

Compaq StorageWorks™

Release Notes

HSJ50 Array Controller Operating Software HSOF Version 5.7

This document summarizes features and characteristics of the HSJ50 array controller operating software Versions 5.7, 5.6, 5.4, 5.3, 5.2, and 5.1 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 Trade-mark 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 Micro-soft 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 Soft-ware, 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 documen-tation accompanying such products. Nothing herein should be construed as constituting a further or additional warranty.

Printed in the U.S.A.

Release Notes – HSJ50 Array Controller Operating Software HSOF Version 5.7
Fifth Edition (May 2000)
EK-HSJ50-RN. E01

Visit our Web Site for the Latest Information

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

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

Release Notes Contents

These release notes cover the following topics:

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

Release Package Contents

The Version 5.7 release package consists of the following:

- A cover letter
- A PCMCIA program card containing HSOF Version 5.7 software
- The HSJ50 documentation set
 - *HSJ50 Array Controller HSOF Version 5.1 Configuration Manual*
 - *HSJ50 Array Controller HSOF Version 5.1 Service Manual*
 - *HSJ50 Array Controller HSOF Version 5.1 CLI Reference Manual*
 - *HSJ50 Array Controller Operating Software Version 5.7 Release Notes*
 - *HSJ50 Array Controller Operating Software (HSOF), Version 5.7 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 HSJ50 Array Controller running HSOF Software Version 5.7. Individuals responsible for configuring, installing, and using the HSJ50 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 HSJ50 Array Controller User's Manuals.

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 V57Z-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 5.7:

- Hardware Support
- Operating System Support
- Host Adapter Support
- PKQdriver Version
- Device Support

2.1 Hardware Support

HSOF Software Version 5.7 supports the following revisions or higher for the RAID Array HSJ50 Controller and associated hardware:

- HSJ50 controller module, hardware revision A
- Version 3 cache module, hardware revision A or B
- BA350–MA controller shelf
- BA350–Sx device shelf
- BA350–Sx wide device shelf with 8-bit or 16-bit personality module (controller operates in 8-bit mode only)
- BA35x–HF power supply
- HS35X–BA Single External Cache Battery SBB (with 1 ECB)
- HS35X–BB Dual External Cache Battery SBB (with 2 ECBs)
- HSSIM–AA 32 MB SIMM Pack for Version 3 Cache module upgrades

2.2 Operating System Support

HSOF Version 5.7 software on HSJ50 controllers is supported by the following operating system versions:

- OpenVMS™ Alpha:
V6.2-1H3, V7.1-1H1, V7.1-1H2, V7.1-1H3, 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 5.7, 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.

- CIXCD-AB
- CIXCD-AC (both for XMI-based systems)
- CIBCA-Bx (for BI-based systems)-VAX
- CIPCA (for PCI-based systems)-Alpha

2.4 PKQdriver Version

All OpenVMS systems, using the StorageWorks Command Console V2.3 Agent and HSOF V5.7 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 V5.7, your system may experience an operating system crash.

2.4.1 Link Date and Image Identification Information

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

Table 2-1 PKQdriver Identification Information

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

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

2.4.2 Identifying Your PKQdriver

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

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

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

You will receive the first data screen.

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

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

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

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

```
Image Identification Information
```

```
image name: "SYS$PKQDRIVER"
```

```
image file identification: "X-6A52"
```

```
image file build identification: "X6TE-
0050130101"
```

```
link date/time: 10-Jan-2000 17:50:12.67
```

```
linker identification: "A11-39"
```

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

In this example:

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

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

2.4.3 Installing the Current PKQdriver

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

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

2.5 Device Support

HSOF Software V5.7 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 ¹	1.05	440C	A01
RZ26N-VA/VW ¹	1.05	446	A01
SWXD3-SF/WF ¹	1.05	446	A01
DS-RZ26N-VZ ¹	1.05	1003	A01
DS-RZ1BB-VW	2.10	LYJO/0656	A01
RZ28-VA/VW ¹	2.10	435E	B03
RZ28B-VA	2.10	0003	A01
RZ28D-VA/VW ¹	2.10	0008	A01
SWXD3-SG/WG ¹	2.10	0008	A01
RZ28M-VA/VW ¹	2.10	0466	A01
DS-RZ28M-VZ ¹	2.10	1003	A01
SWXD3-SH/WH ¹	2.10	0466	A01
RZ74-VA	3.57	T427B	B07
DS-RZ1CB-VW	4.1	LYJO/0656	A01
RZ29B-VA/VW ¹	4.3	0007	B01
SWXD3-SE/WE ¹	4.3	0007	C02/A01

