HSZ40 Array Controller Operating Software

HSOF Version 3.1
Release Notes

Order Number: EK–HSZ40–RN. L01

This document summarizes features and characteristics of the HSZ40 array controller operating software Version 3.1 that are not covered elsewhere in the documentation. These release notes also contain instructions for installing the software.

Software Version: HSOF V3.1

Digital Equipment Corporation
Maynard, Massachusetts

March 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 tregistered 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 1
Intended Audience 2
Hardware and Software Support 2
Hardware Support 2
Operating System Support 3
Host Adapter Support 3
Device Support 5
New Features 10
DILX Enhancement 10
Ability to Install New Tape Drive Microcode 10
Set SCSI Speed for Devices 10
Clarifications 11
Logical Device Event Reports 11
SAVE_CONFIGURATION from Previously Initialized Storagesets 11
Using SAVE_CONFIGURATION to Save the Subsystem's Configuration 11
Considerations for Saving the Subsystem's Configuration 12
Support for Larger Storagesets 12
Faster Storageset Initialization 12
Operating Constraints 13
Limits on the Total Number of Storagesets 13
Restrictions on Moving Storagesets to Previous Software Versions 13
Partitioning Not Supported with SCSI Multiple Bus Failover 14
SCSI Multiple Bus Failover Host Operating System Support 14
CFMENU Constraints 14
CLONE Utility Constraints 15
Low/Failed Battery on Dual-Redundant Configurations with Preferred Paths 15

Unit Problems after Battery Low 15
CLEAR_ERRORS LOST_DATA command 15
Only One Qualifier per CLI Command Recommended..15
Avoiding Problem Situations 16
Device Adds, Moves, and Changes 16
Tape Drive Firmware Revision 16
Running CONFIG or CFMENU Utilities During Backup Operations.................................16
Changing Host Adapter Types 16
CLONE Utility 16
Restarting a Controller..17
Moving Write-Back Cache Modules..17
Documentation Additions and Corrections 18
HSOF Software Installation..19
Required Parts and Tools..20
Nonredundant Configurations 20
Dual-Redundant Configurations..21
Identifying Your HSOF Software Revision Level..22
Order Numbers..23

Appendix A

Possible problem with disks initialized with SAVE_CONFIG under HSOF V2.7 on HSZ40/20/SWXRC Controllers.. A-1
How to Detect if You Have This Problem A-2
Windows NT platforms A-2
Novell NetWare platforms.. A-3
Sun Solaris and SunOS platforms A-3
OpenVMS platforms.. A-4
DIGITAL UNIX platforms A-4
AIX platforms A-5
AIX 4.1.4.. A-5
AIX 3.2.5.. A-5
HP-UX platforms.. A-5
Solution.. A-6

Introduction

These release notes provide information for the HSZ40 array controller operating software Version 3.1[1] not covered elsewhere in the documentation. This document should be used by individuals responsible for configuring, installing, and using the HSZ40 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 3.1 software
- New features—Modifications in the software that add significant functions which the controller can perform
- Changes from Version 3.0—Modifications to operating parameters or functions that do not necessarily add to the functions performed by the controller
- 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
- Corrections and additions to the documentation
- Installation instructions
- Identification of the software revision level

Release Package Contents

The Version 3.1 release package consists of the following:

- A cover letter
- The HSZ40 documentation set:
  - *DIGITAL Storageworks HSZ40 Array Controller*
    HSOF Version 3.1
    Configuration Manual

---

[1] The OpenVMS™ SHOW CLUSTER command and the HSOF software SHOW THIS_CONTROLLER command display the version as “V31Z”.

- *DIGITAL Storageworks HSZ40 Array Controller*
  HSOF Version 3.1
  Service Manual
- *DIGITAL Storageworks HSZ40 Array Controller*
  HSOF Version 3.1
  CLI Reference Manual
- *DIGITAL Storageworks Array Controller*
  HSZ40 Array Controller Operating Software HSOF Version 3.1 Release Notes
- *DIGITAL Storageworks Family Array Controller*
  Operating Software (HSOF), Version 3.1 Software Product Description

- A PCMCIA program card containing HSOF Version 3.1 software

Intended Audience

This document has been prepared for DIGITAL customers who have purchased HSZ40 array controllers and for DIGITAL Multivendor Customer Services personnel responsible for installing and maintaining systems that include HSZ40 array controllers.

