

# *Compaq StorageWorks™*

## *Release Notes*

### **HSD30 Array Controller Operating Software HSOF Version 3.7**

---

*This document summarizes features and characteristics of the HSD30 array controller operating software 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 – HSD30 Array Controller Operating Software HSOF Version 3.7  
Tenth Edition (May 2000)  
EK-HSD30-RN. K01

## 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 HSOE 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 Part Numbers
- 11.0 Order Numbers
- 12.0 Replacing or Adding an HSD Controller or Cache Module
- 13.0 The HSUTIL Utility

## Release Package Contents

The Version 3.7 release package consists of the following:

- A cover letter
- A PCMCIA program card containing HSOE Version 3.7 software
- The HS Family of Array Controllers documentation set
  - *HS Family of Array Controllers User's Guide*
  - *HS Family of Array Controllers Service Manual*
  - *HS Family of Array Controllers CLI Reference Manual*
  - *HSD30 Array Controller Operating Software Version 3.7 Release Notes*
  - *HSD30 Array Controller Operating Software (HSOE), Version 3.7 Software Product Description*
- License keys to enable optional features: write-back cache, disk mirroring, and RAID (each to be purchased separately)

## Intended Audience

These release notes provide information for the HSD30 Array Controller running HSOE Software Version 3.7. Individuals responsible for configuring, installing, and using the HSD30 controller should use this document.

Read this entire document before upgrading the controller's software. The information in these release notes takes precedence over the information in the HS Family of Array Controllers User's Guide.

## 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
- Device Support

### 2.1 Hardware Support

HSOF Software Version 3.7 supports the following revisions or higher for the RAID Array HSD30 Controller and associated hardware:

- HSD30-Ax and HSD30-Cx controller modules
- Cache module Version 1, hardware revision A and Version 2, hardware revision A or B
- BA350-MA controller shelf
- BA350-Sx 8-bit SCSI device shelf
- BA356-Sx wide device shelf with 8-bit or 16-bit personality module (controller operates in 8-bit mode only)
- BA35x-HF power supply

### 2.2 Operating System Support

HSOF Version 3.7 software on HSD30 controllers is supported by the following operating system versions:

- OpenVMS™ Alpha:  
V6.2-1H3, V7.1-1H1, V7.1-1H 2, V7.1-1H 3, V7.1-2, V7.2, and V7.2-1
- OpenVMS VAX:  
V6.2, V7.1, and V7.2

#### NOTE

With the release of HSOF Version 3.7, the StorageWorks Command Console V2.3 Agent is used. An update of the PKQdriver is required for use with this Agent. Please view the following Section, 2.4 “*PKQdriver Version*” for PKQdriver update information.

## 2.3 Host Adapter Support

The following adapters have been tested and are supported in this release.

- KFMSA for XMI on VAX
- KFMSB for XMI on Alpha 7000 series
- KFPSA for PCI on newer Alphas
- KFESA for EISA bus for Alphas
- KFESB for EISA bus for Alphas

## 2.4 PKQdriver Version

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

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 the 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 the Table 2-1 then your driver will need to be updated.

If additional support with PKQdriver installation is needed, please contact your Customer Service Representative.

## 2.5 Device Support

HSOF Software V3.7 supports the 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 <sup>1</sup>	1.05	440C	A01
RZ26N-VA/VW <sup>1</sup>	1.05	446	A01
SWXD3-SF/WF <sup>1</sup>	1.05	446	A01
DS-RZ26N-VZ <sup>1</sup>	1.05	1003	A01
DS-RZ1BB-VW	2.10	LYJO/0656	A01
RZ28-VA/VW <sup>1</sup>	2.10	435E	B03
RZ28B-VA	2.10	0003	A01
RZ28D-VA/VW <sup>1</sup>	2.10	0008	A01
SWXD3-SG/WG <sup>1</sup>	2.10	0008	A01
RZ28M-VA/VW <sup>1</sup>	2.10	0466	A01
DS-RZ28M-VZ <sup>1</sup>	2.10	1003	A01
SWXD3-SH/WH <sup>1</sup>	2.10	0466	A01
RZ74-VA	3.57	T427B	B07
DS-RZ1CB-VW	4.1	LYJO/0656	A01
RZ29B-VA/VW <sup>1</sup>	4.3	0007	B01
SWXD3-SE/WE <sup>1</sup>	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**

<b>Device</b>	<b>Capacity G= Gigabytes T=Terabytes</b>	<b>Minimum Microcode Version</b>	<b>Minimum Hardware Version</b>	<b>Notes</b>
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
DS-TZ820	25/50G	01aj	A01	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
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	V80	A01	3,4,5,7
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 Giga-bytes	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
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

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**

<b>Device</b>	<b>Capacity in Giga-bytes</b>	<b>Minimum Microcode Version</b>	<b>Minimum Hardware Version</b>
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 and 3.2 briefly describe the new items in HSOF Software V3.7.

- New Devices Supported in HSOF Software
- Issue Resolutions

## 3.1 New Devices Supported in HSOF V3.7 Software

### 3.1.1 Disk devices

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

### 3.1.2 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.

## 3.2 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 controller reset when running backup with compare during the reading and writing of EBCDIC formatted tapes on the TSZ07 tape device.
- 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.
- 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.
- A rare occurrence in which RAIDsets were not properly failing over in the event of battery failure.
- In a rare occurrence the HSD30 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 4.4 briefly describe changes that were introduced in the code version previous to HSOF V3.7 Software release that are not covered in other documentation.

- Correction for TILX Incompatibility with Disks
- Action for repeated Reboots Problem Resolution
- Support for Large Drive Formatting
- Support for up to Fourteen 18GB Drives

### 4.1 Correction for TILX Incompatibility with Disks

In previous versions of HSOF software, data integrity problems could infrequently occur if the tape in-line exerciser (TILX) was run on tapes attached to a controller on which disks were also configured. This problem is corrected in HSOF Version 3.4.

### 4.2 Action for Repeated Reboots

Very rarely, a hardware failure might cause a controller to reboot repeatedly, preventing preferred disk devices from completing failover to the dual-redundant partner. 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.

### 4.3 Support for Large-Drive Formatting

This release supports formatting of larger disk devices (9 GB and 18 GB).

## 4.4 Support for up to Fourteen 18 GB Drives

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

## 5.0 Clarifications

Sections 5.1 through 5.14 contain clarification on the following subjects:

- EMU Alarm Clarification
- Write-Back Cache Batteries and Battery Handling
- Logical Device Event Reports
- SHOW on Both Controllers
- Logical Block Address Does Not Match for MEDIUM ERROR
- Adding RAIDsets When Battery Condition is Low
- Drive-Level Event Reports
- Saving the Subsystem Configuration
- RRD42 CD-ROM Reader
- CLI Type Ahead, Recall, and Editing
- Faster Storageset Initialization
- Multiple Disk Drive Formatting
- Ability to Install New Microcode on a Tape Drive and Solid State Disk Drives
- Last Failure Code 010E0110 Normal Occurrence After Upgrade

### 5.1 EMU Alarm Clarification

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 an 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 last 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 apply, 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 HSD30s (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 HSD30s
  - 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 HSD30 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 one 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 HSD30 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.

- 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.

- 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 *HSFAM of Array Controllers Service Manual*, for instructions on removing and replacing the controller module and write-back cache module.

- 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.
- Connect a maintenance terminal to the controller's maintenance terminal port on the front bezel.
- 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
```

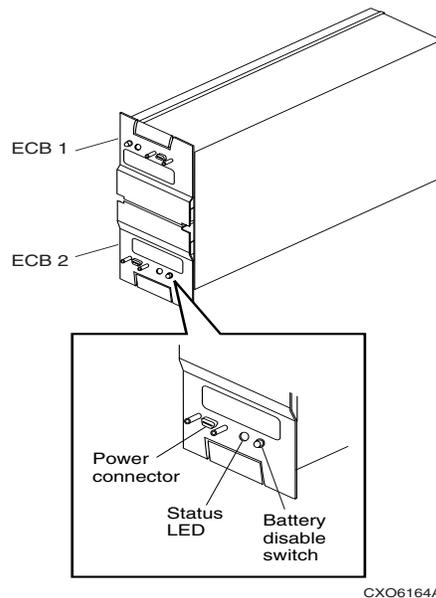
When the shut down completes, the controller or controllers' green reset LED on the operator control panel (OCP) stops blinking.

- Power off the entire storage subsystem.
- 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 two seconds (refer to Figure 5-1). The cache LED will flash once, then shut off.

#### NOTE

The battery is no longer powering the cache.

- To return to normal operation, the cache battery is enabled when the RAID subsystem is powered on.
- 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**

### 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 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.

## 5.4 SHOW on Both Controllers

Some devices, such as tape devices, may not show correctly on one controller if they are preferred to the other controller.

If the expected information is not displayed from one controller, use the `SHOW` command again to view the information on the other controller.

## 5.5 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) displays 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.

## 5.6 Adding RAIDsets When Battery Condition is Low

You can use the `add raidset` and `add mirrorset` commands, regardless of the write-back cache battery condition. However, if the `cache_policy` is set to A and the batteries are low, the controller does not allow access to any RAIDsets or mirrorsets. `Cache_policy A` requires that the cache batteries be fully charged before you can use RAIDsets or mirrorsets.

## 5.7 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.

This is a permanent restriction.

## 5.8 Saving the Subsystem Configuration

The `save_configuration` switch is intended to be used to provide a means to restore nonvolatile (NVRAM) contents when a controller in a single-controller module configuration is replaced. The `save_configuration` switch cannot cross hardware revisions or software versions. `save_configuration` is not available for upgrades of firmware or hardware, and does not perform inter-platform conversions. For example, do not use `save_configuration` to upgrade from HSOF Version 3.2, or from an HSD30 to an HSD50 array controller.

The controller stores the subsystem's configuration—the storagesets you configured, the switches set for each, the HSOF patches you've installed, and so on—in its nonvolatile memory. This means that, if the controller fails in a nonredundant configuration, you would have to reconfigure the subsystem from scratch. (In a dual-redundant configuration, this information is stored by both controllers, which effectively eliminates the risk of losing it.)

### NOTE

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

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