Table 2-2 Supported Disk Drives - continued

Device	Capacity in Gigabytes	Minimum Microcode Version	Minimum Hardware Version
DS-RZ1CF-VA/VW	4.3	0370/0371	A01
DS-RZ1DB-VW	9.1	LYJO/0307	A01
DS-RZ1DF-VA/VW	9.1	0372/1614	A01
DS-RZ40-VA	9.1	LYGO	A01
DS-RZ1DD-VA/VW	9.1	0305/3B07	A01
DS-RZ1EF-VA/VW	18.2	0372/N1H1	A01
DS-RZ1ED-VW	18.2	0306/0305/3B07	A01
DS-RZ1EA-VW	18.2	3B05/ B016	A01
DS-RZ1DA-VW	9.1	3B06/ B016	A01
DS-RZ1FC-VW	36.4	3B07	A01

Table 2-2 Notes:

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

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

NOTE

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

Table 2-3 Supported Tape Drives

Device	Capacity G= Gigabytes T=Terabytes	Minimum Microcode Version	Minimum Hardware Version	Notes
TL812	960/1920G	1.2 robot /CC33 drive	A01	1,2,3,4,5
TL822	5280/10560G	1g4F robot /CC33 drive	A01	1,2,3,4,5
TL826	3520/7040G	1g4F robot /CC33 drive	A01	1,2,3,4,5
DS-TL893	924/1848T	V2A/5A	A01	1,2,3,4,5
DS-TL894	1.68/3.36T	V1.24	A01	1,2,3,4,5
DS-TL895	3.1/6.2T	230	A01	1,2,3,4,5
DS-TL896	6.16/12.32T	V2A/5A	A01	1,2,3,4,5
TZ87-VA	10/20G	930A	A01	3,5,6
TZ87N-VA	10/20G	930A	A01	3,4,5,6
TZ87-TA	10/20G	9514	B02	1,2,3,6
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

Table 2-3 Supported Tape Drives - continued

Device	Capacity G= Gigabytes T=Terabytes	Minimum Microcode Version	Minimum Hardware Version	Notes
DS-TL890	560/1120G	3.23 robot/V55 drive	A01	1,2,3,4,5
DS-TL891	350/700G	3.23 robot/V55 drive	A02	1,2,3,4,5
DS-TL892	350/700G	3.23 robot/V55 drive	A02	1,2,3,4,5
DS-AIT35- VW	35/70G	4.03	A01	5,6,7

Table 2-3 Notes:

1. Requires 0.2 meter SCSI-1 to SCSI-2 transition cable, Compaq internal part number 17-03831-01 for DWZZA-AA, and Compaq part number 17-04367-01 for SSB DWZZA-VA and DWZZB-VW.
2. Requires SWXA2 single-ended to differential SCSI signal converter.
3. Capacity values represent compressed data. The compression factor is device dependent based on individual device algorithms.
4. Cannot read TK50, TK70 or TZ30 formatted tapes.
5. Requires a KZPSA or PMAZC host adapter.
6. Tape device code load is supported.
7. Wide tape devices require a SWXSS-06 shelf with an 8-bit I/O module.

Table 2-4 Supported Solid State Devices

Device	Capacity in Gigabytes	Minimum Microcode Version	Minimum Hardware Version	Notes
EZ31-VW	0.134	V064	A01	2, 3
EZ32-VW	0.268	V064	A01	2, 3
EZ51R-VA	0.10	V096	D01	2, 3
EZ54R-VA	0.42	V109	C02	2, 3
EZ58R-VA	0.85	V110	D01	1, 2, 3
EZ64-VA	0.475	V064	A01	2, 3
EZ64-VW	0.475	V070	A01	2, 3
EZ69-VA	0.950	V064	A01	2, 3
EZ69-VW	0.950	V070	A01	2, 3
EZ454	.536	Y018	A01	2, 3
EZ832	3.2	Y018	A01	2, 3
EZ41	0.134	V012	A01	2, 3
EZ42	0.268	V012	A01	2, 3
EZ51	.107	V109	C02	2, 3
EZ54	.428	V109	C02	2, 3
EZ705	0.536	V012	A01	2, 3
EZ711	1.1	V012	A01	2, 3
EZ716	1.6	V012	A01	2, 3

Table 2-4 Notes:

1. Code load is not supported for these drives
2. Formatting supported for these drives
3. Do not warm-swap solid-state disk drives. Make sure power to the device shelf is turned off before removing or inserting this device.