Hardware and Software Support

This section lists the hardware and software compatible with HSOF Version 3.1 software.

Hardware Support

HSOF Version 3.1 software supports the following revisions for the HSZ40 controllers and associated hardware:

- HSZ40–B*x* and HSZ40–C*x* controller modules
- Version 2 cache module, hardware revision A or B
- BA350–MA controller shelf
- BA350–S*x* 8-bit SCSI device shelf
- BA356-Sx wide device shelf with 8-bit or 16-bit I/O module (controller operates in 8-bit mode only)
- BA35*x*–HF power supply

Operating System Support

HSOF Version 3.1 software on HSZ40 controllers is supported by the following operating system versions, within the limitations described in *DIGITAL Storageworks HSZ40 Array Controller HSOF Version 3.1 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 HSZ40 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) adapter firmware version A10

The following host adapters are supported by the OpenVMS Alpha operating system Version 6.2 and higher for HSZ40 controllers:

- KZTSA (for DEC 3000 systems) adapter firmware version A11 and higher
- KFTIA (for TurboLaser 8200 embedded SCSI) adapter firmware version 2.46 and higher
- PMAZC (for DEC 3000 systems – requires a DWZZ-series signal converter) adapter firmware version 2.1 and higher
- 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) adapter firmware version A10 and higher

The following host adapters are 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 5.1 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]	1.05	440C	A01
RZ26N-VA/VW[1]	1.05	446	A01
SWXD3-SF/WF[1]	1.05	446	A01
DS-RZ26N-VZ[1]	1.05	1003	A01
RZ28-VA/VW[1]	2.10	435E	A01
RZ28B-VA	2.10	0003	A01
RZ28D-VA/VW[1]	2.10	0008	A01
SWXD3-SG/WG[1]	2.10	0008	A01
RZ28M-VA/VW[1]	2.10	0466	A01
DS-RZ28M-VZ[1]	2.10	1003	A01
SWXD3-SH/WH[1]	2.10	0466	A01
RZ29B-VA/VW[1]	4.3	0007	B01
SWXD3-SE/WE[1]	4.3	0007	C02/A01
DS-RZ40-VA	9.1	LYGO	A01
RZ74-VA	3.57	T427B	B07

Note:

1. Wide disk drives require a BA356 shelf with an 8-bit I/O module.

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	3,4,5,6,7
DS-TZ89N-TA	35/70	141F	A01	1,3,4,5,6

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 SSB 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.
7. Wide tape devices require a BA356 shelf with an 8-bit I/O module.

Table 3: Supported Solid State Disk Drives

Device	Capacity in Gigabytes	Minimum Microcode Version	Minimum Hardware Version
EZ31-VW	0.134	V064	A01
EZ32-VW	0.268	V064	A01
EZ51R-VA	0.10	V096	D01
EZ54R-VA	0.42	V096	A01
EZ58R-VA	0.85	V109	C01
EZ64-VA	0.475	V064	A01
EZ64-VW	0.475	V070	A01
EZ69-VA	0.950	V064	A01
EZ69-VW	0.950	V070	A01

Note: Do not warm-swap solid-state disk drives. Make sure power to the device shelf is turned off before removing or inserting this device.

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:

CD–ROM drives are only supported under the DIGITAL UNIX and OpenVMS operating systems.

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

The following list describes changes in the operation of the controller provided by the software. A change is a modification of an operating parameter or function thatwas introduced in software revisions prior to version 5.1, but was not covered in the documentation.

DILX Enhancement

The HSZ40 controller with HSOF Version 3.1 software has a new version Disk In-Line Exerciser (DILX) that acts on disks configured as single-device, RAIDset, stripeset, mirrorset, stripe-mirrorset or partitioned units.