```
INITIALIZE STORAGESET_NAME SAVE_CONFIGURATION
```

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

After initializing a storageset with this switch, the controller keeps the copy up-to-date. For example, every time a new patch or storageset is added to the HSOF software, the controller re-copies the new configuration to all storagesets that were initialized with the `save_configuration` switch.

If a controller in a nonredundant-configuration fails, its replacement automatically searches the devices in the subsystem for a saved configuration. If it finds a saved configuration, it transparently loads it into nonvolatile memory and brings the subsystem online. It is unnecessary to issue an initialize command again after reconfiguring the devices with a new controller.

Compaq does not recommend initializing all of your storagesets with the `save_configuration` switch. Every time you change the subsystem's configuration, the controller writes the new configuration to the storagesets that you initialized with this switch. Too many write operations can adversely affect performance.

If you are upgrading from HSOF Version 3.2 to Version 3.4 in a single-controller configuration, you can refresh the `SAVE_CONFIGURATION` data on your storagesets to include the new software revision by issuing the following command:

```
SET UNIT XXX WRITE_PROTECT
```

[where xxx is any existing unit on the controller]

```
SET UNIT XXX NOWRITE_PROTECT
```

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

## 5.9 RRD42 CD-ROM Reader

Mounting an RRD42 might put a message in the OpenVMS error log. In earlier versions of OpenVMS, the first four digits in Longword 19 of the controller dependent information is 02D1. New versions of OpenVMS report `ASC: 209.(D1) ASCQ: 2.(02)`. The message indicates "Unexpected bus phase," and can be ignored safely.

## 5.10 CLI Type Ahead, Recall, and Editing

With HSOF Version 3.1 or higher, you no longer need to wait for the current command to complete before entering the next command. You can recall any of the last four commands entered, and then edit that command to save time when entering new commands.

Refer to *HS Family of Array Controllers User's Guide* for more information on entering CLI commands.

## 5.11 Faster Storageset Initialization

The initialize command has been optimized in Version 3.1, and runs significantly faster than on previous versions of HSOF software. It takes less than 30 seconds to initialize a large RAIDset.

## 5.12 Multiple Disk Drive Formatting

You can perform a low-level format on up to seven disk drives at once with HSOF Version 3.1 software. See the "Formatting a disk using HSUTIL" section 13.1.2, -in these release notes for instructions on how to format disk drives.

### 5.13 Ability to Install New Microcode on a Tape Drive and Solid State Disk Drives

HSOF Version 3.1 software allows you to use the controller's HSUTIL program to install new microcode to the TZ89 tape drive.

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

### 5.14 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: HSD30 Se-  
rial Number: ZG05051997 Hardware Version: 0000(00) Firmware Version: V37Z(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 6.23 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, 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:

- Updated Storageset Limits
- Hot Swapping Devices
- Maintenance Terminal Port Speeds
- Restrictions on Moving Storagesets to Previous Software Versions
- Initialization Devices
- OVERRIDE\_ONLINE Switch
- Unable to Access Tape Drives After Power Cycle
- CFMENU Constraints
- Device Adds, Moves, and Changes
- Tape Drive Firmware Revision
- Running CONFIG or CFMENU Utilities During Backup or Copy Operations
- Changing Host Adapter Types
- CLONE Utility
- Moving Write-Back Cache Modules
- Set Tape Density When Adding Drive and Changing Cartridge Types
- Set SCSI Speed for Devices
- Configuration Saved on Disk
- Host Tape Density Selection
- Disk Code Load Capability
- Disk Formatting Capability
- Disk Geometries of Mirrorsets and Clones
- New CLEAR\_ERRORS INVALID\_CACHE Functionality
- New Cache Battery Diagnostic

### 6.1 Updated Storageset Limits

The following limits apply to storagesets configured on single controller or dual-redundant controller configurations:

- 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 and/or RAIDsets.
- There can be a maximum of 30 storagesets (RAIDsets, mirrorsets, and/or stripesets).
- There can be no more than 32 physical device members for a unit.
- Each storageset can store a maximum of 256 GB.

**NOTE**

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.

## 6.2 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.

## 6.3 Maintenance Terminal Port Speeds

HSOF Version 3.4 software supports terminal port connections at 4800, 9600, and 19200 bits per second (bps). Connection speeds at 300, 1200, and 2400 bps are not supported.

## 6.4 Restrictions on Moving Stagesets to Previous Software Versions

In HSOF Version 2.5 software, striped mirrorsets reported a different disk geometry than the same nonmirrored stripeset. This has been resolved in Versions 2.7, 3.1, 3.2, and 3.4, such that striped mirrorsets initialized under these versions report a disk geometry identical to the same nonmirrored stripeset.

Due to this improvement, be aware of the following when moving stripesets to previous versions of software:

- Striped mirrorsets created under HSOF Software Versions 2.7, 3.1, or 3.2 cannot be used on previous versions; the metadata is incompatible.
- Any stripeset that is converted to a striped mirrorset using the mirror command under HSOF Software Versions 2.7, 3.1 or 3.2 cannot be used on previous versions. You must unmirror each of the members to return the stripeset to an earlier version.
- Any stripeset clone that is created under Versions 2.7, 3.1 or 3.2 (using the Clone utility) cannot be used with previous software versions.
- Containers initialized with the `SAVE_CONFIGURATION` switch introduced in HSOF Version 2.7 cannot be used with versions prior to Version 2.7.

Striped mirrorsets created under previous software versions and used with Versions 2.7, 3.1 or 3.2 carry forward the same geometry they had under the previous version. If possible, back up all data and reinitialize the stripesets under Versions 2.7, 3.1 or 3.2 to take full advantage of geometry improvements.

## 6.5 Initialization Devices

Controller-attached devices can serve as initialization devices for VAX 7000 and VAX 10000 systems, provided the VAX 7000/10000 console code is at Version 4.0 or higher.

## 6.6 `OVERRIDE_ONLINE` Switch

Attempting to delete tape-based units using the `OVERRIDE_ONLINE` switch results in conditions requiring controller reinitialization. Do not enter the `OVERRIDE_ONLINE` switch when deleting tape-based units.

## 6.7 Unable to Access Tape Drives after Power Cycle

You might not be able to access a tape drive after the power for the drive is turned off and back on. Use the `CLEAR_ERRORS UNKNOWN DEVICE-NAME` command to place the device in a known state.

## 6.8 CFMENU Constraints

CFMENU can only delete single sparesets. Delete sparesets one at a time.

## 6.9 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.

## 6.10 Tape Drive Firmware Revision Upgrade

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.

## 6.11 Running CONFIG or CFMENU Utilities During Backup Operations

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

## 6.12 Changing Host Adapter Types

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

## 6.13 CLONE Utility

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

## 6.14 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 HS Family of Array Controllers User's Guide or the HS Family of Array Controllers 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.

## 6.15 Set Tape Density When Adding Drive and Changing Cartridge Types

The `default_format` switch allows you to specify the tape format for use by the controller for write operations to a tape drive. Also use this switch to tell the controller to allow the host to set the tape format.

Not all tape devices support all formats. The easiest way to determine what formats are supported by a specific device is to enter the `set tape-unit-number default_format=?` command. The controller displays the valid options for the unit.

### NOTE

Host system software must support the density selection for the affected device for this option to operate correctly. Using the `HOST_SELECTED` switch for other devices may yield unpredictable results.

Compaq recommends the use of the `set tape-unit-number default_format=host_selected` command with only TSZ07 tape drives. Set the switch to `host_selected` to allow the host to set the tape format.

When first adding a tape drive to the controller configuration, set the tape density with the `set tape-unit-number default_format=` command before using the OpenVMS initialize command. In dual-redundant controller configurations, set the tape density from both controllers in the pair before the tape drive is mounted. Not setting the tape density might cause the tape to be written at the wrong density, or the OpenVMS initialize command to fail with an invalid format error.

## 6.16 Set SCSI Speed for Devices

HSOF Version 3.1 software makes it possible for the initiator and target to negotiate a faster data rate than is supported by the cable length. HSOF Version 3.1 software has a new CLI command to set a maximum data transfer rate between the controller and any device. Possible settings are 10MHz, 5MHz, or asynchronous mode.

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

