



DIGITAL StorageWorks HSZ50 Array Controller Operating Software

HSOF Version 5.2 Release Notes

Order Number: EK-HSZ50-RN. D01

This document summarizes features and characteristics of the HSZ50 array controller operating software Version 5.2 that are not covered elsewhere in the documentation. These release notes also contain instructions for installing the software and should be retained for future reference.

Software Version: HSOF V5.2

**Digital Equipment Corporation
Maynard, Massachusetts**

December 1997

While Digital Equipment Corporation believes the information included in this manual is correct as of the date of publication, it is subject to change without notice. DIGITAL makes no representations that the interconnection of its products in the manner described in this document will not infringe existing or future patent rights, nor do the descriptions contained in this document imply the granting of licenses to make, use, or sell equipment or software in accordance with the description. No responsibility is assumed for the use or reliability of firmware on equipment not supplied by DIGITAL or its affiliated companies. Possession, use, or copying of the software or firmware described in this documentation is authorized only pursuant to a valid written license from DIGITAL, an authorized sublicensor, or the identified licensor.

Commercial Computer Software, Computer Software Documentation and Technical Data for Commercial Items are licensed to the U.S. Government with DIGITAL's standard commercial license and, when applicable, the rights in DFAR 252.227 -7015, "Technical Data—Commercial Items."

© Digital Equipment Corporation 1997.

Printed in U.S.A.
All rights reserved.

Alpha, CI, DCL, DECconnect, DECserver, DIGITAL, DSSI, HSC, HSJ, HSD, HSZ, MSCP, OpenVMS, StorageWorks, TMSCP, VAX, VAXcluster, VAX 7000, VAX 10000, VMS, VMScluster, and the DIGITAL logo are trademarks of Digital Equipment Corporation.

UNIX is a registered trademark in the United States and other countries, licensed exclusively through X/Open Company, Ltd. IBM AIX is a registered trademark of International Business Machines Corporation. Intel is a registered trademark of Intel Corporation. HP-UX is a registered trademark of Hewlett-Packard Company. Novel Netware is a registered trademark of Novel, Inc. Sun and Solaris are registered trademarks of Sun Microsystems, Inc. Windows NT and NT are trademarks of Microsoft Corporation. Windows is a registered trademark of Microsoft Corporation.

All other trademarks and registered trademarks are the property of their respective holders.

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 not installed and used in accordance with the instruction manual, 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. Restrictions apply to the use of the local-connection port on this series of controllers; failure to observe these restrictions may result in harmful interference. Always disconnect this port as soon as possible after completing the setup operation. Any changes or modifications made to this equipment may void the user's authority to operate the equipment.

Warning!

This is a Class A product. In a domestic environment this product may cause radio interference in which case the user may be required to take adequate measures.

Achtung!

Dieses ist ein Gerät der Funkstörgrenzwertklasse A. In Wohnbereichen können bei Betrieb dieses Gerätes Rundfunkstörungen auftreten, in welchen Fällen der Benutzer für entsprechende Gegenmaßnahmen verantwortlich ist.

Avertissement!

Cet appareil est un appareil de Classe A. Dans un environnement résidentiel cet appareil peut provoquer des brouillages radioélectriques. Dans ce cas, il peut être demandé à l'utilisateur de prendre les mesures appropriées.

Contents

Introduction	1
Topics Covered.....	1
Release Package Contents	2
Intended Audience	2
Hardware and Software Support	3
Hardware Support.....	3
Operating System Support	4
Host Adapter Support	4
Device Support	5
New Features of Version 5.2	9
Summary	9
Items from Previous Release	10
DILX Enhancement	10
Ability to Install New Device Microcode	10
Clarifications	11
Battery Handling	11
Battery Handling and CACHE_UPS	11
Battery Handling and Cache Policy	12
Failed Battery Action	13

SHOW on Both Controllers	13
Logical Block Address Does Not Match for MEDIUM ERROR	14
Adding RAIDsets When Battery Condition is Low	14
Drive-Level Event Reports	14
HSZ50 Recognition by DIGITAL UNIX Versions Prior to Version 4.0.	14
Logical Device Event Reports	14
Using SAVE_CONFIGURATION to Save the Subsystem's Configuration	15
Operating Constraints	17
Maintenance Terminal Port Speeds	17
Limits on the Total Number of Stagesets	17
Partitioning Not Supported With SCSI Multiple Bus Failover	18
SCSI Multiple Bus Failover Host Operating System Support	18
CFMENU Constraints.....	18
CLONE Utility Constraints.....	18
CLEAR_ERRORS LOST_DATA command	18
Only One Qualifier per CLI Command Recommended	19
Avoiding Problem Situations	20
Use of the CACHE_UPS Qualifier	20
Device Adds, Moves, and Changes	20
Tape Drive Firmware Revision Upgrade	21
Running CONFIG or CFMENU Utilities during Backup or Copy Operations	21
Changing Host Adapter Types	21
CLONE Utility	21
Extending the Life of Your Write-Back Cache Batteries	21
Shutting Down a Controller	22
Disabling Write-Back Cache Batteries	23
Restarting the Controller	23
Documentation Additions and Corrections	24
Fault Management Utility (FMU) Codes	24
Executive Services Last Failure	24
Host Interconnect Port Service Last Failure	24
Last Failure Codes No Longer Used	25
Upgrading Array Controller Software	26
Program Card Upgrade (Single Controller Configuration)	26
Program Card Upgrade (Dual-Redundant Configuration)	28

Identifying Your HSOF Software Revision Level	29
Order Numbers	30
Appendix A.....	31
Possible Problem with Disks Initialized with SAVE_CONFIG Under HSOF Version 2.7 on HSZ40/20/SWXRC Controller	31
How to Detect Whether You Have This Problem	33
Windows NT platforms	33
Novell NetWare platforms	33
Sun Solaris and Sun OS Platforms.....	34
OpenVMS platforms	35
DIGITAL UNIX platforms.....	35
AIX platforms	36
AIX 4.1.4:	36
AIX 3.2.5:	36
HP-UX platforms	36
Solution	37

Introduction

These release notes provide information for the HSZ50 array controller operating software Version 5.2¹ not covered elsewhere in the documentation. This document should be used by individuals responsible for configuring, installing, and using the HSZ50 controllers.

Read this entire document before installing or upgrading the software.