Ability to Install New Tape Drive Microcode

HSOF Version 3.1 software allows you to run the controller's HSUTIL utility to install new microcode to the following tape drive:

- TZ89

  HSOF Version 3.1 software also allows you to use the controller's HSUTIL program to install new microcode to the following solid state disk drives:

- EZ51
- EZ54
- EZ64
- EZ69

See *DIGITAL Storageworks HSZ40 Array Controller HSOF Version 3.1 Service Manual* for instructions on how to install new device microcode.

Set SCSI Speed for Devices

HSOF Version 3.1 software makes it possible for the initiator and target to negotiate a faster data rate than is supported by the cable length. HSOF Version 3.1 software has a new CLI command to set a maximum data transfer rate between the controller and any device.

To set the transfer speed, use the SET command as follows:

SET *device-name* **TRANSFER_RATE_REQUESTED**=*rate*

rate can be one of the following:

ASYNCHRONOUS

10MHZ

5MHZ

The default is 10MHZ.

Clarifications

This section presents clarifications on controller behavior in certain situations.

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.

SAVE_CONFIGURATION from Previously Initialized Storagesets

If the storagesets in your configuration were previously initialized and the INITIALIZE container-name SAVE_CONFIGURATION command was issued to save your configuration to disk, it is not necessary to enter an INITIALIZE command again after you have reconfigured your devices with a new controller.

Using SAVE_CONFIGURATION to Save the Subsystem's Configuration

The controller stores the subsystem's configuration—the storagesets you configured, the qualifiers you set for each, the HSOF patches you've installed, and so on—in its non-volatile memory. This means that, if the controller fails in a nonredundant configuration, you would have to reconfigure the subsystem from scratch. (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 configuration subsystem, DIGITAL recommends that you save the subsystem's configuration on at least one of the devices or storagesets with the following command:

```
INITIALIZE storageset_name SAVE_CONFIGURATION
```

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

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 non-volatile memory and brings the subsystem online. It is not necessary to issue an INITIALIZE command again after you have reconfigured your devices with a new controller.

Considerations for Saving the Subsystem's Configuration

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 the storagesets that you initialized with this qualifier. Too many write operations can adversely affect performance.

The SAVE_CONFIGURATION qualifier cannot cross 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 3.0 to Version 3.1, or from an HSZ40 to an HSZ50 array controller.

Support for Larger Storagesets

HSOF Version 3.1 software allows you to create storagesets with capacities of up to 120GB. You can configure drives of any size into storagesets, such as RAIDsets or stripesets. The maximum size of a storageset is 120GB, regardless of the size of the member drives.

Faster Storageset Initialization

The INITIALIZE command was optimized in Version 3.0, and runs significantly faster than on previous versions of HSOF software. It takes less than 30 seconds to initialize a large RAIDset.

Operating Constraints

This section describes the operating constraints for HSOF Version 3.1 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.

Limits on the Total Number of Storagesets

The following limits apply to storagesets 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 storagesets 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 storageset.

Restrictions on Moving Storagesets to Previous Software Versions

In HSOF Software Version 2.5, striped mirrorsets reported a different disk geometry than the same non-mirrored stripeset. This has been resolved in Versions 2.7 , 3.0 and 3.1. Striped mirrorsets initialized under these versions report a disk geometry identical to the same nonmirrored stripeset.

Due to this improvement, observe the following restrictions when moving stripesets to Version 2.5 of software:

- Striped mirrorsets created under HSOF Version 2.7, 3.0 or 3.1 software cannot be used on previous versions; the metadata is incompatible.
- Stripesets that are converted to a striped mirrorset using the MIRROR command under HSOF Version 2.7, 3.0 or 3.1 software cannot be used on previous versions. You must UNMIRROR each of the members to return the stripeset to an earlier version.
- Stripeset clones that are created under HSOF Version 2.7 or 3.1 software with the CLONE utility cannot be used with previous software versions.
- Containers initialized with the SAVE_CONFIGURATION switch cannot be used with previous software versions.