Table 2-5 Supported CD-ROM Readers

Device	Capacity in Gigabytes	Minimum Microcode Version	Minimum Hardware Version
RRD42-VB/VU	0.6	1.1a	A01
RRD43-VA	0.6	0064	A02
RRD44-VA	0.6	3493	A02
RRD45-VA/VU	0.6	1645	A01
RRD46-VA	0.6	1337	A01
RRD47-VA	0.6	1206	A01

Table 2-5 Notes:

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

3.0 New Items

Sections 3.1 and 3.2 briefly describe the new items in HSOF Software V5.7.

- New Devices Supported in HSOF Software
- Issue Resolutions

3.1 New Devices Supported in HSOF V5.7 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 V5.7 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 V5.7 release. The issues addressed by HSOF V5.7 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 of a TKZ61 tape device becoming misconfigured.
- A rare occurrence in which RAIDsets were not properly failing over in the event of battery failure.
- In a rare occurrence the HSJ50 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.
- A rare occurrence of a TZ89 tape device error causing a controller reset with a last failure code of 103360188

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 V5.7 Software release that are not covered in other documentation.

- Correction for TILX Incompatibility with Disks
- Action for repeated Reboots Problem Resolution
- NOR/DDDL Errors
- 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 5.4.

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 5.4 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 NOR/DDL Errors

A small number of NOR/DDL (No Response/Dual Data Link) errors may occur during cluster transitions and other normal cluster events. However, some customers have experienced a NOR/DDL error rate higher than expected. Under certain rare conditions on the CI bus, the HSJ controller CI interface ASIC may become unable to recognize the host's acknowledgment that a packet from the controller has been received. When this condition persists on both paths, the controller closes the virtual circuit with an error log of NOR/DDL. This version of HSOF software corrects this specific problem.

4.4 Support for Large-Drive Formatting

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

4.5 Support for up to Fourteen 18 GB Drives

The maximum storage set 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 5.6 contain clarification on the following subjects:

- EMU Alarm Clarification
- Write-Back Cache Batteries and Battery Handling
- Logical Device Event Reports
- Saving the Subsystem Configuration
- 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 storage set failures.

5.2 Write-Back Cache Batteries and Battery Handling

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

5.2.1 *CACHE_UPS* Switch Use

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

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

When the *CACHE_UPS* switch is enabled, an uninterruptible power supply (UPS) is assumed to be in use to maintain power to the write-back cache module in the event of a power failure. Note that the HSOF software does not have the capability to monitor the remaining capacity of an external UPS. There are many variables in a UPS configuration; some configurations might put data at risk if the UPS is exhausted. If UPS power is lost to the host/storage system without warning, write-back cache data is maintained by the cache batteries only. With *CACHE_UPS* set, if UPS power is 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 applies, and regular periodic battery tests are performed.

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

Table 5-1 Cache Policy Settings

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

NOTE

Write-through cache mode is not a 100% guarantee of user data consistency within storagesets (RAID and mirrorsets). This is because of an inherent delay in the writing of data to the different members of these storagesets. If a power failure occurs between the writing of some members and others (an extremely narrow window), data on the various members will be inconsistent.

Therefore, Cache Policy A (which denies access to RAID/mirrorsets on controllers with low batteries) is the only 100% guarantee of user data consistency. That is why Cache Policy A enforces loss of storageset access when a battery does not pass the battery test.

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

5.2.3 Failed Battery Action

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

- Dual-controller HSJ50s (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 “fail-back” leading to no access to RAIDsets and mirrorsets
- Single-controller HSJ50s
 - 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 HSJ50 array controller configuration with write-back cache are not followed, there is potential for loss of data that may exist on any devices connected with them. In addition, if you are turning off the power to the controller subsystem for any reason for longer than 1 day (such as a holiday, system move, or replacing a bad SCSI host cable), you will need to turn off the external cache batteries (ECBs) to prevent them from discharging. To avoid problems when the controller restarts, always use the following process to shut it down:

CAUTION

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

5.2.4.1 Shutting Down a Controller

1. Use the proper procedures for shutting down the operating system if the host system is also going to be turned off.
2. If the host is not going to be turned off, it is not necessary to shut down the host system, but use the proper operating system procedures to dismount any units that are accessed through the HSJ50 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 RAID-set or mirrorset units, or single devices with write-back caching turned on, set norun on all units.
4. When the dismount and/or the operating system shutdown procedures are complete, invoke the controller shutdown commands on the controllers.

