Printrak, A Motorola Company

MANUFACTURING INSTALLATION PROCEDURE FOR THE HSZ70 RAID CONTROLLER 0603-00017 Rev D

Revised: 2/28/00

Printrak, a Motorola Company
1250 North Tustin Ave.
Anaheim, CA 92807
UNITED STATES OF AMERICA

1.0 HSZ70 HARDWARE CONFIGURATION.

The first step in configuring an HSZ70 RAID controller is setting up the hardware. This section details how to perform this function.

1.1 The components which make up a Controller Module are depicted below.

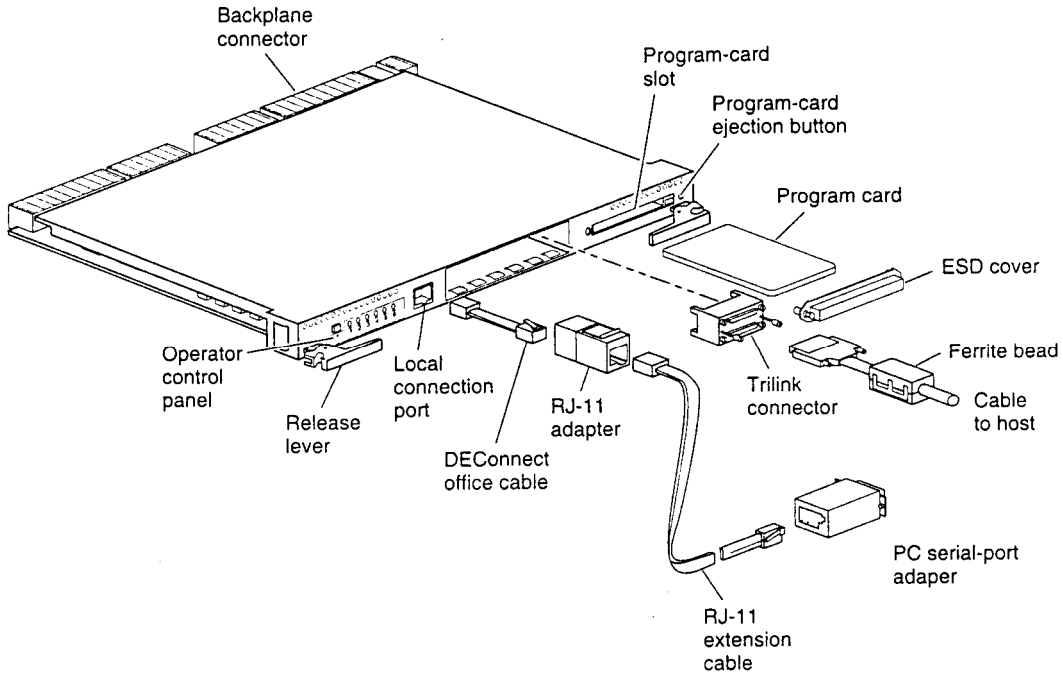

1.2 The Cache Module is depicted below:

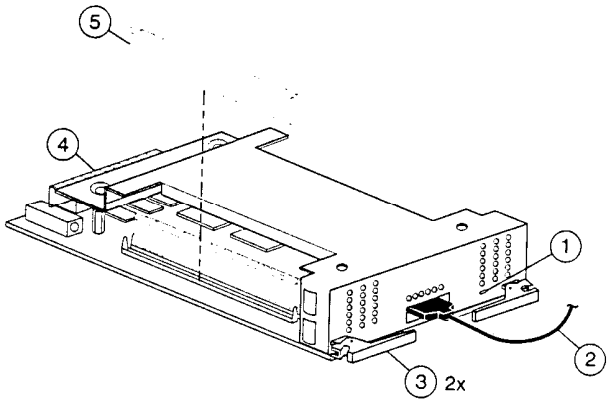

Item	Description	Compaq Part Number
1	Cache-memory power LED	-
2	ECB cable	70-33268-01
3	Retaining lever	-
4	Backplane connector	-
5	SIMM (2 or 4) 64 MByte SIMM pack	54-23391-01 (DS-HSSIM-AB)