Striped mirrorsets created under previous software versions and used with HSOF Version 2.7, 3.0 or 3.1 software carry forward the same disk geometry they had under the previous version. If possible, back up your data and reinitialize the stripesets under Version 2.7, 3.0 or 3.1 to take full advantage of disk geometry improvements.

Partitioning Not Supported with SCSI Multiple Bus Failover

The SCSI multiple bus failover feature in HSOF Version 3.1 software does not support partitioned disks or storagesets. You must delete any existing partitions before enabling multiple bus failover, and you cannot create partitions once the controllers have been configured to operate in multiple bus failover mode.

SCSI Multiple Bus Failover Host Operating System Support

Although the HSZ40 array controller has the capability to support SCSI multiple bus failover, currently, OpenVMS, DIGITAL UNIX, and 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. You must delete sparesets one at a time by responding Y to one member, then respond N to the rest of the spareset's members. CFMENU will delete 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.

Low/Failed Battery on Dual-Redundant Configurations with Preferred Paths

If a cache battery goes low or fails for a controller in a dual-redundant configuration and preferred paths are set, that controller's units may go "AVAILABLE" on the "OTHER" controller and INOPERATIVE on it. This causes a MOUNT VERIFY on the units from OpenVMS. You will need to issue the SET unit-name NOPREFERRED_PATH command which will allow the units to go out of MOUNTVERIFY. Later you can reissue the SET unit-name PREFERRED_PATH= command to the affected units.

Unit Problems after Battery Low

An unusual combination of circumstances may cause a unit to unexpectedly become inoperative or report lost data after a write-back cache battery changes state from "low" to "good." No data has been lost. Reset the controller to correct the unit state.

CLEAR_ERRORS LOST_DATA command

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

Only One Qualifier per CLI Command Recommended

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

```
CLI> SET D102 WRITEBACK_CACHE
CLI> SET D102 NORUN
```

Avoiding Problem Situations

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

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

If you are not using HSUTIL to install your new tape firmware, do the following:

Before installing a new version of tape drive firmware, delete the tape drive completely from the controller configuration. You can add the tape drive back to the configuration after the firmware installation is complete.

Running CONFIG or CFMENU Utilities During Backup Operations

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

Changing Host Adapter Types

If for some reason you change your host adapter and maintain the same node number, you MUST reinitialize your 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.

Restarting a Controller

When you encounter problems with your controller, sometimes you will need to shut it down and restart it. To avoid problems when restarting a controller, always use the following sequence of commands:

1. Shutdown the controller:

   CLI> **SHUTDOWN THIS_CONTROLLER**

2. If you need to remove the PCMCIA card, always press the reset button while ejecting the card.

Moving Write-Back Cache Modules

When moving a write-back cache module to a new environment (that is, different controller, different devices), it is very important that you perform a clean shutdown of the controller and clear any data from the cache. Follow these steps to move a write-back cache module:

1. Shutdown the controller and flush unwritten data from the cache module using the appropriate CLI command:

   CLI> **SHUTDOWN THIS_CONTROLLER**

   or

   CLI> **SHUTDOWN OTHER_CONTROLLER**

2. Wait for the SHUTDOWN command to complete and verify that the controller does not report any errors.
3. Remove the controller and cache module following the instructions in the *DIGITAL Storageworks HSZ40 Array Controller HSOF Version 3.1 Service Manual*.

The SHUTDOWN command flushes cache data to the devices, and you can install the module in another location without problems caused by uncleared cache data.

Documentation Additions and Corrections

Following are additions and corrections to *DIGITAL Storageworks HSZ40 Array Controller HSOF Version 3.1 Configuration Manual*, *DIGITAL Storageworks HSZ40 Array Controller HSOF Version 3.1 Service Manual* and *DIGITAL Storageworks HSZ40 Array Controller HSOF Version 3.1 CLI Reference Manual*.

(none identified)

HSOF Software Installation

Use the procedures described in this section to install the HSOF Version 3.1 software supplied in this kit. The procedure for upgrading your software to Version 3.1 requires that you shutdown and then restart the controllers. This process typically takes less than 5 minutes to accomplish.

