root user or superuser.

- CLI Compatibility
- Using the CLI
- Command Summary
- Using the cache Commands
- Data Service Command `fwcadm`

Reading Fast Write Cache Man Pages

- To read the legacy CLI Fast Write Cache `fwcadm(1M)` man page, type:

```
# man -M /usr/opt/SUNWesm/SUNWnvm/man fwcadm
```

- To read the Fast Write Cache management service man pages, type:

```
# man -M /usr/opt/SUNWesm/SUNWmscmu/man manpagename
```

where *manpagename* is `cache_config`, `cache_health`, `cache_control`, `cache_vol_health`, or `cache_vol_config`.

CLI Compatibility

The Fast Write Cache CLI is compatible with the `SUNWnvm` package-installed legacy CLI. See the `fwcadm(1M)` man page for more information about the `SUNWnvm` command line. See also “Reading Fast Write Cache Man Pages” earlier in this chapter..

Using the CLI

There are three ways to use the CLI to execute commands:

- As part of an existing or new script file which uses the `fwcadm` command (see the `fwcadm(1M)` man page).
- As a single command executed from the Solaris root user prompt or group of commands executed from a script file, such as:

```
# /usr/opt/SUNWesm/bin/esm_cli -e "cache/health -S"
```

or

```
/usr/opt/SUNWesm/bin/esm_cli -e "cache/config -t 256"  
/usr/opt/SUNWesm/bin/esm_cli -e "cache/control -d"  
/usr/opt/SUNWesm/bin/esm_cli -e "cache/control -e"
```

- As a single command executed from the Fast Write Cache shell, such as:

```
[1] sx> health -State
```

Note – If you enter a command incorrectly or an error is generated in the Fast Write Cache shell, the shell returns you to the Solaris operating environment prompt.

▼ To Start the Command Line Interface Shell

1. Become superuser (root).
2. Enter the CLI shell by typing:

```
# /usr/opt/SUNWesm/bin/esm_cli  
[1] sx> cd cache  
[2] sx>
```

Command Summary

You use three commands in the Fast Write Cache CLI shell (`esm_cli`). These three commands correspond to the Fast Write Cache management service GUI tabs described in Chapter 3:

- `config`—use to change cache size, number of threads, and Read/Write options
- `control`—use to enable and disable the cache
- `health`—use to see volume information and status

Note – `cache start` fails if there is insufficient contiguous host memory. The host memory requirement is based on the `cache_mem` parameter in the `/etc/sd.cf` file and the size of the NVRAM memory.

TABLE 4-1 Command Parameters

Command Parameters		Definition
health	-S	Display state and status of cache
	-l	Display all cache properties and their values
	-s	Display data acquisition in seconds
	-o	Display read and write options
	-c	Display cache size
	-b	Display block size
	-u	Display cache memory used
	-r	Display read hits total
	-e	Display read misses total
	-w	Display write hits total
	-i	Display write misses total
	-n	Display the status of the NVRAM cards
	-V	.Display Cache Package version
-h	Provide help on the command	
config	-S	Displays the state and status of the cache
	-c [value]	Return [set] size in MB of cache memory
	-t [value]	Return [set] number of threads
	-o [value]	Return [set] option values Values must be one or two of the following: readcache noreadcache writethru nowritethru
	-h	Provide help on the command
control	-S	Display the state of the cache
	-e	Enable the cache
	-d	Disable the cache
	-h	Provide help on the command
volume health	-S	Display status of volume(s)
	-o	Display readcache,writethru for volume(s)
	-r	Display cache reads for volume(s)
	-w	Display cache writes for volumes(s)
	-e	Display disk I/) reads for volumes(s)
	-i	Display disk I/) writes for volume(s)
	-d	Display dirty write blocks for volume(s)
	-b	Display write blocks to disk for volume(s)
	-f	Display failed write blocks for volume(s)
	-h	Provide help on command
volume config	-S	Display status of volume(s) and options
	-o	Display/set readcache,writethru for volume(s)
	-s	Set sync option on volume(s)
	-p	Set purge option on volume(s)
	-r	Set redevid option on volume(s)
	-h	Provide help on command

Using the cache Commands

The `esm_cli` cache commands may be used as an alternate to the Management Console. See Chapter 3. TABLE 4-1 shows the options available.

You may want to use the command line interface for a single interaction and script invocation, or to change more than one option at the same time.

- **To start the cache command shell, type:**

```
#!/esm_cli
```

The system returns the `sx` prompt, and you type the following:

```
sx>  
sx>ls  
Cache  
sx>cd cache  
sx>ls  
control*  
config*  
health*  
volumes/
```

Command Examples

▼ To start the cache

```
sx>control -e  
control: INFO: cache enabled
```

▼ To stop the cache

```
sx>control -d  
control: INFO: cache disabled
```

▼ To change the number of threads

```
sx>config -t [value] (where value is 512)  
config: INFO: number of threads set to 512
```

Configuration values have been saved until the cache is re-enabled. To commit values to cache, go to the Cache Control CLI to disable and enable the cache. See example 3. above.

▼ To change the cache memory size

```
sx> config -c [value]
```

The value for memory might be [8..64] in 8 MB increments.

▼ Checking the overall health

```
sx> health -S
health: INFO: cache/health
Running:OK
```

▼ Checking the volume health

```
sx> volumes/config -S
config: INFO: cache/volumes
volume /dev/rdisk/c0t1d0s0  status = Online
volume /dev/rdisk/c0t1d0s4  status = Online
```

or

```
sx> volumes/health -S
health: INFO: cache/volumes
volume /dev/rdisk/c0t1d0s0  status = Online
volume /dev/rdisk/c0t1d0s4  status = Online
```

▼ Changing the cache read/write options on all volumes

```
sx> volumes/config -o readcache,nowritethru
config: INFO: cache/volumes
config: INFO: Cache Volume(s) updated.
```

Note that the absence of volume names means that the options are set for all volume and that this sets all the volumes to this caching but the main caching options are not changed.

For example

cache/config -o shows values of readcache,nowritethru

but

cache/volumes/config -o shows value for each volume.

▼ Changing the cache read/write options on a single volume

```
sx> volumes/config -o readcache,nowritethru /dev/rdisk/c1t65d0s0
config: INFO: cache/volumes
config: INFO: Cache Volume(s) updated.
```

▼ List current statistics of the overall health

```
sx> health -l
health: INFO: cache/health
Running:OK
Cache Options = readcache nowritethru
Cache Size = 64
Block Size = 8192
Cache Used = 64
Read Hits Total = 17
Read Misses Total = 5
Write Hits Total = 18799
Write Misses Total = 0
Data Acquisition frequency = 15 seconds
```

▼ Changing the data acquisition interval

```
sx> health -d 20
health: INFO: Data Acquisition frequency = 20 seconds
Data Acquisition frequency = 20 seconds
```

Data Service Command `fwcadm`

`fwcadm` is the administration command for the cache, NVRAM card, and the Storage Volume (SV) driver. `fwcadm` must be specified with one of the following parameters:

- `cache`
- `nvr`
- `volume`

`fwcadm cache` Command

Use the `cache` command to do the following:

- Enable or disable the cache
- View cache statistics
- Destage cache
- Clear the offline state of a failed device
- Re-identify a physical drive that has been replaced

Note – You must be superuser to run all of the `cache` command options except `-s`.