CAUTION

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

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

```
SHUTDOWN OTHER_CONTROLLER
SHUTDOWN THIS_CONTROLLER
```

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

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

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

5.2.5 Disabling the Write-Back Cache Batteries

Refer to Chapter 2 in *HSFAM of Array Controllers Service Manual*, for instructions on removing and replacing the controller module and write-back cache module.

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

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

- Single Controller:

SHUTDOWN THIS_CONTROLLER

- Dual Controllers:

SHUTDOWN OTHER_CONTROLLER

SHUTDOWN THIS_CONTROLLER

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

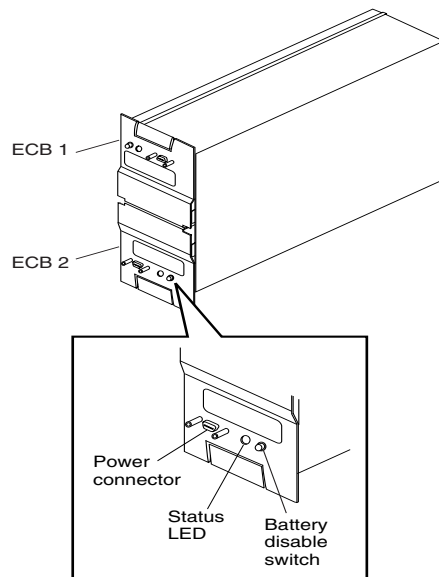
4. Power off the entire storage subsystem.
5. Press the battery disable switch, the small button labeled SHUT OFF next to the flashing status LED on the ECB, and hold it in for approximately 2 seconds (refer to Figure 5-1). The cache LED will flash once, then shut off.

NOTE

The battery is no longer powering the cache.

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

Figure 5-1. External Cache Battery Disable Switch



CXO6164A

5.2.5.1 Restarting the Controller

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

5.3 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 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 5.2, or from an HSJ50 to an HSJ50 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 5.2 to Version 5.4 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.

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

HSOF Version 5.1 software allows you to use the controller's HSUTIL program to install new microcode to the TZ89 tape drive.

HSOF Version 5.1 software also allows you to use the controller's HSUTIL program to install new microcode to solid state disk devices.

5.6 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: HSJ50 Serial Num-  
ber: ZG05051997 Hardware Version: 0000(00) Firmware  
Version: V57Z(00)
```

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

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

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

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

6.0 Operating Constraints

Paragraphs 6.1 through 6.11 describe the operating constraints for HSOE Version 5.7 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 Stageset Limits
- Hot Swapping Devices
- Maintenance Terminal Port Speeds
- CFMENU Constraints
- Changing Host Adapter Types
- CLONE Utility
- Low/Failed Battery on Dual-Redundant Configurations with Preferred Paths
- Device Adds, Moves, and Changes
- Tape Drive Firmware Revision
- Set Tape Density When Adding Drive and Changing Cartridge Types
- Running CONFIG or CFMENU Utilities During Backup or Copy Operations

6.1 Updated Stageset Limits

The following limits apply to stagesets 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 stagesets (RAIDsets, mirrorsets, and/or stripesets).
- There can be no more than 32 physical device members for a unit.
- Each stageset 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 5.4 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 CFMENU Constraints

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

6.5 Changing Host Adapter Types

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

6.6 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.

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

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

6.8 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.

6.9 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.

6.10 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.

6.11 Running CONFIG or CFMENU Utilities During Backup Operations

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

7.0 Avoiding Problem Situations

In certain situations, you may experience unusual controller behavior. Paragraphs 7.1 through 7.6 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
- Fast Shadow Member Eviction
- Only One Switch per CLI Command Recommended
- DILX Enhancement

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

As noted in Section 3.2, in a rare occurrence the HSJ50 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 HSJ50 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 Fast Shadow Member Eviction

An MSCP flag is provided to enable rapid shadow member eviction when a device error is detected. The OpenVMS operating system can set this flag based on SYSGEN parameters in the field SHADOW_SYS_DISK. The MSCP flag is called MD.SER. When set, and an I/O encounters a device error, the I/O is returned as failed without further error recovery. The OpenVMS operating system can then evict a shadowset member, as appropriate.

7.5 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
```

7.6 DILX Enhancement

The HSJ50 controller with HSOF Version 5.1 software has a new Version Disk In-Line Exerciser (DILX) that acts on disks configured as single-device, RAIDset, sripeset, mirrorset, striped-mirrorset, or partitioned units.