Note

HSOF Version 3.1 software waits up to one minute after the controller restarts before presenting a CLI prompt or accepting any commands. This delay does not affect unit availability to the host. The delay provides time for controller's internal configuration operations to process completely before entered commands can change the configuration.

After you insert the HSOF Version 3.1 program cards and press the reset button, allow 60 seconds for the CLI prompt to appear.

DIGITAL recomends dismounting devices attached to the controller before performing the upgrade procedure. If it is not convenient to dismount the devices, perform the upgrade when I/O load is slow.

Caution

If the controller about to be upgraded is running HSOF V2.7 and contains JBOD disks that were initialized using the SAVE_CONFIGURATION option, it is important to see Appendix A BEFORE upgrading this controller to HSOF V3.1.

Required Parts and Tools

You must have the tools listed in Table 5 to upgrade a controller module to Version 3.1.

Table 5 Required Tools for Adding a Second Controller

TOOLS REQUIRED	PURPOSE
Maintenance terminal and cable	To shutdown controllers, restart controllers, and invoke C_SWAP
ESD wrist strap and ESD mat	To protect all equipment against electrostatic discharge
3/32-inch Allen wrench	To loosen the controller mounting screws; to reinstall the controller
5/32-inch Allen wrench	To unlock the SW800-series cabinet
Small flat-head screwdriver	To connect the SCSI cable adapter to the controller

Nonredundant Configurations

You must shutdown and restart the controller during this upgrade. During this time, the units will be unavailable to the host system. Before upgrading the controller software, prepare the host system by dismounting units or by shutting down the system.

Use the following procedure to upgrade the HSOF software in a nonredundant controller:

1. Stop all I/O to the units in your subsystem.
2. Establish a local terminal connection to the controller.
3. Enter the SHUTDOWN command:

   ```
   $ SHUTDOWN THIS_CONTROLLER
   ```

   Wait for the command to complete. When the controller shuts down, the green Reset (//) LED stops flashing and stays on.

   CAUTION

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

4. Remove the ESD shield covering the PCMCIA program card.
5. Press and hold down the reset button while ejecting the program card.

6. Release the reset button.
7. Remove the program card.
8. While holding in the controller reset button, insert the Version 3.1 program card, pressing the card in until the eject button extends outward.
9. Release the reset button.
10. Reinstall the ESD shield.
11. The controller restarts. Communication with the host systems begins as described in the *Digital Storageworks HSZ40 Array Controller HSOF Version 3.1 Configuration Manual*.

Dual-Redundant Configurations

The procedure below requires that you change your dual-redundant configuration into two single controller configurations for the HSOF Version 3.1 software upgrade. You must shutdown and restart both controllers during this upgrade. Units will be unavailable to the host system during this time. Before upgrading the controller software, prepare the host system for this situation by dismounting units or shutting down the system.

Note

DIGITAL does not recommend the use of the SAVE_CONFIGURATION option for dual-redundant configurations.

Use the following procedure to upgrade the HSOF software in a dual-redundant controller:

1. Stop all I/O to the units in your subsystem.
2. Establish a local terminal connection to one of the controllers.
3. Enter the SET NOFAILOVER command to take the controllers out of dual-redundant failover mode.

```
$ SET NOFAILOVER
```

4. Enter the SHUTDOWN command.

   $ SHUTDOWN THIS_CONTROLLER

   Wait for the command to complete. When t he controller shuts down, the green RESET (//) LED stops flashing and stays lit.

   CAUTION

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

5. Remove the electrostatic-discharge (ESD) shield covering the PCMCIA program card.
6. Press and hold down the reset button while ejecting the program card.
7. Release the reset button.
8. Remove the program card.
9. While holding in the controller reset button, insert the V3.1 program card, pressing the card in until the eject button extends outward.
10. Release the reset button.
11. Reinstall the ESD shield.