Syntax

```
cache [purge diskname | sync diskname | redevid diskname]  
      [-d] [-e] [-f file] [ -s ]
```

Parameters

Parameter	Description
-d	<code>fwcadm cache -d</code> disables the cache.
-e	<code>fwcadm cache -e</code> reads the configuration file (the default is <code>/etc/opt/SUNWopt/sd.cf</code>) and enables the storage device cache with those parameters.
-f <i>file</i>	<code>fwcadm cache -f file</code> specifies an alternate configuration file. The default configuration file is <code>/etc/opt/SUNWopt/sd.cf</code> .
-s	<code>fwcadm cache -s</code> displays cache statistics. Press the <code>t</code> key to toggle between two screens. The first screen shows general statistics about the data cache, and the second screen displays total counts.
purge <i>diskname</i>	<code>fwcadm cache purge diskname</code> discards the failed blocks and clears the offline state of the failed device.
sync <i>diskname</i>	<code>fwcadm cache sync diskname</code> destages the failed blocks and clears the offline state of the device.
redev id <i>diskname</i>	<code>fwcadm cache redev id diskname</code> allows the re-identification of a replaced physical disk. This is necessary only if there is pinned data for the device.

`fwcadm cache -s` Display

When `cache -s` is running, you can change the display by selecting various keys.

TABLE 4-2 `fwcadm cache -s` Display Options

Key	Description
+	Increases the screen update delay an additional second.
-	Decreases the screen update delay by a second (minimum delay is 1 second).
B	Toggles between normal and bold types.
b or CTRL-b	Scrolls the display backwards to the previous set of devices currently not in view.
C	Clears the screen and displays the statistics again.
f or CTRL-f	Scrolls the display forward to the next set of devices currently not in view.
M or m	This key does not apply to Fast Write Cache but is designed for use with other Network Storage Data Services.

TABLE 4-2 fwcadm cache -s Display Options (*Continued*)

Key	Description
R	Toggles between normal and reverse video.
T or t	Toggles between regular (per second statistics) and cumulative screens.
Z	Clears the index cache statistics.

Example

In the following example, `fwcadm cache -s` displays the per second general statistics. Pressing the `t` key toggles to the second screen, which displays the cumulative counts.

```
# fwcadm cache -s

SAMPLE 1          *****      Storage Cache          *****      12:01:25
                  disk_io          cache          write_blocks
cd cached_partition reads  writes reads  writes dirty  to disk  failed
-----
0  ...dsk/clt5d0s1    352    320     0     0     0     0     0
1  ...dsk/clt5d0s3    320   1024     0   352     8   128     0
2  ...dsk/clt5d0s4    320    960     0   320     4   120     0
3  ...dsk/clt5d0s5    288     0    64   320   112     0     0
4  ...dsk/clt5d0s6    320     0    32   352   112     0     0
5  ...dsk/clt5d0s7     64     0   256   352    80     0     0
-----
      Kbytes/s total: 1664    2304    352    1696

accesses/s   read/s   write/s   %readh   %writeh
(misses/s) (misses/s)
-----
    126.12    11.01    53.05    17.5     84.1
      ( 52.05 ) (10.01)

<t key is pressed to toggle to the next screen>

SAMPLE 3          *****      Storage Cache (Cumulative)          *****      12:01:28
                  disk_io          cache
cd cached_partition  reads  writes  reads  writes
-----
0  ...dsk/clt5d0s1   170328 1403352    128  581312
1  ...dsk/clt5d0s3   43464 2040406    992 2024694
2  ...dsk/clt5d0s4   23616 183488     576 184160
```

3	...dsk/clt5d0s5	19680	186208	4512	186752
4	...dsk/clt5d0s6	20384	167552	3808	168096
5	...dsk/clt5d0s7	4128	164576	20000	164832
	Kbytes/s total:	281600	4145582	30016	3309846
accesses/s	read/s	write/s	%readh	%writeh	
	(misses/s)	(misses/s)			
-----	-----	-----	-----	-----	-----
90569	938	67892	10.8	82.9	
	(7747)	(13992)			
-----	-----	-----	-----	-----	-----
96128K	8192				
Write blocks available:					

Net 0:	3488	Net 1:	0	Net 2:	0
				Net 3:	0
LRU stats:	Blocks	Requeued	Optimized		
	-----	-----	-----		
	8192	5321	551933		
Total Cache Memory Usage: 2544 Kbytes					
Total Stats Memory Usage: 0 Kbytes					

Note – Fast Write Cache write activity proceeds from *cache writes* to *dirty* to *todisk* to *disk_io* writes.

cd

This field displays the CD cache descriptor number.

cached_partition

This field displays the volume name.

disk_io

These fields display the number of Kbytes read from disk and Kbytes written to disk.

cache

These fields display the number of Kbytes read from cache on a read hit and Kbytes written to cache on a write hit (that is, a fast write).

write_blocks

These fields display the following:

- **dirty** – number of 8Kbyte cache blocks having dirty sectors
- **todisk** – number of 8Kbyte cache blocks being processed to build I/O requests to the real driver
- **failed** – number of 8Kbyte cache blocks held in cache due to an I/O error

accesses

This field displays the total read hits, read misses, write hits, and write misses.

read/(misses)

The first number in this field (read) is the number of times that all data for a read request is in cache.

The second number in this field (misses) is the number of times that some data for read I/O is not in cache.

write/(misses)

The first number in this field (write) is incremented for each cache block allocated in fast write mode.

The second number in this field (misses) is the number of synchronous writes.

%readh

This percentage is obtained by dividing the read number by the total of the read hits and misses. That is, $read\ hits / (read\ hits + read\ misses) = \%readh$.

%writeh

This percentage is obtained by dividing the write number by the total of the write hits and misses. That is, $write\ hits / (write\ hits + write\ misses) = \%writeh$.

cache size

This field displays the size of cache (host memory) plus the write cache (NVRAM).

block size

This field displays the cache block size.

Write blocks available

This field displays the number of free 8Kbyte blocks in NVRAM. NET is the NVRAM card.

LRU stats (Least Recently Used)

- **Blocks** – number of 8Kbyte cache blocks in host memory.
- **Requeued** – number of times the cache blocks were placed at the Most Recently Used (MRU) end of queue when released.
- **Optimized** – number of times blocks were *not* placed at the MRU end of queue when released. (The 25% most recently used blocks in the queue are not requeued.)

Total Cache Memory Usage

This field displays the total cache overhead used in Kbytes.

fwcadm nvram Command

Use the `fwcadm nvram` command to display the NVRAM card status.

Note – You must be superuser to run the `fwcadm nvram` command.

Syntax

```
fwcadm nvram -s
```

Options

Parameter	Description
<code>-s</code>	<code>fwcadm nvram -s</code> displays the status of the NVRAM cards.