```
SET DEVICE-NAME TRANSFER_RATE_REQUESTED=RATE
```

where `rate` can be one of the following:

- ASYNCHRONOUS
- 10MHZ
- 5MHZ

The default is 10MHZ.

Refer to HS Family of Array Controllers User's Guide for more information on setting the data transfer rate according to SCSI cable lengths.

## 6.17 Configuration Saved on Disk

A new CLI option causes the controllers to save a copy of the controller configuration (nonredundant controllers only) on one or more disk drives in the subsystem. The saved information is updated whenever the configuration of the nonredundant controller is changed. Should a nonredundant controller be replaced by another controller with no configuration, the saved configuration is restored from one of the devices to the controller.

The information saved includes the controller settings, device, storageset and unit configurations, and software patches. The section "Using SAVE\_CONFIGURATION to Save the Subsystem's Configuration" further describes this feature. Note that any container initialized with the SAVE\_CONFIGURATION option cannot be used on previous HSD30 array controller HSOF software versions. See the section entitled "Replacing a Controller or Cache Module with C\_SWAP" in these release notes for instructions on how to replace a failed controller.

## 6.18 Host Tape Density Selection

You can now set the density for tape drive writes from the host operating system, and from the controller CLI.

Currently, the TSZ07 is the only tape drive for which you can set tape density from OpenVMS through the HSD controller using `openvms init` and `openvms mount`. Do not use these commands to set tape density on other tape drives. Set tape density on other tape drives through programs that issue QIOs, as some customer applications do.

Inform the controller to allow the host to set tape density with the CLI command:

```
SET UNIT-NAME DEFAULT_FORMAT=HOST_SELECTED
```

When using dual-redundant controllers, tape format changes only appear on the controller that the device is mounted through. `SHOW UNIT` only shows the updated format on the controller that has the tape online.

## 6.19 Disk Code Load Capability

Upload disk firmware to compatible disks in the subsystem using the HSUTIL utility. For more information, see Section 13, "The HSUTIL Utility".

## 6.20 Disk Formatting Capability

Format compatible disks in the subsystem using the HSUTIL utility. For more information, see Section 13, "The HSUTIL Utility".

This section describes the features and user interface of the HSUTIL utility. HSUTIL provides two functions:

Device format-Enables you to perform a basic SCSI format operation on a disk device.

Device code load-Enables you to download device firmware into supported disk devices.

## 6.21 Disk Geometries of Mirrorsets and Clones

In software Version 2.5, mirrored devices reported different geometries than unmirrored devices of the same type. This also affected the clone utility, such that a clone of a unit may have reported a different size to the operating system than the original unit. Software Version 2.7 resolves this issue. In Version 2.7, all mirrorsets and clones report the same disk geometry/capacity as their unmirrored equivalents.

## 6.22 New CLEAR\_ERRORS\_INVALID\_CACHE Functionality

A controller presents an invalid cache error during initialization if it detects a mismatch between the cache information on the controller and the cache information on the cache module. The CLEAR\_ERRORS\_INVALID\_CACHE COMMAND has new switches to tell the controller whether to use the information in the controller module or the information in the cache module.

The new syntax of the command is:

```
CLEAR_ERRORS_INVALID_CACHE CONTROLLER DATA_RETENTION
```

In this command:

- Controller identifies the controller to clear the error on, either this\_controller or other\_controller
- data\_retention specifies whether to keep the cache data and overwrite the controller information (NODESTROY\_UNFLUSHED\_DATA) or keep the controller information and discard the cache data (DESTROY\_UNFLUSHED\_DATA).

Use the NODESTROY\_UNFLUSHED\_DATA parameter in the following situations:

- When you have replaced the controller module
  - The controller's nonvolatile memory (NVMEM) has lost its memory, possibly from an NVMEM battery failure
  - You removed the cache module while there was unflushed data
- For example:

```
CLEAR_ERRORS_INVALID_CACHE THIS_CONTROLLER
```

```
NODESTROY_UNFLUSHED_DATA
```

Use the DESTROY\_UNFLUSHED\_DATA PARAMETER in the following situations:

- When you have replaced the cache module
  - Any other reason not listed above.
- For example:

```
CLEAR_ERRORS_INVALID_CACHE OTHER_CONTROLLER
DESTROY_UNFLUSHED_DATA
```

## 6.23 New Cache Battery Diagnostic

Software Versions 2.7, 3.1, and 3.2 check the condition of the write-back cache batteries every 24 hours. If a battery failure is detected, all RAIDsets and mirrorsets become inoperative on the controller with the failed batteries. In dual-redundant controller configurations, the RAIDsets and mirrorsets will fail over to the other controller.

## 7.0 Avoiding Problem Situations

In certain situations, you may experience unusual controller behavior. Paragraphs 7.1 through 7.12 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
- Using the CLI and CF Menu Simultaneously in a Dual Configuration
- Write-Back Cache Battery Replacement
- Only One Switch per CLI Command Recommended
- UNMIRROR and Device Capacity
- Unit problems After Battery Low
- CLEAR\_ERRORS LOST\_DATA Command
- SET NOPATH Requires Controller Reset
- Device Extended Status
- SET FAILOVER Causes Target to Restart
- Error Log Codes
- Chunk Size

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

As noted in Section 3.2, in a rare occurrence the HSD30 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 HSD30 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 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.

## 7.3 Write-Back Cache (WBC) 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).

## 7.4 Only One Switch per CLI Command Recommended

Certain switches to CLI commands are incompatible. To avoid problems in this area, only specify 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
```

## 7.5 UNMIRROR and Device Capacity

The `unmirror` command is valid only for devices that have a capacity equal to the capacity of the mirrorset. If a mirrorset comprises devices of different capacities, the mirrorset capacity is limited to the size of the smallest member; larger members will have unused capacity. If a member with unused capacity is the last remaining member of a mirrorset, you cannot use the `UNMIRROR` command to change the device back to a single-disk unit, because that change would cause a change in reported disk capacity, possibly confusing the operating system.

## 7.6 Unit Problems after Battery Low

An unusual combination of circumstances could 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.

## 7.7 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.

## 7.8 SET NOPATH Requires Controller Reset

Shut down the controller using the CLI SHUTDOWN THIS\_CONTROLLER command to restart an HSD controller after disabling the host paths with the SET THIS\_CONTROLLER NOPATH command. Failure to do so could result in subsequent unexpected conditions requiring controller reinitialization. The shut down allows devices preferred to this controller to become visible on the other controller.

You do not need to reset an HSD controller after enabling paths.

## 7.9 Device Extended Status

HSD controllers with software Versions 2.7, 3.1, and 3.2 transfer additional device error log information to the host when the information is made available by the device.

## 7.10 SET FAILOVER Causes Target to Restart

The set\_failover copy= command now causes the target controller to restart after the copy completes. For example, the command cli> set failover copy=this\_controller causes the "other controller" to restart after the configuration information is copied from "this controller".

The set\_failover copy= command now copies license information as well as the controller configuration.

## 7.11 Error Log Codes

Refer to revision D01 of the *HS Family of Array Controllers Service Manual* for lists of error log codes.

## 7.12 Chunk Size

Please add the following information regarding "Chunk Size" to your *HS Family of Array Controllers User's Guide*

### 7.12.1 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)

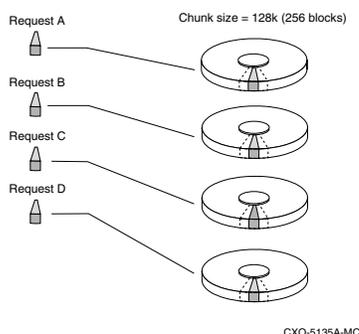
**CAUTION**

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.

**7.12.2 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 7-1.

**Figure 7-1 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 7-1 shows a few examples of chunk size selection).

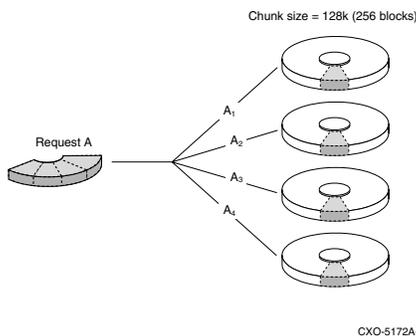
**Table 7-1 Example Chunk Size**

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

### 7.12.3 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 7–2.

**Figure 7-2. Chunk Size Smaller than the Request Size**



CXO-5172A-MC

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

### 7.12.4 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

## 8.0 Documentation Additions and Corrections

The following sections describe corrections to:

- *HS Family of Array Controllers User’s Guide*
- *HS Family of Array Controllers Service Manual*
- *HS Family of Array Controllers CLI Reference Manual*

### 8.1 HS Family of Array Controllers User’s Guide

#### 8.1.1 Normal or Normalizing Member

In the HS Family of Array Controllers User’s Guide on page 5-28, the second paragraph should read: “Mirrorsets that have at least one NORMAL or NORMALIZING member other than the one being removed automatically adjust to the removal of the device.”

