

Compaq StorageWorks™

Release Notes

HSZ40 Array Controller Operating Software HSOF Version 3.7

This document summarizes features, characteristics, and known limitations of the HSZ40 Array Controller operating Versions 3.7, 3.6, 3.4, 3.3, 3.2, and 3.1 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.

Notice

© 2000 Compaq Computer Corporation.

COMPAQ, the Compaq logo, and StorageWorks Registered in U. S. Patent and Trademark Office. OpenVMS is a trademark and/or service mark of Compaq Information Technologies Group, L.P.

Microsoft, MS-DOS, Windows, and Windows NT are registered trademarks of Microsoft Corporation in the United States and/or other countries.

UNIX is a registered trademark of The Open Group.

All other product names mentioned herein may be trademarks or registered trademarks of their respective companies.

Confidential computer software. Valid license from Compaq required for possession, use or copying. Consistent with FAR 12.211 and 12.212, Commercial Computer Software, Computer Software Documentation, and Technical Data for Commercial Items are licensed to the U.S. Government under vendor's standard commercial license.

Compaq shall not be liable for technical or editorial errors or omissions contained herein. The information in this document is subject to change without notice.

The information in this publication is subject to change without notice and is provided "AS IS" WITHOUT WARRANTY OF ANY KIND. THE ENTIRE RISK ARISING OUT OF THE USE OF THIS INFORMATION REMAINS WITH RECIPIENT. IN NO EVENT SHALL COMPAQ BE LIABLE FOR ANY DIRECT, CONSEQUENTIAL, INCIDENTAL, SPECIAL, PUNITIVE OR OTHER DAMAGES WHATSOEVER (INCLUDING WITHOUT LIMITATION, DAMAGES FOR LOSS OF BUSINESS PROFITS, BUSINESS INTERRUPTION OR LOSS OF BUSINESS INFORMATION), EVEN IF COMPAQ HAS BEEN ADVISED OF THE POSSIBILITY OF SUCH DAMAGES. THE FOREGOING SHALL APPLY REGARDLESS OF THE NEGLIGENCE OR OTHER FAULT OF EITHER PARTY AND REGARDLESS OF WHETHER SUCH LIABILITY SOUNDS IN CONTRACT, NEGLIGENCE, TORT, OR ANY OTHER THEORY OF LEGAL LIABILITY, AND NOTWITHSTANDING ANY FAILURE OF ESSENTIAL PURPOSE OF ANY LIMITED REMEDY.

The limited warranties for Compaq products are exclusively set forth in the documentation accompanying such products. Nothing herein should be construed as constituting a further or additional warranty.

Printed in the U.S.A.

Release Notes – HSZ40 Array Controller Operating Software
HSOF Version 3.7
Fourteenth Edition (April 2000)
EK-HSZ40-RN. P01

Visit our Web Site for the Latest Information

Visit our web site for the latest technical tips and documentation. We can be found in the technical area of our web page:

<http://www.compaq.com/storage/>

Release Notes Contents

These release notes cover the following topics:

- 1.0 Identifying Your HSOF Revision Level
- 2.0 Hardware and Software Support
- 3.0 New Items
- 4.0 Features From Previous Releases
- 5.0 Clarifications
- 6.0 Operating Constraints
- 7.0 Avoiding Problem Situations
- 8.0 Documentation Additions and Corrections
- 9.0 Software Installation
- 10.0 Possible Problem Identified
- 11.0 Order Numbers

Release Package Contents

The Version 3.7 Release Package contains the following:

- A PCMCIA program card containing HSOF software Version 3.7
- A CD-ROM containing the following:
 - .txt file describing the contents of the CDROM
 - .pdf files of documentation
 - Adobe Acrobat Reader
- The HSZ40 Array Controller documentation set:
 - *HSZ40 Array Controller HSOF Version 3.1 CLI Reference Manual*
 - *HSZ40 Array Controller HSOF Version 3.1 Service Manual*
 - *HSZ40 Array Controller HSOF Version 3.1 Configuration Manual*
 - *HSZ40 Solutions Software Version 3.7 Software Product Description*
 - *Release Notes – HSZ40 Array Controller Operating Software HSOF Version 3.7*

Intended Audience

This document has been prepared for Compaq customers who currently use and maintain the HSZ40 Array Controller.

1.0 Identifying Your HSOF Revision Level

The release package you received includes a PCMCIA program card containing the new HSOF software.

You can identify your software revision level by entering the `SHOW this_controller` command at the Command Line Interpreter (CLI) prompt. The resulting display lists the software revision level that is currently installed. When you have completed the installation process for the new version HSOF software, you should see the following displayed: `Version V37Z-x` (where `x` represents the number of patches associated with this version of software that are correctly installed).

2.0 Hardware and Software Support

Sections 2.1 through 2.5 list the hardware and software compatible with HSOF Software Version 3.7:

- Hardware Support
- Operating System Support
- Host Adapter Support
- PKQdriver Version for Compaq OpenVMS™ Systems
- Device Support

2.1 Hardware Support

HSOF Software Version 3.7 supports the following revisions or higher for your controller and associated hardware:

- HSZ40-Bx and HSZ40-Cx controller modules
- Cache module Version 2, hardware revision A or B
- BA350-MA controller shelf
- BA350-Sx 8-bit SCSI device shelf
- BA356-Sx wide shelf with 8-bit or 16-bit personality module (controller operates in 8-bit mode only)
- BA35x-HF power supply

2.2 Operating System Support

HSOF Software Version 3.7 on HSZ40 controllers is supported by the following operating system versions:

- Compaq Tru64™ UNIX:
V4.0d, V4.0e, V4.0f, V5.0, and V5.0a
- OpenVMS Alpha:
V6.2-1H3, V7.1-1H1, V7.1-1H2, V7.1-1H3, V7.1-2, V7.2, and V7.2-1

NOTE

If you are operating with any Compaq Tru64 UNIX versions prior to V5.0, it is important to install HSOF software V3.7 on your controllers **before** upgrading to Compaq Tru64 UNIX versions 5.0 or higher.

2.3 Host Adapter Support

The following host adapters are supported by the Compaq Tru64 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 1000/2000/2100/4000/4100/8200/8400) adapter firmware Version A10
- KZPBA (for DIGITAL AlphaServers 1000/2000/2100/4000/4100/8200/8400) adapter firmware Version A10 and higher

The following host adapters are supported by the OpenVMS operating system for HSZ40 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 1000/2000/2100/4000/4100/8200/8400) adapter firmware Version A10
- KZPBA (for DIGITAL AlphaServers 1000/2000/2100/4000/4100/8200/8400) adapter firmware Version A10 and higher

2.4 PKQdriver Version for Compaq OpenVMS Systems

Compaq recommends updating your PKQdriver version when using HSOF V3.7 Software on OpenVMS systems. Sections 2.4.1 through 2.4.3 describe the following:

- Link Date and Image Identification Information
- Identifying Your PKQdriver
- Installing the Current PKQdriver

CAUTION

All OpenVMS systems, using the StorageWorks Command Console V2.3 Agent and HSOF V3.7 Software, **MUST** update their PKQdriver.

If you do not use the proper PKQdriver when using SWCC V2.3 Agents with HSOF Software V3.7, your system may experience an operating system crash.

2.4.1 Link Date and Image Identification Information

Table 2-1 provides Link Dates and Image Identification information for each OpenVMS version. Please refer to this table when identifying your current PKQdriver.

Table 2-1 PKQdriver Identification Information

OpenVMS Version	Link Date	Image Identification Number
V6.2-1H*	16-Nov-1999	4U4A3A3A3A8
V7.1-1H*	11-Nov-1999	6A13A13A6A11
V7.1-2	11-Nov-1999	6A13A13A6A11
V7.2 (Alpha)	11-Nov-1999	6A52
V7.2-1	10-Jan-2000	6A52

To identify your PKQdriver, observe both the Link Date and the Image Identification information. If your driver, with one of the above Image Identification numbers, does not have the corresponding Link Date or later you will need to upgrade your driver.

2.4.2 Identifying Your PKQdriver

The most important factor in identifying your current PKQdriver is the Image Identification number you will see and then its corresponding Link Date. Please use Table 2-1 to view the Image Identification number for each OpenVMS version and its corresponding Link Date.

1. To identify your PKQdriver and obtain the Image Identification and Link Date information, issue the following command:

```
$ ANALYZE/IMAGE/INTERACTIVE
  SYS$LOADABLE_IMAGES:SYS$PKQDRIVER.EXE.
```

You will receive the first data screen.

2. Press enter three times until you come to the fourth data screen.

The fourth data screen contains the Image Identification and Link Date information.

Below is an example of the fourth data screen showing the Image Identification and Link Date information.

Press RETURN to continue, or enter a period (.) for next file:

```
Image Identification Information
  image name: "SYS$PKQDRIVER"
  image file identification: "X-6A52"
  image file build identification: "X6TE-0050130101"
  link date/time: 10-Jan-2000 17:50:12.67
  linker identification: "A11-39"
```

Press RETURN to continue, or enter a period (.) for next file:

In this example:

- The Image Identification number is 6A52
- The Link Date is 10-Jan-2000

If the above screen did not show the 10-Jan-2000 Link Date or later for the Image Identification number 6A52 for an OpenVMS V7.2-1 system, then this driver would need to be updated.

2.4.3 Installing the Current PKQdriver

If your driver does not have the Image Identification number and corresponding Link Date or later that were provided in Table 2-1, then your driver will need to be updated.

For additional support with PKQdriver installation, please contact your Customer Service Representative.

2.5 Device Support

HSOF Software V3.7 supports the following devices described in Tables 2-2 through 2-5:

- Supported Disk Drives
- Supported Tape Drives
- Supported Solid State Devices
- Supported CD-ROM Readers

Table 2-2 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
DS-RZ1BB-VW	2.10	LYJO/0656	A01
RZ28-VA/VW ¹	2.10	435E	B03
RZ28B-VA	2.10	0003	A01
RZ28D-VA/VW ¹	2.10	0008	A01

Table 2-2 Supported Disk Drives – continued

Device	Capacity in Gigabytes	Minimum Microcode Version	Minimum Hardware Version
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
RZ74-VA	3.57	T427B	B07
DS-RZ1CB-VW	4.1	LYJO/0656	A01
RZ29B-VA/VW ¹	4.3	0007	B01
SWXD3-SE/WE ¹	4.3	0007	C02/A01
DS-RZ1CF-VA/VW	4.3	0370/0371	A01
DS-RZ1DB-VW	9.1	LYJO/0307	A01
DS-RZ1DF-VA/VW	9.1	0372/1614	A01
DS-RZ40-VA	9.1	LYGO	A01
DS-RZ1DD-VA/VW	9.1	0305/3B07	A01
DS-RZ1EF-VA/VW	18.2	0372/N1H1	A01
DS-RZ1ED-VW	18.2	0306/0305/3B07	A01
DS-RZ1EA-VW	18.2	3B05/ B016	A01
DS-RZ1DA-VW	9.1	3B06/ B016	A01
DS-RZ1FC-VW	36.4	3B07	A01

Table 2-2 Notes:

1. Wide disk drives require a SWXSS-06 shelf.

All drive “VW” models require DS-SWXSS-06 wide device shelves.

NOTE

Any accidental use of the ADD TAPE command will result in conditions requiring reinitialization of the controller.