Example

```
# fwcadm nvram -s
```

Device	Size	Inst	Parent	Status	Dirty	Battery	UCE	CE	Mirror
<code>nvram0</code>	32MB	#0	<code>sbus #0</code>	Active	1	Good	0	0	<code>nvram1</code>
<code>nvram1</code>	32MB	#1	<code>sbus #1</code>	Active	1	Good	0	0	<code>nvram0</code>

Device

This field displays the entry that is in the `/dev` directory.

Size

This field displays the total amount of memory in Mbytes detected by the NVRAM driver.

Inst

This field displays the instance number assigned by the system to the device during configuration. This is useful in matching a specified card to output (for example, `prtconf(1M)`).

Parent

This field displays the name and instance number of the I/O bus on which the NVRAM card resides. This is useful in matching a specified card and bus combination to output (for example, `prtconf(1M)`).

Status

This field displays the current state of the device.

TABLE 4-3 Device State Description

State	Description
Inactive	Card is not in use.
Active	Card is in use.
Disabled	Card has a problem.

Dirty

This field displays the status of the dirty bit. The dirty bit controls the state of the LED on the SBus NVRAM card. The dirty bit is set when the cache is enabled whether or not there is any dirty data on the card. It remains set if the cache was *not* cleanly shut down (see “Unclean and Clean Shutdowns” in Chapter 6.). The dirty bit is not the best indicator of whether unwritten data is in the cache, as it keeps driver-specific information. Use `fwcadm cache -s` to display the cache statistics for dirty blocks.

This bit is useful for error recovery when the cache has an unclean shutdown. The state of this bit is maintained during a loss of power to the NVRAM card (that is, it is nonvolatile). During recovery, any dirty blocks on the NVRAM card are automatically destaged to disk. (See “Fast Write Cache Modes” in Appendix A.)

Note – Only the SBus NVRAM card has an LED tied to the dirty bit register. No other NVRAM cards for Fast Write Cache have this LED, but they do have a dirty bit.

TABLE 4-4 SBus NVRAMBit Description

Value	Description
0	No valuable data is on the card.
1	Valuable data may be on the card.

Note – On the SBus NVRAM card, the system is the power source for the LED and *not* the batteries on the card. For the LED to illuminate, the card must be installed in a system with power turned on. There is only one way to check for data on the SBus NVRAM card that is not in a system: Plug the card into a system, power it on, and immediately press L1-A or Stop-A to halt the boot process. The LED illuminates if the cache manager software left the dirty bit set due to a previous unclean shutdown.

TABLE 4-5 Dirty LED States

State	Reasons
ON	Cache is enabled (regardless of any dirty blocks in the cache). Unclean shutdown occurred while the cache was enabled.
OFF	Cache was disabled by a clean shutdown. NVRAM card is not installed. NVRAM card is installed, but the system power is off.

Note – Only the SBus NVRAM card has a Dirty LED. Future NVRAM cards will all have a Dirty bit and no Dirty LED.



Caution – Do not attempt to alter the state of the dirty bit, as data may be lost.

Battery

This field displays the state of the three on-board lithium batteries on the SBus NVRAM card or the four batteries on the PCI Bus NVRAM card.

TABLE 4-6 Battery Status Description

Status	Description
Good	All batteries are good.
Bad	Low voltage in one or more batteries.
Disabled	SBus NVRAM: At least one battery is disabled. More than one battery may be disabled. The battery is not connected to the array; either its fuse is open or the battery enable jumper is OFF. If all three batteries are installed, and the three battery enable jumpers are installed, the battery fuse is open. PCI bus NVRAM: Enable/disable switch is off.

If a card has a Bad or Disabled battery status, `fwcadm nvram -s` displays a WARNING message and specifies which battery on the card has the problem. For example:

Device	Size	Inst	Parent	Status	Dirty	Battery	UCE	CE	Mirror
nvram0	32MB	#0	sbus #0	Disabled	0	Disabled	0	0	nvram1
WARNING: Battery 2 is disabled									
nvram1	32MB	#1	sbus #0	Inactive	0	Good	0	0	nvram0

In this example for the SBus NVRAM card, the message indicates that Battery 2 is disabled on `nvram0`. The possible causes and solutions are:

- Battery enable jumper is not installed (install the jumper)
- Battery is not installed (replace the card)
- Fuse is open (replace the card)

UCE (Uncorrectable Errors)

This field displays the number of uncorrectable ECC errors detected. These errors are considered fatal; if they occur, the card is unusable, and you must replace it. The data is not lost if this type of error is detected on one of the cards, as the NVRAM cards are mirrored.

CE (Correctable Errors)

This field displays the number of correctable ECC errors detected. This is not a fatal error condition. The NVRAM driver has a memory scrubber that scans for these errors and attempts to fix them. If the scrubber cannot fix the error, the error is uncorrectable; you must replace the card.

Mirror

This field displays the mirror device that is configured with this device. If the device is not mirrored, it specifies `none`. The mirror contains the same data as this device. Mirroring prevents data loss if a UCE error occurs or an NVRAM card fails during recovery.

fwcadm volume Command

Use the `fwcadm volume` command to administer the Storage Volume (SV) driver. Use it to control the SV driver by providing facilities to enable and disable the SV driver for specified devices, and to dynamically reconfigure the system.

Note – You must be superuser to run the `fwcadm volume` command.

Syntax

```
fwcadm volume { -d | -e | -f file | -r | -s }
```

Options

Parameter	Description
-d	<code>fwcadm volume -d</code> disables the SV devices specified in the configuration file.
-e	<code>fwcadm volume -e</code> reads the configuration file (the default is <code>/etc/opt/SUNWspsv/sv.cf</code>) and enables the specified SV devices.
-f <i>file</i>	<code>fwcadm volume -f</code> specifies an alternate configuration file. The default configuration file is <code>/etc/opt/SUNWspsv/sv.cf</code> .
-r	<code>fwcadm volume -r</code> reconfigures the SV subsystem. It compares the contents of the configuration file to the state of the running system, and then enables and disables devices to reconfigure the running system as specified in the configuration file.
-s	<code>fwcadm volume -s</code> displays the current state of the SV subsystem.

sd.cf Configuration File

You can specify the following parameters in the `/etc/opt/SUNWscm/sd.cf` configuration file.

Parameter	Description
<code>cache_mem</code> <i>net size_in_mbytes</i>	Amount of host memory to be allocated. For Fast Write Cache, <code>cache_mem</code> should be no less than the size of one NVRAM card. The <i>net</i> parameter is a reserved value and must be 0
<code>threads</code> <i>num</i>	Specifies the number of flush threads to create. One is needed for each cached volume. Threads in excess of the number of cached volumes are not used.
<code>write_cache_mem</code> <i>megabytes</i>	Specifies the amount of memory reserved for write cache. Ignored for Fast Write Cache.

sv.cf Configuration File

The `/etc/opt/SUNWspv/sv.cf` file specifies the SV devices to be enabled, disabled, or reconfigured by the `fwcadm volume` command. It contains the raw devices and access modes.

The following are examples of `sv.cf` entries:

```
# Example 1: Layer SV onto /dev/rdisk/clt0d0s0, access through
# the SCM cache.
/dev/rdisk/clt0d0s0      cache

# Example 2: Layer SV onto /dev/rdisk/clt2d0s0, access the device
# directly.
/dev/rdisk/clt2d0s0      raw

# Example 3: Layer SV onto /dev/rdisk/clt4d0s0, use the default
# access mode.
/dev/rdisk/clt4d0s0
```

Specifying a raw device enables both the raw and block devices. For example, specifying `/dev/rdisk/clt0d0s0 cache` in Example 1 above also enables `/dev/dsk/clt0d0s0`.