8.0 Documentation Additions and Corrections

The following sections describe corrections to:

- *HSJ50 Array Controller HSOF Version 5.1 Configuration Manual*
- *HSJ50 Array Controller HSOF Version 5.1 CLI Reference Manual*
- *HSJ50 Array Controller HSOF Version 5.1 Service Manual*

8.1 HSJ50 Array Controller HSOF V5.1 Configuration Manual

8.1.1 Chunk Size

Referring to the section titled *Initialize Switches* in *HSJ50 Array Controller HSOF Version 5.1 Configuration Manual*, pages 3-18 through 3-20 the information regarding “Chunk Size” should be replaced with the following information.

The information on maximum chunk size is incorrect and should be removed.

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

8.1.1.1 Specifying Chunk Size

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

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

CAUTION

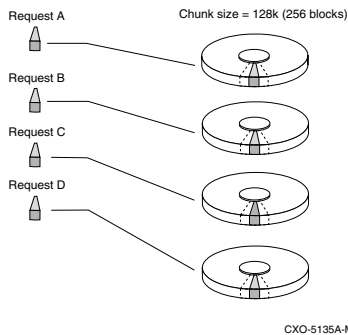
The default chunk size should be used with caution. It is imperative that you analyze the I/O transfers being used with your application to determine the proper chunk size. An improper value used as a chunk size may cause serious performance problems with the subsystem.

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

8.1.1.2 Increasing the Request Rate

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

Figure 8-1 Chunk Size Larger than the Request Size



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

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

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


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

Table 8-1 Example Chunk Size

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

8.1.1.3 Increasing the Data Transfer Rate

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

Figure 8-2. Chunk Size Smaller than the Request Size

CXO-5172A-MC

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

8.1.1.4 Increasing Sequential Write Performance

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

8.2 HSJ50 Array Controller HSOF V5.1 CLI Reference Manual

8.2.1 Chunk Size

Referring to the section titled *Initialize, Switches* in the *HSJ50 Array Controller HSOF Version 5.1 CLI Reference Manual*, page 2-38 the information regarding “Chunk Size” should be replaced. Please refer to the above section of these Release Notes (*Section 8.1.1, Chunk Size*) for the correct information.

8.2.2 SET FAILEDSET Autospare Not Valid Switch for HSJ50 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 HSJ50 controller.

8.2.3 MIRROR_CACHE Not Valid Switch on SET THIS_CONTROLLER and SET OTHER_CONTROLLER COMMANDS

The SET THIS_CONTROLLER and SET OTHER_CONTROLLER commands in the *HSJ50 Array Controller HSOF Version 5.1 CLI Reference Manual* contain a description of the MIRROR_CACHE switch. MIRROR_CACHE is not a supported switch for the HSJ50 array controller.

8.3 HSJ50 Array Controller Service HSOF V5.1 Manual

8.3.1 Software Installation

Please see section 9.0 *Software Installation* in this document. Section 9.0 provides the specific procedures needed to install your HSOF V5.7 Software. The information in this section can be used in conjunction with the information found in the *HSJ50 Array Controller HSOF Version 5.1 Service Manual*, Chapter 3 “Installing”, pages 3-3 through 3-21.

8.3.2 Replacing a Controller with C_SWAP

The procedure in *HSJ50 Array Controller 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.

9.0 Software Installation

This section provides the specific procedures needed to install your HSO V5.7 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 HSO V5.7 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 HSO software in a non-redundant (single) controller:

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

3. Shutdown the controller with the following command:

SHUTDOWN THIS_CONTROLLER

CAUTION

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

CAUTION

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

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

NOTE

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

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

NOTE

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

NOTE

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

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

NOTE

You can identify the version of the HSOE 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 HSJ50 HSOE 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 5.7 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 V5.7 From:	Shutdown Upgrade to V5.7 From:
Versions 5.4, 5.5, and 5.6	Versions 5.0, 5.1, 5.2, and 5.3

As shown in the above table, Rolling Upgrades to Version 5.7 are supported from Versions 5.4, 5.5, and 5.6.