Table 2-3 Supported Tape Drives

Device	Capacity G= Gigabytes T= Terabytes	Minimum Microcode Version	Minimum Hardware Version	Notes
TL812	960/1920G	1.2 robot /CC33 drive	A01	1,2,3,4,5
TL822	5280/10560G	1g4F robot/CC33 drive	A01	1,2,3,4,5
TL826	3520/7040G	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-TL895	3.1/6.2T	230	A01	1,2,3,4,5
DS-TL896	6.16/12.32T	V2A/5A	A01	1,2,3,4,5
TZ87-VA	10/20G	930A	A01	3,5,6
TZ87N-VA	10/20G	930A	A01	3,4,5,6
TZ87-TA	10/20G	9514	B02	1,2,3,6
TZ875-NT	50/100G	930A	A01	1,3,5,6
TZ875-TA	50/100G	930A	A01	1,3,5,6
TZ877-AE/AF	70/140G	930A	A01	1,3,5,6
TZ88N-VA	20/40G	CC33	A01	3,4,5,6
DS-TZS20-VA	25/50G	01aj	A01	3,4,5,6
TZ885-NT/NE	100/200G	CC33	A01	1,3,4,5,6
TZ887- NT/NE	140/280G	CC33	A01	1,3,4,5,6
DS-TZ89N-VW	35/70G	80F	A01	3,4,5,6
DS-TZ89N-TA	35/70G	141F	A01	1,3,4,5,6
DS-TL890	560/1120G	3.23 robot/V55 drive	A01	1,2,3,4,5
DS-TL891	350/700G	3.23 robot/V55 drive	A02	1,2,3,4,5
DS-TL892	350/700G	3.23 robot/V55 drive	A02	1,2,3,4,5
DS-AIT35-VW	35/70G	4.03	A01	5,6,7

Table 2-3 Notes:

1. Requires 0.2 meter SCSI-1 to SCSI-2 transition cable, Compaq internal part number 17-03831-01 for DWZZA-AA, and Compaq part number 17-04367-01 for SSB DWZZA-VA and DWZZB-VW.
2. Requires SWXA2 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 SWXSS-06 shelf with an 8-bit I/O module

Table 2-4 Supported Solid State Devices

Device	Capacity in Gigabytes	Minimum Microcode Version	Minimum Hardware Version	Notes
EZ31-VW	0.134	V064	A01	2, 3
EZ32-VW	0.268	V064	A01	2, 3
EZ51R-VA	0.10	V096	D01	2, 3
EZ54R-VA	0.42	V109	C02	2, 3
EZ58R-VA	0.85	V110	D01	1, 2, 3
EZ64-VA	0.475	V064	A01	2, 3
EZ64-VW	0.475	V070	A01	2, 3
EZ69-VA	0.950	V064	A01	2, 3
EZ69-VW	0.950	V070	A01	2, 3
EZ454	.536	Y018	A01	2, 3
EZ832	3.2	Y018	A01	2, 3
EZ41	0.134	V012	A01	2, 3
EZ42	0.268	V012	A01	2, 3
EZ51	.107	V109	C02	2, 3

Table 2-4 Supported Solid State Devices - continued

Device	Capacity in Gigabytes	Minimum Microcode Version	Minimum Hardware Version	Notes
EZ54	.428	V109	C02	2, 3
EZ705	0.536	V012	A01	2, 3
EZ711	1.1	V012	A01	2, 3
EZ716	1.6	V012	A01	2, 3

Table 2-4 Notes:

1. Code load is not supported for these drives.
2. Formatting supported for these drives.
3. 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 2-5 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
RRD46-VA	0.6	1337	A01
RRD47-VA	0.6	1206	A01

Table 2-5 Notes:

Do not warm-swap CD-ROM drives. Make sure power to the device shelf is turned off before removing or inserting this device.

3.0 New Items

Sections 3.1 through 3.3 briefly describe the new items in HSOF Software V3.7.

- StorageWorks Command Console Support in a Cluster Environment
- New Devices Supported in HSOF Software
- Issue Resolutions

3.1 StorageWorks Command Console Support in a Cluster Environment

StorageWorks Command Console V2.3 Agents are able to access the same controller from multiple hosts when used with HSOF Software V3.7. This allows the use of SWCC from multiple hosts in a cluster environment.

3.2 New Devices Supported in HSOF V3.7 Software

3.2.1 Disk devices

- DS-RZ1EA-VW with 18GB
- DS-RZ1DA-VW with 9GB
- DS-RZ1DD-VW with 9GB
- DS-RZ1FC-VW with 36GB

3.2.2 Tape Devices

- DS-AIT35-VW with 35/70GB

NOTE

Although some of the above listed Devices may have been supported in previous releases, all of these Devices have been classified as “new” for the HSOF V3.7 Software Release. Please see Section 2.5 for additional information on Device Support.

3.3 Issue Resolutions

This section addresses issues that occurred in a previous release of the HSOF software and have been resolved in the HSOF V3.7 release. The issues addressed by HSOF V3.7 Software are as follows:

- A rare occurrence of a disk device responding to TEST UNIT READY but not to any subsequent commands
- A rare occurrence of the controller experiencing a reset with the last failure code of 43130100 when using scripts that perform the configuration/de-configuration of storagesets
- A rare occurrence in which an unusual drive error caused RAIDsets to suspend activity on the controller. This situation could occur when a drive became busy and as a result, would tie up the port
- An occasional problem with failover on a dual-redundant controller pair during certain unusual cache battery states.
- An occasional issue of a disk entering a mount verification timeout state when changing the preferred path on a stripeset.
- A rare occurrence of a drive error causing an “online with lost data” message and a spareset lockup or I/O failure when a drive was being added to the storageset
- For Compaq Tru64 UNIX V4.0 systems, a rare occurrence in which multiple Bad Block Replacements in the same data transfer from a mirrorset cause a Unit Attention message to be sent to the host. This may lead to Compaq Tru64 UNIX V4.0x perception of an I/O failure, resulting in various host effects. The correction for this problem requires installation of both HSOF Software V3.7 and a patch kit to Compaq Tru64 UNIX V4.0x. See Section 7.1, *Compaq Tru64 UNIX Patch Kits for Issue Resolution*, to obtain information on Compaq Tru64 V4.0 versions and their corresponding patch kits
- For OpenVMS systems, a rare occurrence of a medium offline error generated for the OpenVMS loader when a mount command was issued immediately after a dismount or when a tape was mounted in a magazine loader that was set to the automatic mode and the backup spanned from one tape volume to the next

- In a rare occurrence the HSZ40 controller may not be able to use the HSUTIL program to perform a “code load” firmware upgrade to TZ87/88/89 tape drives or any associated Automated Tape Library. In such a case the HSUTIL program refuses to perform the upgrade because it incorrectly determines that there is a tape present in the target tape drive when there is none. Please see Section 7.1 *HSUTIL Program Failure During a Code Load Firmware Upgrade* for a description of this issue and directions for resolution.
- A rare occurrence of a drive reporting “Not Ready/Becoming Ready” where the drive was unable to recover, triggering a controller reset

4.0 Features from Previous Releases

Sections 4.1 through 4.5 briefly describe changes that were introduced in the code version previous to HSOF V3.7 Software release that are not covered in other documentation

- Host-Level I/O Failure Problem Resolution
- Support for TZ887
- Action for repeated Reboots Problem Resolution
- Support for Large Drive Formatting
- Support for up to Fourteen 18GB Drives

4.1 Host-Level I/O Failure Problem Resolution

The V3.4 release addressed the problem in which host-level I/O failures occur due to a storageset unit not-ready condition. This condition is caused by disk devices becoming spontaneously not ready or not responding to controller read or write commands.

4.2 Support for TZ887

A few sites reported that the tape and loader unit would become misconfigured when a TZ887 was added using the CLI ADD command. This is corrected in HSOF Version 3.4. See Table 2-3, “Supported Tape Drives,” in the “Device Support” section for information including capacity, minimum microcode version, and minimum hardware version.

4.3 Action for Repeated Reboot Problem Resolution

Very rarely, a hardware failure might cause a controller to reboot repeatedly, preventing preferred disk devices from completing failover to the other controller. HSOF Version 3.4 contains a modification so that if the same bugcheck repeats itself three times in a ten minute period, the controller will shutdown with no restart to permit the disk devices to failover to the working controller in a dual-redundant pair.

4.4 Support for Large Drive Formatting

The V3.4 release supported formatting of large disk devices (9 GB and 18 GB).

4.5 Support for up to Fourteen 18 GB Drives

The maximum storageset size has been increased to 256 GB so that the controller can support up to fourteen 18 GB drives.

5.0 Clarifications

Sections 5.1 through 5.9 contain clarification on the following subjects:

- EMU Alarm
- Write-Back Cache Batteries and Battery Handling
- Using SHOW on Both Controllers
- Logical Block Address Does Not Match for MEDIUM ERROR
- Adding RAIDsets When Battery Condition is Low
- Drive-Level Event Reports
- Logical Device Event Reports
- Saving the Subsystem Configuration
- Last Failure Code 010E0110 Normal Occurrence After Upgrade

5.1 EMU Alarm

The environmental monitoring unit (EMU) alarm sounds for individual device failures, but does not sound for storageset failures.

5.2 Write-Back Cache Batteries and 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. The following paragraphs clarify battery-related topics.

5.2.1 *CACHE_UPS* Switch Use

If the SET *controller CACHE_UPS* switch is activated, the only action taken in the event of battery failure is to send an error message.

The *CACHE_UPS* is not intended to replace batteries. The switch is intended to control the impact that the current battery state has on cache policy. Battery state is still reported and if the power fails and UPS is broken or exhausted the batteries will still maintain cache contents.

When the *CACHE_UPS* switch is enabled, 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 external UPS. There are many variables in a UPS configuration; some configurations might put data at risk if the UPS is exhausted. 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, if UPS power is lost and if the cache batteries have failed, 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.

5.2.2 Cache Policy

The setting chosen for SET controller *CACHE_POLICY* (to either 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. The battery is tested by the software every four minutes. Full caching operations 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 the:

- Battery becomes fully charged or
 - 10 hours elapse without the battery becoming fully charged.
- In the latter 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.

Table 5-1 summarizes the effect of both cache policy settings on access to RAIDsets and mirrorsets when a battery is “low” during the initial test period.

Table 5-1 Cache Policy Settings

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 storagesets (RAID and mirrorsets). This is because of an inherent delay in the writing of data to the different members of these storagesets. 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/mirrorsets on controllers with low batteries) is the only 100% guarantee of user data consistency. That is why Cache Policy A enforces loss of storageset access when a battery 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.

5.2.3 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. The following failed battery action is taken when a failed battery is detected during either initial or periodic testing:

- Dual-controller HSZ40s (dual batteries on each cache module and the partner controller 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 RAIDsets and mirrorsets
- Single-controller HSZ40s
 - No shut down
 - No RAIDset or mirrorset access permitted
 - Individual disk units accessed in write-through mode

5.2.4 Extending the Life of the External Cache Batteries

Occasionally, circumstances will require you to shut down your controller and restart it. If the correct steps for turning off power to an HSZ40 array controller configuration with write-back cache are not followed, there is potential for loss of data that may exist on any devices connected with them. In addition, if you are turning off the power to the controller subsystem for any reason for longer than 1 day (such as a holiday, system move, or replacing a bad SCSI host cable), you will need to turn off the external cache batteries (ECBs) 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, could result in permanent and irreversible degradation to the batteries.

5.2.4.1 Shutting Down a Controller

1. Use the proper procedures for shutting down the operating system if the host system is also going to be turned off.
2. If the host is not going to be turned off, it is not necessary to shut down the host system, but use the proper operating system procedures to dismount any units that are accessed through the HSZ40 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.
4. When the dismount and/or the operating system shutdown procedures are complete, invoke the controller SHUTDOWN commands on the controllers.

CAUTION

Do not turn off the power to the controller subsystem until all shutdown 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.

The green reset LED stops blinking when the SHUTDOWN command completes.

6. When the controller SHUTDOWN command successfully completes, 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).

5.2.5 Disabling the Write-Back Cache Batteries

Refer to Chapter 2 in *HSZ40 Array Controller HSOF Version 3.1 Service Manual*, for instructions on removing and replacing the controller module and write-back cache module.

1. Stop all host activity and dismount all device units from the host system. Device service is interrupted for the duration of the service cycle. You can shut down the host operating system to accomplish this.
2. Connect a maintenance terminal to the controller's maintenance terminal port on the front bezel.

3. Take the controller out of service by issuing the one of the following two CLI command steps.

- Single Controller:

SHUTDOWN THIS_CONTROLLER

- Dual Controllers:

SHUTDOWN OTHER_CONTROLLER

SHUTDOWN THIS_CONTROLLER

When the shut down completes, the controller or controllers' green reset LED on the operator control panel (OCP) stops blinking.


4. Power off the entire storage subsystem.
5. Press the battery disable switch, the small button labeled SHUT OFF next to the flashing status LED on the ECB, and hold it in for approximately 2 seconds (refer to Figure 1). The cache LED will flash once, then shut off.

NOTE

The battery is no longer powering the cache.

6. To return to normal operation, the cache battery is enabled when the RAID subsystem is powered on.
7. Restart the Controller
 - **PCMCIA Card Removed:** Hold down the reset button while inserting the PCMCIA Card. When you release the reset button, the controller restarts.
 - **PCMCIA Card Not Removed:** Press the reset button and the controller restarts.

Figure 5-1 External Cache Battery Disable Switch



CXO6164A

5.2.5.1 Restarting the Controller

1. Before you restore power to the subsystem, remove the controller module, and the write-back cache module. Remove the battery disable jumper and replace it so it is only covering one pin, then reassemble the unit.
2. 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.
3. If you did not remove the PCMCIA card, press the reset button and the controller restarts.

5.3 Using 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 by using the `SHOW this_controller` command, use the `SHOW other_controller` command to view the information.

5.4 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.

5.5 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 `SET controller 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 3.1 CLI Reference Manual* for additional information about `CACHE_POLICY`.

5.6 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.

5.7 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.

5.8 Saving Your Configuration

The subsystem configuration should be saved to another location other than NVMEM. If the configuration is saved to another location, then the configuration will not have to be re-entered when the controller module is replaced. The method of ensuring the current subsystem configuration is saved differs depending upon whether your subsystem is a single controller (Section 5.8.1) or a dual-redundant controller configuration (Section 5.8.2).