1.3 Cache Memory Configurations.

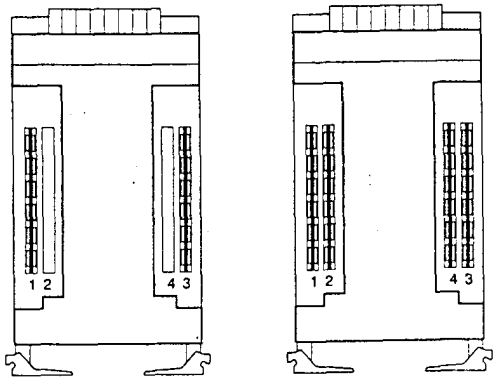

1.4 Inserting SIMMs into the Cache Module.

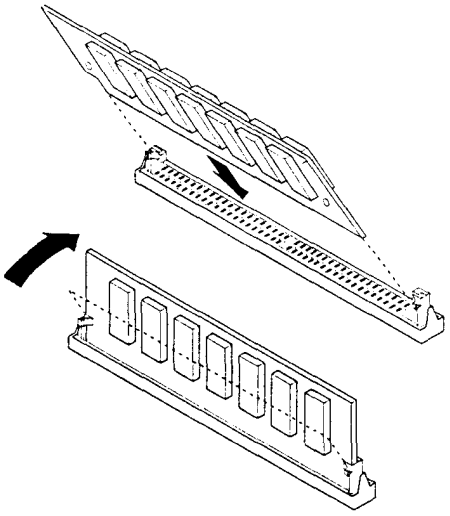

1.5 Installing the Controller and Cache Modules into the Controller Shelf.

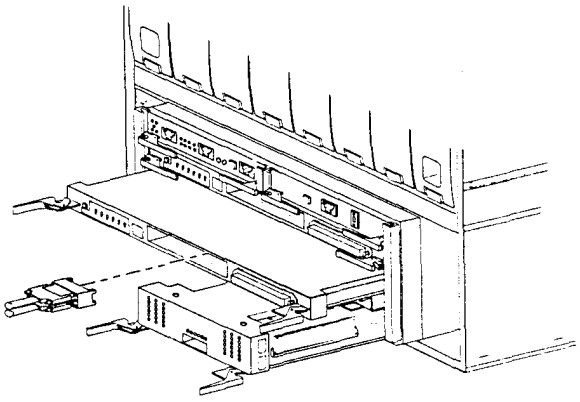

1.6 Connecting cables to an HSZ70 (dual-redundant configuration is shown).

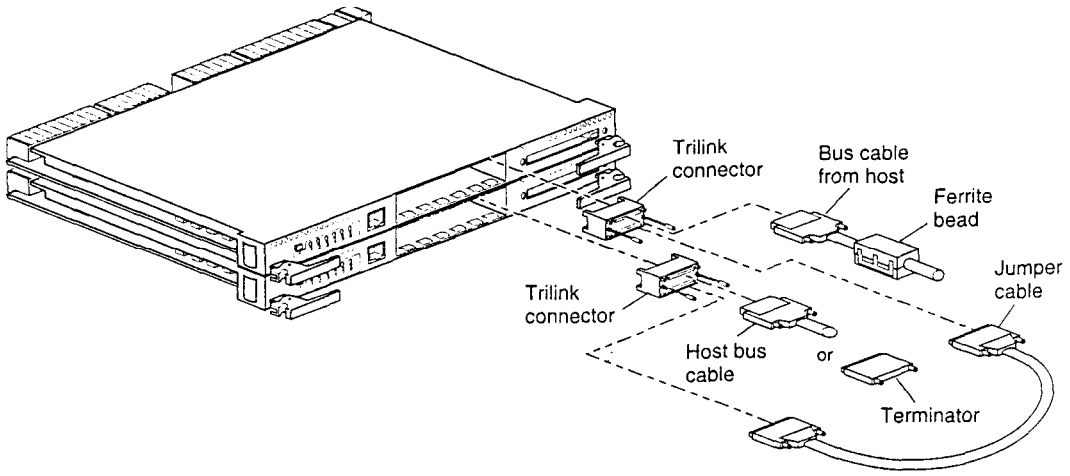

1.7 Connecting an HSZ70 to a VT (dual-redundant configuration shown - the VT must be connected to the controller identified as "This Controller" in configuration).

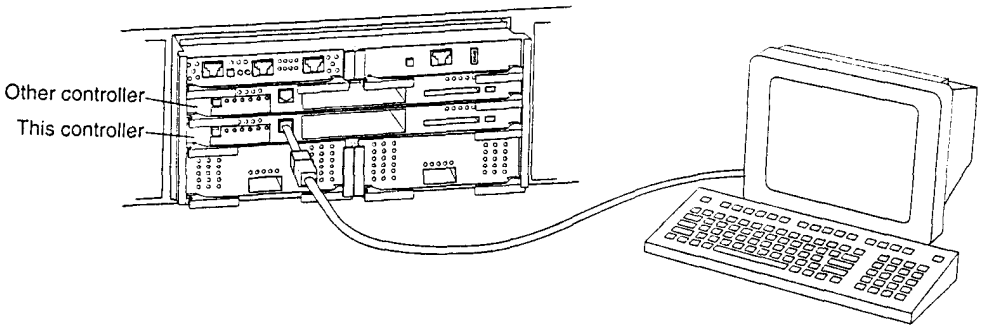

1.8 Detail of Operator Control Panel (OCP). Explanation of the features is provided in Section 6.

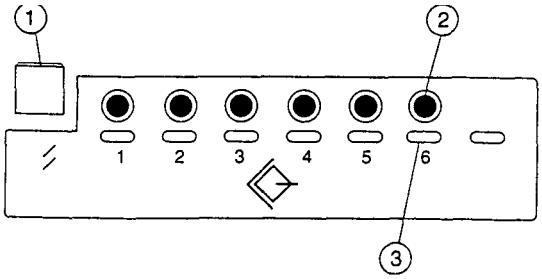

Item	Description
1	Reset button
2	Port Quiesce Buttons
3	Device Port LEDs

2.0 STORAGEWORKS HARDWARE CONFIGURATION.

StorageWorks shelves are used with the HSZ70 RAID controller. The following precautions are recommended to reduce damage to RAID components. Bent pins on a StorageWorks shelf (especially in the sockets where the Personality Module is installed) can permanently damage some or all of the devices on that shelf. ***Bent pins have been known to blow out every disk drive on a shelf****, so read these precautions completely prior to working on a StorageWorks shelf.*

Shelf Inspection: *When installing a Personality Module (PM) in a shelf, use a flashlight to inspect the socket where it will be inserted for any sign of bent pins prior to PM installation. Also inspect the pins where any tape drives will be installed (for the tan DLT tape drives, this will be the pins to the extreme left side of the tape drive, for the blue DLT tape drives, this will be the pins dead center behind the DLT). The pins may only be slightly bent, so move the flashlight around to get reflections from several angles. DO NOT ATTEMPT TO STRAIGHTEN ANY BENT OR RECESSED PINS - REPLACE THE SHELF.*

Personality Modules: *When installing a PM in a DEC StorageWorks shelf, gently slide the PM in until some resistance is met, then apply additional pressure until the catch at the bottom of the PM engages. Give it one last push to ensure that the connector is firmly seated. When removing them, it may be necessary to release the catch at the bottom in order to get the PM out of the slot. This can be done by removing the devices in slots 0 and 1, then slipping a flat-bladed screwdriver under the edge of the catch to release it.*

Disk Drives: *When installing disk drives in DEC StorageWorks shelves, gently slide the disk drive into place until the catches at the top and bottom of the disk drive engage. Give the disk drive one last firm push to ensure that the connector is firmly seated. When power is applied, verify that the green LED comes on for about 30 seconds, then goes off.*

Tape Drives: *When installing tape drives in DEC StorageWorks shelves, gently slide the tape drive into place until the catches at the top and bottom of the tape drive engage. Give the tape drive one last firm push to ensure that the connector is firmly seated. When power is applied, verify that the activity LEDs cycle normally for a few seconds. The LEDs on a 4-mm DAT should all go out (unless a tape is present) and DLTs should have the Operate Lever LED on or the LEDs indicating the size of the installed tape should be lit.*

Connecting Cables: *Inspect all cables prior to attaching them to a DEC StorageWorks shelf or associated controllers/computers. Note any bent or recessed pins. It is advisable to view the connector on the cable from several directions, as not all pin problems appear from one angle only. When attaching cables, ensure that the screws are not cross-threaded and that one side of the connector is not cocked at an extreme angle during attachment (this can also bend or recess pins during cinching down of the screws). DO NOT ATTEMPT TO FIX ANY BENT OR RECESSED PINS - REPLACE THE CABLE.*

2.1 StorageWorks shelf slot designation (slot 7 always has to be a power supply; slot 6 can be a power supply, a disk drive, or a tape drive; slots 0 through 5 can be disk or tape drives).

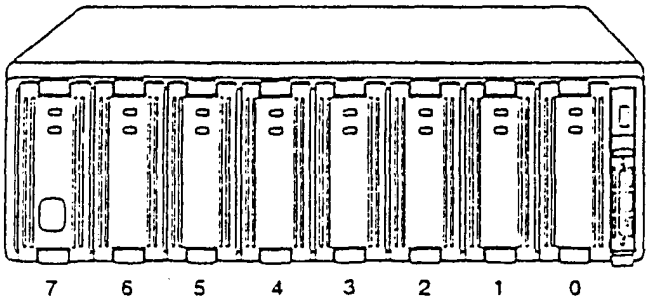

2.2 If tape drives are sharing the shelf with disk drives, the *Shelf OK* jumper will have to be removed to prevent false errors on the power supply LEDs. To access this jumper, Blower 1 must be removed from the rear of the shelf. This can be done as shown in the next two figures. First, the safety screw must be removed, then the mounting tabs must be pressed together to remove Blower 1 as shown.