Topics Covered

These release notes cover the following topics:

- Hardware and software supported by HSOF Version 5.2Z-*x* software (where *x* represents number of patches associated with this version of software that are correctly installed)
- New features— Modifications in the software that add significant functionality to the controller
- Items from previous releases
- Clarifications—Explanations of controller behavior in certain situations
- Operating constraints—Limitations placed on the operation of the controller by the nature of its design
- Avoiding problem situations—Information to help you avoid and recover from unusual controller behavior in certain situations, under specific conditions
- Corrections and additions to the documentation
- HSOF software installation/upgrade instructions
- Identification of the software revision level
- Order numbers

¹ The OpenVMS™ SHOW CLUSTER command and the HSOF software SHOW THIS_CONTROLLER command display the version as “V52Z”.

Release Package Contents

The Version 5.2 release package consists of the following:

- A cover letter
- The HSZ50 documentation set which includes many documents produced for HSOF Version 5.1, all still applicable to HSOF Version 5.2:
 - *DIGITAL Storageworks HSZ50 Array Controller
HSOF Version 5.1
Configuration Manual*
 - *DIGITAL Storageworks HSZ50 Array Controller
HSOF Version 5.1
Service Manual*
 - *DIGITAL Storageworks HSZ50 Array Controller
HSOF Version 5.1
CLI Reference Manual*
 - *DIGITAL Storageworks HSZ50 Array Controller
Operating Software HSOF Version 5.2 Release Notes*
 - *DIGITAL Storageworks Family Array Controller
Operating Software (HSOF), Version 5.2 Software Product Description*
- A PCMCIA program card containing HSOF software Version 5.2

Intended Audience

This document has been prepared for DIGITAL customers who have purchased HSZ50 array controllers and for DIGITAL Multivendor Customer Services personnel responsible for installing and maintaining systems that include HSZ50 array controllers.

Hardware and Software Support

This section lists the hardware and software compatible with HSOF Version 5.2 software.

Hardware Support

HSOF Version 5.2 software supports the following revisions for the HSZ50 controllers and associated hardware:

- HSZ50 controller modules, hardware revision A or B
- Version 3 cache module, hardware revision A or B
- BA350-MA controller shelf
- BA350-Sx device shelf
- BA356-Sx wide device shelf with 8-bit or 16-bit personality module (controller operates in 8-bit mode only)
- BA35x-HF power supply
- HS35X-BA Single External Cache Battery SBB (with 1 ECB)
- HS35X-BB Dual External Cache Battery SBB (with 2 ECBs)
- HSSIM-AA 32 MB SIMM Pack for Version 3 Cache module upgrades

Operating System Support

HSOF Version 5.2 software on HSZ50 controllers is supported by the following operating system Versions, within the limitations described in *DIGITAL Storageworks HSZ50 Array Controller, HSOF Version 5.1 Configuration Manual*:

- DIGITAL UNIX™ Versions 3.2c, 3.2d, 3.2g, and 4.0b.
- DIGITAL UNIX™ Versions 4.0 and 4.0a with the following patches:
 - V4.0 requires patch OSF 400-114.
 - V4.0a requires patch OSF 405-034.
- OpenVMS Alpha Version 6.2, 7.0, and 7.1.
- Windows NT® Server Version 3.51 and 4.0.

Host Adapter Support

The following host adapters are supported by the DIGITAL UNIX operating system for HSZ50 controllers:

- KZTSA (for DEC 3000™ systems) adapter firmware Version A11
- PMAZC (for DEC 3000 systems – Requires a DWZZ-series signal converter) adapter firmware Version 2.0.
- KZMSA (for DEC 7000™ and DEC 10000™ systems and for DIGITAL AlphaServers 8200 and 8400 – Requires a DWZZ-series signal converter.)
- KZPSA (for DIGITAL AlphaServers 1000/2000/2100/8200/8400/4000/4100) adapter firmware Version A10 and higher

The following host adapters are supported by the OpenVMS Alpha operating system for HSZ50 controllers:

- KZTSA (for DEC 3000 systems) adapter firmware Version A11.
- KFTIA (for TurboLaser 8200 embedded SCSI) adapter firmware Version 2.46.
- PMAZC (for DEC 3000 systems – requires a DWZZ-series signal converter) adapter firmware Version 2.1.
- KZMSA (for DEC 7000 and DEC 10000 systems – Requires a DWZZ-series signal converter.)
- KZPAA (for DIGITAL AlphaServers 1000/2000/2100 systems.)

- KZPSA (for DIGITAL AlphaServers 1000/2000/2100/8200/8400/4060/4100) adapter firmware Version A10 and higher

The following host adapter is supported by the Windows NT operating system Version 3.51 for HSZ40 controllers:

- KZPSA (for DIGITAL AlphaServers 400/1000/2000/2100) adapter firmware version A10 and higher

Device Support

HSOF Version 5.2 Software supports the devices listed in Table 1 through Table 4 at the indicated hardware and microcode levels or higher.

Table 1: Supported Disk Drives

Device	Capacity in Gigabytes	Minimum Microcode Version	Minimum Hardware Version
RZ25-VA	0.43	0900	B01
RZ26-VA	1.05	T392	D02
RZ26L-VA/VW	1.05	440C	A01
RZ26N-VA/VW	1.05	446	A01
SWXD3-SF/WF	1.05	446	A01
DS-RZ1BB-VW	2.1	LYJO/0656	A01
DS-RZ1CB-VW	4.3	LYJO/0656	A01
DS-RZ1DB-VW	9.1	LYJO/0307	A01
DS-RZ26N-VZ	1.05	1003	A01
RZ28-VA/VW	2.10	435E	A01
RZ28B-VA	2.10	0003	A01
RZ28D-VA/VW	2.10	0008	A01
SWXD3-SG/WG	2.10	0008	A01
RZ28M-VA/VW	2.10	0466	A01
DS-RZ28M-VZ	2.10	1003	A01
SWXD3-SH/WH	2.10	0466	A01
RZ29B-VA/VW	4.3	0007	B01
SWXD3-SE/WE	4.3	0007	C02/A01
DS-RZ40-VA	9.1	LYGO	A01
RZ74-VA	3.57	T427B	B07