## 8.1.2 Changing Raidset Characteristics

In the HS Family of Array Controllers User's Guide on page 6-14, the second sentence of section 6.5.5, "Changing RAIDset Characteristics," should begin, "When a RAIDset loses a member..."

## 8.2 HS Family of Array Controllers Service Manual

### 8.2.1 Software Installation

Please see section 9.0 *Software Installation* in this document. Section 9.0 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 *HS Family of Array Controllers Service Manual, Chapter 4 "Installation"*, pages 4-27 through 4-31.

### 8.2.2 Replacing a Controller with C\_SWAP

The procedure in *HS Family of Array Controllers 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.

### 8.2.3 Fault Management Utility (FMU) Codes:

Add the following codes:

#### 8.2.3.1 Executive Services Last Failure:

- 01192390-A processor interrupt was generated by the CACHEA Dynamic Ram controller and Arbitration engine (DRAB) with an indication that the CACHE backup battery has failed or is low (needs charging).
- 011A2390-A processor interrupt was generated by the CACHEB Dynamic Ram controller and Arbitration engine (DRAB) with an indication that the CACHE backup battery has failed or is low (needs charging).

#### 8.2.3.2 Host Interconnect Port Service Last Failure:

- 42332080-Receive\_main found destination address in the rcv packet does not match node address.

#### 8.2.3.3 Last Failure Codes No Longer Used:

- 010B2380-A processor interrupt was generated by the CACHEA Dynamic Ram controller and Arbitration engine (DRAB) with an indication that the CACHE backup battery has failed or is low (needs charging).
- 010C2380-A processor interrupt was generated by the CACHEB Dynamic Ram controller and Arbitration engine (DRAB) with an indication that the CACHE backup battery has failed or is low (needs charging).

## 8.3 HS Family of Array Controllers CLI Reference Manual

### 8.3.1 SET FAILEDSET Autospare Not Valid Switch for HSD30 Controllers

The *HS Family of Array Controllers CLI Reference Manual* incorrectly identifies AUTOSPARE as a valid switch on the SET FAILEDSET command. AUTOSPARE is not a valid switch for the HSD30 controller.

### 8.3.2 CLI Command Time Out

The time out duration for CLI commands has been extended from 10 minutes to 30 minutes. The time out is the length of time that one controller in a dual-redundant configuration waits for the other controller to complete a command.

### 8.3.3 Cache and Battery Information

The following commands now display more accurate information about the cache and battery state, even when the cache is disabled:

```
SHOW THIS_CONTROLLER
SHOW THIS_CONTROLLER FULL
SHOW OTHER_CONTROLLER
SHOW OTHER_CONTROLLER FULL
```

### 8.3.4 NORECONSTRUCT Not Allowed

The noreconstruct option for RAIDsets is no longer allowed. The valid options are:

```
RECONSTRUCT=NORMAL
RECONSTRUCT=FAST
```

### 8.3.5 New Switches for CLEAR\_ERRORS INVALID\_CACHE

The CLEAR\_ERRORS INVALID\_CACHE command has new switches to tell the controller whether to use the cache information in the controller module or the cache information in the cache module.

```
CLEAR_ERRORS INVALID_CACHE THIS_CONTROLLER NODESTROY_UNFLUSHED_DATA
CLEAR_ERRORS INVALID_CACHE OTHER_CONTROLLER NODESTROY_UNFLUSHED_DATA
CLEAR_ERRORS INVALID_CACHE THIS_CONTROLLER DESTROY_UNFLUSHED_DATA
CLEAR_ERRORS INVALID_CACHE OTHER_CONTROLLER DESTROY_UNFLUSHED_DATA
```

### 8.3.6 Customer Data Not Lost

The documentation for the following commands includes warnings that customer data may be lost. Customer data is not lost when using these commands.

`CLEAR_ERRORS LOST_DATA`

`DELETE UNIT OVERRIDE ONLINE`

`RESTART OTHER_CONTROLLER OVERRIDE ONLINE`

`RESTART THIS_CONTROLLER OVERRIDE ONLINE`

`SELFTTEST OTHER_CONTROLLER OVERRIDE ONLINE`

`SELFTTEST THIS_CONTROLLER OVERRIDE ONLINE`

`SHUTDOWN OTHER_CONTROLLER OVERRIDE ONLINE`

`SHUTDOWN THIS_CONTROLLER OVERRIDE ONLINE`

### 8.3.7 New Tape Formats

The values for default\_format used with the set unit command have been expanded to include the following choices:

**Table 8-1 New Tape Formats**

host_selected	tz88_compression
1600bpi_9track	tz88_nocompression
6250bpi_9track	3490e_compression
tz85	3490e_nocompression
tz86	3480_compression
tz87_compression	3480_nocompression
tz87_compression	tz89_compression
tz87_nocompression	tz89_nocompression
dat_dds_compression	ait_1_compression
dat_dds_nocompression	ait_1_nocompression

Valid values depend upon the specific device and the media loaded. To list valid options for a given unit or media enter the following command:

```
SET UNIT-NUMBER DEFAULT_FORMAT=?
```

The SHOW UNIT-NAME FULL command shows the default tape format and whether the format was set from the host or from the controller CLI. There is a new entry at the end of the switches section that reads:

```
DEFAULT_FORMAT= XXX (ZZZ)
```

where xxx is one of the formats listed above, and zzz is either CLI\_SELECTED or HOST\_SELECTED.

### 8.3.8 Add Loader Correction

The example for creating a command disk on page 7-102 in the User's Guide and page 4-22 in the Service Manual incorrectly shows the loader being added as a unit rather than the passthrough container. The correct example should be:

```
ADD TAPE TAPE320 3 2 0

ADD UNIT T322 TAPE320

ADD LOADER LDR321 3 2 1

ADD PASSTHROUGH PASS1 LDR321

ADD UNIT D332 PASS1
```

## 9.0 Software Installation

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 *HS Family of Array Controllers Service Manual, Chapter 4 "Installation"*, pages 4-27 through 4-31.

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

The following section describes two types of controller upgrades:

- Single Controller Upgrade
- Dual-Redundant Controller Upgrade

### 9.1 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.

- 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.

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

- 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.

- Replace the ESD cover on the controller, pushing the locking pins to secure them in place.
- 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 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 HSD30 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.

### 9.2.1 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.

### 9.3 Dual-Redundant Controller Shutdown Upgrade

If your HSOE 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
```

## 10.0 Part Numbers

Table 10-1 contains controller options that support the current HSOF version.

**Table 10-1 Part Numbers**

Option	Description
HSD30-BA/CA	StorageWorks HSD30 Array Controller with no cache module, 6 SCSI-2 device ports
HSD30-BD/CD	StorageWorks HSD30 Array Controller with a 16 MB read cache module, 6 SCSI-2 device ports
HSD30-BF/CF	StorageWorks HSD30 Array Controller with a 32 MB read cache module, 6 SCSI-2 device ports
HSD30-XD	16 MB read cache module
HSD30-XF	32 MB read cache module
HSD30-YX	Write-back cache option kit -- contains two onboard cache batteries, one battery bracket, one write-back cache license

## 11.0 Order Numbers

Table 11-1 contains the software and documentation kit order number.

**Table 11-1 Order Numbers**

Option	Description
QA-2YJAC-HS	HSOF Software and Documentation

## 12.0 Replacing or Adding an HSD Controller or Cache Module

There are two procedures for replacing a failed controller or cache module in a dual-redundant configuration: an online controller swap method (C\_SWAP) in which the surviving controller continues to process I/O, and an off-line method in which you must shut down both controllers.

### 12.1 Replacing a Controller or Cache Module with C\_SWAP

The online (C\_SWAP) method described in these release notes is valid for HSOF software Version 2.7 and higher.

To replace a controller or cache module using C\_SWAP, two power supplies are required in the controller shelf in the SW500 and the SW800 cabinet.

The following sections guide you through the online method using the C\_SWAP utility.

#### 12.1.1 Required Parts and Tools

You must have the tools listed in Table 12-1 to replace a controller or cache module.

**Table 12-1 Required Parts and Tools**

Tools Required	Purpose
Maintenance terminal and cable	To shut down controllers, restart controllers, and invoke c_swap
ESD wrist strap and ESD mat	To protect all equipment against electrostatic discharge
3/32-inch Allen wrench	To loosen the controller mounting screws; to reinstall the controller
5/32-inch Allen wrench	To unlock the SW800-series cabinet
Small flat-head screwdriver	To connect the CI cable to the controller

**CAUTION**

Before invoking the C\_SWAP utility, terminate all other running utilities and disable all other terminals.

**12.1.2 Preparing the Subsystem Configuration (Online C\_SWAP Method)**

Complete the following steps to prepare your subsystem for the replacement.

1. Have the replacement modules at hand. The modules should be factory-fresh, or should have been shut down cleanly (with the SHUTDOWN command) in their last application.
2. Connect a maintenance terminal to the controller that you are not replacing.
3. If the controller to be replaced is still functioning, use the shutdown command to take it out of service. When using this command, do not specify any command switches.