5.8.1 Single Controller Configurations

If the controller in a single controller subsystem requires replacement, the configuration and all pertinent code patches stored in nonvolatile memory (NVMEM) would be lost upon installing a replacement controller module. The CLI command: `INITIALIZE container-name SAVE_CONFIGURATION` provides the means to save the NVRAM contents to another location so that it may be retrieved when needed.

The CLI command does this by saving a copy of the configuration file to all members of the container specified in the command. Refer to the *HSZ40 Array Controller HSOF V3.1 CLI Reference Manual* for details concerning the INITIALIZE command; see also the description in the *HSZ40 Array Controller HSOF V3.1 Configuration Manual*, Chapter 2, “Planning Storagesets”.

The use of the CLI command: `INITIALIZE container-name SAVE_CONFIGURATION` saves you from having to enter the subsystem configuration from scratch if the controller in a single-controller configuration requires replacement.

The *SAVE_CONFIGURATION* switch should **not** be used when upgrading the HSOF software on a single-controller configuration. The configuration data that is stored in NVMEM is not overwritten when the new software is loaded (see the *HSZ40 Array Controller HSOF V3.1 Configuration Manual*, Chapter 2, “*Planning Stagesets*”, for additional information).

NOTE

The *SAVE_CONFIGURATION* switch should not be used when upgrading your hardware, and will not perform inter-platform conversions. For example, you cannot use the *SAVE_CONFIGURATION* switch to upgrade an HSZ40 to an HSZ70 Array Controller. Configuration information cannot be retrieved from stagesets on other HSx controllers (for example, HSD, HSJ, or prior HSOF software versions operating on HSZ controllers).

NOTE

Compaq recommends that the *SAVE_CONFIGURATION* switch only be used for non-redundant (single) controller configurations. To save the configuration information for dual-redundant configurations, use the CLI command:
SET FAILOVER *COPY=controller* (see Paragraph 5.8.2).

5.8.1.1 Initializing Using the SAVE_CONFIGURATION Switch

When initializing a container using the *SAVE_CONFIGURATION* switch, the controller copies the subsystem configuration to all members of the specified container. If you use the switch for a multi-device container, (a stripeset for example), the complete information is stored on **each** device in the container. Be aware that a capacity reduction of 256 KB applies to each member of the container being initialized using the *SAVE_CONFIGURATION* switch.

After initializing a container using the *SAVE_CONFIGURATION* switch, the controller keeps the copy up-to-date. Each time you change the subsystem configuration or add a patch to the HSOF software, the controller copies the new configuration to all containers that were initialized with the *SAVE_CONFIGURATION* switch.

Compaq does not recommend initializing **all** of your containers using the *SAVE_CONFIGURATION* switch (initializing all of the containers causes too many configuration-write operations). Too many write operations are unnecessary and can adversely affect system performance.

5.8.1.2 Replacing a Controller

There are three possibilities that occur regarding the saving of a subsystem configuration when replacing a controller:

- Controller is “factory initialized” and no prior subsystem configuration data
- Controller is not “factory initialized” and has no prior subsystem configuration data
- Controller has some prior subsystem configuration data in it

With any of these possibilities, a local connection should be made to the controller and a printed copy of the configuration as shown by the CLI command: *SHOW disk* should be available before replacing the controller.

5.8.1.2.1 Controller Factory Initialized (No Configuration Data)

With this situation, the new controller searches the devices in the subsystem for a saved configuration:

- Upon finding a saved configuration, it loads it into the new controller's NVMEM and brings the subsystem on-line (it is not necessary to issue an INITIALIZE command after installing a new controller).

If a local connection was made to the maintenance port of the controller, then the CLI will display the events going on: "Looking for Configuration", "Restarting Configuration", etc.

- If a saved configuration is not found, no configuration data is loaded into NVMEM and the subsystem is not brought on-line. Subsystem configuration data will have to be entered into the controller. The methods of entering this data is by way of the StorageWorks Command Console (SWCC) or by way of the CLI command: RUN *config* through a maintenance terminal (RUN *config* brings up the configuration utility) Refer to the *HSZ40 Array Controller HSOF 3.1 Configuration Manual*, Chapter 4 for additional details.

If a local connection was made to the maintenance port of the controller, then the CLI will display the events going on: "Looking for Configuration", "No Configuration Found", etc.

5.8.1.2.2 Controller Not Factory Initialized (No Configuration Data)

With this situation, the new controller does not search the devices in the subsystem for a saved configuration and the controller is not brought on-line. The CLI command: SET this_controller INITIAL_CONFIGURATION must be entered to allow the configuration saved on a storageset to be brought in.

5.8.1.2.3 Controller Has Prior Configuration Data

With this situation, the controller software displays many errors in finding the disks it assumes is part of the configuration. The CLI command: SET this_controller INITIAL_CONFIGURATION must be entered to clear the incorrect configuration from NVMEM and bring in the proper configuration from the container that has it stored. See the *HSZ40 Array Controller HSOF V3.1 CLI Reference Manual* for additional details.

CAUTION

Do **not** enter any other SET *this_controller* command other than a SET *this_controller INITIAL_CONFIGURATION*. Entering another type of SET *this_controller* command may change the configuration data stored in the controller NVMEM. When this happens, the software initiates an update of the saved configuration to the container(s) in which the configuration was saved, thereby destroying the proper saved configuration

5.8.1.3 Upgrading HSOF Software

When upgrading HSOF software, you can refresh the configuration data on your storagesets to include the new software revision by issuing the following command:

```
SET UNIT XXX WRITE_PROTECT
SET UNIT XXX NOWRITE_PROTECT
```

[where xxx is any existing unit on the controller]

These commands change the contents of NVMEM, causing the updated configuration information to be automatically written to the save configuration area of all containers initialized with this option.

5.8.2 Dual-Redundant Controller Configurations

Subsystem configuration and code patch information is stored in the NVMEM of both controllers of a dual-redundant configuration. If one of the two controller modules in a dual-redundant configuration needs replacing, no configuration information is lost because the other controller has it. There is no need to use the CLI command: INITIALIZE *container-name SAVE_CONFIGURATION*; in fact, its use is discouraged.

The method of saving the configuration in dual-redundant controller configurations is by using the CLI command: SET FAILOVER COPY=*controller* (refer to *HSZ40 Array Controller HSOF V3.1 CLI Reference Manual* for additional information). The act of placing the two controllers into failover mode with this command ensures that the configuration in the primary controller (the one that had the configuration placed in it) is copied to the companion controller in the dual-redundant pair.

NOTE

Ensure you know which controller has the good configuration information before entering this command. The device configuration from the controller specified in the CLI command: SET FAILOVER COPY=*controller* overwrites the information on the companion controller.

5.9 Last Failure Code 010E0110 Normal Occurrence After Upgrade

When updating firmware on the controllers, the following message may appear suggesting that there is a problem with the controller or that there was an issue during the update process. The below message is normal and is not cause for alarm.

```
Last fail code: 010E0110
```

```
Press " ?" at any time for help.
```

```
%CER-HSZ_2> --13-JAN-1946 04:32:20 (time not set)--  
System Information Page-and Last Failure entries reset  
to default settings. Last Failure Entry: 1. Flags:  
00000002 Template: 1.(01) Description: Last Failure  
Event Power On Time: 1. Years, 127. Days, 3. Hours, 27.  
Minutes, 5. Seconds Controller Model: HSZ40 Serial  
Number: ZG05051997 Hardware Version: 0000(00) Firmware  
Version: V57Z(00)
```

```
Instance Code: 0102030A Description:
```

```
An unrecoverable firmware inconsistency was detected or  
an intentional restart or shutdown of controller  
operation was requested.
```

This is a normal occurrence for the first boot following manufacture of the controller module and during the transition from one firmware version to another if and only if the format of the System Information Page (SIP) is different between the two versions.

If this event is reported at any other time, follow the recommended repair action associated with this Last Failure Code.

6.0 Operating Constraints

Paragraphs 6.1 through 6.11 describe the operating constraints for HSOF Version 3.7 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 or host-related information, 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:

- Dependency Between HSOF V3.7 and Compaq Tru64 UNIX V5.0 and Higher
- SCSI Target ID Restriction for Compaq Tru64 UNIX Versions Prior to V5.0
- CFMENU Constraints
- Low/Failed Battery on Dual-Redundant Configurations with Preferred Paths
- Device Adds, Moves, and Changes
- Tape Drive Firmware Revision
- Running CONFIG or CFMENU Utilities During Backup or Copy Operations
- Changing Host Adapter Types
- Moving Write-Back Cache Modules
- Limits on the Total Number of Stagesets
- CLONE Utility Constraints

6.1 Dependency Between HSOF V3.7 and Compaq Tru64 UNIX V5.0 and Higher

If you are operating with any Compaq Tru64 UNIX versions prior to V5.0, it is important to install HSOF software V3.7 on your controllers **before** upgrading to Compaq Tru64 UNIX versions 5.0 or higher.

6.2 SCSI Target ID Restriction for Compaq Tru64 UNIX Versions Prior to V5.0

The HSZ40 controllers can support host SCSI Target IDs between 0 and 15. However, Compaq Tru64 UNIX versions prior to V5.0 recognize SCSI Target IDs in the range 0 to 7 only.

Therefore, in Compaq Tru64 UNIX versions prior to V5.0 the useable host SCSI IDs for the HSZ40 array controller subsystems must be assigned within the range of 0 to 7. Please note that SCSI ID 7 is typically assigned to the host SCSI adapter.

NOTE

This restriction does not exist for Compaq Tru64 UNIX V5.0 and higher.

6.3 CFMENU Constraints

CFMENU can only delete single sparesets. Delete sparesets one at a time. Delete spareset members one at a time by responding Y (yes) to one member, then responding N (no) to the rest of the spareset's members. CFMENU deletes the spareset member.

You can also delete spareset member 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.

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

If an external cache battery goes low or fails for a controller in a dual-redundant configuration and preferred paths have been set, that controller's units may go available on the "OTHER" controller and become inoperative on the preferred controller. 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 mount verify. Later you can reissue the `SET unit-name NOPREFERRED_PATH` command to the affected units.

6.5 Device Adds, Moves, and Changes

The controller maintains configuration maps of device types and locations and uses those maps for communicating 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.

6.6 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 re-add the tape drive after the firmware installation is complete. This prevents interrupts to the drive during the code loading operation.
- After the firmware is installed, refer to the HSZ40 Array Controller HSOF Version 3.1 Configuration Manual, HSZ40 Array Controller HSOF Version 3.1 Service Manual, HSZ40 Array Controller HSOF Version 3.1 CLI Reference Manual for adding the tape drive back to the configuration.

6.7 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.

6.8 Changing Host Adapter Types

If for some reason you change your host adapter and maintain the same node number, you **MUST** restart your controller or controllers.

6.9 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 shut down of the controller and clear any data from the cache. Follow these steps to move a write-back cache module:

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

SHUTDOWN THIS_CONTROLLER

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 *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.

6.10 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 256GB per stageset

NOTE

On Compaq Tru64 UNIX, when building stripesets with drives of capacity 4 GB or greater, the largest stripeset that can be built contains 11 members.

On OpenVMS, when building RAIDsets with drives of capacity 4GB or greater, if you are using OpenVMS 6.2, the largest RAIDset that can be built contains 12 members. For OpenVMS V7.0 this limitation does not exist.

6.11 CLONE Utility Constraints

The CLONE utility cannot be used with partitioned units. 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.

7.0 Avoiding Problem Situations

In certain situations, you may experience unusual controller behavior. Paragraphs 7.1 through 7.10 present information to help you avoid such situations and to recover from them if they occur:

- HSUTIL Program Failure During a Code Load Firmware Upgrade
- Compaq Tru64 UNIX Patch Kits for Issue Resolution
- Unit Problems after Battery Low
- CLEAR_ERRORS LOST_DATA Command
- Only One Switch per CLI Command Recommended
- Hot Swapping Devices
- Write-Back Cache Battery Replacement
- Restoring Power after Emergency Cab Shutdown
- Using the CLI and CF Menu Simultaneously in a Dual Configuration
- Changing Preferred Access Device on Compaq Tru64 UNIX

7.1 **HSUTIL Program Failure During a Code Load Firmware Upgrade**

As noted in Section 3.3, in a rare occurrence the HSZ40 controller may not be able to use the HSUTIL program to perform a “code load” firmware upgrade to TZ87/88/89 tape drives or any associated Automated Tape Library. In such a case the HSUTIL program refuses to perform the upgrade because it incorrectly determines that there is a tape present in the target tape drive when there is none.

The HSZ40 controller has the capability to perform firmware upgrade “code loads” to tape drives included in its storage subsystem configuration. This code load capability is manually accessed via the HSUTIL program in the controller. The code load procedure requires that there be no tape in the drive. In rare circumstances, HSUTIL in these controllers incorrectly determines that there is a tape in the target drive, and refuses to perform the firmware upgrade.