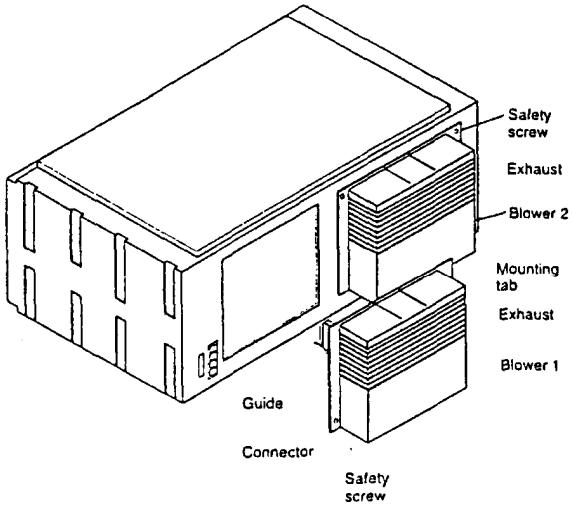

2.3 The *Shelf OK* jumper is the next to the jumper socket behind Slot 0. The jumper must be removed to clear the error condition on the power supplies (it is recommended that the jumper be placed on one jumper pin to keep it available in case it is ever needed).

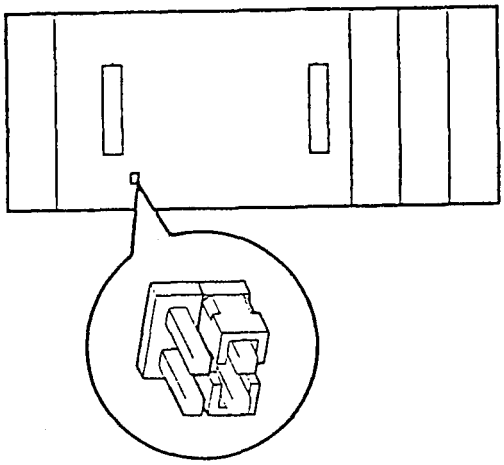

2.4 If this shelf is meant to be used for two cables, it must have a split-backplane terminator installed in the right-side socket (viewed from the rear). There may be an EMI (electro-magnetic interference) shield installed over this socket, which will need to be removed. This is identified in the figure below as item number 2 (the socket is identified as item number 1). The shield may need to be pried off by pushing a screwdriver blade into first the top clip until it pops loose, then the bottom clip.

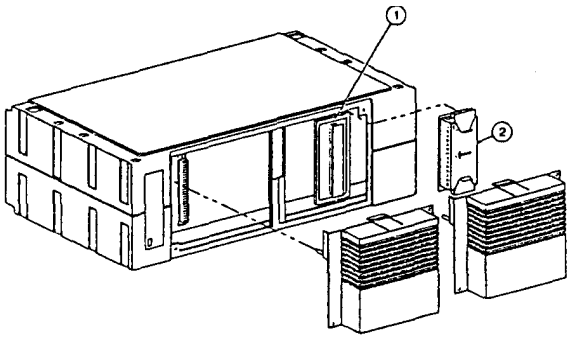

2.5 A terminator can be distinguished from a jumper by the integrated circuit chips soldered to the board. A terminator appears similar to the figure shown below:

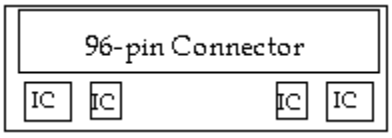

2.6 If this shelf is meant to be used for only one cable, it must have a single-backplane jumper installed in the right-side socket (viewed from the rear). A jumper looks like the figure shown below:

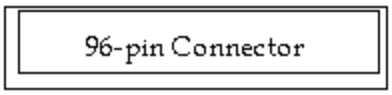

2.7 StorageWorks shelf Power Supply module in StorageWorks Building Block (SBB).

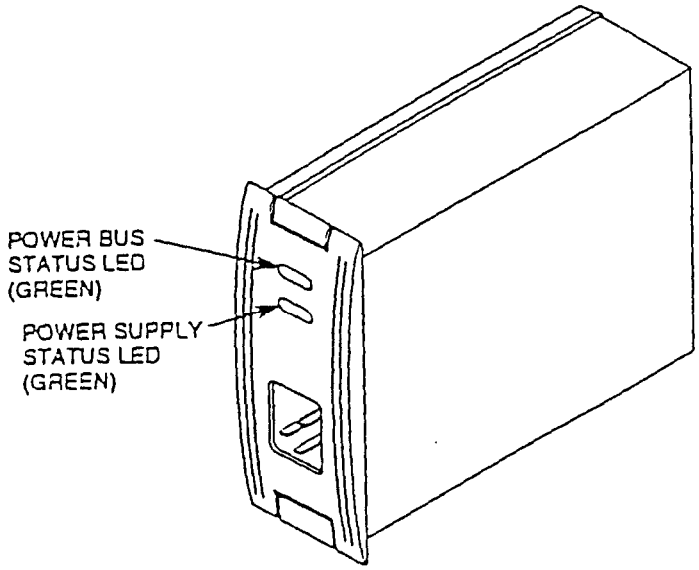

Status LED	State	Indication
Shelf (upper) PS (lower)	On On	Normal operation.
Shelf (upper) PS (lower)	Off Off	Fault status, shelf and power supply fault.
Shelf (upper) PS (lower)	Off On	Fault status, shelf fault but no power supply fault. Replace shelf blower.

2.8 StorageWorks disk drive in SBB (9.1 GByte and 18.2 GByte sizes).

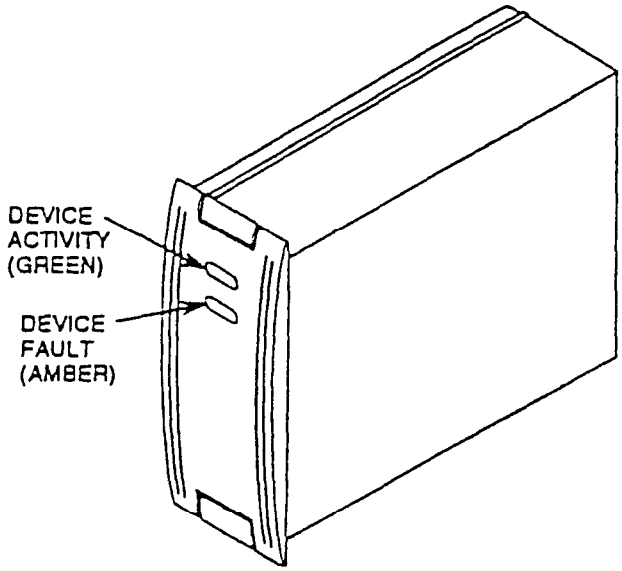

LED	Status	Indication
Device activity Device fault	On Off	SBB is operating normally.
Device activity Device fault	Flashing Off	SBB is operating normally.
Device activity Device fault	Off Off	SBB is operating normally. The SBB is inactive, and there is no fault.
Device activity Device fault	On On	Fault status SBB is probably not responding to control signals. Digital recommends that you replace the SBB.
Device activity Device fault	Off On	Fault status SBB is inactive and spun down. Digital recommends that you replace the SBB.
Device activity Device fault	On Flashing	Fault status SBB is active and is spinning down because of the fault.
Device activity Device fault	Off Flashing	Fault status SBB has been identified by the controller as failed. Digital recommends that you replace the SBB.