```
SHUTDOWN OTHER_CONTROLLER
```

**CAUTION**

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

When the controller halts, the green Reset (//) LED stops flashing and stays lit.

4. Take the operating controller out of dual-redundant (failover) mode.

```
SET NOFAILOVER
```

5. With a small flat-head screwdriver, loosen the captive screws that secure the trilink cable connector to the controller. Do not remove the connector (see Figure 5-1, Location of Write Back Cache Module Batteries).
6. Loosen the four captive retaining screws on the controller's front bezel. Do not remove the module yet.
7. Start the C\_SWAP program.

```
RUN C_SWAP
```

**12.1.3 Removing the Modules (Online C\_SWAP Method)**

Use the following steps to remove the modules.

1. When the controller prompts you, answer the question:  
Do you wish to remove other HSD30?
2. Enter "YES"
3. Answer the question:  
Will its cache module also be removed Y/N?
4. Enter "YES" if the cache module requires service or replacement.

5. Wait for the following text to be displayed at the console:

Port 1 quiesced

Port 2 quiesced

Port 3 quiesced

Port 4 quiesced

Port 5 quiesced

Port 6 quiesced

All ports quiesced

Remove the other controller (the one WITHOUT a blinking green LED) within 5 minutes.

Time remaining 4 minutes 50 seconds.

**NOTE**

If the cache module is not being removed, the time allowed to remove the controller will be two minutes.

6. Place the ESD wrist strap around your wrist. Ensure that the strap fits snugly around your wrist.
7. Attach or clip the other end of the ESD wrist strap to the cabinet grounding stud or a convenient cabinet grounding point (nonpainted surface).
8. Unsnap and remove the program card ESD shield on the controller you are removing (see Figure 12-1).
9. Remove the program card by pushing the eject button next to the card (see Figure 12-2), Removing and Installing the PCMCIA Card).
10. Pull the card out and save it for use in the replacement controller module.
11. Remove the CI host cable from the controller to be replaced.
12. Slide the module out of the shelf (noting which rails the module was seated in) and place it on an ESD mat (see Removing and Installing the Controller).
13. Wait for the following text to be displayed at the operating controller's console:

Port 1 restarted

Port 2 restarted

Port 3 restarted

Port 4 restarted

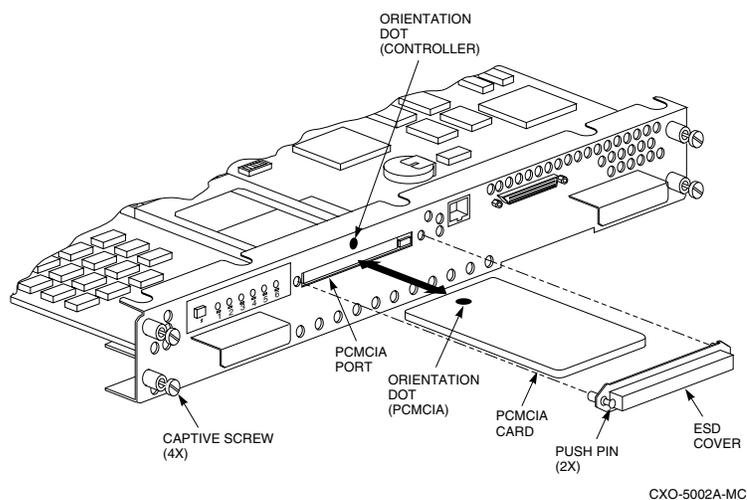
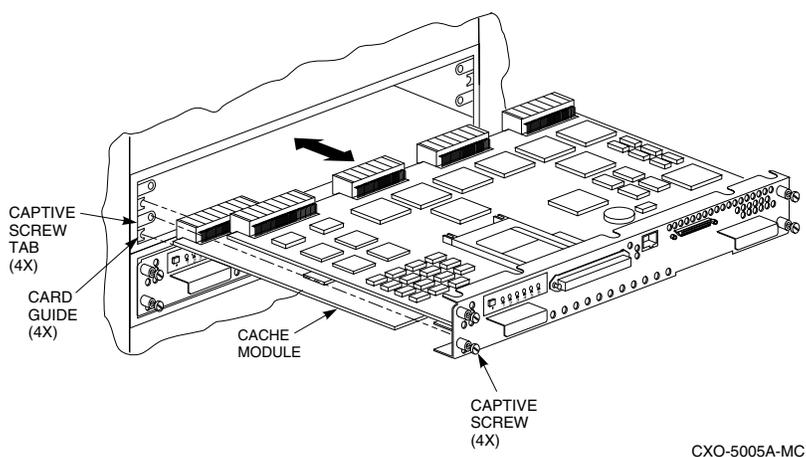
Port 5 restarted

Port 6 restarted

**NOTE**

You can remove the cache module before or after port activity has restarted.

14. If necessary, remove the cache module.

**Figure 12-1 Removing and Installing the PCMCIA Card****Figure 12-2. Removing and Installing the Controller**

### 12.1.4 Installing the New Hardware (Online C\_SWAP Method)

Use these steps to install the new hardware modules.

1. When the controller prompts you, answer the question:  
Do you have a replacement controller available?
2. If you have a replacement module ready, press "Y" for yes.
3. Answer the question:  
\*\*\*Sequence to INSERT the other HSD has begun.\*\*\*  
Do you wish to INSERT the other controller Y/N?
4. Press "Y" for yes.

5. Wait for the following text to appear on the operating controller's console:

```
Attempting to quiesce all ports.
```

```
Port 1 quiesced
```

```
Port 2 quiesced
```

```
Port 3 quiesced
```

```
Port 4 quiesced
```

```
Port 5 quiesced
```

```
Port 6 quiesced
```

```
All ports quiesced.
```

Insert the other HSD WITHOUT its program card, and press Return.

6. If necessary, install the cache module. Slide the module straight in along the rails and then push firmly to seat it in the backplane.
7. Install the controller module. Slide the module straight in along the rails and then push firmly to seat it in the backplane.

**CAUTION**

Do not overtighten the controller's front panel captive screws.  
Damage to the controller PC board or front panel may result.

8. Tighten the four front panel captive screws.
9. Press Return on the operating controller's console.
10. Wait for the following text to be displayed on the operating controller's console:

```
Port 1 restarted
```

```
Port 2 restarted
```

```
Port 3 restarted
```

```
Port 4 restarted
```

```
Port 5 restarted
```

```
Port 6 restarted
```

```
Controller Warm Swap terminated.
```

```
The configuration has two controllers
```

```
To restart the other HSD30
```

```
1) Enter the command RESTART OTHER_CONTROLLER
```

```
2) Press and hold in the Reset (//) button while inserting the program card.
```

```
3) Release Reset (//); the controller will initialize.
```

```
4) Configure new controller by referring to controller's user guide
```

### 12.1.5 Restarting the Subsystem (Online C\_SWAP Method)

Use the following steps to restart the subsystem.

1. Set the original controller to enable the new controller:  
**RESTART\_OTHER\_CONTROLLER**
2. Connect the maintenance terminal to the newly installed controller.
3. Press and hold the Reset (//) button on the new controller while inserting the program card from the replaced controller.
4. Release the Reset button to initialize the controller. Wait for the CLI prompt to display at the terminal. Ignore the "Controllers misconfigured" message that displays.
5. If the new controller reports an invalid cache error, enter one of the following CLI commands to clear the error:

If you have replaced the cache board, enter the following command:

```
CLEAR_ERRORS_INVALID_CACHE_THIS_CONTROLLER  
DESTROY_UNFLUSHED_DATA
```

Otherwise enter:

```
CLEAR_ERRORS_INVALID_CACHE_THIS_CONTROLLER  
NODESTROY_UNFLUSHED_DATA
```

6. Ensure that the new controller is not in dual-redundant mode by entering the following command on the new controller.

```
SET_NOFAILOVER
```

If the controller was already in nonredundant mode, a message is displayed indicating that.

7. Enter the following command from the new controller CLI to put the controllers into dual-redundant mode:

```
SET_FAILOVER_COPY=OTHER_CONTROLLER
```

The new controller restarts.

8. Verify that all host information is correct.

```
SHOW_THIS_CONTROLLER
```

Modify any incorrect information.

9. If you changed any information, restart the new controller.