You have likely encountered this problem if the following occurs:

- You are attempting to do a firmware upgrade to a TZ87, TZ88, TZ89, or any associated Automated Tape Library
- You are using program HSUTIL in the controller
- There is no tape in the target tape drive
- You receive a HSUTIL dialog similar to the following:

```
After you answer the next question, the code load will start.
```

```
Do you want to continue (y/n) [n] ? y
```

```
HSUTIL is doing code load. Please be patient.
```

```
Error at PTL 3.4.0. Sense key:05 ASC/Q:24/8D.
```

```
The tape cartridge must be removed to update the tape drive firmware.
```

```
The cartridge is in the drive. Please unload.
```

```
HSUTIL - Normal Termination at 03-APR-2000 14:36:06
```

The recommended work around is to use the built-in firmware upgrade capability of the tape device itself. This is done by using a specially prepared tape containing an image of the firmware upgrade code, and manipulating the tape drive and/or Automated Tape Library loader controls to activate the built-in firmware upgrade function of the tape drive or loader.

Please contact your service representative for further instructions on performing this procedure.

7.2 Compaq Tru64 UNIX Patch Kits for Issue Resolution

As noted in Section 3.3, a rare issue related to Bad Block Replacements in data transfers on mirrorsets is corrected by the application of **both** HSOF Software V3.7 and a patch kit to Compaq Tru64 UNIX V4.0x operating systems. The appropriate patch kit is identified in Table 7-1 below:

Table 7-1 Patch Kits for Compaq Tru64 UNIX V4.0x

Compaq Tru64 UNIX Version	Patch Kit
V4.0d	6
V4.0e	4
V4.0f	3

Please use your usual update procedure to obtain the above listed patch kits.

NOTE

The correction in these patch kits is already present in Compaq Tru64 UNIX V5.0 and higher.

7.3 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.

7.4 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.

7.5 Only One Switch per CLI Command Recommended

Certain switches within CLI commands are incompatible. To avoid problems, specify only one switch 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
```

7.6 Hot Swapping Devices

When hot swapping a device, wait 90 seconds after you restore power before enabling the bus, issuing CLI commands to the controller, or allowing any activity to the controller.

7.7 Write-Back Cache Battery Replacement

Please observe the two year posted replacement period for the write-back cache batteries as specified in the *Instructions for Write-Back Cache Battery Service Label* (FRU Insert EK-WBCIN-FI rev B01).

7.8 Restoring Power After Emergency Cab Shutdown

This section provides the procedures to restore power to each of the extended cabs after the master cab has been powered off by an “emergency cab shutdown”. Examples causing this type of shutdown would be fan loss or power-supply loss, etc.

It is recommended that the extended cabs be powered off and powered back on prior to the reboot of the master cab. The steps to follow are:

1. Power off remaining (extended) cabs.
2. Power up extended cabs.
3. Power up Master Cab.

There are two procedures that may be used to accomplish the above three steps. Compaq recommends the first procedure.

7.8.1 Using AC switch

1. Turn off AC switch after EMU alarm sounds on extended cabs.
2. Turn on AC switch on extended cabs.
3. Turn on AC switch on Master Cab.

7.8.2 Using EMU/PVA

1. On extended cabs, wait for EMU alarm to sound, turn off EMU alarm, then press and hold PVA switch until the cab shuts down.
2. Press EMU button on extended cabs.
3. Press EMU button on Master Cab.

7.9 Using the CLI and CF Menu Simultaneously in a Dual Configuration

When in a dual configuration do not issue a simultaneous CLI command on one controller and a CFMENU command on the other controller. This can cause a deadlock situation between the two controllers while each waits for a CLI lock.

7.10 Changing the Preferred Access Device on Compaq Tru64 UNIX

To change the preferred access device for the SWCC UNIX Agent, perform one of the following procedures:

- Directly edit the storage.ini file which is kept in the /usr/opt/SWCC500/etc directory. It is an ASCII file, and any ASCII editor may be used

Or

- Delete/rename the existing storage.ini file. When the Agent is restarted, a new storage.ini file is created with default names and monitor intervals

Or

- Delete the existing storage.ini file and run the swcc_config utility. Select option 2 (Modify Storage Subsystem Information) and select the add option. This will rescan the device and will use the current access device

Once the storage.ini file has been changed, the user must stop and restart the Agent in order for the changes to take effect.

8.0 Documentation Additions and Corrections

The following sections describe corrections to:

8.1 *Array Controller HSOF Version 3.1 Configuration Manual*

8.2 *Array Controller HSOF Version 3.1 CLI Reference Manual.*

8.3 *Array Controller HSOF Version 3.1 Service Manual*

8.1 Array Controller HSOF V3.1 Configuration Manual

8.1.1 **Chunk Size**

Referring to the section titled *Initialize Switches* in the *HSZ40 Array Controller HSOF Version 3.1 Configuration Manual*, pages 3-22 through 3-24 the information regarding “Chunk Size” should be replaced with the following information.

The information on maximum chunk size was incorrect and has been removed.

In addition, cautionary information on using the default chunk size as well as information on calculating chunk size, increasing the request rate, increasing the data transfer rate and increasing sequential write performance have been included and are to be added to the documents referenced.

8.1.1.1 Specifying Chunk Size

Specify the chunk size of the data to be stored to control the stripesize used in RAIDsets and stripesets:

- *CHUNKSIZE=DEFAULT*
lets the controller set the chunk size based on the number of disk drives (d) in a stripeset or RAIDset.
If $d \leq 9$ then chunk size = 256. or 128 kilobytes (K)
If $d > 9$ then chunk size = 128. or 64 kilobytes (K)


CAUTION

The default chunk size should be used with caution. It is imperative that you analyze the I/O transfers being used with your application to determine the proper chunk size. An improper value used as a chunk size may cause serious performance problems with the subsystem.

- *CHUNKSIZE=n*
lets you specify a chunk size in blocks. The relationship between chunk size and request size determines whether striping increases the request rate or the data-transfer rate.

8.1.1.2 Increasing the Request Rate

A large chunk size (relative to the average request size) increases the request rate by allowing multiple disk drives to respond to multiple requests. If one disk drive contains all of the data for one request, then the other disk drives in the storageset are available to handle other requests. Thus, in principle, separate I/O requests can be handled in parallel, thereby increasing the request rate. This concept is shown in Figure 8-1.

Figure 8-1 Chunk Size Larger than the Request Size

Applications such as interactive transaction processing, office automation, and file services for general timesharing tend to require high I/O request rates.

Large chunk sizes also tend to increase the performance of random reads and writes. It is recommended that you use a chunk size of 10 to 20 times the average request size, rounded to the closest prime number. In general, a chunk size of 239 works well for Compaq Tru64 UNIX systems with a transfer size of 16 sectors; 113 works well for OpenVMS systems with a transfer size of 8 sectors.

To calculate the chunk size that should be used for your subsystem, you first must analyze the types of requests that are being made to the subsystem:

- Lots of parallel I/O that use a small area of disk should use a chunk size of 10 times the average transfer request rate.
- Random I/Os that are scattered over all the areas of the disks should use a chunk size of 20 times the average transfer request rate.
- If you don't know, then you should use a chunk size of 15 times the average transfer request rate.
- If you have mostly sequential reads or writes (like those needed to work with large graphic files), then make the chunk size a small number (i.e. 17 sectors in Table 8-1 shows a few examples of chunk size selection.


Table 8-1 Example Chunk Size

Transfer Size (KB)	High Locality	Unknown Locality	Low Locality
2	41	59	790
4	79	113	163
8	157	239	317

8.1.1.3 Increasing the Data Transfer Rate

A small chunk size relative to the average request size increases the data transfer rate by allowing multiple disk drives to participate in one I/O request. This concept is shown in Figure 8–2.

Figure 8-2 Chunk Size Smaller than the Request Size



Applications such as CAD, image processing, data collection and reduction, and sequential file processing tend to require high data-transfer rates.

8.1.1.4 Increasing Sequential Write Performance

Sequential write and read requests on stripesets (or striped mirrorsets), should use a small chunk size relative to the I/O size to increase the performance. A chunk size of 17 generally works well

8.2 Array Controller HSOF V3.1 CLI Reference Manual

8.2.1 *Chunk Size*

Referring to the section titled *Initialize Switches*, in the *HSZ40 Array Controller HSOF Version 3.1 CLI Reference Manual*, page 2-38 the information regarding “Chunk Size” should be replaced. Please refer to the above section of these Release Notes (*Section 8.1.1, Chunk Size*) for the correct information.

8.2.2 *SET unit-number*

Referring to the section titled *Set unit-number* in the *HSZ40 Array Controller HSOF Version 3.1 CLI Reference Manual* on page 2-80, the below example was included in error. Please strike through this example.

To set unit T47 to 1600 bpi

```
CLI> SET T47 DEFAULT_FORMAT=1600BPI_9TRACK
```

8.3 Array Controller HSOF V3.1 Service Manual

8.3.1 *Software Installation*

Please see section 9.0, *Software Installation* for software installation procedures.

Section 9.1 provides the specific procedures needed to install your HSOF V3.7 Software on all **Non-Compaq Tru64 UNIX systems**.

Section 9.2 provides the specific procedures needed to install your HSOF V3.7 Software on all **Compaq Tru64 UNIX systems**.

The information in these sections can be used in conjunction with the information found in the *HSZ40 Array Controller HSOF V3.1 Service Manual, Chapter 3 “Installing and Upgrading”*, pages 3-3 through 3-21.

8.3.2 HSUTIL Procedure

Please insert the following procedure in the *HSZ40 Array Controller HSOF V3.1 Service Manual, Chapter 3 “Installing and Upgrading”, page 3-36* after the section “*Considerations for Installing New device Firmware*”.

Use the below procedure to upgrade firmware with HSUTIL.

1. Connect a local terminal to the controller that accesses the device you want to upgrade.
2. Configure a single-disk unit.

NOTE

In the steps that follow, you will copy the firmware image to this unit, then use HSUTIL to distribute it to the devices you are upgrading. This unit must be a newly initialized disk with no label or file structure to ensure that the firmware image resides in contiguous blocks starting from LBN 0 or another known LBN. Additionally, write-back caching must be disabled (see *HSZ40 Array Controller HSOF Version 3.1 CLI Reference Manual*). See the *HSZ40 Array Controller HSOF Version 3.1 Configuration Manual* for instructions on configuring a single-disk unit. See the Firmware Release Notes for details and exceptions to the code load process.

3. Copy the firmware image to the single-disk unit that you configured in Step 2. The firmware image must begin at a known LBN (usually 0) and must be contiguous. See the documentation that accompanied your host’s operating system for instructions on copying firmware images to a disk drive.

CAUTION

You must quiesce the host load before running HSUTIL or damage to the storage device can occur.

4. Start HSUTIL with the following command:

```
HSZ> RUN HSUTIL
```

HSUTIL displays the following:

```
*** Available functions are:
```

- 0. EXIT
- 1. FORMAT_DISK
- 2. DEVICE_CODE_LOAD_DISK
- 3. DEVICE_CODE_LOAD_TAPE

5. Enter 2 or 3 from the HSUTIL menu.

Example: 2

6. Choose the single-disk unit as the source disk for the download. You will see:

```
Available single device units on this controller include:
Unit   Associated Device   SCSI Product ID   Current Device
Rev 0   DISK10500             RZ1BB-CS (C) DEC  0844
```

The HSUTIL Menu will ask:

```
Which unit is the code to be loaded FROM?
```

Example: 0

7. Enter the starting LBN of the firmware image (Usually LBN 0)
or
Press Enter for the default [LBN 0]

Example: [LBN 0]

8. Enter the product ID of the device you want to upgrade. This ID corresponds to the product information that's reported in the "Type" column when you issue the SHOW DISK command. The HSUTIL menu will ask the following:

```
What is the SCSI PRODUCT ID of the device that you want code
load TO?
```

Example: rz1bb-cs

HSUTIL lists all devices that correspond to the product ID you entered:

```
Unattached devices on this controller include:
```

Device	SCSI Product ID	Current Device Rev
DISK20500	RZ1BB-CS (C) DEC	0844
DISK10400	RZ1BB-CS (C) DEC	0844

9. Enter the disk or tape name of the device you want to upgrade.

Example: disk10400

10. If the vendor's method of updating the code image has changed for this device, it may be necessary to treat this as an unsupported device. (See the Firmware Release Notes for device support information.) The HSUTIL menu will ask the following:

```
Do you want to process this device as an unsupported device
(y/n)
```

Example: n

11. Confirm or enter the byte count of the firmware image. This can be found in the Release Notes that were shipped with the Firmware.