Table 2: Supported Tape Drives

Device	Capacity in Gigabytes	Minimum Microcode Version	Minimum Hardware Version	Notes
TL812	960/1920	1.2 robot /CC33 drive	A01	1,2,3,4,5
TL822	5280/10560	1g4F robot /CC33 drive	A01	1,2,3,4,5
TL826	3520/7040	1g4F robot /CC33 drive	A01	1,2,3,4,5
DS-TL893	924/1848T	V2A/5A	A01	1,2,3,4,5
DS-TL894	1.68/3.36T	V1.24	A01	1,2,3,4,5
DS-TL896	6.16/12.32T	V2A/5A	A01	1,2,3,4,5
TZ87-VA	10/20	930A	A01	3,5,6
TZ87N-VA	10/20	930A	A01	3,4,5,6
TZ87-TA	10/20	9514	B02	1,2,3,6
TZ875-NT	50/100	930A	A01	1,3,5,6
TZ875-TA	50/100	930A	A01	1,3,5,6
TZ877-AE/AF	70/140	930A	A01	1,3,5,6
TZ88N-VA	20/40	CC33	A01	3,4,5,6
TZ885-NT/NE	100/200	CC33	A01	1,3,4,5,6
TZ887- NT/NE	140/280	CC33	A01	1,3,4,5,6
DS-TZ89N-VW	35/70	141F	A01	3,4,5,6
DS-TZ89N-TA	35/70	141F	A01	1,3,4,5,6
DS-TL890	560/1120	3.23 robot/V55 drive	A01	1,2,3,4,5
DS-TL891	350/700	3.23 robot/V55 drive	A02	1,2,3,4,5
DS-TL892	350/700	3.23 robot/V55 drive	A02	1,2,3,4,5

Notes:

Tape devices are not supported on Windows NT™, NOVELL™, Sun Solaris™, HP-UX™ or IBM AIX™ systems.

In the Capacity column, T = Terabytes

1. Requires 0.2 meter SCSI-1 to SCSI-2 transition cable, DIGITAL internal part number 17-03831-01 for DWZZA-AA, and DIGITAL part number 17-04367-01 for Storage building blocks (SBB) DWZZA-VA and DWZZB-VW.
2. Requires DWZZA/DWZZB single-ended to differential SCSI signal converter.

3. Capacity values represent compressed data. The compression factor is device dependent based on individual device algorithms.
4. Cannot read TK50, TK70 or TZ30 formatted tapes.
5. Requires a KZPSA or PMAZC host adapter.
6. Tape device code load is supported using HSUTIL.

Table 3: Supported Solid State Disk Drives

Device	Capacity in Gigabytes	Minimum Microcode Version	Minimum Hardware Version	HSUTIL Microcode Load Supported?
EZ31-VW	0.134	V064	A01	Yes, See Note 2
EZ32-VW	0.268	V064	A01	Yes, See Note 2
EZ51R-VA	0.10	V096	D01	Yes
EZ54R-VA	0.42	V096	A01	Yes
EZ58R-VA	0.85	V109	C01	No
EZ64-VA	0.475	V064	A01	Yes
EZ64-VW	0.475	V070	A01	Yes
EZ69-VA	0.950	V064	A01	Yes
EZ69-VW	0.950	V070	A01	Yes

Notes:

1. Do not warm-swap solid-state disk drives. Make sure power to the device shelf is turned off before removing or inserting this device. This applies to all solid-state disk drives.
2. The EZ31 and EZ32 are supported as devices in HSOF Version 5.1, but device code load to this drive with HSUTIL is not supported with an entry in the internal HSUTIL table. Code load to the EZ31 and EZ32 may be performed using the "unsupported device" code load feature of HSUTIL. The steps are as follows:
 - a. answer the "unsupported device" question: yes
 - b. total size: 1024
 - c. single write buffer: no
 - d. buffer size: 8
 - e. download microcode and SAVE: yes
 - f. bytes reversed: no

Table 4: Supported CD-ROM Readers

Device	Capacity in Gigabytes	Minimum Microcode Version	Minimum Hardware Version
RRD42-VB/VU	0.6	1.1a	A01
RRD43-VA	0.6	0064	A02
RRD44-VA	0.6	3493	A02
RRD45-VA/VU	0.6	1645	A01

Notes:

1. CD-ROM drives are only supported under the DIGITAL UNIX and OpenVMS operating systems.
2. Do not warm-swap CD-ROM drives. Make sure power to the device shelf is turned off before removing or inserting this device.

New Features of Version 5.2

Summary

HSOF Version 5.2 provides the following:

- Corrects problems that were identified after the release of HSOF Version 5.1.
- Incorporates all patches issued to HSOF Version 5.1.
- Improves internal battery test algorithms to resolve some anomalies in Version 5.1 battery testing.
- Provides support for the following new devices:
 - DS-RZ1BB-VW
 - DS-RZ1CB-VW
 - DS-RZ1DB-VW
 - DS-TL890
 - DS-TL891
 - DS-TL892

Items from Previous Release

The following list briefly describes features introduced in HSOF Version 5.1 that are not described in the HSOF 5.1 documentation set.

DILX Enhancement

The Disk In-Line Exerciser (DILX) acts on disks configured as single-device, RAIDset, stripeset, mirrorset, stripe-mirrorset or partitioned units.

Ability to Install New Device Microcode

The HSUTIL program can be used to install new microcode to tape drive TZ89 and solid state disk drives EZ51, EZ54, EZ64, or EZ69.

Clarifications

This section presents clarifications on controller behavior in certain situations.

Battery Handling

Diagnostic testing is performed on the write-back cache batteries when the controller first initializes, and then periodically after the batteries have become fully charged. This section clarifies battery-related topics.

Battery Handling and CACHE_UPS

If the CACHE_UPS qualifier is activated, the following behavior is true at all times:

- no battery checks are performed
- no battery condition is reported or messages sent
- no action is taken in the event of battery failure

When the CACHE_UPS qualifier is enabled, the controller ignores the condition of the write-back cache batteries because an uninterruptible power supply (UPS) is assumed to be in use to maintain power to the write-back cache module in the event of a power failure. Note that the HSOF software does not have the capability to monitor the remaining capacity of an UPS. There are many variables in an UPS configuration; some configurations might put data at risk if the UPS is exhausted. For example, if UPS power is lost to the host/storage system without warning, write-back cache data is maintained by the cache batteries only. With CACHE_UPS set, the state of the batteries is unknown and data could be at risk.

Caution

Use of this setting without a functional UPS in place could result in data loss if power is interrupted and batteries have failed.

Battery Handling and Cache Policy

The setting chosen for Cache Policy (A or B) affects access to RAIDsets and Mirrorsets during the initial test only, and determines the cache mode the controller uses during initial test only.