Releases prior to HSOF Software Version 5.4 must upgrade to HSOF Software Version 5.7 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 HSOF Software version requires a Dual-Redundant Controller Shutdown Upgrade, follow the below procedures.

NOTE

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

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

3. Enter the SET NOFAILOVER command to take the controllers out of dual-redundant (failover) mode.

```
$ SET NOFAILOVER
```

4. Shutdown the controller with the following command:

```
SHUTDOWN THIS_CONTROLLER
```

CAUTION

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

CAUTION

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

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

NOTE

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

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

NOTE

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

NOTE

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

8. Replace the ESD cover on the controller, pushing the locking pins to secure them in place.
9. Mount the storage units on the host. Resume I/O activity to the controller.
10. Repeat 5 through 9 for the second controller.
11. After you complete all of the above steps on both controllers, put the controllers back into dual-redundant (failover) mode by entering the following command:

```
$ SET FAILOVER COPY=CONFIGURATION_SOURCE
```


10.0 Order Numbers

Table 10-1 contains the software and documentation kit order number.

Table 10-1 Order Numbers

Option	Description
HSJ50-AF, HSJ50-AH, HSJ50-AJ	StorageWorks HSJ50 Array Controller with a 128 MB cache module, 36/42 SCSI-2 device connections, 1 external cache battery in a SBB, one 2-meter ECB cable
HSJ52-AF	Two StorageWorks HSJ50 Array Controllers with 32 MB cache modules, 72/84 SCSI-2 device connections, 2 external cache batteries in a SBB, two 2-meter ECB cables
HSJ52-AH	Two StorageWorks HSJ50 Array Controllers with 64 MB cache modules, 72/84 SCSI-2 device connections, 2 external cache batteries in a SBB, two 2-meter ECB cables
HSJ52-AJ	Two StorageWorks HSJ50 Array Controllers with 128 MB cache modules, 72/84 SCSI-2 device connections, 2 external cache batteries in a SBB, two 2-meter ECB cables
HSSIM-AA	32 MB SIMM pack for cache modules upgrades
HS35X-BA	One replacement external cache battery in a SBB
HS35X-BB	Two replacement external cache batteries in a SBB

NOTE

One prerequisite QB-5C5AA-SA software kit must be ordered for each controller ordered (two kits for HSJ52 configurations).

11.0 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.

11.1 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.

11.1.1 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.

11.1.2 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
```

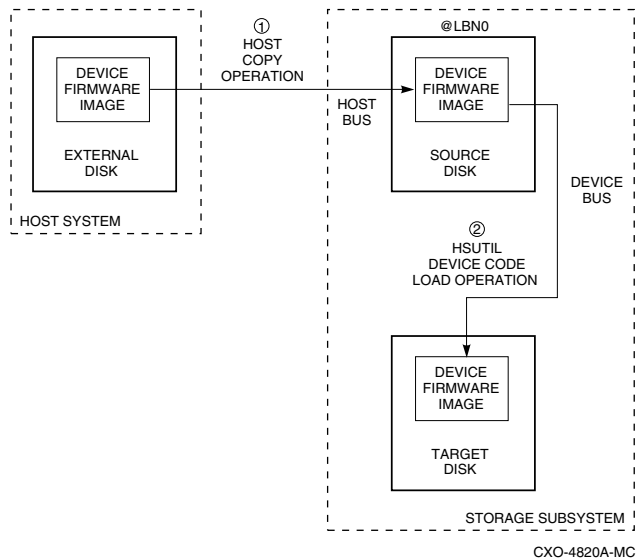
```
Format of DISK150 has successfully finished.
```

```
HSUTIL - Normal Termination at: 13-JAN-1996 05:22:33
```

DISK150 is now formatted and usable.

11.2 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 11-1. The Device Code Load Process

11.2.1 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.

11.2.2 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.

11.2.3 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.

11.2.4 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.

11.2.5 Code Loading a Device Using HSUTIL

Example 4 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.

3. 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
```

4. 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
```

5. 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
```

6. 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
```

7. 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
```

8. Select the target device from the list:

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

9. 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
```

10. 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
```

11. 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
```

11.2.6 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.

11.2.7 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 Table 11-1:

Table 11-1 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