```
What is the TOTAL SIZE of the code image in bytes (min:65537,
max:524288)
```

Example: 265216

12. Confirm the download.
13. Some disk firmware releases require that you reformat the disk after upgrading its firmware. See the documentation that accompanied the firmware to determine if you need to reformat the device.

8.3.3 C_SWAP Procedure

The procedure in the *HSZ40 Array Controller HSOF Version 3.1 Service Manual* for restarting the subsystem after replacing a controller with C_SWAP should contain the following note:

NOTE

The following error message is normal. Please continue with controller configuration:

```
Other controller not responding - RESET signal
asserted.
```

8.3.4 Fault Management Utility Instance Codes

Add the following Fault Management Utility (FMU) Instance Codes:

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

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

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

9.0 Software Installation

9.1 Software Installation for Non-Compaq Tru64 UNIX Systems

This section provides the specific procedures needed to install your HSOF V3.7 Software. The information in this section can be used in conjunction with the information found in the *HSZ40 Array Controller HSOF V3.1 Service Manual, Chapter 3 “Installing and Upgrading”*, pages 3-3 through 3-21.

Prior to installing HSOF V3.7 Software, please note the following:

CAUTION

If JBOD disks that were initialized using the `SAVE_CONFIGURATION` switch on a HSZ40 controller running HSOF Version 2.7 have been moved onto your HSZ40 controller, please contact your service representative before upgrading your HSZ40 controller to HSOF Version 3.7 software.

The following section describes two types of controller upgrades:

- Single Controller Upgrade
- Dual-Redundant Controller Upgrade

9.1.1 Single Controller Upgrade

A Single Controller Upgrade is synonymous with a Shutdown Upgrade. A Shutdown Upgrade is one in which the storage device is taken off-line for the period of time that the software is being installed. Use the following procedure to upgrade the HSOF software in a non-redundant (single) controller:

1. Halt all host I/O activity to the controller and prepare the host system for a controller shutdown using the appropriate procedures for your operating system.

2. Establish a local connection to the controller through the maintenance port of the controller.
3. Shutdown the controller with the following command:
SHUTDOWN THIS_CONTROLLER

CAUTION

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

CAUTION

Before continuing, wait until the Reset button on the controller stops flashing (lit continuously). This indicates that unwritten data is flushed from the cache module.

4. Remove the PCMCIA program card ESD cover from the controller.
5. While pressing and holding the controller Reset button, press the program card eject button to eject the program card from the controller. Release the button after the card is ejected.

NOTE

The controller will initialize and the Reset button will flash at an approximate once per second rate when ready for operation.

6. Press and hold the Reset button while inserting the new program card. Release the button after the new card is inserted.

NOTE

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.

NOTE

The controller will initialize and the Reset button will flash at an approximate once per second rate when ready for operation.

7. Replace the ESD cover on the controller, pushing the locking pins to secure them in place.
8. Mount the storage units on the host. Resume I/O activity to the controller.

NOTE

You can identify the version of the HSOF software that is currently loaded in the controller by entering a `SHOW THIS_CONTROLLER` command at the CLI prompt.

9.1.2 Dual-Redundant Upgrade

Before you can begin the Dual-Redundant upgrade you will first need to determine which one of the following two upgrade procedures you will perform:

- Dual-Redundant Controller Rolling Upgrade

Or

- Dual-Redundant Controller Shutdown Upgrade

Below is a description of each of the two procedures and Table 9-1 shows which procedure is necessary for the HSZ40 HSOF Software versions.

Dual-Redundant Controller Rolling Upgrade:

The Dual-Redundant Controller Rolling Upgrade procedure is one in which the disk subsystem does not have to be shut down to upgrade the software. Using the failover capabilities of the controllers, only one controller at a time needs to be shut down, allowing the host system to have disk drive activity from the disk subsystem being upgraded. When performing rolling upgrades, it is required that the most recent software patches (if any), be installed on both the lower and higher code versions.

Dual-Redundant Controller Shutdown Upgrade:

The Dual-Redundant Controller Shutdown Upgrade procedure requires that you change your dual-redundant configuration into two single controller configurations for the HSOF Version 3.7 Software upgrade. You must shut down 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.

Table 9-1 HSOF Software Versions and Supported Upgrade Procedures

Rolling Upgrade to V3.7 From:	Shutdown Upgrade to V3.7 From:
Versions 3.4, 3.5, and 3.6	Versions 3.0, 3.1, 3.2, and 3.3

As shown in the above table, Rolling Upgrades to Version 3.7 are supported from Versions 3.4, 3.5, and 3.6.

Releases prior to HSOF Software Version 3.4 must upgrade to HSOF Software Version 3.7 by performing the shutdown upgrade procedure.

9.1.2.1 Dual-Redundant Controller Rolling Upgrade

If your HSOF Software version supports a Dual-Redundant Controller Rolling Upgrade, follow the below procedures.

1. Establish a local connection to the controller through the maintenance port of one of the controllers (“this controller”).
2. Wait for the CLI prompt, and then shut down “other controller” by using the following command:

```
SHUTDOWN OTHER_CONTROLLER
```

CAUTION

Before continuing, wait until the Reset button stops flashing and stays lit continuously. This indicates that unwritten data is flushed from the cache module.

3. Remove the program card ESD cover from “other controller”.
4. Press and hold the controller Reset button while ejecting the old program card. Release the button after the card is ejected.
5. Press and hold the Reset button while inserting the new program card in “other controller” Release the button after the new card is inserted.

NOTE

Wait for the following to happen before continuing:
The Reset button flashes at an approximate once per second rate
and
The maintenance terminal displays a message that finishes with
other controller restarted.

6. Read the status of “other controller” by issuing the following command:
`SHOW OTHER_CONTROLLER .`

NOTE

When the **SHOW OTHER_CONTROLLER** command verifies that the
other controller restarted,
it is safe to continue .

7. Replace the ESD cover on “other controller” and push the locking pins to secure it.
8. First pass through this procedure?
YES — Move the maintenance port cable from “this controller”
to “other controller” maintenance port and repeat
procedure from step 2 to step 8 to replace the other
program card.

NOTE

“Other controller” switches with “this controller”—the controller
that the maintenance port cable was plugged into originally—
 (“this controller”) becomes “other controller” for the second pass
through steps 2–8.

NO — You have completed this procedure.

9.1.2.2 Dual-Redundant Controller Shutdown Upgrade

If your HSOF Software version requires a Dual-Redundant Controller Shutdown Upgrade, follow the below procedures.

NOTE

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

1. Halt all host I/O activity to the controller and prepare the host system for a controller shutdown using the appropriate procedures for your operating system.
2. Establish a local connection to the controller through the maintenance port of the controller.
3. Enter the SET NOFAILOVER command to take the controllers out of dual-redundant (failover) mode.
\$ SET NOFAILOVER
4. Shutdown the controller with the following command:
SHUTDOWN THIS_CONTROLLER

CAUTION

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

CAUTION

Before continuing, wait until the Reset button on the controller stops flashing (lit continuously). This indicates that unwritten data is flushed from the cache module.

5. Remove the PCMCIA program card ESD cover from the controller.
6. While pressing and holding the controller Reset button, press the program card eject button to eject the program card from the controller. Release the button after the card is ejected.

NOTE

The controller will initialize and the Reset button will flash at an approximate once per second rate when ready for operation.

7. Press and hold the Reset button while inserting the new program card. Release the button after the new card is inserted.

NOTE

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.

NOTE

The controller will initialize and the Reset button will flash at an approximate once per second rate when ready for operation.

8. Replace the ESD cover on the controller, pushing the locking pins to secure them in place.
9. Mount the storage units on the host. Resume I/O activity to the controller.
10. Repeat 5 through 9 for the second controller.
11. 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

9.2 Software Installation for Compaq Tru64 UNIX Systems

CAUTION

If you are operating with any Compaq Tru64 UNIX versions prior to V5.0, it is important to install HSOF software V3.7 on your controllers **before** upgrading to Compaq Tru64 UNIX versions 5.0 or higher.

This section provides the specific procedures needed to install your HSOF V3.7 Software. The information in this section can be used in conjunction with the information found in the *HSZ40 Array Controller HSOF V3.1 Service Manual, Chapter 3 “Installing and Upgrading”*, pages 3-3 through 3-21.

Prior to installing HSOF V3.7 Software, please note the following:

CAUTION

If JBOD disks that were initialized using the *SAVE_CONFIGURATION* switch on a HSZ40 controller running HSOF Version 2.7 have been moved onto your HSZ40 controller, please contact your service representative before upgrading your HSZ40 controller to HSOF Version 3.7 software.

The following section describes two types of controller upgrades:

- Single Controller Upgrade
- Dual-Redundant Controller Upgrade

9.2.1 Single Controller Upgrade

A Single Controller Upgrade is synonymous with a Shutdown Upgrade. A Shutdown Upgrade is one in which the storage device is taken off-line for the period of time that the software is being installed. Use the following procedure to upgrade the HSOF software in a non-redundant (single) controller:

1. Halt all host I/O activity to the controller and prepare the host system for a controller shutdown using the appropriate procedures for your operating system.
2. Establish a local connection to the controller through the maintenance port of the controller.
3. Shutdown the controller with the following command:

SHUTDOWN THIS_CONTROLLER

CAUTION

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

CAUTION

Before continuing, wait until the Reset button on the controller stops flashing (lit continuously). This indicates that unwritten data is flushed from the cache module.

4. Remove the PCMCIA program card ESD cover from the controller.
5. While pressing and holding the controller Reset button, press the program card eject button to eject the program card from the controller. Release the button after the card is ejected.

NOTE

The controller will initialize and the Reset button will flash at an approximate once per second rate when ready for operation.

6. Press and hold the Reset button while inserting the new program card. Release the button after the new card is inserted.

NOTE

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.

NOTE

The controller will initialize and the Reset button will flash at an approximate once per second rate when ready for operation.

7. Replace the ESD cover on the controller, pushing the locking pins to secure them in place.
8. Mount the storage units on the host. Resume I/O activity to the controller.

NOTE

You can identify the version of the HSOF software that is currently loaded in the controller by entering a **SHOW THIS_CONTROLLER** command at the CLI prompt.

9.2.2 Dual-Redundant Upgrade

Before you can begin the Dual-Redundant upgrade you will first need to determine which one of the following two upgrade procedures you will perform:

- Dual-Redundant Controller Rolling Upgrade
- Or
- Dual-Redundant Controller Shutdown Upgrade

Below is a description of each of the two procedures and Table 9-2 shows which procedure is necessary for the HSZ40 HSOF Software versions.

Dual-Redundant Controller Rolling Upgrade:

The Dual-Redundant Controller Rolling Upgrade procedure is one in which the disk subsystem does not have to be shut down to upgrade the software. Using the failover capabilities of the controllers, only one controller at a time needs to be shut down, allowing the host system to have disk drive activity from the disk subsystem being upgraded. When performing rolling upgrades, it is required that the most recent software patches (if any), be installed on both the lower and higher code versions.

Dual-Redundant Controller Shutdown Upgrade:

The Dual-Redundant Controller Shutdown Upgrade procedure requires that you change your dual-redundant configuration into two single controller configurations for the HSOF Version 3.7 Software upgrade. You must shut down 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.

Table 9-2 HSOF Software Versions and Supported Upgrade Procedures

Rolling Upgrade to V3.7 From:	Shutdown Upgrade to V3.7 From:
Versions 3.4, 3.5, and 3.6	Versions 3.0, 3.1, 53.2, and 3.3

As shown in the above table, Rolling Upgrades to Version 3.7 are supported from Versions 3.4, 3.5, and 3.6.

Releases prior to HSOF Software Version 3.4 must upgrade to HSOF Software Version 3.7 by performing the shutdown upgrade procedure.

9.2.2.1 Dual-Redundant Controller Rolling Upgrade

If your HSOF Software version supports a Dual-Redundant Controller Rolling Upgrade, follow the below procedures.

CAUTION

Prior to installing HSOF V3.7 software, it is important to first determine whether the two controllers in your dual-redundant configurations have the same or different product ID's, this will determine the upgrade procedure you will use to continue.

To view the product ID's on your dual-redundant controllers, establish a local connection to the controller through the maintenance port of one of the controllers, (which becomes "this controller"). Enter the following commands and observe the resulting displays:

`SHOW THIS_CONTROLLER and SHOW OTHER_CONTROLLER`

The product ID is shown in the first characters of the display, prior to the hardware serial number. Most controllers have the product ID of "HSZ40" but a very few controllers may have the product ID of "HSZ40 (C) DEC."

- IF the Product ID's of the two HSZ40 controllers in a dual-redundant pair are the same, use the procedure for *Same Product ID's for Dual-Redundant Rolling Upgrade* (Section 9.2.2.1.1).
- IF the Product ID's of the two HSZ40 controllers in a dual-redundant pair are different, and IF your current Compaq Tru64 UNIX operating system version is V4.0x or lower, use the procedure for *Different Product ID's and Compaq Tru64 UNIX V4.0x for Dual-Redundant Rolling Upgrade* (Section 9.2.2.1.2).
- IF the Product ID's of the two HSZ40 controllers in a dual-redundant pair are **different, and** IF your current Compaq Tru64 UNIX operating system version is **V5.0**, use the procedure below for *Different Product ID's and Compaq Tru64 UNIX V5.0 for Dual-Redundant Rolling Upgrade* (Section 9.2.2.1.3).

9.2.2.1.1 Procedure for Same Product ID's for Dual-Redundant Rolling Upgrade

1. Establish a local connection to the controller through the maintenance port of one of the controllers (“this controller”).
2. Wait for the CLI prompt, and then shut down “other controller” by using the following command:

```
SHUTDOWN OTHER_CONTROLLER
```

CAUTION

Before continuing, wait until the Reset button stops flashing and stays lit continuously. This indicates that unwritten data is flushed from the cache module.

3. Remove the program card ESD cover from “other controller”.
4. Press and hold the controller Reset button while ejecting the old program card. Release the button after the card is ejected.
5. Press and hold the Reset button while inserting the new program card in “other controller” Release the button after the new card is inserted.

NOTE

Wait for the following to happen before continuing:
The Reset button flashes at an approximate once per second rate and
The maintenance terminal displays a message that finishes with
other controller restarted.

6. Read the status of “other controller” by issuing the following command:
SHOW OTHER_CONTROLLER.

NOTE

When the **SHOW OTHER_CONTROLLER** command verifies that the **other controller restarted**, it is safe to continue .

7. Replace the ESD cover on “other controller” and push the locking pins to secure it.
8. First pass through this procedure?
 - YES — Move the maintenance port cable from “this controller” to “other controller” maintenance port and repeat procedure from step 2 to step 8 to replace the other program card.

NOTE

“Other controller” switches with “this controller”—the controller that the maintenance port cable was plugged into originally— (“this controller”) becomes “other controller” for the second pass through steps 2–8.

NO — You have completed this procedure.

9.2.2.1.2 Procedure for Different Product ID’s and Compaq Tru64 UNIX V4.0x for Dual-Redundant Rolling Upgrade

1. Establish a local connection to the controller through the maintenance port of one of the controllers (“this controller”).
2. Wait for the CLI prompt, and then shut down “other controller” by using the following command:
`SHUTDOWN OTHER_CONTROLLER`

CAUTION

Before continuing, wait until the Reset button stops flashing and stays lit continuously. This indicates that unwritten data is flushed from the cache module.

3. Remove the program card ESD cover from “other controller”.
4. Press and hold the controller Reset button while ejecting the old program card. Release the button after the card is ejected.

5. Press and hold the Reset button while inserting the new program card in “other controller” Release the button after the new card is inserted.

NOTE

Wait for the following to happen before continuing:
The Reset button flashes at an approximate once per second rate
and
The maintenance terminal displays a message that finishes with
other controller restarted.

6. Read the status of “other controller” by issuing the following command:
SHOW OTHER_CONTROLLER.

NOTE

When the **SHOW OTHER_CONTROLLER** command verifies that the
other controller restarted,
it is safe to continue .

7. Replace the ESD cover on “other controller” and push the locking pins to secure it.
8. When “other controller” has completely rebooted, move the local connection to the maintenance port of **other** (newly upgraded) controller, which will cause the newly upgraded controller to become “this controller.”
Issue the **SHOW THIS_CONTROLLER** command.
The resulting display will show that the controllers are misconfigured due to the differing product ID’s.
9. Issue the **SET NOFAILOVER** command to shutdown the “other controller” still containing old software. This will cause the “other controller” to shut down and devices to fail over to “this controller”.
10. Observing all notes and caution statements in steps 2 through 7 remove the old program card, insert the new one, and reset “other controller” to start it.

11. When the second controller has restarted as described in Step 6, issue the **SET FAILOVER COPY=controller** command.

Either **THIS** or **OTHER** may be used as the source in the **COPY=** parameter since both controllers have already been updated.

9.2.2.1.3 Different Product ID's and Compaq Tru64 UNIX V5.0 for Dual-Redundant Rolling Upgrade

When your HSZ40 controllers with different product ID's are configured in a system using Compaq Tru64 UNIX V5.0, it is necessary to perform some steps involving the operating system before and after controller upgrade to HSOF V3.7.

1. Determine the **master** controller by following the sub-steps below:
 - a. Use the following command to view details of the devices attached to the host.

```
# hwmgr -view dev -cat disk
HWID: Device Name      Mfg  Model          Location
-----
27: /dev/disk/floppy0c  3.5in floppy  fdi0-unit-0
34: /dev/disk/dsk0c    DEC  RZ28 (C) DEC  bus-0-targ-1-lun-0
35: /dev/disk/dsk1c    DEC  RZ26L (C) DEC  bus-0-targ-2-lun-0
36: /dev/disk/cdrom0c  DEC  RRD44 (C) DEC  bus-0-targ-5-lun-0
37: /dev/disk/dsk2c    DEC  HSZ40          bus-1-targ-3-lun-0
38: /dev/disk/dsk3c    DEC  HSZ40          bus-1-targ-3-lun-1
39: /dev/disk/dsk4c    DEC  HSZ40 (C) DEC  bus-1-targ-6-lun-0
40: /dev/disk/dsk5c    DEC  HSZ40 (C) DEC  bus-1-targ-6-lun-1
```

In the above example dsk2 (D300 on the controller) and dsk3 (D301 on the controller) are reporting the Product ID of "HSZ40" while dsk4 (D600 on the controller) and dsk5 (D601 on the controller) are reporting the Product ID of "HSZ40 (C) DEC" in the Model column.

NOTE

If all the units configured on the dual controller pair are “ONLINE” to only one of the controllers (i.e. the other controller is in stand by mode), then all the devices will show only one Product ID. The controller which has the units “ONLINE” is the **master** controller. Go to Step 2 and perform Steps 2 through 13. It is not necessary in this case to perform Step 14.

- b. Use the command below to determine which device the root file system is mounted on.

```
# df
Filesystem      512-blocks    Used   Avail Capacity  Mounted on
/dev/disk/dsk4a  338542    276436   28250    91%    /
/dev/disk/dsk4g  3389096    520056  2530130   18%    /usr
/proc           0          0         0    100%    /proc
```

In the example above the root file system is mounted on dsk4.

- c. Shut the system down to single-user mode. This will prevent background processes from touching the disks while this procedure is in progress.

```
# shutdown now
```

- d. If the root file system is not mounted on a unit configured on the controller pair, choose either of the controllers to be the "master" controller in the upgrade operation below. Go to step 2.

If the root file system is mounted on a unit configured on the controller pair, then do the following in order to determine which is the master controller.

- e. Connect the CLI interface to one of the controllers and perform the SHOW <UNIT> command for the unit with the root device (D600 in case of the example).

```

HSZ> show d600
      LUN                               Uses
-----
D600                                DISK600
...
      State:
        ONLINE to this controller
        Not reserved
        PREFERRED_PATH = THIS_CONTROLLER
...

```

If the output shows that the unit is “ONLINE” to **this** controller (as is the case in this example), then **THIS_CONTROLLER** is the **master** controller.

If the output shows that the unit is “ONLINE” to the **other** controller, then **OTHER_CONTROLLER** is the **master** controller.

2. Establish a local connection to the **master** controller (determined in Step 1) through the maintenance port of that controller, which now becomes “this controller”.
3. Wait for the CLI prompt, and then shut down “other controller” by using the following command:
SHUTDOWN OTHER_CONTROLLER

CAUTION

Before continuing, wait until the Reset button stops flashing and stays lit continuously. This indicates that unwritten data is flushed from the cache module.

4. Remove the program card ESD cover from “other controller”.
5. Press and hold the controller Reset button while ejecting the old program card. Release the button after the card is ejected.
6. Press and hold the Reset button while inserting the new program card in “other controller” Release the button after the new card is inserted.

NOTE

Wait for the following to happen before continuing:
The Reset button flashes at an approximate once per second rate
and
The maintenance terminal displays a message that finishes with
other controller restarted.

7. Read the status of “other controller” by issuing the following command:
SHOW OTHER_CONTROLLER.

NOTE

When the **SHOW OTHER_CONTROLLER** command verifies that the
other controller restarted,
it is safe to continue .

8. Replace the ESD cover on “other controller” and push the locking pins to secure it.
9. When “other controller” has completely rebooted, move the local connection to the maintenance port of **other** (newly upgraded) controller, which will cause the newly upgraded controller to become “this controller.”
Issue the **SHOW THIS_CONTROLLER** command.
The resulting display will show that the controllers are misconfigured due to the differing product ID’s.
10. Issue the **SET NOFAILOVER** command to shutdown the “other controller” still containing old software. This will cause the “other controller” to shut down and devices to fail over to “this controller”.
11. Observing all notes and caution statements in steps 2 through 7 remove the old program card, insert the new one, and reset “other controller” to start it.
12. When the second controller has restarted as described in Step 6, issue the **SET FAILOVER COPY=controller** command.

CAUTION

Be sure to use **OTHER** as the qualifier in the **COPY=** command to properly match the Product ID's.

13. Follow the commands in the example below to correct the device files:
 - a. Scan the SCSI interfaces to update the operating system's view of device changes and create new devices as necessary.

NOTE

Depending on the size of your configuration, the scan may take several minutes to complete. The presence of tape devices will further increase the delay to complete the scan. For this reason, you may wish to use the **-bus** qualifier to specify the bus you want to scan. The correct bus number can be determined by examining the "location" field of the "hwmgr -view devices" output.

```
# hwmgr -view devices
HWID: Device Name   Mfg  Model          Location
-----
39: /dev/disk/dsk4c DEC  HSZ40  (C) DEC bus-1-targ-6-lun-0
```

Issue the command for the appropriate bus (1 in this example).

```
# hwmgr -scan scsi -bus 1
```

- b. In order to understand which devices need to be redirected to the newly created devices, examine the current list of devices prior to rebooting the system:

```
# hwmgr -view devices
HWID: Device Name      Mfg  Model              Location
-----
4: /dev/kevm
27: /dev/disk/floppy0c  3.5in floppy  fdi0-unit-0
34: /dev/disk/dsk0c    DEC  RZ28 (C) DEC  bus-0-targ-1-lun-0
35: /dev/disk/dsk1c    DEC  RZ26L (C) DEC  bus-0-targ-2-lun-0
36: /dev/disk/cdrom0c  DEC  RRD44 (C) DEC  bus-0-targ-5-lun-0
37: /dev/disk/dsk2c    DEC  HSZ40              bus-1-targ-3-lun-0
38: /dev/disk/dsk3c    DEC  HSZ40              bus-1-targ-3-lun-1
39: /dev/disk/dsk4c    DEC  HSZ40 (C) DEC  bus-1-targ-6-lun-0
40: /dev/disk/dsk5c    DEC  HSZ40 (C) DEC  bus-1-targ-6-lun-1
41: /dev/cport/scp0    HSZ40          bus-1-targ-3-lun-0
42: /dev/cport/scp1    HSZ40 (C) DEC  bus-1-targ-6-lun-0
44: /dev/disk/dsk6c    DEC  HSZ40 (C) DEC  bus-1-targ-3-lun-0
45: /dev/disk/dsk7c    DEC  HSZ40 (C) DEC  bus-1-targ-3-lun-1
```

Note that you can ignore the control port "scp0" device.

We will need to redirect the following device names (with invalid PID's):

```
37: /dev/disk/dsk2c    DEC  HSZ40              bus-1-targ-3-lun-0
38: /dev/disk/dsk3c    DEC  HSZ40              bus-1-targ-3-lun-1
```

To the following new device names (containing valid PID's):

```
44: /dev/disk/dsk6c    DEC  HSZ40 (C) DEC  bus-1-targ-3-lun-0
45: /dev/disk/dsk7c    DEC  HSZ40 (C) DEC  bus-1-targ-3-lun-1
```

Remember the following:

- HWID 37 will be redirected to 44
- HWID 38 will be redirected to 45

- c. At this point, you should reboot the system (ONLY TO SINGLE-USER MODE).

Then mount the root file system to enable writing to the disk:

```
# shutdown -h now
.
.
.
P00>>> boot -flag s dkb600
.
.
.
Starting secondary cpu 1
INIT: SINGLE-USER MODE
# mountroot
.
.
.
```

- d. Examine the output of "hwmgr -view devices" and "hwmgr -show scsi" (The scsi DID output will be necessary to execute the hwmgr -redirect commands).