Initial Test begins when the controller is turned on and continues until the batteries are charged, or up to ten hours. Initial testing includes a recharge of the battery after potential discharge during shut down. Every four minutes, the software tests the battery. Full caching operations can begin when the batteries are fully charged.

Cache Policy A is the default setting. Cache Policy A and B apply only to a “low” battery situation during controller initialization until either:

1. the battery becomes fully charged or
2. 10 hours elapse without the battery becoming fully charged (in this case, failed battery action is taken, as described in the “Failed Battery Action”. section of these release notes.)

Once a battery has become fully charged, Cache Policy A and B no longer applies, and regular periodic battery tests are performed.

The following table summarizes the effect of both cache policy settings on access to RAIDsets and mirrorsets when a battery is “low” during the initial test period.

Cache Policy	RAID/Mirrorset access	Cache Mode
Cache Policy A	no access	write-through on individual disk units (also called JBOD)
Cache Policy B	access	write-through on RAID/mirrorsets and individual disk units

Note

Write-through cache mode is not a 100% guarantee of user data consistency within storage sets (RAID and mirror sets). This is because in these storage sets, there is an inherent delay in the writing of data to the different members. If a power failure occurs between the writing of some members and others (an extremely narrow window), data on the various members will be inconsistent. Therefore, Cache Policy A (which denies access to RAID/mirror sets on controllers with low batteries) is the only 100% guarantee of user data consistency. That is why Cache Policy A enforces loss of storage set access in the event of a battery which does not pass the battery test.

HSOF software makes every attempt to notify the user that a battery problem (not good or open circuit) has occurred by printing a message to the console and posting an error to the host errorlog. However, not all operating systems present errorlogs reliably.

Failed Battery Action

If a battery passes the initial test, cache policy no longer applies, and periodic testing begins. During periodic testing, the battery is tested every 24 hours. When a failed battery is detected during either initial or periodic testing, the following failed battery action is taken:

Dual-controlled HSZ50s (partner is running and has good batteries):

- controller performs controlled shut down
- all units failover to partner
- no auto-reboot after shut down
- manual restart before battery replacement causes units to “failback” leading to no access to RAID sets and mirror sets

Single-controller HSZ50:

- no shut down
- no RAID set or mirror set access permitted
- individual disk units accessed in write-through mode

SHOW on Both Controllers

Some devices, such as tape devices, might not show correctly on one controller if they are preferred to the other controller. If the expected information is not displayed from one controller, use the SHOW command again to view the information on the other controller.

Logical Block Address Does Not Match for MEDIUM ERROR

The Logical Block Address (LBA) is a number used to identify a block on a mass storage device. The LBA that appears in the Information field of an Event Log (%EVL) display when the Sense Key field is MEDIUM ERROR does not agree with the LBA number contained in the Information field of the extended sense data that is sent to the host system. It is one less than the number in the extended sense data.

Adding RAIDsets When Battery Condition is Low

You can use the ADD RAIDSET and ADD MIRRORSET commands, regardless of the write-back cache battery condition. However, if the CACHE_POLICY is set to A and the batteries are low, the controller does not allow access to any RAIDsets or mirrorsets. CACHE_POLICY A requires that the cache batteries be fully charged before you can use RAIDsets or mirrorsets. Refer to the *HSOF Array Controller Version 5.1 CLI Reference Manual* for additional information about CACHE_POLICY.

Drive-Level Event Reports

Under certain circumstances the "devtype" (device type) field contained in drive-level event reports %EVL displays as well as error logs, will contain the value 1F(hexadecimal). When "devtype" is 1F the content of the "device identification" and "device serial number" fields are undefined and should be ignored.

HSZ50 Recognition by DIGITAL UNIX Versions Prior to Version 4.0.

For the DIGITAL UNIX operating system to recognize the HSZ50 array controller, the "/sys/data/cam_data.c" must be modified. Duplicate the section for the HSZ40 entry. Change HSZ4 to HSZ5. Rebuild the kernel and copy it to the root. Restart the operating system.

Logical Device Event Reports

Under certain circumstances the Port, Target, and LUN fields contained in event reports associated with storageset logical devices, Event Log displays, as well as sense data responses, will be set to 255 (decimal) instead of the Port, Target, and LUN of the first physical device in the storageset. In addition, when the Port, Target, LUN are set to 255, the Device Type is set to 0 (magnetic disk device), and the Device Product ID and Device Firmware Revision Level fields will be ASCII space filled.

Using SAVE_CONFIGURATION to Save the Subsystem's Configuration

The SAVE_CONFIGURATION qualifier is intended to be used to provide a means to restore nonvolatile memory (NVRAM) contents when a controller in a single-controller module configuration is replaced. The SAVE_CONFIGURATION qualifier cannot be used across hardware revisions or software Versions. SAVE_CONFIGURATION is not available for upgrades of firmware or hardware, and will not perform inter-platform conversions. For example, you cannot use SAVE_CONFIGURATION to upgrade from HSOF Version 5.1 to Version 5.2, or from an HSZ40 to an HSZ50 array controller.

The controller stores the subsystem's configuration information such as —the storagesets you configured, the qualifiers set for each, the HSOF patches you've installed, and so on, in its nonvolatile memory. This means that, if the controller in a nonredundant configuration should fail, you would have to reconfigure the subsystem from scratch if you haven't set the SAVE_CONFIGURATION qualifier. In a dual-redundant configuration, this information is stored by both controllers, which effectively eliminates the risk of losing it.

Note

DIGITAL recommends that the SAVE_CONFIGURATION qualifier only be used for nonredundant controller configurations. To save the configuration information for dual-redundant configurations use the SET FAILOVER COPY=command.

If you are operating a nonredundant controller configuration subsystem, DIGITAL recommends saving the subsystem's configuration on at least one of the devices or storagesets with the following command:

```
INITIALIZE storageset_name SAVE_CONFIGURATION
```

When initializing a device or storageset with the SAVE_CONFIGURATION qualifier, the controller copies the subsystem's configuration to the disk drives. If you use the qualifier for a multi-device storageset, such as a stripeset, the complete information is stored on each device in the storageset. The capacity of a device that is initialized with the SAVE_CONFIGURATION qualifier is reduced by 256KB. This capacity reduction applies to each member of a storageset.

After initializing a storageset with this qualifier, the controller keeps the copy up-to-date. For example, every time you add a new storageset or patch to the HSOF software, the controller re-copies the new configuration to all storagesets that were initialized with the SAVE_CONFIGURATION qualifier.