2.9 To remove an SBB from a StorageWorks shelf, grasp the top and bottom tabs and squeeze together to release the SBB, then pull the SBB out of the shelf.

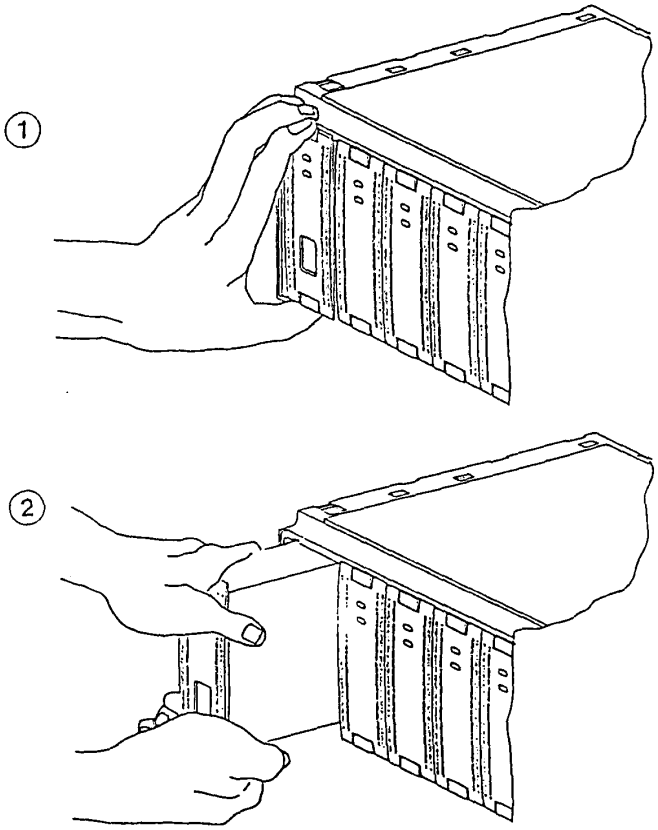

2.10 SCSI addresses are generally assigned by the slot the SCSI device is installed in. If there is the need to change the SCSI ID of a device, the table below can be used to set the DIP switches on the SCSI device to the proper setting.

	Switch Number							
Address	1	2	3	4	5	6	7	8
0	OFF	OFF	OFF	OFF	OFF	OFF	OFF	OFF
1	ON	OFF	OFF	OFF	OFF	OFF	OFF	OFF
2	OFF	ON	OFF	OFF	OFF	OFF	OFF	OFF
3	ON	ON	OFF	OFF	OFF	OFF	OFF	OFF
4	OFF	OFF	ON	OFF	OFF	OFF	OFF	OFF
5	ON	OFF	ON	OFF	OFF	OFF	OFF	OFF
6	OFF	ON	ON	OFF	OFF	OFF	OFF	OFF
7*	ON	ON	ON	OFF	OFF	OFF	OFF	OFF
Automatic†	OFF	OFF	OFF	ON	ON	ON	OFF	OFF

* Normally reserved for the host.

† Default setting: address is defined by the shelf connector.

2.11 Older StorageWorks shelves use the BA35X-MH Personality Module illustrated below for cable connection.

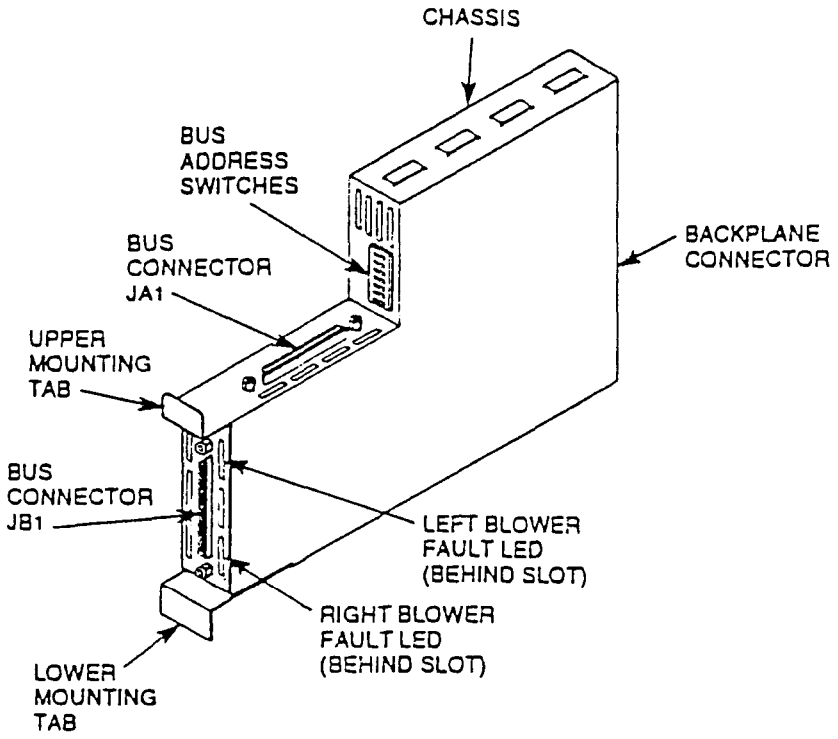

LED	Status	Indication
Blower 2†	Off	Normal status. Blowers are operational.
Blower 1	Off	
Power Supply Shelf Status	On	
Blower 2	On	Overtemp condition. Both blowers to high speed.
Blower 1	On	
Power Supply Shelf Status	On	
Blower 2	On	Blower 2 has failed. Blower 1 to high speed.
Blower 1	Off	
Power Supply Shelf Status	Off	
Blower 2	Off	Blower 1 has failed. Blower 2 to high speed.
Blower 1	On	
Power Supply Shelf Status	Off	
Blower 2	On	
Blower 1	On	Both blowers failed.
Power Supply Shelf Status	Off	

2.12 Newer StorageWorks shelves use the BA35X-FA Personality Module for single-backplane (one connector) configurations and the BA35X-FB Personality Module for split-backplane (two connector) configurations. These are shown below.