```
# hwmgr -view dev -cat disk
HWID: Device Name      Mfg  Model          Location
-----
27: /dev/disk/floppy0c  3.5in floppy  fdi0-unit-0
34: /dev/disk/dsk0c    DEC  RZ28   (C) DEC bus-0-targ-1-lun-0
35: /dev/disk/dsk1c    DEC  RZ26L (C) DEC bus-0-targ-2-lun-0
36: /dev/disk/cdrom0c  DEC  RRD44 (C) DEC bus-0-targ-5-lun-0
39: /dev/disk/dsk4c    DEC  HSZ40 (C) DEC bus-1-targ-6-lun-0
40: /dev/disk/dsk5c    DEC  HSZ40 (C) DEC bus-1-targ-6-lun-1
44: /dev/disk/dsk6c    DEC  HSZ40 (C) DEC bus-1-targ-3-lun-0
45: /dev/disk/dsk7c    DEC  HSZ40 (C) DEC bus-1-targ-3-lun-1
```

```
# hwmgr -show scsi

      SCSI          DEVICE DEVICE  DRIVER NUM DEVICE FIRST
HWID:DEVICE HOSTNAM TYPE   SUBTYPE OWNER  PATH  FILE   VALID
-----
34:  0      ajkitt  disk   none    0     1   dsk0 [0/1/0]
35:  1      ajkitt  disk   none    0     1   dsk1 [0/2/0]
36:  2      ajkitt  cdrom  none    0     1  cdrom0 [0/5/0]
37:  3      ajkitt  disk   none    0     1   (null)
38:  4      ajkitt  disk   none    0     1   (null)
39:  5      ajkitt  disk   none    2     1   dsk4 [1/6/0]
40:  6      ajkitt  disk   none    0     1   dsk5 [1/6/1]
44:  7      ajkitt  disk   none    0     1   dsk6 [1/3/0]
45:  8      ajkitt  disk   none    0     1   dsk7 [1/3/1]
```

Note the following correspondence, and recall our intentions:

HWID = SCSI DID

37 = 3

38 = 4

44 = 7

45 = 8

REDIRECT

HWID SCSI DID

37 to 44 3 to 7

38 to 45 4 to 8

e. The redirection is accomplished by the following hwmgr commands:

```
# hwmgr -redirect scsi -src 3 -dest 7
hwmgr: Redirect operation was successful
```

```
# hwmgr -redirect scsi -src 4 -dest 8
hwmgr: Redirect operation was successful
```

Final result, and proof that all devices are reachable:

```
# hwmgr -view devices
HWID: Device Name      Mfg  Model          Location
-----
 4: /dev/kevm
27: /dev/disk/floppy0c  3.5in floppy  fdi0-unit-0
34: /dev/disk/dsk0c    DEC  RZ28  (C) DEC bus-0-targ-1-lun-0
35: /dev/disk/dsk1c    DEC  RZ26L (C) DEC bus-0-targ-2-lun-0
36: /dev/disk/cdrom0c  DEC  RRD44 (C) DEC bus-0-targ-5-lun-0
37: /dev/disk/dsk2c    DEC  HSZ40 (C) DEC bus-1-targ-3-lun-0
38: /dev/disk/dsk3c    DEC  HSZ40 (C) DEC bus-1-targ-3-lun-1
39: /dev/disk/dsk4c    DEC  HSZ40 (C) DEC bus-1-targ-6-lun-0
40: /dev/disk/dsk5c    DEC  HSZ40 (C) DEC bus-1-targ-6-lun-1
42: /dev/cport/scpl    HSZ40 (C) DEC bus-1-targ-3-lun-0
```

```
# hwmgr -show scsi
```

SCSI HWID:	SCSI DEVICE	HOSTNAM	DEVICE TYPE	DEVICE SUBTYPE	DRIVER OWNER	NUM PATH	DEVICE FILE	FIRST VALID
34:	0	ajkitt	disk	none	0	1	dsk0	[0/1/0]
35:	1	ajkitt	disk	none	0	1	dsk1	[0/2/0]
36:	2	ajkitt	cdrom	none	0	1	cdrom0	[0/5/0]
37:	3	ajkitt	disk	none	0	1	dsk2	[1/3/0]
38:	4	ajkitt	disk	none	0	1	dsk3	[1/3/1]
39:	5	ajkitt	disk	none	2	1	dsk4	[1/6/0]
40:	6	ajkitt	disk	none	0	1	dsk5	[1/6/1]

```
# mount /usr
```

```
# file /dev/rdisk/dsk*c
```

```
/dev/rdisk/dsk0c:      character special (19/22) SCSI #0  
"RZ28" disk #0
```

```
(SCSI ID #1) (SCSI LUN #0)
```

```
/dev/rdisk/dsk1c:      character special (19/38) SCSI #0  
"RZ26L" disk #1
```

```
(SCSI ID #2) (SCSI LUN #0)
```

```
/dev/rdisk/dsk2c:      character special (19/70) SCSI #1  
"HSZ40" disk #3
```

```
(SCSI ID #3) (SCSI LUN #0)
```

```
/dev/rdisk/dsk3c:      character special (19/86) SCSI #1  
"HSZ40" disk #4
```

```
(SCSI ID #3) (SCSI LUN #1)
```

```
/dev/rdisk/dsk4c:      character special (19/102) SCSI #1  
"HSZ40" disk #5
```

```
(SCSI ID #6) (SCSI LUN #0)
```

```
/dev/rdisk/dsk5c:      character special (19/118) SCSI #1  
"HSZ40" disk #6
```

```
(SCSI ID #6) (SCSI LUN #1)
```

- f. The system is now ready for multi-user mode:

```
# ^D
INIT: New run level: 3
.
.
.
<<end of example>>
```

14. The controller software upgrade to HSOF V3.7 is now complete.

9.2.2.2 Dual-Redundant Controller Shutdown Upgrade

If your HSOF Software version requires a Dual-Redundant Controller Shutdown Upgrade, follow the below procedures.

NOTE

Compaq does not recommend the use of the INITIALIZE *SAVE_CONFIGURATION* option for dual-redundant configurations.

CAUTION

Prior to installing HSOF V3.7 software, it is important to first determine whether the two controllers in your dual-redundant configurations have the same or different product ID's, this will determine the upgrade procedure you will use to continue.

To view the product ID's on your dual-redundant controllers, establish a local connection to the controller through the maintenance port of one of the controllers, (which becomes "this controller"). Enter the following commands and observe the resulting displays:

SHOW THIS_CONTROLLER and SHOW OTHER_CONTROLLER

The product ID is shown in the first characters of the display, prior to the hardware serial number. Most controllers have the product ID of "HSZ40" but a very few controllers may have the product ID of "HSZ40 (C) DEC."

- IF the Product ID's of the two HSZ40 controllers in a dual-redundant pair are the same, use the procedure for *Same Product ID's for Dual-Redundant Rolling Upgrade* (Section 9.2.2.2.1).
- IF the Product ID's of the two HSZ40 controllers in a dual-redundant pair are different, and IF your current Compaq Tru64 UNIX operating system version is V4.0x or lower, use the procedure for *Different Product ID's and Compaq Tru64 UNIX V4.0x for Dual-Redundant Rolling Upgrade* (Section 9.2.2.2.2).
- IF the Product ID's of the two HSZ40 controllers in a dual-redundant pair are **different, and** IF your current Compaq Tru64 UNIX operating system version is **V5.0**, use the procedure below for *Different Product ID's and Compaq Tru64 UNIX V5.0 for Dual-Redundant Rolling Upgrade* (Section 9.2.2.2.3).

9.2.2.2.1 Same Product ID's for Dual-Redundant Shutdown Upgrade

1. Halt all host I/O activity to the controller and prepare the host system for a controller shutdown using the appropriate procedures for your operating system.
2. Establish a local connection to the controller through the maintenance port of the controller.
3. Enter the SET NOFAILOVER command to take the controllers out of dual-redundant (failover) mode.

\$ SET NOFAILOVER

4. Shutdown the controller with the following command:

SHUTDOWN THIS_CONTROLLER

CAUTION

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

CAUTION

Before continuing, wait until the Reset button on the controller stops flashing (lit continuously). This indicates that unwritten data is flushed from the cache module.

5. Remove the PCMCIA program card ESD cover from the controller.
6. While pressing and holding the controller Reset button, press the program card eject button to eject the program card from the controller. Release the button after the card is ejected.

NOTE

The controller will initialize and the Reset button will flash at an approximate once per second rate when ready for operation.

7. Press and hold the Reset button while inserting the new program card. Release the button after the new card is inserted.

NOTE

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.

NOTE

The controller will initialize and the Reset button will flash at an approximate once per second rate when ready for operation.

8. Replace the ESD cover on the controller, pushing the locking pins to secure them in place.
9. Mount the storage units on the host. Resume I/O activity to the controller.
10. Repeat 5 through 9 for the second controller.
11. 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
```

9.2.2.2.2 Different Product ID's and Compaq Tru64 UNIX V4.0x for Dual-Redundant Shutdown Upgrade

1. Halt all host I/O activity to the controller and prepare the host system for a controller shutdown using the appropriate procedures for your operating system.
2. Establish a local connection to the controller through the maintenance port of the controller.
3. Enter the SET NOFAILOVER command to take the controllers out of dual-redundant (failover) mode.

```
$ SET NOFAILOVER
```

4. Shutdown the controller with the following command:

```
SHUTDOWN THIS_CONTROLLER
```

CAUTION

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

CAUTION

Before continuing, wait until the Reset button on the controller stops flashing (lit continuously). This indicates that unwritten data is flushed from the cache module.

5. Remove the PCMCIA program card ESD cover from the controller.
6. While pressing and holding the controller Reset button, press the program card eject button to eject the program card from the controller. Release the button after the card is ejected.

NOTE

The controller will initialize and the Reset button will flash at an approximate once per second rate when ready for operation.

7. Press and hold the Reset button while inserting the new program card. Release the button after the new card is inserted.

NOTE

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.

NOTE

The controller will initialize and the Reset button will flash at an approximate once per second rate when ready for operation.

8. Replace the ESD cover on the controller, pushing the locking pins to secure them in place.
9. Mount the storage units on the host. Resume I/O activity to the controller.
10. Repeat steps 5 through 9 for the second controller.
11. 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

9.2.2.2.3 Different Product ID's and Compaq Tru64 UNIX V5.0x for Dual-Redundant Shutdown Upgrade

When your HSZ40 controllers with different Product ID's are configured in a system using Compaq Tru64 UNIX V5.0, it is necessary to perform some steps involving the operating system before and after controller upgrade to HSOF V3.7.

1. Determine the **master** controller by following the sub-steps below:
 - a. Use the following command to view details of the devices attached to the host.

```
# hwmgr -view dev -cat disk
HWID: Device Name      Mfg  Model          Location
-----
27: /dev/disk/floppy0c  3.5in floppy  fdi0-unit-0
34: /dev/disk/dsk0c    DEC  RZ28 (C) DEC  bus-0-targ-1-lun-0
35: /dev/disk/dsk1c    DEC  RZ26L (C) DEC  bus-0-targ-2-lun-0
36: /dev/disk/cdrom0c  DEC  RRD44 (C) DEC  bus-0-targ-5-lun-0
37: /dev/disk/dsk2c    DEC  HSZ40          bus-1-targ-3-lun-0
38: /dev/disk/dsk3c    DEC  HSZ40          bus-1-targ-3-lun-1
39: /dev/disk/dsk4c    DEC  HSZ40 (C) DEC  bus-1-targ-6-lun-0
40: /dev/disk/dsk5c    DEC  HSZ40 (C) DEC  bus-1-targ-6-lun-1
```

In the above example dsk2 (D300 on the controller) and dsk3 (D301 on the controller) are reporting the Product ID of "HSZ40" while dsk4 (D600 on the controller) and dsk5 (D601 on the controller) are reporting the Product ID of "HSZ40 (C) DEC" in the Model column.

NOTE

If all the units configured on the dual controller pair are “ONLINE” to only one of the controllers (i.e. the other controller is in stand by mode), then all the devices will show only one Product ID. The controller which has the units “ONLINE” is the **master** controller. Go to Step 2 and perform Steps 2 through 12. It is not necessary in this case to perform Step 13.

- b. Use the command below to determine which device the root file system is mounted on.

```
# df
Filesystem      512-blocks    Used    Avail Capacity  Mounted on
/dev/disk/dsk4a  338542    276436    28250    91%    /
/dev/disk/dsk4g  3389096    520056  2530130    18%    /usr
/proc            0          0         0    100%    /proc
```

In the example above the root file system is mounted on dsk4.

- c. Shut the system down to single-user mode. This will prevent background processes from touching the disks while this procedure is in progress.

```
# shutdown now
```

- d. If the root file system is not mounted on a unit configured on the controller pair, choose **either** of the controllers to be the “**master**” controller in the upgrade operation below. Go to step 2.

If the root file system is mounted on a unit configured on the controller pair, then do the following in order to determine which is the **master** controller.

- e. Connect the CLI interface to one of the controllers and perform **SHOW <UNIT>** command for the unit with the root device (D600 in case of the example).

```
HSZ> show d600
      LUN                               Uses
-----
D600                                DISK600
...
      State:
      ONLINE to this controller
      Not reserved
      PREFERRED_PATH = THIS_CONTROLLER
...
```

If the output shows that the unit is “ONLINE” to **this** controller (as is the case in this example), then **THIS_CONTROLLER** is the **master** controller.