    The controller restarts. Communication with the host system begins as described in *Digital Storageworks HSZ40 Array Controller HSOF Version 3.1 Configuration Manual.*
12. Repeat all of the previous steps for the second controller.
13. After you complete all of the above steps on both controllers, put the controllers back into dual-redundant (failover) mode by entering the following command:

    $ **SET FAILOVER COPY=CONFIGURATION−SOURCE**

Identifying Your HSOF Software Revision Level

You can identify HSOF Version 3.1 software by entering the SHOW THIS_CONTROLLER command at the CLI prompt. The resulting display lists the software revision level as version “V31Z–*x*” (where *x* represents the number of patches associated with this version of software).

Order Numbers

The following table contains order numbers for controller options and preconfigured options.

HSZ40-BA/CA	StorageWorks HSZ40 Array Controller with no cache module, 6 SCSI-2 device ports
HSZ40-BD/CD	StorageWorks HSZ40 Array Controller with a 16 MB read cache module, 6 SCSI-2 device ports
HSZ40-BF/CF	StorageWorks HSZ40 Array Controller with a 32 MB read cache module, 6 SCSI-2 device ports
HSZ40-XD	16 MB read cache module
HSZ40-XF	32 MB read cache module
HSZ40-YX	Write-back cache option kit -- contains two onboard cache batteries, one battery bracket, one write-back cache license
QA-2YJAC-HS	HSOF Software and Documentation

Appendix A

Possible problem with disks initialized with SAVE_CONFIG under HSOF V2.7 on HSZ40/20/SWXRC Controllers

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

1. Use disks in "JBOD" configuration (that is, disks which are not members of controller-based storagesets such as RAIDsets and mirrorsets)
2. Initialized disks under HSOF V27Z using the SAVE_CONFIG switch 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 V30Z, V31Z, V50Z, or V51Z.

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 V27Z 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 will use 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 config 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 config data).

Various configuration events will cause the controller to write the config 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 controller which has this problem is moved to a controller running HSOF V31Z 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 will take upon receiving this error will vary depending on the platform, but may include rendering the entire file system or database inaccessible.

How to Detect if 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, wait until shut down is complete

b. Restart the HSZ 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 Admin' 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. Make sure that you never use the last 1MB space, leave it unpartitioned.

h. If 'f' is false, there is no unpartitioned space at the end of the disk, then the very last 196 Blocks (100KB) on the drive are at risk for the problem described in this appendix. See the "Solution" section below for the recovery procedure.

Novell NetWare 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.

NetWare reserves 2% of the space at the end of each disk for bad block replacement. 500 blocks (256KB) at the end of this 2% space will be 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 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 below for the recovery procedure.

Sun Solaris and SunOS 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 will 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 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 <unitname>, substituting the actual name of the unit in question for <unitname>.

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 below for the recovery procedure.

OpenVMS platforms

The problem described in this appendix affects single-disks units in JBOD configuration with SAVE_CONFIG data stored on them. If you have not moved JBOD disk units with SAVE_CONFIG data saved on them to the HSZ40 controller being upgraded to V3.1, 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 "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:
   ```

a. Compare the total block size obtained from the "show device" command with the block size capacity obtained in step 'c.'

b. If the reported sizes are different, this disk is at risk for the problem. See the "Solution" section below for the recovery procedure.

DIGITAL UNIX 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 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 "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:
```

a. Compare the sectors/unit output from disklabel command with the block size capacity obtained in step 'c.'

b. If the reported sizes are different, this disk is at risk for the problem. See the "Solution" section below for the recovery procedure.

AIX 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 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:

AIX 4.1.4

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

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

a. If this operation results in a read error as shown below, the disk is at risk for the problem. See the "Solution" section below 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 SAVE_CONFIG to initialize JBOD disks under HSOF V27Z, 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 above, use the steps below to resolve the discrepancy in controller/operating system views of the disk at the earliest opportunity.

DIGITAL recommends that the recovery process described below be performed BEFORE moving the disk to a controller running V3.1. Any files which may have been written in the SAVE_CONFIG area will be accessible to the operating system after the restore process; however, any such 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.