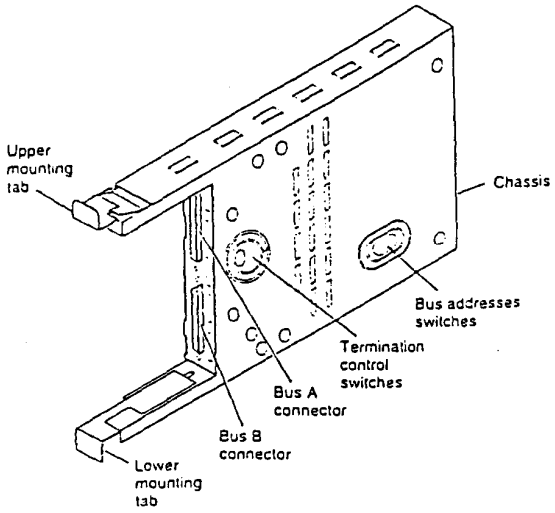

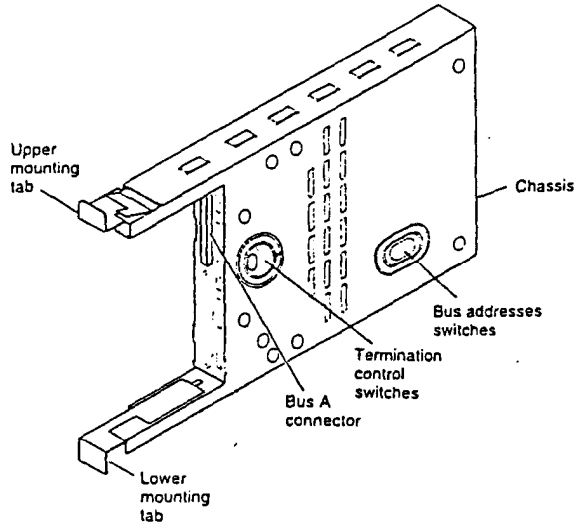

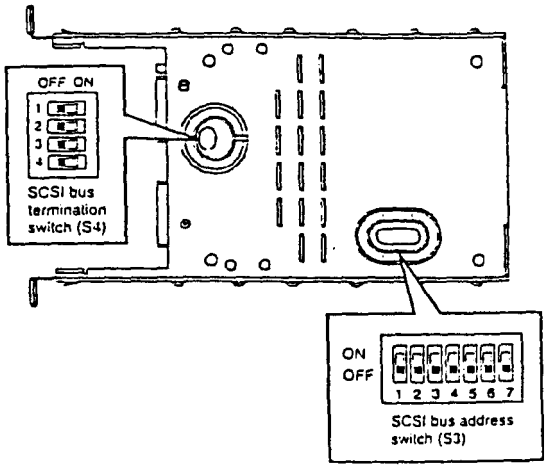

2.13 All HSZ70 controllers come with a backup battery for the cache module. A dual-controller cache battery configuration is shown.

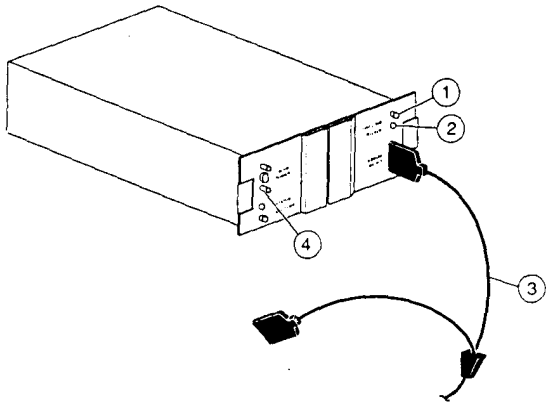

Item	Description	Part No.
	ECB for dual-redundant configurations	DS–HS35X–BD
	ECB for single configurations	DS–HS35X–BC
1	Shut off button	—
2	Status LED	—
3	ECB cable	70–33268–01
4	VHDCI male port for second battery	—

2.14 Gently slide the cache battery module into any unused slot on the StorageWorks shelves as shown here.

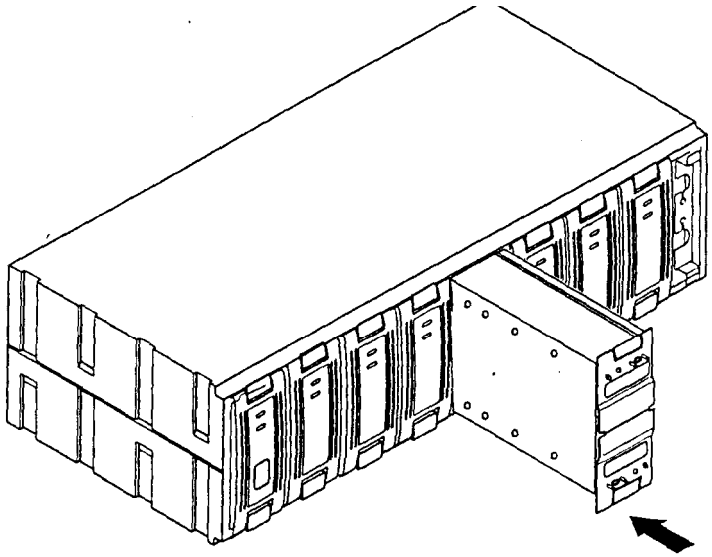

2.15 The cache battery must be shut down manually whenever the RAID is powered down. ***FAILURE TO PROPERLY SHUTDOWN A CACHE BATTERY CAN KILL THAT BATTERY***. The procedure to do this is as follows:

2.15.1 Unmount the RAID volumes with the *umount* command using the format **umount -A** and verify that the RAID volumes have been unmounted by using the **df** command.

2.15.2 If the RAID controller is hooked up to the computer through the serial port, use the ***tip cli*** command to access the RAID controller (see Section 3 for details on how to set this up). If there is no serial connection between the controller and the computer, connect a VT or laptop setup for VT100 emulation to the controller serial port.

2.15.3 Enter the command ***shutdown this*** to halt the controller.

2.15.4 Power down the RAID cabinets, controller shelf included.

2.15.5 Find the small button labeled *SHUT OFF* next to the flashing status LED on the cache battery. Press this battery disable switch and hold it in until the LED stops flashing (watch the LED for 10-15 seconds after releasing the button to confirm that the LED is no longer flashing).

2.15.6 The battery is now disabled and is no longer powering the cache.

2.15.7 To start the battery again, simply power up the RAID cabinets. The cache battery is enabled when the HSZ70 controller is powered on. NOTE: The RAID will not work properly until the battery LED stops flashing. The controller will be accessible, but will only work with a reduced instruction set. If the battery never charges up or is listed as *FAILED*, refer to Section 8 (Troubleshooting) for information on running the RAID without the battery.

2.16 Connect the HSZ70 controller SCSI port to the KZPSA Wide SCSI card inside the computer controlling this RAID with a 68-pin Wide SCSI cable. This will allow the computer to access the RAID drives in Section 5.

3.0 CONNECTIONS TO AN ALPHA COMPUTER.

If this RAID is connected by SCSI cable to a computer running Digital UNIX, connection of a serial cable between the RS-232 port of the computer and the HSZ70 will allow the user to monitor and configure the RAID without the need of a RAID PC or VT monitor. The following steps detail how to accomplish this.

3.1 Connect a serial cable from the COM2 port on the computer to the serial port on the HSZ70. NOTE: There is a cable included with the HSZ70 controller which can be used. If this cable is not available, it can be ordered (DEC P/N 17-04074-01, Printrak P/N 0904-00312).