Note – A line containing a `cache` designator must be specified for Fast Write Cache. The `raw` designator is used by other data services to target I/O through the SV module and around the cache.

Normal Operations

This chapter describes how to perform the following Fast Write Cache system administration procedures during normal operating conditions.

- Removing Fast Write Cache Hardware
- Disabling the Batteries on the NVRAM Card
- Selecting the File Systems and Volumes to Cache
- Removable Disk Packs

Removing Fast Write Cache Hardware

The NVRAM card is like a piece of your disk. If the system experiences a power outage, a system panic, or a machine abort, the Fast Write Cache driver cannot destage cached data to the disk. This type of *unclean* shutdown leaves valid data on the NVRAM card. In this state, removing the card is like removing a piece of the disk.



Caution – Removing an NVRAM card that contains valid data may result in loss of data. If you must move a card containing valid data, see Chapter 6.

▼ To Remove an NVRAM Card

1. Type the following commands to ensure that there is no valid data on the NVRAM card:

```
# fwcadm volume -d
# fwcadm cache -d
```

2. If any disk errors were noted, fix these prior to continuing, and then repeat step 1.
3. Shut down the system by following the instructions that came with your operating system.
4. Turn off the power to the machine.
5. Remove the NVRAM card by following the instructions that came with your system.
6. Put the NVRAM card in an antistatic bag.
You can safely move the card to a different system. If you must move a card containing valid data, see Chapter 6.
7. Reassemble the system.
8. Turn on the power on the system and reboot.

Disabling the Batteries on the NVRAM Card

The NVRAM card has permanently mounted lithium batteries that maintain the nonvolatile memory cache in case of power loss. The SBus NVRAM Card has three batteries that last at least five years in a running system or in storage with the batteries disabled. The batteries last approximately 18 months in a powered down system or in storage with the batteries enabled. The PCI Bus NVRAM Card has four batteries that last at least ten years in a running system and approximately 34 months in a powered down system or in storage with the batteries enabled.

The Fast Write Cache driver monitors the state of the batteries, and a console message is displayed when the battery state changes. When the battery falls below nominal voltage levels, the Fast Write Cache driver disables its functions, and all subsequent disk requests are passed directly to the disks.

If a battery fails, call your authorized Sun service provider.



Caution – Do not attempt to replace the batteries yourself. The batteries are soldered to the card and are sensitive to heat. There is a danger of explosion if the batteries are incorrectly replaced.

If you disable the batteries on the NVRAM card, the memory is lost and its contents are cleared.

▼ To Disable the Batteries on the SBus NVRAM Cards

1. Remove the NVRAM card according to the instructions in “Removing Fast Write Cache Hardware”.
2. Move each jumper so that only one jumper pin is inserted in each jumper

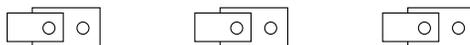


FIGURE 5-1 Jumpers and Jumper Pins With Batteries Disabled

Note – If you must completely clear the memory on the card, ensure that the batteries are disabled for at least 48 hours.

▼ To Disable the Batteries on the PCI Bus NVRAM Cards

1. **Remove the NVRAM card according to the instructions in “Removing Fast Write Cache Hardware”.**

2. **Look for the Enable and Disable switch.**

The PCI NVRAM holds four batteries. There is one switch to enable and disable all the batteries.

3. **Move the switch from Disable to Enable. Make sure that you pull or push the switch from the middle of the switch bar to ensure that the switch enables or disables all batteries.**

When you unpack the card, the switch is set to Disable.

- For NFS™ servers, cache all exported file systems
- On databases, cache only the logs. Do not cache indexes unless the write load to index is small
- The current NVRAM mirroring mechanism causes overall throughput to be one-half the performance of one card (54 Mbytes per second on SBus, 45 Mbytes per second on PCI). If you are writing at higher rates than this to existing volumes, then you may consider not caching these volumes.

Selecting the File Systems and Volumes to Cache

When choosing file systems to cache, consider that good candidates for acceleration receive synchronous write requests (for example, a UNIX® File System (UFS) mounted with the `force directio` set).

For Network File System (NFS) servers, cache all exported filesystems except those that are read-only. On database cache only the logs, not indexes, unless the write load to the index is small; otherwise it will thrash the cache.

If you are not getting the write performance you want, you may have too much data going through the cache; you might have to cut back on the number of cached devices. To see if your cache is being over-utilized, check the number of write blocks available on the Health Panel (FIGURE 3-11). If the count is under 500 for a sustained period, then the write cache is saturated.

Because the NVRAM cards are mirrored, the overall throughput is half the performance of one card. Thus, maximum throughput on Sbus is 54 Mbytes per second; on PCI, it is 45 Mbytes per second. If you are writing at higher throughputs to existing volumes, then do not cache these volumes.

To recover disk data stored in an NVRAM card after a system crash, the data must be restored to disk before file systems are mounted or applications write to raw volumes. Because of this process, the following file systems must not be cached.

File System	Reason the File System Should Not be Cached
/	The driver is not loaded until after the root file system (/) is mounted. If the NVRAM card contains dirty data, the data cannot be recovered until after the driver is loaded. Recovering data after the root file system is mounted may corrupt the root file system.
/usr	The commands that start the recovery process reside on this partition. As above, caching this partition could result in data being recovered from the cache after the file system was mounted

Note – Because the cache is applied to volumes that sit above the cache, recovery occurs after the Volume Manager is started. As a result, caching Volume Manager subdisks is not allowed.

▼ To Enable Caching for a File System

1. Add the raw disk entry representing the device on which the file system will be located.

For example, to cache the /mnt1 file system, do the following:

- a. Determine the mount point for the file system. For example,

```
# grep mnt1 /etc/vstab
/dev/dsk/c1d0s2 /dev/rdisk/c1d0s2 /mnt1 ufs 1 yes -
```

- b. Add the raw disk name to the /etc/opt/SUNWspv/sv.cf file. For example, vi the /etc/opt/SUNWspv/sv.cf file and add the following line:

```
/dev/rdisk/c1d0s2 cache
```

2. Type the following reconfiguration command to enable the volume through the cache:

```
# fwadm volume -r
```

3. If the cache is not enabled, type the following command to enable the cache:

```
# fwadm cache -e
```

▼ To Enable Caching for a Volume

1. Add the volume name to the /etc/opt/SUNWspv/sv.cf file. For example, vi the /etc/opt/SUNWspv/sv.cf file and add the following line:

```
/dev/vx/rdsk/v0101 cache
```

2. Type the following reconfiguration command to enable the volume through the cache:

```
# fwadm volume -r
```

3. If the cache is not enabled, type the following command to enable the cache:

```
# fwcadm cache -e
```

▼ To Display Cache Statistics

- Type the following command to display cache statistics:

```
# fwcadm cache -s
```

Removable Disk Packs

Before removing a disk pack, perform the following procedure.