DIGITAL does not recommend initializing **all** of your storagesets with the SAVE_CONFIGURATION qualifier. Every time you change the subsystem's configuration, the controller writes the new configuration to all the storagesets that you initialized with this qualifier. Too many write operations can adversely affect system performance.

If a controller in a nonredundant-configuration fails, its replacement automatically searches the devices in the subsystem for a saved configuration. If it finds a saved configuration, it transparently loads it into the controller's nonvolatile memory and brings the subsystem online. It is not necessary to issue an INITIALIZE command again after reconfiguring devices with a new controller.

If you are upgrading from HSOF Version 5.1 to Version 5.2 in a single-controller configuration, you can refresh the SAVE_CONFIG data on your storagesets to include the new software revision by issuing the following command:

```
SET UNIT xxx WRITE_PROTECT
```

[where xxx is any existing unit on the controller]

```
SET UNIT xxx NOWRITE_PROTECT
```

These commands, which have no net effect, serve to change the contents of nonvolatile memory, causing the updated SAVE_CONFIG information to be automatically written to the SAVE_CONFIG area of all units which have been initialized with this option.

Operating Constraints

This section describes the operating constraints for HSOF Version 5.2 software. An operating constraint is defined by the parameters within which the controller is designed to operate. Constraints of other system components, such as host adapters, may also be described in this section. Keep these constraints in mind to avoid problems and to obtain the maximum performance available from your controller.

Maintenance Terminal Port Speeds

HSOF Version 5.2 Software supports terminal port connections of 4800, 9600, and 19200 bits per second (bps). Connections speeds at 300, 1200, and 2400 bps are not supported.

Limits on the Total Number of Stagesets

The following limits apply to stagesets configured on a single controller or dual-redundant controller configuration:

- A mirrorset can have a maximum of 6 members
- A stripeset or RAIDset can have a maximum of 14 members
- There can be a maximum of 20 mirrorsets or RAIDsets or both
- There can be a maximum of 30 stagesets including RAIDsets, mirrorsets, and stripesets
- There can be a maximum of 32 physical device members total for a unit
- There can be a maximum of 4 partitions per disk or stageset
- There can be a maximum of 120GB per stageset

Partitioning Not Supported With SCSI Multiple Bus Failover

The SCSI multiple bus failover feature in HSOF Version 5.2 does not support partitioned disks or storagesets. Delete any existing partitions before enabling multiple bus failover. You cannot create partitions once configured for multiple bus mode.

SCSI Multiple Bus Failover Host Operating System Support

Although the HSZ50 array controller has the capability to support SCSI multiple bus failover, currently, OpenVMS, DIGITAL UNIX, and the Windows NT Server operating systems do not support this feature. Check your operating system's Software Product Description and release notes to determine whether your operating system provides support for this feature.

CFMENU Constraints

CFMENU can only delete one spareset member at a time. Delete spareset members one at a time by responding "Y" (yes) to one member, and "N" (no) to the remaining spareset's members. CFMENU deletes the spareset member.

You can also delete spareset members using the CLI command `DELETE SPARESET disk_name`. While you still have to delete each member individually with separate CLI commands for each one, this is a faster method for deleting multiple spareset members.

CLONE Utility Constraints

The CLONE utility cannot be used with partitioned units.

CLEAR_ERRORS LOST_DATA command

When entering the `CLEAR_ERRORS LOST_DATA` command for a RAIDset-based unit, always enter the command through the preferred path on the controller that owns the unit. Entering the command from the companion controller in a dual-redundant configuration does not clear the lost data.

Only One Qualifier per CLI Command Recommended

Certain qualifiers to CLI commands are incompatible. To avoid problems, specify only one qualifier per CLI command. For example, to set both `WRITEBACK_CACHE` and `NORUN` on unit D102, use the following two commands:

```
SET D102 WRITEBACK_CACHE
```

```
SET D102 NORUN
```

Avoiding Problem Situations

In certain situations, you could experience unusual controller behavior. This section presents information to help you avoid such situations and to recover from them if they occur.

Use of the `CACHE_UPS` Qualifier

When the `CACHE_UPS` qualifier on the `SET THIS_CONTROLLER` command is enabled, the controller ignores the condition of the write-back cache batteries because an uninterruptible power supply (UPS) is assumed to be in use to maintain power to the write-back cache module in the event of a power failure. Note that HSOF software does not have the capability to monitor the remaining capacity of an UPS. There are many variables in an UPS configuration; some configurations may put data at risk if the UPS is exhausted. For example, if UPS power is lost to the entire host/storage system without warning, write-back cache data is maintained by the cache batteries only.

Warning

Use of this feature without an UPS could result in data loss if power is interrupted and batteries have failed.

Device Adds, Moves, and Changes

The controller maintains configuration maps of device types and locations. It uses the location maps to communicate with devices. If you add, move, or change devices while the controller is powered off and without changing the controller configuration first, the controller is not able to work with the changed devices when it returns to service.

Tape Drive Firmware Revision Upgrade

If you are not using HSUTIL to install new tape firmware, do the following: Before installing a new version of tape drive firmware, delete the tape drive completely from the controller configuration. This prevents interrupts to the drive during the code loading operation. After the firmware is installed, refer to:

- *DIGITAL Storageworks HSZ50 Array Controller HSOF Version 5.1 Configuration Manual*,
- *DIGITAL Storageworks HSZ50 Array Controller HSOF Version 5.1 Service Manual*, and the
- *DIGITAL Storageworks HSZ50 Array Controller HSOF Version 5.1 CLI Reference Manual* for adding the tape drive back to the configuration.

Running CONFIG or CFMENU Utilities during Backup or Copy Operations

Do not run the CONFIG or CFMENU utilities during a tape backup or copy operation.

Changing Host Adapter Types

If you change the host adapter and maintain the same node number, restart the controller(s).

CLONE Utility

When running the CLONE utility in a dual-redundant controller configuration, do not issue any CLI commands to devices that are involved in the clone operation.

Extending the Life of Your Write-Back Cache Batteries

Occasionally, circumstances require shutting down your controller and restarting it. If the correct steps for turning off the power to an HSZ50 array controller configuration with write-back cache are not followed, there is a potential for loss of data that might exist on any devices connected to them. In addition, turning off the power to the controller subsystem for any reason for longer than one day (such as a long holiday, system move, replacing a bad SCSI host cable, and so forth), turn off the external cache batteries (ECB's) to prevent them from discharging. To avoid problems when the controller restarts, always use the following process to shut it down:

Caution

Allowing the write-back cache batteries to discharge completely, and leaving them in a discharged state for any length of time may result in permanent and irreversible degradation to the batteries.