3.2 Login on the computer attached to the RAID as **root** with a password of **Superuser**

3.3 Click on the ***Applications*** button at the top of the screen.

3.4 Click on ***xterm*** to open an xterm window.

3.5 Enter the command **vi /etc/remote** to access the /etc/remote file.

3.6 If the two lines below do not exist anywhere in this file, go to the last entry (which should be *system2400*) and insert them:
cli:dv=/dev/tty01:pa=none:br#9600:
:el=^D:oe=^D

3.7 Hit the **Esc** key to enter command mode.

3.8 Enter the command **ZZ** (that's shift-zz) to save and exit the file.

3.9 Enter the command **cd /dev** to switch to the /dev directory.

3.10 Hold down the **Ctrl** button and click and hold the right mouse button. Make sure the screen size is set to *Default* (or *Small* for systems with a Raritan switchbox connecting several computers to one monitor). If the screen size is not set properly, scroll down the menu and highlight the correct size.

3.11 Enter the command **tip cli** to activate the utility. The computer will beep once and the message *connected* will be displayed.

3.12 Press the **Enter** key again and note the appearance of the *HSZ>* prompt.

3.13 To confirm functionality of the interface, enter the command **run cfmenu** to display the raid drive status. After perusing the disks to confirm that the controller interface is fully functional, press the **Enter** key to exit the *cfmenu* screen.

3.14 To terminate this utility, press the **~** key (**Shift-`**) followed by **Ctrl-d**

4.0 RAID CONFIGURATION CONSIDERATIONS.

Before the RAID can be configured, the details of raidsets, mirrorsets, hot spares and other factors must be worked out. This section gives some pointers on how to divide up the disk drives on a RAID for maximum integrity and efficiency.

- Spread the disks from a raidset across all of the available channels in order to minimize disk crashes if a channel dies.
- Spread the disks from a raidset across all of the available shelves in order to minimize loss of data if a shelf dies.
- Allocate all of the parity disks and the hot spare to the same channel (preferably a different channel from the one used by the data disks).
- Arrange the disks so that there is an 8 GByte block of usable space allocated for each volume.
- There are a maximum of six channels available with HSZ70 controllers, with a maximum of 7 disks per channel, or 42 disks total.
- It is advisable to always have a hot spare in case a disk in a raidset or mirrorset dies. This will automatically replace the failed disk.
- An example of a 31 disk RAID is given on the next page. This RAID is divided into five raidsets of six disks each, plus a hot spare. Note the formula used to determine total RAID capacity. Even though there is no dedicated parity disk with the RAID 3/5 protocol (the parity being striped across each disk in the raidset), for ease of calculation, one disk has been designated the parity disk for each raidset. There is also a factor of 0.95 used to account for losses due to saved-configuration data stored on each disk. (If mirrorsets are being used, the total amount of capacity of the mirrorset is divided by the number of disks in the mirrorset to find out what the usable space equals).

  31 disks total minus 1 hot spare = 30 disks for raidsets.
  30 disks divided into 5 raidsets composed of 25 data disks and 5 parity disks.
  5 data disks X 18.2 GByte data disks X 0.95 = 86.45 GBytes of usable space per raidset.
  86.45 GBytes X 5 raidsets = 432.25 GBytes total disk space.

  A graphical representation of this configuration is shown on the next page. The five raidsets are R1, R2, R3, R4, and R5. The SCSI ID of each disk is taken from the slot on the StorageWorks shelf in which the disk is placed, so raidset R1 consists of the six drives with SCSI ID "0" which are spread across all six channels. R2 consists of the six disks with SCSI ID "1", R3 consists of the six disks with SCSI ID "2", etc. Note how the raidsets are divided among the shelves and channels. The loss of a single shelf or a single channel may degrade RAID operation, but no data is lost. The hot spare can be put on any shelf, but in this case it was placed on the shelf with all of the parity disks.

Channel 1

P S	P S		R5 DD 10400	R4 DD 10300	R3 DD 10200	R2 DD 10100	R1 DD 10000	P M
	6	5	4	3	2	1	0	

Channel 2

P S	P S		R5 DD 20400	R4 DD 20300	R3 DD 20200	R2 DD 20100	R1 DD 20000	P M
	6	5	4	3	2	1	0	

Channel 3

P S	P S		R5 DD 30400	R4 DD 30300	R3 DD 30200	R2 DD 30100	R1 DD 30000	P M
	6	5	4	3	2	1	0	

Channel 4

P S	P S		R5 DD 40400	R4 DD 40300	R3 DD 40200	R2 DD 40100	R1 DD 40000	P M
	6	5	4	3	2	1	0	

Channel 5

P S	P S		R5 DD 50400	R4 DD 50300	R3 DD 50200	R2 DD 50100	R1 DD 50000	P M
	6	5	4	3	2	1	0	

Channel 6

P S	P S	HSP 60500	R5 PD 60400	R4 PD 60300	R3 PD 60200	R2 PD 60100	R1 PD 60000	P M
	6	5	4	3	2	1	0	

PS = Power Supply
DD = Data Disk
PD = Parity Disk
PM = Personality Module
HSP = Hot Spare

The raidsets in the RAID on the previous page are organized as follows:

R1 is composed of disks Disk10000, Disk20000, Disk30000, Disk40000, Disk50000, and Disk60000
R2 is composed of disks Disk10100, Disk20100, Disk30100, Disk40100, Disk50100, and Disk60100
R3 is composed of disks Disk10200, Disk20200, Disk30200, Disk40200, Disk50200, and Disk60200
R4 is composed of disks Disk10300, Disk20300, Disk30300, Disk40300, Disk50300, and Disk60300
R5 is composed of disks Disk10400, Disk20400, Disk30400, Disk40400, Disk50400, and Disk60400
The Hot Spare is located at Disk60500.

- A variation on the previous RAID configuration is the one shown on the next page. The differences here are that only 18 disks are used, a mirrorset has been added, and split-backplane shelves allow two channels to share each shelf so that only three shelves are needed (since the raidsets now consist of only three disks each). Loss of either the Channel 1/2 or Channel 3/4 shelf will degrade four of the five raidsets (the hot spare on the Channel 5/6 shelf will kick in to replace the lost disk on one of the raidsets), and will result in the mirrorset being degraded, but not destoryed. Loss of the Channel 5/6 shelf will degrade all five raidsets, but will leave the mirrorset fully operational. The loss of any remaining raidset disk will destroy that raidset and force the user to reload the replacement disks from scratch when a new disk and Channel 5/6 shelf are installed. The various capacities of the raidsets and mirrorset are shown below:

Raidset capacity

18 disks total minus 1 hot spare minus 2 mirrorset disks = 15 disks for raidsets.
15 disks divided into 5 raidsets composed of 10 data disks and 5 parity disks.
2 data disks X 18.2 GByte data disks X 0.95 = 34.58 GBytes of usable space per raidset.
34.58 GBytes X 5 raidsets = 172.9 GBytes total disk space in raidsets.