If the output shows that the unit is “ONLINE” to the **other** controller, then **OTHER_CONTROLLER** is the **master** controller.

2. Halt all host I/O activity to the controller and prepare the host system for a controller shutdown using the appropriate procedures for your operating system.
3. Establish a local connection to the controller through the maintenance port of the controller.
4. Enter the SET NOFAILOVER command to take the controllers out of dual-redundant (failover) mode.
\$ SET NOFAILOVER
5. Shutdown the controller with the following command:
SHUTDOWN THIS_CONTROLLER

CAUTION

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

CAUTION

Before continuing, wait until the Reset button on the controller stops flashing (lit continuously). This indicates that unwritten data is flushed from the cache module.

6. Remove the PCMCIA program card ESD cover from the controller.

7. While pressing and holding the controller Reset button, press the program card eject button to eject the program card from the controller. Release the button after the card is ejected.

NOTE

The controller will initialize and the Reset button will flash at an approximate once per second rate when ready for operation.

8. Press and hold the Reset button while inserting the new program card. Release the button after the new card is inserted.

NOTE

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.

NOTE

The controller will initialize and the Reset button will flash at an approximate once per second rate when ready for operation.

9. Replace the ESD cover on the controller, pushing the locking pins to secure them in place.
10. Mount the storage units on the host. Resume I/O activity to the controller.
11. Repeat steps 5 through 9 for the second controller.
12. After you complete all of the above steps on both controllers, select the “Master” as the CONFIGURATION_SOURCE and put the controllers back into dual-redundant (failover) mode by entering the following command:
\$ SET FAILOVER COPY=CONFIGURATION_SOURCE

13. Follow the commands in the example below to correct the device files.
 - a. Scan the SCSI interfaces to update the operating system's view of device changes and then create new devices as necessary.

NOTE

Depending on the size of your configuration, the scan may take several minutes to complete. The presence of tape devices will further increase the delay to complete the scan. For this reason, you may wish to use the -bus qualifier to specify the bus you want to scan. The correct bus number can be determined by examining the "location" field of the "hwmgr -view devices" output.

```
# hwmgr -view devices
HWID: Device Name   Mfg  Model           Location
-----
39: /dev/disk/dsk4c DEC  HSZ40  (C) DEC bus-1-targ-6-lun-0
```

Issue the command for the appropriate bus (1 in this example).

```
# hwmgr -scan scsi -bus 1
```

- b. In order to understand which devices need to be redirected to the newly created devices, examine the current list of devices prior to rebooting the system:

```
# hwmgr -view devices
HWID: Device Name      Mfg  Model          Location
-----
4: /dev/kevm
27: /dev/disk/floppy0c  3.5in floppy  fdi0-unit-0
34: /dev/disk/dsk0c    DEC  RZ28 (C) DEC  bus-0-targ-1-lun-0
35: /dev/disk/dsk1c    DEC  RZ26L (C) DEC  bus-0-targ-2-lun-0
36: /dev/disk/cdrom0c  DEC  RRD44 (C) DEC  bus-0-targ-5-lun-0
37: /dev/disk/dsk2c    DEC  HSZ40          bus-1-targ-3-lun-0
38: /dev/disk/dsk3c    DEC  HSZ40          bus-1-targ-3-lun-1
39: /dev/disk/dsk4c    DEC  HSZ40 (C) DEC  bus-1-targ-6-lun-0
40: /dev/disk/dsk5c    DEC  HSZ40 (C) DEC  bus-1-targ-6-lun-1
41: /dev/cport/scp0    HSZ40          bus-1-targ-3-lun-0
42: /dev/cport/scp1    HSZ40 (C) DEC  bus-1-targ-6-lun-0
44: /dev/disk/dsk6c    DEC  HSZ40 (C) DEC  bus-1-targ-3-lun-0
45: /dev/disk/dsk7c    DEC  HSZ40 (C) DEC  bus-1-targ-3-lun-1
```

Note that you can ignore the control port "scp0" device.

We will need to redirect the following device names (with invalid PID's):

```
37: /dev/disk/dsk2c    DEC  HSZ40          bus-1-targ-3-lun-0
38: /dev/disk/dsk3c    DEC  HSZ40          bus-1-targ-3-lun-1
```

To the following new device names (containing valid PID's):

```
44: /dev/disk/dsk6c    DEC  HSZ40 (C) DEC  bus-1-targ-3-lun-0
45: /dev/disk/dsk7c    DEC  HSZ40 (C) DEC  bus-1-targ-3-lun-1
```

Remember the following:

- HWID 37 will be redirected to 44
 - HWID 38 will be redirected to 45
- c. At this point, you should reboot the system (ONLY TO SINGLE-USER MODE).
Then mount the root file system to enable writing to the disk:

```
# shutdown -h now
.
.
.
P00>>> boot -flag s dkb600
.
.
.
Starting secondary cpu 1
INIT: SINGLE-USER MODE
# mountroot
.
.
.
```

- d. Examine the output of "hwmgr -view devices" and "hwmgr -show scsi" (The scsi DID output will be necessary to execute the hwmgr -redirect commands).

```
# hwmgr -view dev -cat disk
HWID: Device Name      Mfg  Model      Location
-----
27: /dev/disk/floppy0c  3.5in floppy fdi0-unit-0
34: /dev/disk/dsk0c    DEC  RZ28  (C) DEC bus-0-targ-1-lun-0
35: /dev/disk/dsk1c    DEC  RZ26L (C) DEC bus-0-targ-2-lun-0
36: /dev/disk/cdrom0c  DEC  RRD44 (C) DEC bus-0-targ-5-lun-0
39: /dev/disk/dsk4c    DEC  HSZ40 (C) DEC bus-1-targ-6-lun-0
40: /dev/disk/dsk5c    DEC  HSZ40 (C) DEC bus-1-targ-6-lun-1
44: /dev/disk/dsk6c    DEC  HSZ40 (C) DEC bus-1-targ-3-lun-0
45: /dev/disk/dsk7c    DEC  HSZ40 (C) DEC bus-1-targ-3-lun-1
```

```
# hwmgr -show scsi
      SCSI          DEVICE DEVICE  DRIVER NUM DEVICE FIRST
HWID:DEVICE HOSTNAM TYPE   SUBTYPE OWNER PATH FILE  VALID
-----
34:  0      ajkitt  disk   none    0     1   dsk0 [0/1/0]
35:  1      ajkitt  disk   none    0     1   dsk1 [0/2/0]
36:  2      ajkitt  cdrom  none    0     1  cdrom0 [0/5/0]
37:  3      ajkitt  disk   none    0     1   (null)
38:  4      ajkitt  disk   none    0     1   (null)
39:  5      ajkitt  disk   none    2     1   dsk4 [1/6/0]
40:  6      ajkitt  disk   none    0     1   dsk5 [1/6/1]
44:  7      ajkitt  disk   none    0     1   dsk6 [1/3/0]
45:  8      ajkitt  disk   none    0     1   dsk7 [1/3/1]
```

Note the following correspondence, and recall our intentions:

HWID = SCSI DID

37 = 3

38 = 4

44 = 7

45 = 8

REDIRECT

HWID SCSI DID

37 to 44 3 to 7

38 to 45 4 to 8

e. The redirection is accomplished by the following hwmgr commands:

```
# hwmgr -redirect scsi -src 3 -dest 7
```

```
hwmgr: Redirect operation was successful
```

```
# hwmgr -redirect scsi -src 4 -dest 8
```

```
hwmgr: Redirect operation was successful
```


Final result, and proof that all devices are reachable:

```
# hwmgr -view devices
```

```
HWID: Device Name      Mfg  Model          Location
-----
4: /dev/kevm
27: /dev/disk/floppy0c  3.5in floppy  fdi0-unit-0
34: /dev/disk/dsk0c    DEC  RZ28 (C) DEC bus-0-targ-1-lun-0
35: /dev/disk/dsk1c    DEC  RZ26L (C) DEC bus-0-targ-2-lun-0
36: /dev/disk/cdrom0c  DEC  RRD44 (C) DEC bus-0-targ-5-lun-0
37: /dev/disk/dsk2c    DEC  HSZ40 (C) DEC bus-1-targ-3-lun-0
38: /dev/disk/dsk3c    DEC  HSZ40 (C) DEC bus-1-targ-3-lun-1
39: /dev/disk/dsk4c    DEC  HSZ40 (C) DEC bus-1-targ-6-lun-0
40: /dev/disk/dsk5c    DEC  HSZ40 (C) DEC bus-1-targ-6-lun-1
42: /dev/cport/scpl    HSZ40 (C) DEC bus-1-targ-3-lun-0
```

```
# hwmgr -show scsi
```

```
SCSI          DEVICE DEVICE DRIVER NUM DEVICE FIRST
HWID: DEVICE HOSTNAM TYPE  SUBTYPE OWNER PATH FILE  VALID
-----
34:  0      ajkitt  disk   none    0    1   dsk0 [0/1/0]
35:  1      ajkitt  disk   none    0    1   dsk1 [0/2/0]
36:  2      ajkitt  cdrom  none    0    1  cdrom0 [0/5/0]
37:  3      ajkitt  disk   none    0    1   dsk2 [1/3/0]
38:  4      ajkitt  disk   none    0    1   dsk3 [1/3/1]
39:  5      ajkitt  disk   none    2    1   dsk4 [1/6/0]
40:  6      ajkitt  disk   none    0    1   dsk5 [1/6/1]
```

```
# mount /usr
# file /dev/rdisk/dsk*c
/dev/rdisk/dsk0c:      character special (19/22) SCSI #0
"RZ28" disk #0
(SCSI ID #1) (SCSI LUN #0)
/dev/rdisk/dsk1c:      character special (19/38) SCSI #0
"RZ26L" disk #1
(SCSI ID #2) (SCSI LUN #0)
/dev/rdisk/dsk2c:      character special (19/70) SCSI #1
"HSZ40" disk #3
(SCSI ID #3) (SCSI LUN #0)
/dev/rdisk/dsk3c:      character special (19/86) SCSI #1
"HSZ40" disk #4
(SCSI ID #3) (SCSI LUN #1)
/dev/rdisk/dsk4c:      character special (19/102) SCSI #1
"HSZ40" disk #5
(SCSI ID #6) (SCSI LUN #0)
/dev/rdisk/dsk5c:      character special (19/118) SCSI #1
"HSZ40" disk #6
(SCSI ID #6) (SCSI LUN #1)
```

f. The system is now ready for multi-user mode:

```
# ^D
INIT: New run level: 3
.
.
.
<<end of example>>
```

14. The controller software upgrade is now complete. Mount the storage units on the host.

10.0 Possible Problem Identified

Read this section if you initialized a JBOD disk using the *SAVE_CONFIGURATION* switch on an HSZ40 controller running HSOF Version 2.7, and then moved the disk to your HSZ40 controller subsystem.

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 INITIALIZE *SAVE_CONFIGURATION* command **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 the *SAVE_CONFIGURATION* switch and the platform operating system under HSOF Versions 30Z, and above AND 50Z, and above.

If a disk built as described above is moved to an HSZ40 controller running HSOF Version 5.0 and above, problems could occur.

When a disk being used in a JBOD configuration is initialized with the *SAVE_CONFIGURATION* switch, 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 INITIALIZE *SAVE_CONFIGURATION* 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 the INITIALIZE *container-name SAVE_CONFIGURATION* command tries to restore the configuration (unless another drive contains valid configuration data).

Various configuration events cause the controller to write the configuration data to save configuration 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 configuration 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.

10.1 How to Detect Whether You Have This Problem

The problem described in this appendix affects single-disks units in JBOD configurations with save configuration data stored on them. If you have not moved JBOD disk units with save configuration data saved on them to the HSZ40 controller being upgraded to V3.7, 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:

1. At the controller prompt, type SHOW DISKnnn (where nnn is the JBOD disk in question).
2. Look for a "Configuration being backed up on this container" message.
3. Record the block size capacity displayed by the controller.
4. From the Compaq Tru64 UNIX prompt 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 HSZ20
```

```
# disklabel -r /dev/rrza18c
```

```
# /dev/rrza18c:
```

5. Compare the total block size obtained from the "show device" command with the block size capacity obtained in step 3.
6. 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.

10.2 Solution

If you are using the `INITIALIZE SAVE_CONFIGURATION` command to initialize JBOD disks under HSOE Version 27Z, be sure to initialize the disk with the platform file system before rebooting the controller.

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.

Compaq recommends that the recovery process described here be performed before moving the disk to a controller running Version 3.7. Any files written in the save configuration 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.

1. Back up the unit that contains the save `SAVE_CONFIG` information.
2. Unmount the file system(s) contained on that unit.
3. Delete the unit from the configuration in the controller.
4. Initialize the container from the controller without `SAVE_CONFIG`.
5. Add the unit back into the configuration.
6. Initialize and restore unit from backup.

11.0 Order Numbers

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

Table 11-1 Order Numbers

Option	Description
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	16MB 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