Shutting Down a Controller

1. Use the proper procedures for shutting down your operating system if the host system is also going to be turned off.
2. If the host is not going to be turned off, a shut down of the host system is not necessary, but, use the proper operating system procedures to dismount any units that are accessed through the HSZ50 array controllers.
3. If the controller configuration contains any devices or storagesets that are write-back cache enabled, and the system is going to be turned off for an extended length of time, the batteries on the write-back cache modules drain, causing the data in the cache modules to be lost. If you have RAIDset or mirrorset units, or single devices with write-back caching turned on, Set NORUN on all units before shutting down your controller.
4. When the dismount and/or the operating system shut down procedures are complete, invoke the controller SHUTDOWN commands on the controllers.

Caution

Do not turn off the power to the controller subsystem until all shut down procedures have successfully completed.

5. If you have a dual-redundant controller configuration, shut down each controller one at a time, using the following commands:

```
SHUTDOWN OTHER_CONTROLLER
```

```
SHUTDOWN THIS_CONTROLLER
```

If you have a non-redundant controller configuration, only the SHUTDOWN THIS_CONTROLLER command is necessary.


When the SHUTDOWN command successfully completes, the green reset LED stops blinking.

6. Only when the controller SHUTDOWN command has successfully completed, turn off power to the controller subsystem (or just the controller shelf, if you are not shutting down the entire system, by unplugging the power supplies in the controller shelf).

Disabling Write-Back Cache Batteries

Refer to Chapter 2 in *DIGITAL StorageWorks HSZ50 Array Controller HSOF Version 5.1 Service Manual*, for instructions on removing and replacing the controller module and write-back cache module.

Figure 1 External Cache Battery Disable Switch



1. Press the battery disable switch (one for each battery), which is the small button labeled SHUT OFF next to the flashing status LED on the ECB and hold it in for approximately 2 seconds. The cache LED flashes once, then shuts off.
2. The battery is no longer powering the cache.
3. To return to normal operation, the cache battery is enabled when the RAID subsystem is powered on.

Restarting the Controller

1. If you removed the PCMCIA card, hold down the reset button while inserting the PCMCIA card. When you release the reset button, the controller restarts.
2. If you did not remove the PCMCIA card, press the reset button and the controller restarts.

Documentation Additions and Corrections

Following are additions and corrections to:

- *DIGITAL Storageworks HSZ50 Array Controller HSOF Version 5.1 Configuration Manual*
- *DIGITAL Storageworks HSZ50 Array Controller HSOF Version 5.1 Service Manual*
- *DIGITAL Storageworks HSZ50 Array Controller HSOF Version 5.1 CLI Reference Manual.*

In the *DIGITAL Storageworks HSZ50 Array Controller HSOF Version 5.1 CLI Reference Manual* on page 2-81, the second example was included in error. Please strike through that example.

Fault Management Utility (FMU) Codes

Add the following codes:

Executive Services Last Failure

01192390—A processor interrupt was generated by the CACHE A Dynamic Ram controller and ArBitration engine (DRAB) with an indication that the CACHE backup battery has failed or is low (needs charging).

011A2390—A processor interrupt was generated by the CACHE B Dynamic Ram controller and ArBitration engine (DRAB) with an indication that the CACHE backup battery has failed or is low (needs charging).

Host Interconnect Port Service Last Failure

42332080—Receive_main found destination address in the rcv packet does not match node address.

Last Failure Codes No Longer Used

010B2380—A processor interrupt was generated by the CACHEA Dynamic Ram controller and Arbitration engine (DRAB) with an indication that the CACHE backup battery has failed or is low (needs charging).

010C2380—A processor interrupt was generated by the CACHEB Dynamic Ram controller and Arbitration engine (DRAB) with an indication that the CACHE backup battery has failed or is low (needs charging).

Upgrading Array Controller Software

This section describes how to upgrade controller software using the program card for both the single and dual-redundant controller configurations. An HSZ50 operating software upgrade causes a brief interruption in service to attached devices.

Caution

If JBOD disks that were previously initialized using the SAVE_CONFIG option on an HSZ40 controller running HSOF Version 2.7 have been moved onto your HSZ50 controller subsystem, see Appendix A of these release notes **before** upgrading your HSZ50 controller to HSOF Version 5.2.

Program Card Upgrade (Single Controller Configuration)

Use the following procedure to upgrade controller software in a single controller configuration:

1. Halt all I/O activity to the controller using the appropriate procedures for the operating system.
2. Connect a maintenance terminal to the controller.
3. Take the controller out of service:

```
HSZ50> SHUTDOWN THIS_CONTROLLER
```

To ensure the controller has shut down cleanly, check for the following indications on the controller's OCP:

- The Reset (//) light is lit continuously.
- Port lights 1, 2, and 3 are also lit continuously.

Caution

Failure to shut down the controller in this step could result in problems with RAIDsets and cache when the controller is reset.

4. Unsnap and remove the program card ESD cover.
5. While pushing and holding down the operator control panel (OCP) Reset (//) button on the controller, eject and remove the program card.
6. While pushing and holding down the OCP Reset (//) button on the controller, insert the new program card.
7. Release the Reset (//) button to initialize the controller.

After restart, the controller will not allow you to enter any CLI commands for approximately 60 seconds. This delay does not affect unit availability to the host. The delay provides time for the controller internal operations to complete before user commands can change configuration.

When the controller initializes correctly, the green reset indicator flashes once every second.

8. Replace the ESD cover over the program card by pushing the two plastic push pins on each side of the shield.
9. Resume I/O activity to the controller.

Program Card Upgrade (Dual-Redundant Configuration)

In a dual-redundant configuration, the software in both controllers must be upgraded simultaneously. Use the following procedure to upgrade controller software in a dual-redundant configuration:

1. Follow steps 1-4 of the single controller configuration procedure for each controller.
2. Follow steps 5-7 of the single controller configuration procedure for both controllers simultaneously.

The two controllers reinitialize concurrently and synchronize with each other. The timing of this procedure is not critical, however the program cards should be removed and inserted at approximately the same time (within a few seconds).