Mirrorset capacity

2 18.2 GByte disks total divided by 2 (mirrorset division factor) X 0.95 = 17.29 GBytes total

The disks in the RAID are organized as follows (a graphical representation of this configuration appears on the next page):

M1 (mirrorset #1) is composed of disks Disk20500 and Disk40500
R1 is composed of disks Disk10000, Disk30000, and Disk50000
R2 is composed of disks Disk20100, Disk40100, and Disk60100
R3 is composed of disks Disk10200, Disk30200, and Disk50200
R4 is composed of disks Disk20300, Disk40300, and Disk60300
R5 is composed of disks Disk10400, Disk30400, and Disk50400
The Hot Spare is Disk60500.

Channel 1/2

P S	P S	M1 20500	R5 DD 10400	R4 DD 20300	R3 DD 10200	R2 DD 20100	R1 DD 10000	P M
	6	5	4	3	2	1	0	

Channel 3/4

P S	P S	M1 40500	R5 DD 30400	R4 DD 40300	R3 DD 30200	R2 DD 40100	R1 DD 30000	P M
	6	5	4	3	2	1	0	

Channel 5/6

P S	P S	HSP 60500	R5 PD 50400	R4 PD 60300	R3 PD 50200	R2 PD 60100	R1 PD 50000	P M
	6	5	4	3	2	1	0	

PS = Power Supply
DD = Data Disk
PD = Parity Disk
PM = Personality Module
HSP = Hot Spare

5.0 RAID CONFIGURATION.

Once the RAID hardware has been assembled, disk configuration can begin. This section details how to arrange and initialize the various raidsets, mirrosets, and hot spares.

5.1 Attach the HSZ70 controller to either a Video Terminal (VT) or to the computer per the instructions in Section 3.

5.2 Power up the DEC StorageWorks RAID disk drive shelves.

5.3 Power up the HSZ70 controller shelf.

5.4 Check the green LED on the cache battery module and confirm that it is not flashing. If it is, the cache battery is still charging and the RAID will not work properly until this is complete. Once the light stops flashing, move to the VT and, at the *HSZ*> prompt, enter the command **sho this** to display the controller status, which should appear similar to the one shown below. Examine the display and confirm that the entry ***Battery is GOOD*** appears. This must appear before continuing with the configuration procedure. If the message ***Battery has FAILED*** appears, the cache battery must be replaced with a good unit. If this RAID needs to be brought up despite having a bad cache battery, refer to Section 8 (Troubleshooting) Problem #5 for two options in dealing with this problem until a good cache battery can be procured for this RAID.

Controller:
HSZ70 2G93904673 Firmware V73Z-0, Hardware H04
Not configured for dual-redundancy
Device Port SCSI address 7
Time: NOT SET
Host port:
SCSI target (s) (0, 1, 2, 3)
Preferred target (s) (0, 1, 2, 3)
TRANFSER_RATE_REQUESTED = 20 MHZ
Host Functionality Mode = A
Allocation Class 0
Command Console LUN is disabled

Cache:
64 megabyte write cache, version 4
Cache is GOOD
Battery is GOOD
No unflushed data in cache
CACHE_FLUSH_TIMER = DEFAULT (10 seconds)
NOCACHE_UPS

Mirrored Cache:
Not enabled

5.5 Enter the command **set this time=*dd-mmm-yyyy:hh:mm:ss*** to set the time on this controller (for dual controllers, repeat this command substituting *other* for *this*).

where:
dd is the day
mmm is the month
yyyy is the year
hh is the hour
mm is the minute
ss is the second

As an example, to set the time for November 1st, 1999 at 10:30 AM, the command is **set this time=01-Nov-1999:10:30:00**

5.6 If this system has two HSZ70 controllers in a dual-redundant configuration, skip to step 5.9. Otherwise, proceed to step 5.7.

5.7 Confirm that the *SCSI target(s)* under the *Host port* heading are set as shown in step 5.4, with SCSI IDs of 0, 1, 2, and 3. If not, from the *HSZ>* prompt enter the command **set this id=(0,1,2,3)** to setup the correct SCSI allocation for the HSZ70, then enter the command **restart this** for the changes to take effect.

5.8 Confirm that the *Preferred target(s)* under the *Host port* heading are set as shown in step 5.4, with SCSI IDs of 0, 1, 2, and 3. If they are, skip to step 5.16. If not, enter the command **set this preferred=(0,1,2,3)** to set the preferred SCSI busses, then skip to step 5.15.

5.9 Enter the command **set failover copy = this** to activate dual-redundancy.

5.10 Enter the command **set this id=(0,1,2,3)** to allow the main controller to access all the SCSI devices.

5.11 Enter the command **set other id=(0,1,2,3)** to allow the secondary controller to access all the SCSI devices in case the primary controller fails.

5.12 Enter the command **set this preferred=(0,1)** to channel half the SCSI traffic through the main controller.

5.13 Enter the command **set other preferred=(2,3)** to channel the other half of the SCSI traffic through the secondary controller.

5.14 Enter the command **restart other** to restart the secondary controller so the changes take effect.

5.15 Enter the command **restart this** to restart the primary controller so the changes take effect.

5.16 From the *HSZ>* prompt enter the command **run cfmenu** to activate the configuration GUI. If this is a previously-configured RAID, a screen similar to the one pictured on the following page will appear:

```
---------------------------------------------------------------- CFMENU  Configuration Menu Utility -------------------------------------------------------------------
MAIN MENU:                  |Unconfig'd |   Config'd     Device        Product      Stor. set Stor. set Chnk  Trn   In-  Re-         P W W
1. Add/delete devices       | Dev. PTLs |    PTLs        Name            ID          Name    Typ/Sz   size  sp.  it'd duc  Unit    T  P  B
2. Add/delete mirrorsets    |---------- |   -------   -----------  ---------------  -------- -------- ----- ----- ---- ---- -----   -- -- --
3. Add/delete stripesets    |           |raid5:10000  DISK10000  BB01811C9C          R1     RAD/6    64          Y    N    D0       N  Y
4. Add/delete raidsets/     |           |      20100  DISK20100  BB01811C9C          "       "       "           "    "    "        "  "
   sparesets/failedsets     |           |      30000  DISK30000  BB01811C9C          "       "       "           "    "    "        "  "
5. Partition Processing     |           |      40100  DISK40100  BB01811C9C          "       "       "           "    "    "        "  "
6. Initialize device        |           |      50000  DISK50000  BB01811C9C          "       "       "           "    "    "        "  "
   and/or storagesets       |           |      60000  DISK60000  BB01811C9C          "       "       "           "    "    "        "  "
7. Add/delete units         |           |      10200  DISK10200  BB01811C9C          R2     RAD/6    64          Y    N    D1       N  Y
8. Setup terminal           |           |      20300  DISK20300  BB01811C9C          "       "       "           "    "    "        "  "
9. Exit CFMENU              |           |      30200  DISK30200  BB01811C9C          "       "       "           "    "    "        "  "
                            |           |      40300  DISK40300  BB01811C9C          "       "       "           "    "    "        "  "
                            |           |      50100  DISK50100  BB01811C9C          "       "       "           "    "    "        "  "
                            |           |      60100  DISK60100  BB01811C9C          "       "       "           "    "    "        "  "
D=Scroll down U=Scroll up   |           |      10400  DISK10400  BB01811C9C          R3     RAD/6    64          Y    N    D2       N  Y
```