```
RESTART_THIS_CONTROLLER
```

10. After the controller has initialized, reconnect the host cable to the new controller and tighten the captive screws.

11. Enable CI paths by entering the following CLI commands:

```
SET_THIS_CONTROLLER_PATH_A
```

```
SET_THIS_CONTROLLER_PATH_B
```

### 12.1.6 Replacing a Controller or Cache Module with the Offline Method

If you choose not to use the C-SWAP procedure to upgrade HSOF software, follow the steps in this section to replace an HSD30 controller that has write-back cache. C\_SWAP can be used on Version 2.7 and higher.

Before beginning the procedure, write down the following information:

- The name and ID number for the controller you are replacing
- The configurations of all RAID, stripe and/or mirrorsets
- Which devices are part of any RAID, stripe, or mirrorset, and all unit numbers, device names, and array names

In the event that a problem occurs during this process, or a mistake is made in following these procedures, you will need this information for later manual restoration.

- You must shut down both controllers during part of this procedure, making the devices unavailable to the host operating system. Before starting this procedure, prepare your host system for all I/O associated with this controller to stop. OpenVMS shows the devices in mount verify state, and applications that use those devices stall until the controllers are back in service.
- Compaq recommends setting the mount verify time-out (MVTIMEOUT) to be at least 600 seconds (10 minutes) to allow time to complete this procedure.
- If the system uses host-based shadowsets that have members behind these controllers, increase the shadow member time-out (SHADOW\_MBR\_TM) to a value equal to the mount verify time-out (at least 600 seconds).
- If you do not have available a cache module which has already had its memory cleared, you will experience a five-minute delay in executing these procedures while you await the clearing of the memory on the installed cache module.

### 12.1.7 Required Parts and Tools

You must have the tools listed in Table 12-2 to replace a controller module.

**Table 12-2 Required Parts and Tools**

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

### 12.1.8 Preparing the Subsystem Configuration (Offline Method)

To replace a controller module follow these steps:

1. Have a replacement HSD controller available. The replacement controller should be a unit direct from the factory or a unit that has been SHUTDOWN cleanly since its last application in accordance with the SHUTDOWN procedure identified in the HS Family of Array Controllers User's Guide, Revision D01, Appendix B, Page 125.

#### NOTE

Have all license keys for optional firmware available.

2. Connect a maintenance terminal to the controller you are not replacing.
3. Note which units are preferred to each controller, then prefer all units to the surviving controller.

```
SET UNIT-NUMBER PREFERRED_PATH=THIS_CONTROLLER
```

where unit-number is any unit preferred to the controller that is being replaced. Repeat the command for each unit that must be set to the surviving controller.

4. If the controller to be replaced is still functioning, use the SHUTDOWN command to take it out of service.

```
SHUTDOWN OTHER_CONTROLLER
```

When the controller halts, the green Reset (//) LED stops flashing and stays lit.

5. Take the operating controller out of dual-redundant (failover) mode.

```
SET NOFAILOVER
```

6. Shut down the operating controller.

```
SHUTDOWN THIS_CONTROLLER OVERRIDE_ONLINE
```

The shutdown of the controller could take up to five minutes to flush all data from cache. When the controller halts, the green Reset (//) LED stops flashing and stays lit.

### 12.1.9 Removing the Modules (Offline Method)

To remove the modules complete these steps.

1. Place an ESD wrist strap around your wrist. Ensure that the strap fits snugly.
2. Attach or clip the end of the ESD wrist strap to the cabinet grounding stud or a convenient cabinet grounding point (non-painted surface).
3. Unsnap and remove the program card ESD cover on the controller you are removing.
4. Press and hold the reset button for the controller you are removing.
5. Remove the program card on the controller by pushing the eject button next to the card. Pull the card out and save it for use in the replacement controller module.
6. Disconnect the CI Bus cable from the front of the controller you are removing.
7. Loosen the four captive retaining screws on the front bezel of the controller you are removing.
8. Remove the controller module by pulling it straight out from the shelf backplane.

- Slide the module out of the shelf (noting which rails the module was seated in) and place it on an ESD mat.
- Remove the write-back cache module, if present.
- If the cache module has not failed and you do not have another module that has had the memory cleared, you must clear (discharge) the memory on the current cache module. Locate the battery disable jumper on the cache module. It should be installed on one contact pin. Place the jumper on both pins and leave it there for five minutes. Put the jumper back on one pin to keep from losing it.

### **12.1.10 Installing the New Hardware (Offline Method)**

Complete these steps to install the new hardware modules.

- Install a new cache module or reinstall the original cache module in its original position. Slide the module straight in along the rails and then push firmly to seat it in the backplane.
- Slide the replacement controller module into the shelf without the program card installed. Slide the module straight into the shelf and then push firmly to seat it in the backplane.

#### **CAUTION**

Do not overtighten the controller's front panel captive screws.  
Damage to the controller PC board or front panel could result.

- Tighten the four front panel captive screws.

### **12.1.11 Restarting the Subsystem (Offline Method)**

Restart the subsystem using the following steps.

- Press the Reset (//) button on the original controller, and wait for the CLI prompt to appear.
- Tell the original controller to enable the new controller with the command  
**RESTART OTHER\_CONTROLLER**
- Connect a maintenance terminal to the newly-installed controller.
- Press and hold the Reset (//) button on the new controller while inserting the program card.
- Release the Reset button to initialize the controller. Wait for the CLI prompt to appear at the terminal. Ignore the "Controllers misconfigured" message, that displays.
- Ensure that the new controller is not in dual-redundant mode by typing the following command:

**SET NOFAILOVER**

If the controller was already in non-redundant mode, a message is displayed indicating that.

- Ensure the new controller cannot communicate with any host system yet by typing the following command:

**SET THIS\_CONTROLLER NOPATH\_A**

**SET THIS\_CONTROLLER NOPATH\_B**

The new controller must not be able to communicate with any host system because its cluster parameters are undefined.

- Reconnect the CI Bus cable to the new controller.

9. Configure the host interconnect parameters by typing the following commands:

```
SET THIS_CONTROLLER SCS_NODENAME = (STRING ENCLOSED IN QUOTES)
```

```
SET THIS_CONTROLLER ID = (DECIMAL VALUE)
```

10. Type the following command from the new controller CLI to put the controllers in dual-redundant mode.

```
SET FAILOVER COPY=OTHER_CONTROLLER
```

11. Reinitialize the new controller by typing the command:

```
RESTART THIS_CONTROLLER
```

12. Turn on the host interconnect path(s) by typing the following commands:

```
SET THIS_CONTROLLER PATH_A
```

```
SET THIS_CONTROLLER PATH_B
```

13. Restore the preferred paths for each unit to their original setting using the command

```
SET UNIT-NUMBER PREFERRED_PATH=THIS_CONTROLLER
```

where unit-number is any unit to be serviced through the new controller. Repeat this command for each unit that you want to change.

## 12.2 Using C\_SWAP to Add a Second HSD Controller

You can add a second controller to a single controller configuration to create a dual-redundant configuration. There are two procedures for adding a second controller: an online method using the C\_SWAP utility in which the existing controller continues to process I/O, and an offline method in which you must shut down the existing controller. The online (C\_SWAP) method described in these release notes is valid for HSOF software Version 2.7 and higher.

To add a controller module using the online (C\_SWAP) method, your system must have two power supplies in the controller shelf of the SW500 and SW800 cabinet. If your controller shelf has only one power supply, you must shut down the operating controller to add a second controller. Refer to the HS Family of Array Controllers User's Guide, Rev. D01, Chapter 5, or the HS Family of Array Controllers Service Manual, Rev D01, Chapter 7 for this offline method of adding a second controller.

The following steps guide you through the online method using the C\_SWAP utility.

### 12.2.1 Required Parts and Tools

You must have the tools listed in the Table 12-3 to add a second controller module.

**Table 12-3 Required parts and Tools**

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

#### CAUTION

Before starting the C\_SWAP utility, terminate all other running utilities and disable all other terminals.

### 12.2.2 Preparing the Subsystem Configuration (Online C\_SWAP Method)

Use these steps to prepare the subsystem for the second controller.

Have the new module at hand. The module should be factory-fresh or should have been shutdown cleanly (with the shutdown command) in its last application.

1. Connect a maintenance terminal to the existing controller.
2. Show the configuration of the existing controller:

**SHOW THIS\_CONTROLLER**

The controller will display the following information:

Controller:

```
HSD30 (C) DEC CX20100000 Firmware V02.7,(1) Hardware F01
Configured for dual-redundancy with ZG03231995
In dual-redundant configuration
SCSI address 7
Time: 18 DEC-1995 16:32:54
```

Host port:

```
Node name: HSDA2, valid CI node 21, 32 max nodes
System ID 42001015F525
Path A is on
Path B is on
MSCP allocation class 3
TMSCP allocation class 3
CI_ARBITRATION = ASYNCHRONOUS
MAXIMUM_HOSTS = 5
```

Cache:

```
32 megabyte writeback cache, version 2
Cache is GOOD
Battery is good
No unflushed data in cache
CACHE_FLUSH_TIMER = 1 (seconds)
CACHE_POLICY = B
Host Functionality Mode = A
```

3. Record the size and type of the cache module (if any) that is installed with the first controller. Ensure that the cache module for the new controller is the same size and type.
4. Start the C\_SWAP program:

```
RUN C_SWAP
```

### **12.2.3 Installing the Second Controller and Cache Modules**

Follow these steps to add the second controller and cache.

1. When the controller prompts you, answer the question:  
Do you have a replacement HSD readily available? [N] ? y
2. If you have a replacement module available, enter "YES."
3. Answer the question:  
\*\*\*Sequence to INSERT the other HSD has begun.\*\*\*  
Do you wish to INSERT the other HSD y/n?
4. Press "Y" for yes.

5. Wait for the following text to appear on the operating controller's console:  
Attempting to quiesce all ports.  
Port 1 quiesced  
Port 2 quiesced  
Port 3 quiesced  
Port 4 quiesced  
Port 5 quiesced  
Port 6 quiesced  
All ports quiesced.
6. Insert the other HSD controller WITHOUT its program card, and press Return.
7. If you are installing a new cache module, slide it straight in along the rails and then push firmly to seat it in the backplane.
8. Check the new controller to ensure that the program card is not installed. If it is installed, take the card out by removing the ESD shield and then pressing the eject button next to the card.
9. Install the controller module. Slide the module straight in along the rails and then push firmly to seat it in the backplane.
10. Tighten the four front panel captive screws.

**CAUTION**

Do not overtighten the controller's front panel captive screws.  
Damage to the controller PC board or front panel may result.

11. Press Return on the operating controller's console.
12. Wait for the following text to be displayed on the operating controller's console:  
Port 1 restarted  
Port 2 restarted  
Port 3 restarted  
Port 4 restarted  
Port 5 restarted  
Port 6 restarted Controller Warm Swap terminated.  
The configuration has two controllers  
To restart the other HSD30  
1) Enter the command RESTART OTHER\_CONTROLLER  
2) Press and hold in the Reset (//) button while inserting the program card.  
3) Release Reset (//); the controller will initialize.  
4) Configure new controller by referring to controller's user guide

### 12.2.4 Restarting the Subsystem (Online C\_SWAP Method)

Restart the subsystem using these steps.

1. Set the original controller to allow the new controller to start:

```
RESTART OTHER_CONTROLLER
```

2. Connect the maintenance terminal to the newly installed controller.
3. Press and hold the Reset (//) button on the new controller while inserting the program card from the new controller.
4. Release the Reset button to initialize the controller. Wait for the CLI prompt to appear at the terminal. Ignore the "Controllers misconfigured" message, which is displayed.
5. If the new controller reports an invalid cache error, enter the following CLI command to clear the error:

```
CLEAR_ERRORS INVALID_CACHE THIS_CONTROLLER DESTROY_UNFLUSHED_DATA
```

6. Ensure that the new controller is not in dual-redundant mode by entering the following command on the new controller:

```
SET NOFAILOVER
```

If the controller was already in nonredundant mode, a message is displayed indicating that.

7. Enter the following command from the new controller CLI to put the controllers into dual-redundant mode:

```
SET FAILOVER COPY=OTHER_CONTROLLER
```

The new controller initializes.

8. Verify that all host settings are correct:

```
SHOW THIS_CONTROLLER
```

9. Modify any host information.
10. After the modifications have been made, restart the new controller:

```
RESTART THIS_CONTROLLER
```

11. After the controller has initialized, connect the CI cable to the new controller and tighten the captive screws.
12. Enable CI paths by entering the following CLI commands:

```
SET THIS_CONTROLLER PATH_A
```

```
SET THIS_CONTROLLER PATH_B
```

## 13.0 The HSUTIL Utility

This section describes the features and user interface of the HSUTIL utility. HSUTIL provides two functions:

- Device format-Enables you to perform a basic SCSI format operation on a disk device.
- Device code load-Enables you to download device firmware into supported disk devices.

### 13.1 The Device Format Function

The device format function enables you to perform a SCSI format operation on a disk. Once started, the HSUTIL utility performs the operation without intervention.

### 13.1.1 Device Format Special Considerations

Be aware of the following special considerations when using the HSUTIL utility to format disks:

- HSUTIL formats a single target device per session.
- HSUTIL does not format devices configured as units, stripesets, mirrorsets, RAIDsets, sparesets, or failedsets. The destination device must be an unattached device.
- If a power failure or SCSI bus reset occurs during a format operation, the device could become unusable until a successful format operation is completed. To minimize the possibility of this problem, Compaq recommends that you prevent activity to all devices on the same port as the device being formatted.
- You can use the Ctrl/C or Ctrl/Y keys to exit HSUTIL before the program begins issuing SCSI commands to perform a format. Once the program begins issuing SCSI commands, however, the exit key combinations are disabled.
- Medium formatting is accomplished with no initiator control over defect management. The destination device determines the degree of defect management to be performed.
- Special considerations for dual-redundant configurations:
  - While HSUTIL is active, do not invoke any CLI command or run any local program that might reference the target device from the other controller. Results are unpredictable.
  - Do not reinitialize either controller during a format operation. The resulting SCSI bus reset may render the target device unusable until after a successful format operation is completed.

### 13.1.2 Formatting a Disk Using HSUTIL

This section demonstrates the use of the HSUTIL utility to format a disk. In this example, you are formatting target disk DISK150.

1. Make a backup of any user data on the disk to be formatted.
2. If possible, prevent I/O operations to all other devices on the same device SCSI bus to reduce the possibility of a SCSI bus reset.
3. Invoke the HSUTIL utility and select the format operation by entering "1" at the prompt:

```
RUN HSUTIL
```

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

```
0. EXIT
```

```
1. FORMAT
```

```
2. CODE_LOAD
```

```
Enter the number of the function you wish to perform (0:2) [2] ?1
```

The program reports the unattached units available for formatting:

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

Device	SCSI	Product	ID	Current	Device	Rev
DISK150	RZ26L	(C) DEC		T392		
DISK250	RZ28	(C) DEC		441C		

4. Select DISK150 as the device to be formatted:  
 Enter the device to format ? DISK150

The program displays a number of informational messages and offers you the option of exiting before the format operation begins:

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

The format operation begins, finishing a number of minutes later:

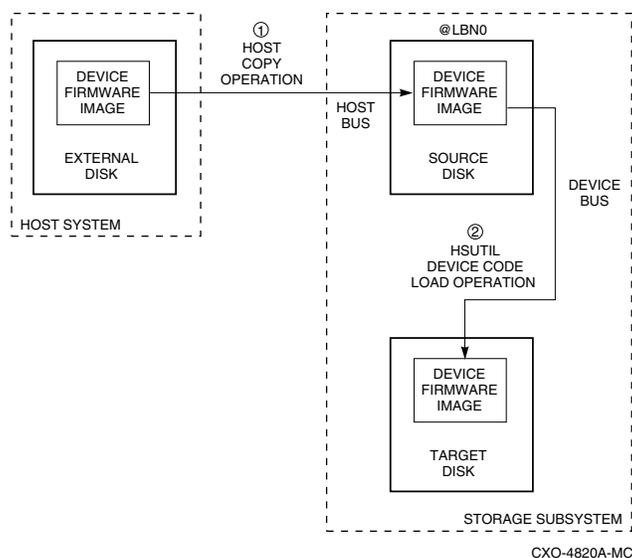
```
HSUTIL Started at: 13-JAN-1996 04:49:48
Format of DISK150 has successfully finished.
HSUTIL - Normal Termination at: 13-JAN-1996 05:22:33
```

DISK150 is now formatted and usable.

## 13.2 The Device Code Load Function

The device code load function allows you to load new firmware into supported devices. The new firmware image must be at a specific LBN on a source disk on the same controller. The source and target disks must be in the configuration on the controller from which HSUTIL is invoked. The source disk can be a raw disk device configured as a unit with no file system or label on it, or it could be a normal file system disk. The device code load process takes place in two steps.

**Figure 13-1. The Device Code Load Process**



### **13.2.1 Preparing the Source Disk**

Copy the new device firmware image to a source disk in the storage subsystem. You cannot load the firmware image from a location external to the controller.

The device firmware image must be located in contiguous blocks at a known logical block number (LBN) on the source disk. Prepare for the code load operation by copying the device firmware image to the source disk at a known LBN location.

One way to prepare the firmware image for downloading is to copy it to a raw source disk containing no data. A raw source disk has no label and no formatted file structure on it. Using a raw disk as the source disk puts the firmware image in contiguous blocks at a specific LBN (usually LBN 0) on the source disk.

The details of copying the firmware image to the source disk from the host disk are specific to the host operating system. Refer to the sample Host Copy Scripts below for examples of copying a firmware image in OpenVMS.

### **13.2.2 Code Loading the Target Disk**

Once the firmware image is located in contiguous blocks at a known LBN on the source disk, use HSUTIL to download it into the target disk.

### **13.2.3 Sample OpenVMS Host Copy Script**

If you are installing new firmware on only one device type, such as RZ29B, and be sure to select the correct code for the code load, then load the code first onto a spare drive at LBN: 0. These commands show how to copy firmware for a single device type.

From an OpenVMS account you would type these commands:

```
$ MOUNT/FOREIGN $2$DUA300:
$ COPY/CONTIGUOUS NODE::LCA:[FIRMWARE.RZ2X]RZ29B_DEC_0014.LOD -
TO $2$DUA300:
$ DISMOUNT/NOUNLOAD $2$DUA300:
```

Drive \$2\$DUA300: must be installed but not mounted on a shelf visible to the controller you are going to connect to. You are now ready to update code in an RZ29B using the code starting at LBN: 0 on DUA300:.

These commands show how to copy firmware for multiple device types.

If you are updating more than one type of supported device, you can copy the code for each device to a spare disk drive or any drive that can be moved to other controllers with little interference. In this example you can use any single mounted disk drive and would continue until you have all the software you need.