▼ To Remove Disk Packs

1. You must deconfigure devices from the Fast Write Cache before removing the disks by removing the entry for the device from the `/etc/opt/SUNWspsv/sv.cf` file.
2. Type the following command to reconfigure the system as specified in the `sv.cf` file:

```
# fwcadm volume -r
```

3. Use the `umount` command if there are file systems mounted on the disks.
4. Physically remove the disk pack.

Abnormal Operations

This chapter describes how to manage Fast Write Cache under abnormal operating conditions. Abnormal conditions include cases where the system did not shut down cleanly, or where a disk accelerated by Fast Write Cache encounters errors or failure.

- Unclean and Clean Shutdowns
- Moving an NVRAM Card Containing Data
- Handling Disk Failures

Unclean and Clean Shutdowns

An *unclean* shutdown results if one of the following occurs before Fast Write Cache has been shut down with the `fwcadm cache -d` command:

- Power failure
- Hardware failure
- L1-A, BREAK, or Stop-A key sequences are pressed
- `reboot(1M)`, `halt(1M)`, or `vadmin(1M)` command is used
- Shutdown of any type when there is trapped data because of a disk failure

Note – After an unclean shutdown, Fast Write Cache may contain valuable data that has not been written to disk. Fast Write Cache attempts to destage the cache after you reboot the system or start the cache with the `fwcadm cache -e` command.

A *clean* shutdown results when the system is properly halted according to the instructions that come with the operating system (refer to the `shutdown(1M)` man page). You can also shut down the Fast Write Cache subsystem using the `fwcadm cache -d` command prior to shutting down the system.

Booting a Different Kernel

Fast Write Cache uses device path names to identify data blocks. If you reconfigure your machine after an unclean shutdown, you may have cached data blocks destaged to the wrong device. This could happen under the following conditions:

- A kernel is installed that has different disk device name mappings than the kernel last used with Fast Write Cache.
- A disk controller is added or removed.

Cache `redevi`d

In some cases, if the cache manager software detects a mismatch of the World Wide Name unique device id, it will not destage blocks to a cached volume on recovery after an unclean shutdown. This can happen if all of these conditions occur concurrently:

1. A partition not managed by a volume manager is enabled for cacheing.
2. The device on which the partition resides has a World Wide Name.

3. An unclean shutdown occurs.
4. The device on which the partition resides is replaced and the system is subsequently rebooted with the reconfiguration option which makes the new World Wide Name known to the operating system.

If this happens, the `fwcadm cache redevid` command must be used to cause the cache manager software to store the new World Wide Name. A subsequent `sync` or `purge` operation must then be performed before any I/O requests will be handled for the device.

Moving an NVRAM Card Containing Data



Caution – Moving cards containing valid data can result in loss of data. If you must move NVRAM cards from one system to another, shutdown both systems cleanly, and then remove the cards from the first system and install them in the second system.

Moving an NVRAM Card Without Losing Data

Data *should not* be lost when you do the following:

- Move NVRAM cards within the same machine and reboot.
 - Handle cards properly
 - Install all the original cards in the machine prior to bootup
- Add new NVRAM cards to a machine containing NVRAM cards with valid data.
- Change the system CPU or ID prom and all NVRAM cards remain in the system.
- Remove NVRAM cards from the machine and reboot when the remaining cards in the system form a valid mirror (complete primary or secondary), and an uncorrectable error does *not* occur on the mirror during data recovery.

Moving an NVRAM Card and Losing Data

Data *is* lost when you do the following:

- Remove NVRAM cards from the machine and reboot when the remaining cards in the system do not form a valid mirror (incomplete primary or secondary).

- Remove NVRAM cards from the machine and reboot when the remaining cards in the system form a valid mirror (complete primary or secondary) and an uncorrectable error occurs on the mirror during data recover.
- Move an NVRAM card to another machine. Data on that card is lost.
- Remove an NVRAM card from the machine and disable its batteries.

Note – If you must completely clear the memory on the card, ensure that the batteries are disabled for at least 48 hours.

Handling Disk Failures

When Fast Write Cache is caching disk blocks, data written by an application may not be written to disk immediately. If a disk on a cached volume fails, the system does not notice the failure until Fast Write Cache tries to destage the blocks. When this happens, Fast Write Cache marks the device as being offline. With the exception of read hits, Fast Write Cache does not accept I/O requests for the device until the successful disk synchronization operation has been performed on the failed device.

Rebooting the system also causes Fast Write Cache to destage its cache.

Temporary disk failures are those that can be fixed without major repairs, such as a disk being offline or write-protected.

Serious disk failures, such as a head crash, involve significant repair work and may result in data loss.

Note – Use the `fwcadm cache sync` and `fwcadm cache purge` commands only when the device is quiescent (that is, when all application I/O activity to the device has stopped and the cache has attempted to destage all pending writes). If the Cache Manager detects a World Wide Name mismatch because a physical drive has been replaced, you must use the `fwcadm cache redevicid` command before the `fwcadm cache sync` command or each `purge` command.

Recovering Trapped Data After a Disk Failure

When a disk fails and data blocks are held in cache, correct the condition that caused the failure or repair or replace the disk. When the fault has been corrected, recover the trapped data by issuing the following command, which destages the failed blocks and clears the offline state of the device:

```
# fwcadm cache sync diskname
```

The offline state of a device is not persistent across system reboots or invocations of the cache. Thus, it may not be necessary to issue the `fwcadm cache sync diskname` command if the cache has been stopped and restarted since the disk has been repaired.

Clearing the Offline Device State

Use the following command to clear the offline state of the failed device. This command releases all failed blocks for a device.



Caution – Use this command with caution because data is lost.

```
# fwcadm cache purge diskname
```


Theory of Operation

This appendix provides the operational theory of the Sun Fast Write Cache product. This information is background material to increase your understanding of the system administration issues discussed in Chapter 5 and Chapter 6.

- Introduction
- Buffer Management
- Fast Write Cache Modes
- Recovery
- Hints for Safe Operation

Introduction

Fast Write Cache is implemented as a UNIX device driver, using nonvolatile memory to cache synchronous write requests. As the cache fills, older data is written asynchronously to the underlying device driver. Fast Write Cache works as a layer between other disk drivers and the rest of the UNIX kernel. Stubs replace the original driver's entry points in the device switch tables. Whenever Fast Write Cache performs actual I/O (for example, when its cache must be destaged), it uses the real device driver routines.

Fast Write Cache requires minimal day-to-day administration. However, the system administrator must be familiar with how the `fwcadmin` command works. This command is described in Chapter 4 and Chapter 5.

If you encounter error conditions, see Chapter 6 and Appendix B.

Buffer Management

Except for recovery procedures after an unclean shutdown and initial configuration at start up, data is never read from NVRAM. That is, all NVRAM accesses during normal I/O operation are writes only and the cache software maintains a consistent image of the data in host memory. This method avoids excessive bus traffic to and from the offboard NVRAM at the expense of host memory usage.

When an application does a write to a cached volume, a buffer is allocated from NVRAM and the data is copied into the NVRAM buffer from host memory. The buffer is marked *dirty* and the application receives notification of the write completion. Sometime later a flush thread, after possibly coalescing small individual writes into larger writes, calls the underlying device driver to write data blocks from host memory to disk. After the data has been successfully written to disk, the associated NVRAM data buffers are made available for re-use.