3. Replace the ESD covers over both program cards.
4. Resume I/O activity to the controllers.

Identifying Your HSOF Software Revision Level

You can identify HSOF Version 5.2 Software by entering the `SHOW THIS_CONTROLLER` command at the CLI prompt. The resulting display lists the software revision level as Version “V52Z-*x*” (where *x* represents the number of patches associated with this Version of software that are correctly installed).

Order Numbers

The following table lists order numbers for controller options and preconfigured subsystems.

HSZ50-AF	StorageWorks HSZ50 Array Controller with a 32MB cache module, 32 SCSI-2 device connections, 1 external cache battery in a SBB, one 2-meter ECB cable
HSZ50-AH	StorageWorks HSZ50 Array Controller with a 64 MB cache module, 32 SCSI-2 device connections, 1 external cache battery in a SBB, one 2-meter ECB cable
HSZ50-AJ	StorageWorks HSZ50 Array Controller with a 128 MB cache module, 32 SCSI-2 device connections, 1 external cache battery in a SBB, one 2-meter ECB cable
HSZ52-AF	Two StorageWorks HSZ50 Array Controllers with 32 MB cache modules, 64 SCSI-2 device connections, 2 external cache batteries in a SBB, two 2-meter ECB cables
HSZ52-AH	Two StorageWorks HSZ50 Array Controllers with 64 MB cache modules, 64 SCSI-2 device connections, 2 external cache batteries in a SBB, two 2-meter ECB cables
HSZ52-AJ	Two StorageWorks HSZ50 Array Controllers with 128 MB cache modules, 64 SCSI-2 device connections, 2 external cache batteries in a SBB, two 2-meter ECB cables
HSZ54-AJ	Four StorageWorks HSZ50 Array Controllers with 128 MB cache modules, 144/168 SCSI-2 device connections, 4 external cache batteries in 2 SBBs, four 2-meter ECB cables
HSSIM-AA	One 32 MB SIMM pack for cache module upgrades
HS35X-BA	One replacement external cache battery (ECB) in a SBB
HS35X-BB	Two replacement external cache batteries in a SBB

Note:

One prerequisite QB-5CJAA-SA software kit is required for each controller ordered (two kits for HSZ52 configurations, four kits for HSZ54 configurations).

Appendix A

Possible Problem with Disks Initialized with SAVE_CONFIG Under HSOF Version 2.7 on HSZ40/20/SWXRC Controller

If you initialized a JBOD disk using the `SAVE_CONFIGURATION` switch on an HSZ40 controller which was running HSOF Version 2.7, and then moved the disk to your HSZ50 controller subsystem, read this appendix.

There is a remote possibility that some disks attached to HSZ40/20/SWXRC controllers and the solution products containing them (RA410, SC4200/4600, etc.) might have a problem in the structure of the on-disk file system. Systems which might be affected are those which:

1. Use disks in JBOD configurations (that is, disks which are not members of controller-based storagesets such as RAIDsets and mirrorsets)
2. Initialized disks under HSOF Version 2.7 software using the `SAVE_CONFIG` qualifier AND rebooted the controller BEFORE initializing the disk under the operating system.

Note that the problem does not occur if the file system was built on the disk before the controller was rebooted. Also, the problem does not occur when disks are initialized using SAVE_CONFIG and the platform operating system under HSOF Versions 30Z, 31z, 32z, 50z, 51Z or 52Z.

If a disk built as described above is moved to an HSZ50 controller running HSOF Version 5.0, Version 5.1, or Version 5.2 problems could occur.

Note that all 2GB and 4GB drives on Windows NT platforms are not exposed to this potential problem. Drives on other platforms meeting the above criteria have a small risk of exposure; see the "How to Detect" section of this appendix for procedures to determine whether a disk is exposed.

When a disk being used in a JBOD configuration is initialized with SAVE_CONFIG, the last 500 blocks on the disk are allocated by the controller to store the configuration data. If the controller running HSOF Version 27Z software is rebooted **before** the disk is initialized by the platform operating system, the controller fails to remember the reduction in disk size and reports the unreduced disk capacity to the operating system. When the operating system subsequently builds the file system, the blocks which SAVE_CONFIG uses to update the configuration data are also included in the file system disk space, creating a potential for both the operating system and the controller to write to the last 500 blocks on disk.

If the file system subsequently overwrites configuration data, the controller recognizes that the data is invalid configuration data and ignores it. In this case, controller parameters must be manually re-entered when SAVE_CONFIG tries to restore the configuration (unless another drive contains valid configuration data).

Various configuration events cause the controller to write the configuration data to the SAVE_CONFIG area. If the controller overwrites file system data, the results vary depending on the platform operating system and the application.

If a disk which has this problem is moved to a controller running HSOF Version 51Z or Version 52Z software before the differing file system and controller view of the disk capacity is resolved, and the file system tries to access the SAVE_CONFIG area, the controller returns an error to the operating system. The action that the operating system takes upon receiving this error varies depending on the platform, but could include rendering the entire file system or database inaccessible.

How to Detect Whether You Have This Problem

Windows NT platforms

As previously noted, 2GB and 4GB drives on Windows NT platforms are not exposed to the problem described in this appendix. This problem affects 1GB single-disks units in JBOD configuration with SAVE_CONFIG data stored on them. If you are not using 1GB JBOD disk units with SAVE_CONFIG data saved on them, do not proceed any further. Your system is **not** at risk.

Use the following procedure to check a JBOD 1GB drive with SAVE_CONFIG data saved on it to determine whether it is exposed:

- a. Shut down the host computer and wait until shut down is complete.
- b. Restart the HSZ50 controller(s) by pressing the heart-beat button(s) (Green reset button.)
- c. Wait a minute, then start the host computer.
- d. After the host reboots, start up "Disk Administrator".
- e. Determine which drive on "Disk Administrator" corresponds to the 1GB JBOD disk to be checked.
- f. Check if the JBOD has a 1MB or greater unpartitioned space at the end of disk.
- g. If "f" is true, the disk does **not** have the problem described in this appendix. Ensure the last 1MB space is unused, leaving it unpartitioned.
- h. If "f" is false, there is no unpartitioned space at the end of the disk. The last 196 Blocks (100KB) on the drive are at risk for the problem described in this appendix. See the "Solution" section of this appendix for the recovery procedure.

Novell NetWare platforms

The problem described in this appendix affects single-disks units in JBOD configurations with SAVE_CONFIG data stored on them. If you are not using JBOD disk units with SAVE_CONFIG data saved on them, do not proceed any further. Your system is **not** at risk.