```
$ COPY/CONT NODE::LCA:[FIRMWARE.RZ2X]RZ29B_DEC_0014.LOD -
TO $2$DUA300:[000000]
$ COPY/CONT NODE::LCA:[FIRMWARE.RZ2X]RZ28P4_442C_DEC.FUP -
TO $2$DUA300:[000000]
$
```

Use the dump command to find the starting LBN of each firmware image:

```
$ DUMP/HEAD/BLOCK=COUNT=0 $2$DUA300:[000000]*.FUP
$ DUMP/HEAD/BLOCK=COUNT=0 $2$DUA300:[000000]*.LOD
```

The information below was edited to reduce the printout but each line is from the above dump command.

```
$2$DUA300:[000000]RZ26L_442D_DEC.FUP LBN: 8
$2$DUA300:[000000]RZ26_T392A_DEC.FUP LBN: 520
$2$DUA300:[000000]RZ28P4_442C_DEC.FUP LBN: 1036
$2$DUA300:[000000]RZ28_442D_DEC2104.FUP LBN: 1548
$2$DUA300:[000000]RZ29B_DEC_0014.LOD LBN: 2060
$ DISMOUNT/NOUNLOAD $2$DUA300:
```

Drive \$2\$DUA300: must be installed but not mounted on a shelf visible to the controller you are going to connect to. You are now ready to update code in different drives of choice with their code starting at LBN: \* on DUA300: as printed out above.

### 13.2.4 Device Code Load Special Considerations

Be aware of the following special considerations when using the HSUTIL utility to download firmware to a disk device:

#### CAUTION

HSUTIL has been tested to perform the code load operation properly only with supported devices. Use HSUTIL to code load unsupported devices at your own risk.

- Device compatibility with firmware code loading using the SCSI Write Buffer command varies greatly from vendor to vendor. While HSUTIL does check to see if a target disk is an unsupported device, it does not prevent you from continuing the code load operation on the device. Any unsupported device could be rendered unusable until it is factory-initialized, if the code load is unsuccessful.

#### CAUTION

Delete all non-disk devices from the controller configuration that are on the same shelf as the target disk, and prevent bus activity to all other disks on the same shelf before downloading firmware to the target disk.

- If a power failure or SCSI bus reset occurs during a code load operation, permanent damage to the device can result. To minimize the possibility of rendering a device unusable, Compaq recommends halting activity on other devices on the same port during the code load process.
- HSUTIL processes a single target device per session.
- HSUTIL does not load code to devices configured as units, stripesets, mirrorsets, RAIDsets, sparesets, or failedsets. The destination device must be an unattached device.
- The source device must be configured as a single-device unit.

- The device firmware image must be located on the source disk in contiguous blocks at a known LBN location.
- Once HSUTIL has allocated the source disk for the code load operation, the source disk is unavailable to other subsystem operations.
- Some device firmware releases require a format operation after a code load operation to make the device usable. In this case, you must rerun HSUTIL to perform the format operation. See the release notes for your device firmware version to determine if you must format the device after code loading. If you do not have release notes, contact the device vendor to verify whether a format is required after a firmware update.
- Use the Ctrl/C or Ctrl/Y keys to exit HSUTIL before the program begins issuing SCSI commands to perform a code load. Once the program begins issuing SCSI commands, however, the exit key combinations are disabled.
- In dual-redundant configurations, shut down the controller that you are not using for the code load operation and eject the program card. Reinstall the program card to restart the controller after you have finished the code load.

### 13.2.5 Code Loading a Device Using HSUTIL

The example below demonstrates the use of the HSUTIL utility to download device firmware. In this case, your device's firmware release documentation specifies that it does not require a subsequent format. In the example, you are changing the firmware in device DISK150 from revision 0007 to 0014. The source disk is DISK250.

The details of copying the device firmware image from a location external to the controller to the source disk are specific to the host operating system. In this example, the firmware image is already at LBN 0 of DISK250.

1. Back up all data on the target device using your operating system's backup procedures.
2. Use the show disk command to report the available devices and the firmware level of each. In the example below, DISK150 is the target device and DISK250 (Unit D200) is the source device. Note that the firmware level of DISK150 is reported as "0007" and that it is an unattached device:

```
SHOW DISK FULL
Name          Type      Port  Targ  Lun          Used by
DISK150       disk      1     5     0
              RZ29B    (C) DEC          0007
DISK250       disk      2     5     0           D200
              RZ28    (C) DEC          442D
Switches:
NOTTRANSPORTABLE
Size: 8356756 blocks
```

If the target device is configured as part of a higher-level container, such as a unit, you must delete that part of the configuration. The target device must be an unattached disk drive.

- Invoke the HSUTIL utility and select the code load function by entering option "2" at the prompt:

```
RUN HSUTIL
*** Available functions are:
  0. EXIT
  1. FORMAT
  2. CODE_LOAD
Enter the number of the function you wish to perform (0:2) [2] ?2
```

- The program displays the single-device units available as source disks for code loading:

```
Single device units on this controller include:
  Unit   Associated Device   SCSI Product ID
  200    DISK250                RZ28
```

- Select Unit 200 (DISK250) as the source disk:

```
Which unit is the code to be loaded FROM ? 200
Because you have previously copied the firmware image to LBN 0 on the source
disk, you can accept the default starting LBN of the device firmware image by
pressing Return at the prompt:
What is the starting LBN of the code on the unit where the code is to be loaded
FROM [0] ? Return
```

- Enter the disk model that is the target of the code load. Enter this information exactly as it was reported in the show command, including any spaces or special characters: (You may, however, use only as many characters as are necessary to uniquely identify the device, as shown in the following example.)

```
What is the SCSI PRODUCT ID of the disk that you want code load TO ? RZ29B
```

- The program lists all unattached disk drives of the model specified. In this example, DISK150 is the only unattached device with the desired product ID:

```
Unattached devices on this controller include:
  Device   SCSI Product ID   Current Device Rev
  DISK150  RZ29B             (C) DEC           0007
```

- Select the target device from the list:

```
Which device is the code to be loaded TO ? DISK150
```

- The program displays a number of informational messages and offers you the option of exiting before the format operation begins:

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

- The program begins the device code load operation, finishing within two minutes:

```
Device code has been successfully downloaded to DISK150
```

```
HSUTIL - Normal Termination at: 13-JAN-1996 05:53:03
```

- Enter the show disk command for disk150 to verify that the device code load operation was successful. Note that the program reports a new firmware revision level, such as "0014." The device is now available for use with its new firmware.

```
SHOW DISK150
```

```
Name          Type      Port Targ  Lun          Used by
-----
DISK150       disk      1    5    0
              RZ29B   (C) DEC          0014
```

### **13.2.6 Code Loading a Device Using HSUTIL (Format Required)**

Some devices and some revisions of device firmware require that you format the device after installing new firmware. The Code Load portion of HSUTIL detects when a device format is required, and displays a message on the console indicating this. Follow the instructions for "Formatting a Disk" in this document. You cannot use the device until it is formatted.

Ensure data is backed up from the device before starting the code load.

### 13.3 Abort Codes

In the event that HSUTIL terminates before it successfully completes a format or code load operation, it reports one of the abort codes shown in the Table 13-1:

**Table 13-1 Abort Codes**

Abort Code	Description
1	FAO returned either FAO_BAD_FORMAT or FAO_OVERFLOW
2	Bad return from TS\$READ_TERMINAL_DATA
3	TS\$READ_TERMINAL_DATA returned either an ABORTED or INVALID_BYTE_COUNT
4	User requested an abort via ^Y or ^C
5	An error occurred on a SCSI command
6	Can't find the pub, device is probably missing