The host memory buffer of the data remains available to satisfy read hits, until it is re-used according to a least recently used allocation scheme.

Note – Because the amount of host memory that backs the NVRAM cache is determined by the assignment of the `cache_mem` parameter in the `sd.cf` file, it is possible to have a much larger pool of host buffers to satisfy read hits than write buffers in NVRAM.

Fast Write Cache Modes

The Fast Write Cache software is always in one of two modes:

- Write through – write completion is returned when the data has been written to disk.
- Fast write – write completion is returned when the data is on the NVRAM card.

Note – The terms fast write, write back, and write behind are synonymous.

During normal operation, the cache is in fast write mode. The cache goes into write through mode during recovery or when an error is detected on an NVRAM card.

The goal at bootup is to successfully destage any dirty Fast Write Cache buffers. If there are dirty buffers, I/O is blocked for all devices being recovered. The cache is in write through mode for all other devices until all dirty buffers have been queued for destaging, at which time the cache goes into fast write mode. It is not recommended that the cache be disabled if there are dirty buffers in NVRAM. However, to allow a graceful shutdown for hardware repair, the Fast Write Cache software permits a shutdown if there are dirty buffers that cannot be destaged to disk but are safely stored in NVRAM. If an error has been detected on an NVRAM card, and there are dirty buffers that cannot be destaged to disk, any attempt to shutdown the cache fails.

When an I/O failure occurs, the cache mode does not change but the device is marked offline and the dirty buffers are placed on the failed queue for the device. When the problem that caused the failure is corrected, the administrator can clear the offline state of the device and synchronize or purge the data by using the `fwcadm cache -s` command or the `fwcadm cache -p` command.

Note – If a disk was replaced and the cache manager detects a changed World Wide Name, you must issue the `fwcadm cache redevid` command before data can be synchronized or new I/O can be started to the new disk.

Recovery

At startup, the cache software reads a private header on the NVRAM card. If it finds that an unclean shutdown has occurred, it reads in all metadata (device names and dirty bits) from NVRAM, opens all previously cached devices, and searches the dirty bit maps for dirty blocks. If any dirty blocks are found, the data blocks are read into host memory and queued for destaging. At this point, normal I/O operation can proceed while the dirty blocks are destaged in the background.

Hints for Safe Operation

- In general, devices must be Storage Volume (SV) disabled using the `svadm -d` or `svadm -r` commands before doing any operations requiring a device to be quiescent. Some procedures (for example, an A5000 firmware download) will not complete successfully unless the appropriate devices are SV disabled.
- The volume manager must always be layered below the Fast Write Cache. Do not cache the physical disk drives that make up a storage volume.
- Do not cache the root (/) and /usr file systems (or any file systems that come up before Fast Write Cache) because to recover disk data stored in an NVRAM card after a system crash, the data must be restored to disk before file systems are mounted or applications are written to raw volumes.
- If the NVRAM card contains dirty data, do not reconfigure your disks or move the card to another system until the dirty data is destaged to disk.
- If a disk was replaced and the cache manager detects a changed World Wide Name, you must issue the `fwcadmin cache redevid` command before data can be synchronized or new I/O can be started to the new disk. If the replaced disk belongs to a mirrored or Raid 5 volume and the volume is cached, then nothing needs to be done to the FWC. The volume management software will handle the disk replacement correctly. Note that bit replacement is handled by volume management if the disk belongs to a mirrored or RAID5 volume.

Error Messages

This appendix lists some of the error messages that may be displayed on your system when the NVRAM hardware and driver are installed. The appendix also describes causes and solutions for the errors.

TABLE B-1 Known Error Conditions

Message	Cause	Solution
WARNING: nvmemX: can't map control registers WARNING: nvmemX: can't find size WARNING: nvmemX: can't map memory	NVRAM instance <i>X</i> or the system may have a problem. The card cannot be initialized and will not be used by the system.	Replace the NVRAM card or report the problem to Sun.
WARNING: nvmemX: card did not initialize properly, base = <i>W</i> at <i>Y</i> = <i>Z</i>	NVRAM instance <i>X</i> cannot initialize location <i>Y</i> with zeros. The card will not be used by the system. <i>W</i> is the base address of the card, <i>Y</i> is the location that failed, and <i>Z</i> are the contents of location <i>Y</i> .	Replace the NVRAM card.
WARNING: nvmemX: card detects <i>Y</i> errors after writing/reading zeros	NVRAM instance <i>X</i> is getting <i>Y</i> check bit errors after initializing the memory with zeros. The card will not be used by the system.	Replace the NVRAM card.
WARNING: nvmemX: Batteries are disconnected	The batteries for NVRAM instance <i>X</i> are disabled.	Refer to Chapter " for information on how to enable the batteries.

TABLE B-1 Known Error Conditions (Continued)

Message	Cause	Solution
WARNING: nvmemX: Batteries are running low	The batteries on NVRAM instance <i>X</i> are low. If Fast Write Cache is enabled, it destages all dirty blocks on the card and passes any new I/O directly to the disk.	Replace the NVRAM card.
WARNING: nvmemX: Uncorrectable ECC error detected	NVRAM instance <i>X</i> has an uncorrectable ECC error. If Fast Write Cache is enabled, it destages all dirty blocks on the card and passes any new I/O directly to the disk. If the uncorrectable ECC error occurs during recovery, the data is read from the mirror.	Replace the NVRAM card.
NOTICE: nvmemX: memory scrubber detected an Uncorrectable ECC error Address range <i>Y</i> to <i>Z</i>	The memory scrubber detected an Uncorrectable ECC error within the address range <i>Y</i> to <i>Z</i> on NVRAM instance <i>X</i> . The Uncorrectable ECC error WARNING will follow.	Replace the NVRAM card.
WARNING: nvmemX: Correctable ECC error limit <i>Y</i> reached	The number of Correctable ECC errors detected has reached the error limit <i>Y</i> (the default is 10000). NVRAM instance <i>X</i> may have a problem. If Fast Write Cache is enabled, it destages all dirty blocks on the card and passes any new I/O directly to the disk.	Replace the NVRAM card.
WARNING: nvmem: not all mirrors enabled, expected <i>X</i> actual <i>Y</i>	The number of mirrored pairs is less than expected.	Report the problem to Sun.
WARNING: nvmemX: nvram_init, can't start scrubber thread	A software or system resource problem was detected (possibly out of memory).	Report the problem to Sun.