NetWare reserves 2% of the space at the end of each disk for bad block replacement. The 500 blocks (256KB) at the end of this 2% space are exposed to the problem described in this appendix. A 2% space is larger than is generally needed for replacing bad blocks. For example, reserve space on a 4GB, 2GB, and 1GB disk is 80MB, 40MB, and 20MB respectively. The probability of a bad block being replaced in the last 256KB of this reserve space is very small; however, it is possible. Use the following procedure to check a disk in a JBOD configuration to determine whether it is exposed:

- a. NWSERVER> load install
- b. Open "disk options"
- c. Open "Modify disk partition and Hot Fix"
- d. Select disk drive
- e. Choose "Change Hot Fix"
- f. Record "Redirection Area", this is the BadBlock size.
- g. Calculate 2% of the disk
- h. If BadBlock size is less than (2% - 256KB) then the disk is **not** affected.
- i. If the BadBlock size is greater than (2% - 256KB) then the disk **is** at risk. See the "Solution" section of this appendix for the recovery procedure.

Sun Solaris and Sun OS Platforms

The problem described in this appendix affects single-disks units in JBOD configuration with SAVE_CONFIG data stored on them. If you are not using JBOD disk units with SAVE_CONFIG data saved on them, do not proceed any further. Your system is **not** at risk.

If you followed the *Installation Guide*, you are not at risk. This is due to the fact that the default partition layout reserves the last two cylinders for diagnostic purposes. The 500 blocks in question always reside within those two diagnostic cylinders.

If you changed the default partition layout, **and** allocated the two diagnostic cylinders to a partition, you may be at risk.

If disks in your system are at risk of this problem, use the following procedure to check a disk in a JBOD configuration to determine whether it is exposed:

- a. Use the GUI to display the number of blocks on the unit.
Do this by selecting the LUN in question, and then choosing LUN parameters from the pull-down menu. Write down this number.
- b. Use the TIP command (or an RS-232 terminal) to connect to the controller CLI. If you have problems or questions, this command is documented in the *Installation Guide*.

- c. Use the CLI command `SHOW <unit-name>`, substituting the actual name of the unit in question for `<unit-name>`.
- d. If the GUI and the CLI report different sizes for the same unit, you are at risk for the problem. See the "Solution" section of this appendix for the recovery procedure.

OpenVMS platforms

The problem described in this appendix affects single-disks units in JBOD configurations with `SAVE_CONFIG` data stored on them. If you have not moved JBOD disk units with `SAVE_CONFIG` data saved on them to the HSZ50 controller being upgraded to V5.2, do not proceed any further. Your system is **not** at risk.

If disks in your system are at risk for this problem, use the following procedure to check a disk in JBOD configuration to determine whether it is exposed:

- a. At the controller prompt, type `SHOW DISKnnn` (where `nnn` is the JBOD disk in question).
- b. Look for a "Configuration being backed up on this container" message.
- c. Record the block size capacity displayed by the controller.
- d. From the OpenVMS prompt on one of the hosts, mount the disk in question and type the command:

```
$ SHOW DEVICE/FULL DKA200 :
```
- e. Compare the total block size obtained from the "show device" command with the block size capacity obtained in step 'c.'
- f. If the reported sizes are different, this disk **is** at risk for the problem. See the "Solution" section of this appendix for the recovery procedure.

DIGITAL UNIX platforms

The problem described in this appendix affects single-disks units in JBOD configurations with `SAVE_CONFIG` data stored on them. If you are not using JBOD disk units with `SAVE_CONFIG` data saved on them, do not proceed any further. Your system is **not** at risk.

If disks in your system are at risk of this problem, use the following procedure to check a disk in JBOD configuration to determine whether it is exposed:

- a. At the controller prompt, type `SHOW DISKnnn` (where `nnn` is the JBOD disk in question).
- b. Look for the "Configuration being backed up on this container" message.
- c. Record the block size capacity displayed by the controller.

- d. From the DIGITAL UNIX on one of the hosts, type the following commands (rrza18c is used in the following example as the device in question):

```
# disklabel -rw /dev/rrza18c HSZ40  
# disklabel -r /dev/rrza18c # /dev/rrza18c:
```
- e. Compare the sectors/unit output from disklabel command with the block size capacity obtained in step 'c.'
- f. If the reported sizes are different, this disk **is** at risk for the problem. See the "Solution" section of this appendix for the recovery procedure.

AIX platforms

The problem described in this appendix affects single-disks units in JBOD configurations with SAVE_CONFIG data stored on them. If you are not using JBOD disk units with SAVE_CONFIG data saved on them, do not proceed any further. Your system is **not** at risk.

If disks in your system are at risk of this problem, use the following procedure to check a disk in a JBOD configuration to determine whether it is exposed:

AIX 4.1.4:

- a. Sum the raw device as shown in the following command:

```
sum -r /dev/rhdiskN
```

- b. If this operation results in a read error as shown below, the disk **is** at risk for the problem. See the "Solution" section of this appendix for the recovery procedure.

```
sum: read error on /dev/rhdiskN
```

AIX 3.2.5:

Disks on systems which have the risk factors described above should be regarded as at risk for the problem described in this appendix.

HP-UX platforms

The problem described in this appendix affects single-disks units in JBOD configuration with SAVE_CONFIG data stored on them. If you are not using JBOD disk units with SAVE_CONFIG data saved on them, do not proceed any further. Your system is **not** at risk.

Disks on systems which have the risk factors described above should be regarded as at risk for the problem described in this appendix.

Solution

1. If you are using the SAVE_CONFIGURATION qualifier to initialize JBOD disks under HSOF Version 27Z, be sure to initialize the disk with the platform file system **before** rebooting the controller.
2. If you have the risk factors for the problem as described in the “Symptom” and “Detection” sections of this appendix, use the steps below to resolve the discrepancy in controller/operating system views of the disk at your earliest opportunity.

DIGITAL recommends that the recovery process described here be performed before moving the disk to a controller running Version 5.2. Any files written in the SAVE_CONFIG area are accessible to the operating system after the restore process. However, those files are suspect and should be carefully examined to ensure that the data they contain is correct, or restored from a previous backup.

- a. Back up the unit that contains SAVE_CONFIG information.
- b. Unmount the file system(s) contained on that unit.
- c. Delete the unit from the configuration in the controller.
- d. Initialize the container from the controller without SAVE_CONFIG.
- e. Add the unit back into the configuration.
- f. Initialize and restore unit from backup.