TABLE B-1 Known Error Conditions (Continued)

Message	Cause	Solution
WARNING: nvmemX: nvrām_bcopy, write exceeds memory range	A software failure was detected; trying to write past the end of NVRAM.	Report the problem to Sun.
WARNING: nvmemX: nvrām_bcopy, read exceeds memory range	A software failure was detected; trying to read past the end of NVRAM.	Report the problem to Sun.
WARNING: nvmem: nv_init, config changed, initializing memory	The card configuration has changed since the last boot time. Any new cards are zeroed out and labeled internally, and old cards are handled accordingly.	This message is informational only and requires no action unless NVRAM configuration changes were not made to the system. In that case, report the problem to Sun.
WARNING: nvmem: nv_init, primary X can't find secondary mirror card instance Y	During initialization, the primary mirror instance X could not find its secondary mirror NVRAM card instance Y. The card may have been removed or it has a hardware problem. If the primary mirror is dirty, the driver uses it to recover data when Fast Write Cache is started. Since the primary does not have a secondary mirror, Fast Write Cache destages all dirty blocks on the card and passes any new I/O directly to the disk.	If NVRAM configuration changes were not made, and there were no previous error messages about NVRAM instance Y, report the problem to Sun.
WARNING: nvmem: nv_init, primary X can't find secondary mirror instance Y	During initialization, the primary mirror instance X could not find its secondary mirror instance Y. The secondary mirror card may have been replaced. The driver initializes the new card as a secondary mirror and resynchronizes it with primary X.	If NVRAM configuration changes were not made, report the problem to Sun.

TABLE B-1 Known Error Conditions *(Continued)*

Message	Cause	Solution
WARNING: nvmem: nv_init, Inconsistent NVRAM mirror configuration, mirror disabled	A configuration failure has been detected. Fast Write Cache will not be enabled.	Report the problem to Sun.
WARNING: nvmem: nvram_init, mirror not configured	The mirror is not configured because of inconsistencies found or insufficient number of working NVRAM cards detected. Fast Write Cache will not be enabled.	Replace the failed NVRAM card and ensure that a minimum of two cards is installed.
WARNING: nvmem: nv_init, non- matching card sizes, mirror disabled	The mirror is not configured because of mismatched NVRAM card sizes. Fast Write Cache will not be enabled.	If the cards are the same model, report the problem to Sun. Otherwise, install the same size NVRAM cards.

SunVTS Diagnostic Software

This appendix contains an overview of the SunVTS diagnostic tool and describes how to use SunVTS with NVRAM cards. This information is for experienced SunVTS users.

This appendix is organized as follows:

- SunVTS Overview
- Preparing to Run SunVTS
- Running SunVTS

SunVTS Overview

The SunVTS software executes multiple diagnostic hardware tests from a single user interface and is used to verify the configuration and functionality of most hardware controllers and devices. SunVTS operates primarily from the OpenWindows™ and Common Desktop Environment (CDE) user interfaces, enabling test parameters to be set quickly and easily while a diagnostic test operation is being performed.

The NVRAM SunVTS (`SUNWvtsnp`) tests the memory and error detection and correction (EDC) functionality of the NVRAM card. This test runs in 32-bit and 64-bit operating environments.

Because `SUNWvtsnp` is not a stand-alone package, the appropriate SunVTS must be installed.

Solaris Version	SunVTS version
Solaris 2.6	SUNWvts 2.1.3
Solaris 7	SUNWvts 3.0
Solaris 7	SUNWvts 3.1

The Solaris 2.6 version of the `SUNWvtsnp` package contains a 32-bit `nvtest`.

The Solaris 7 versions of the `SUNWvtsnp` package contain both 32-bit and 64-bit `nvtest` binaries. Package dependency is checked during installation, and the `nvtest` binaries are installed in the appropriate directories. The 64-bit `nvtest` is not installed unless the `SUNWvtsx` (64-bit SunVTS) package is already installed.

The SunVTS documents are listed in TABLE C-1.

TABLE C-1 SunVTS Documentation

Title	Part Number	Description
<i>SunVTS 2.1 User's Guide</i> (Solaris 2.6)	802-7299	Describes the SunVTS environment, including starting and controlling various user interfaces.
<i>SunVTS 2.1.3 Test Reference Manual</i> (Solaris 2.6)	805-4163	Describes each SunVTS test; provides various test options and command-line arguments.
<i>SunVTS 2.1 Quick Reference</i> (Solaris 2.6)	802-7301	Quick reference for starting and using SunVTS.

TABLE C-1 SunVTS Documentation

Title	Part Number	Description
<i>SunVTS 3.0 User's Guide</i> (Solaris 7)	805-4442	Describes the SunVTS environment, including starting and controlling various user interfaces.
<i>SunVTS 3.0 Test Reference Manual</i> (Solaris 7)	805-4443	Describes each SunVTS test; provides various test options and command-line arguments.
<i>SunVTS 3.0 Programmer's Guide</i> (Solaris 7)	805-7338	Describes how to create, integrate, and package diagnostic tests for the SunVTS environment.
<i>SunVTS 3.0 Quick Reference</i> (Solaris 7)	805-4444	Quick-reference for starting and using SunVTS.
<i>SunVTS 3.1 Test Reference Manual</i> (Solaris 7)	805-7407	Describes each SunVTS test; provides various test options and command-line arguments.
<i>SunVTS 3.1 Programmer's Guide</i> (Solaris 7)	806-0452	Describes how to create, integrate, and package diagnostic tests for the SunVTS environment.
<i>SunVTS 3.1 Quick Reference</i> (Solaris 7)	805-7408	Quick-reference for starting and using SunVTS.

Preparing to Run SunVTS

The SunVTS test for NVRAM cards is a destructive test. It destroys any data currently on the cards. There are several other SunVTS tests that are also destructive. Therefore, run the NVRAM test only after disabling the Sun StorEdge™ Fast Write Cache driver.

To prevent data corruption, the whole system should be inactive and preferably not used at all during the test, as these procedures greatly reduce the speed of all data file systems on the system.

▼ To Prepare to Run SunVTS

The Fast Write Cache driver must be disabled before testing the NVRAM devices.

- **Disable the Fast Write Cache driver by typing the following command as superuser:**

```
# fwcadm cache -d
```

Running SunVTS

The following procedures describe how to invoke SunVTS locally and remotely.

▼ To Invoke SunVTS Locally (OpenWindows™)

1. **To invoke SunVTS locally on a machine running the OpenWindows environment, type the following as superuser:**

```
# cd /opt/SUNWvts/bin
# ./vtsk
# ./vtsui.ol
```

2. **Make the following selections from the SunVTS menus:**
 - a. **Select** Select Tests...
 - b. **Select Test Mode** Functional Test.
 - c. **Select Test Set** None.
 - d. **Select Test Intervention** Enabled.
 - e. **Scroll down to** Other Devices **and select the** nvramX **devices, for the NVRAM devices you have. (Use the** fwcadm nvram -s **command to display your devices.)**
 - f. **Select Start to start the test.**

3. After the test finishes, enable Fast Write Cache by typing the following command as superuser:

```
# fwcadm cache -e
```

▼ To Invoke SunVTS Locally (CDE)

1. To invoke SunVTS locally on a machine running CDE, type the following as superuser:

```
# cd /opt/SUNWvts/bin  
# ./sunvts
```

2. Make the following selections from the SunVTS menus:

- a. Select `None` from `Select Devices`.
- b. Enable intervention from `Select Devices`.
- c. Select `Functional Test` from `Select Mode`.
- d. Scroll down to `Other Devices` and select the `nvrAmX` devices, for the NVRAM devices you have. (Use the `fwcadm nvrAm -s` command to display your devices.)
- e. Select `Start` to start the test.

3. After the test finishes, enable Fast Write Cache by typing the following command as superuser:

```
# fwcadm cache -e
```

▼ To Invoke SunVTS Remotely

1. To invoke SunVTS remotely on a machine running either the CDE or OpenWindows environment, type the following on the machine that runs SunVTS:

```
% /usr/openwin/bin/xhost nfs_machine
```

2. On the machine to be tested, type the following as superuser:

For CDE:

```
# cd /opt/SUNWvts/bin
# ./sunvts -display display_machine:0.0
```

For OpenWindows:

```
# cd /opt/SUNWvts/bin
# ./vtsk
# ./vtsui.ol -display display_machine:0.0
```

3. Make the following selections from the SunVTS menus:

- a. **Select** Select Tests.
- b. **Select** Functional Test **from Test Mode.**
- c. **Select** None **from Test Set.**
- d. **Select** Enabled **from Test Intervention.**
- e. **Scroll down to the** Other Devices **and select the nvramX devices, for the NVRAM devices you have. (Use the fwcadm nvram -s command to display your devices.)**
- f. **Select** Start **to start the test.**

Note – If any tests show errors in the card, have Sun Enterprise Services replace the card.

4. After the test finishes, enable Fast Write Cache by typing the following command as superuser:

```
# fwcadm cache -e
```

Glossary

alert	Message sent to identify a problem or impending problem.
cache	Buffer storage that is faster than main storage.
clean buffer	Buffer containing a valid disk block image that has been destaged to disk.
clean shutdown	System is halted according to the instructions that come with the operating system.
destage	To move data from the cache to the disk.
dirty buffer	Buffer containing a disk block image that has not been destaged to disk.
disk	Host view of the logical storage device. Disks are also referred to as volumes.
invalid buffer	Buffer that does not contain a disk block image.
nonvolatile memory	Memory that remains intact even if the system loses power.
NVRAM card	Single-wide SBus or PCI memory card with battery back-up error detection and correction (EDC) static memory.
mirroring	Process of maintaining two or more identical images of a designated disk volume.
node	See <i>subsystem</i> .
RAID	Redundant Array of Independent Disks.
SCM	Storage Cache Management.
Storage Cache Management	Provides quick access to recently used data or data that is physically close to recently accessed data.
Storage Volume driver	Enables standard system call access to a disk device to be redirected into the storage software, allowing standard applications to use the storage features.

- stub** New pointer and the function to which it points. Stubs are used to redirect or intercept the function calls away from the original driver over to the new driver (new layer). This can be done by saving the original driver entry pointers and replacing them with the new driver pointers. When the new driver goes away, it restores the original pointers. This enables the new driver to add and remove itself without affecting the driver above or below it.
- subsystem** A combination of storage hardware (channel adapters, device adapters, CPUs, cache memory) and software that is independently powered and configured. Subsystems are sometimes referred to as nodes.
- SV** Storage Volume driver.
- unclean shutdown** A shutdown that may leave data in the Fast Write Cache.
- volume** See *disk*.

Index

A

abnormal operations, 6-1
alert, Glossary-1

B

batteries
 disabling, 5-3
 enabling, 2-4, 2-5, 2-7
battery status, 4-19
boot time
 new kernel, 6-2
buffer
 dirty, Glossary-1
 management, A-2

C

cache, Glossary-1
 modes, A-3
 statistics, 5-7
card
 batteries, 2-4, 2-5, 2-7
 installation, 2-8, 2-10
 status, 2-9, 2-11
clean
 buffer, Glossary-1
 shutdown, 6-2, Glossary-1
clearing offline device state, 6-5
commands
 fwcadm, 4-1
compliance, 2-2

configuration file
 sd.cf, 4-22
 sv.cf, 4-23
correctable ECC errors, 4-20
correction capability, 2-2

D

data recovery, 6-5
destage, Glossary-1
detection of errors, 2-2
device
 mirror, 4-20
dirty
 bit, 4-17
 buffer, Glossary-1
disabling batteries, 5-3
discharge precautions, 2-3
disks, Glossary-1
 errors, 6-4
 removing packs, 5-7
displaying
 cache statistics, 5-7
documentation
 SunVTS, C-2

E

electrostatic discharge precautions, 2-3
EMC compliance, 2-2
enabling caching
 file system, 5-6

- volume, 5-6
- enabling the NVRAM card batteries, 2-4, 2-5, 2-7
- error detection, 2-2
- errors
 - conditions, B-1
 - correctable, 4-20
 - disk, 6-4
 - uncorrectable, 4-19

F

- Fast Write Cache hardware
 - removing, 5-2
- features, 2-2
- file systems
 - should be cached, 5-4
 - should not be cached, 5-5
- fwcadm command, 4-1
 - fwcadm nvadmin, 4-16
 - fwcadm volume, 4-21

H

- halting the system, 6-2
- hardware, removing, 5-2
- hints
 - safe operation, A-4

I

- installation
 - NVRAM card, 2-8, 2-10
 - verification, 2-9, 2-11
- invalid buffer, Glossary-1
- invoking SunVTS
 - locally (CDE), C-5
 - locally (Open Windows), C-4
 - remotely, C-5

M

- mirror device, 4-20
- mirroring, Glossary-1
- modes, cache, A-3
- moving an NVRAM card, 5-2, 6-3

N

- node, see subsystem, Glossary-1
- nonvolatile memory, Glossary-1
- normal operations, 5-1
- NVRAM
 - card, Glossary-1
 - card batteries, 2-4, 2-5, 2-7
 - card installation, 2-8, 2-10
 - moving, 5-2
 - moving a card, 6-3
 - removing, 5-2
 - status, 4-16

O

- offline state, 6-5
- operations
 - normal, 5-1
- overview
 - SunVTS, C-2

P

- performance capability, 1-1
- power failures, 6-2
- precautions
 - discharge, 2-3
- preparing to run SunVTS, C-3

R

- RAID, Glossary-1
- recovering trapped data, 6-5
- recovery, A-4
- removing
 - disk packs, 5-7
 - Fast Write Cache hardware, 5-2
 - NVRAM card, 5-2
- running SunVTS, C-4

S

- safe operation
 - hints, A-4
- SCM, Glossary-1

- sd.cf configuration file, 4-22
- selecting
 - file systems, 5-4
 - volumes, 5-4
- shutdowns
 - clean, 6-2, Glossary-1
 - unclean, 6-2, Glossary-2
- status, 2-9, 2-11
 - batteries, 4-19
 - NVRAM boards, 4-16
- Storage Cache Management, Glossary-1
- Storage Volume (SV), 4-21
- Storage Volume (SV) driver, Glossary-1
- stub, Glossary-2
- subsystem, Glossary-2
- SunVTS, C-1
 - documentation, C-2
 - invoking locally (CDE), C-5
 - invoking locally (Open Windows), C-4
 - invoking remotely, C-5
 - overview, C-2
 - running, C-4
- sv.cf configuration file, 4-23
- SV, see Storage Volume driver, Glossary-2

U

- unclean
 - shutdown, 6-2, Glossary-2
- uncorrectable ECC errors, 4-19

V

- verifying the installation, 2-9, 2-11
- volumes
 - see disk, Glossary-2