If there are too many disks to appear on one screen, type in a **d** and hit the **Enter** key to scroll down to the next screen. A sample screen is shown below:

```
---------------------------------------------------------------- CFMENU  Configuration Menu Utility -------------------------------------------------------------------
MAIN MENU:                  |Unconfig'd |   Config'd     Device        Product      Stor. set Stor. set Chnk  Trn   In-  Re-         P W W
1. Add/delete devices       | Dev. PTLs |    PTLs        Name            ID          Name    Typ/Sz   size  sp.  it'd duc  Unit    T  P  B
2. Add/delete mirrorsets    |---------- |   -------   -----------  ---------------  -------- -------- ----- ----- ---- ---- -----   -- -- --
3. Add/delete stripesets    |           |raid5:10200  DISK10200  BB01811C9C          R2     RAD/6    64          Y    N    D1       N  Y
4. Add/delete raidsets/     |           |      20300  DISK20300  BB01811C9C          "       "       "           "    "    "        "  "
   sparesets/failedsets     |           |      30200  DISK30200  BB01811C9C          "       "       "           "    "    "        "  "
5. Partition Processing     |           |      40300  DISK40300  BB01811C9C          "       "       "           "    "    "        "  "
6. Initialize device        |           |      50100  DISK50100  BB01811C9C          "       "       "           "    "    "        "  "
   and/or storagesets       |           |      60100  DISK60100  BB01811C9C          "       "       "           "    "    "        "  "
7. Add/delete units         |           |      10400  DISK10400  BB01811C9C          R3     RAD/6    64          Y    N    D2       N  Y
8. Setup terminal           |           |      20500  DISK20500  BB01811C9C          "       "       "           "    "    "        "  "
9. Exit CFMENU              |           |      30400  DISK30400  BB01811C9C          "       "       "           "    "    "        "  "
                            |           |      40500  DISK40500  BB01811C9C          "       "       "           "    "    "        "  "
                            |           |      50200  DISK50200  BB01811C9C          "       "       "           "    "    "        "  "
                            |           |      60200  DISK60200  BB01811C9C          "       "       "           "    "    "        "  "
D=Scroll down U=Scroll up   |           |spare:60300  DISK60300  BB01811C9C
```

To get back to the first screen, type in a **u** and hit the **Enter** key. See Section 10 for an explanation of each of the categories.

If this RAID was never configured, it will appear similar to the following:

```
---------------------------------------------------------------- CFMENU  Configuration Menu Utility -------------------------------------------------------------------
MAIN MENU:                  | Unconfig'd |   Config'd  Device       Product         Stor. set Stor. set Chnk  Trn   In-  Re-         P W W
1. Add/delete devices       |  Dev. PTLs |     PTLs    Name           ID             Name    Typ/Sz   size  sp.  it'd duc  Unit    T  P  B
2. Add/delete mirrorsets    | ---------- |   -------- --------- ------------------  -------- -------- ----- ----- ---- ---- ----- -- -- --
3. Add/delete stripesets    | DISK10200  |
4. Add/delete raidsets/     | DISK20300  |
   sparesets/failedsets     | DISK20100  |
5. Partition Processing     | DISK40300  |
6. Initialize device        | DISK50100  |
   and/or storagesets       | DISK60100  |
7. Add/delete units         | DISK10400  |
8. Setup terminal           | DISK20500  |
9. Exit CFMENU              | DISK30400  |
                            | DISK40500  |
                            | DISK50200  |
                            | DISK60200  |
D=Scroll down U=Scroll up   | DISK60300  |
```

5.17 When the menu on the previous page appears, verify that all of the disk drives physically located in the shelves show up on the GUI at the correct position and reflect the correct type of disk. Enter **d** to scroll down a screen and **u** to scroll up if some of the disks fail to appear.

5.18 If this RAID has never been configured (there are no existing raidsets, mirrorsets, hot spares or devices on this RAID which need to be deleted) enter a **1** (*Add/delete devices*) and proceed to step 5.35. If the current configuration needs to be changed, continue on to step 5.19.

5.19 If there are no units on this RAID which need to be deleted (these are identified in the *Unit* field on the *cfmenu* screen as *D0, D1,* and *D2*), proceed to step 5.24. Otherwise, enter **7** (*Add/delete units)* to proceed to the *Add/delete units* menu.

5.20 Enter **2** (*Delete a unit (eligible units marked by *))* to select units to delete.

5.21 Enter **y** at the prompt *Delete unit D0 (y/n/q) [n]?*

5.22 Repeat step 5.21 for any remaining units which need to be deleted.

5.23 Hit the **Enter** key to return to the MAIN MENU.

5.24 If there are no raidsets on this RAID which need to be deleted (these are identified in the *Stor. set Name* field on the *cfmenu* screen as *R1, R2,* and *R3*) proceed to step 5.30. Otherwise, enter **4** (*Add/delete raidsets/sparesets/failedsets)*

5.25 Enter **3** (*Delete all unbounded raidsets (marked by *))*. This will delete only the raidsets which do not have a unit number, which preserves any existing raidsets if these new disks are being added to an existing RAID.

5.26 If the hot spare is not going to be changed, proceed to step 5.30. Otherwise, enter **4** (*Add/delete device in SPARESET or FAILEDSET (submenu)).*

5.27 Enter **2** (*Remove a device from the SPARESET).*

5.28 Enter **y** when prompted *Remove device DISK... from SPARESET (y/n/q) [n]?*

5.29 Hit the **Enter** key **twice** to return to the MAIN MENU.

5.30 If there are no mirrorsets which need to be deleted, go to step 5.33. Otherwise, enter **2** *(Add/delete mirrorsets).*

5.31 Enter **3** *(Delete all unbounded mirrorsets (marked by *))* to delete all mirrorsets which were removed from units in step 5.21 and 5.22.

5.32 Hit the **Enter** key to return to the MAIN MENU.

5.33 Enter **1** (*Add/delete devices)*

5.34 Enter **4** (*Delete all unbounded devices (marked with *)).* At this point, all of the disks which do not still belong to raidsets or mirrorsets will appear in the *Unconfig'd Dev. PTLs* column. This will also delete devices which no longer physically reside on the RAID (such as disks which have been removed and no replacement disk was installed in that slot - a useful means of clearing empty slots off the display).

5.35 If there are no disks in the *Unconfig'd Dev PTLs* column, go to step 5.38. Otherwise, enter **3** (*Add all devices from list of PTLs not configured (marked with ^))* to begin adding devices listed in the *Unconfig'd Dev. PTLs* column back into the *Config'd PTLs* column so they can be built into raidsets and mirrorsets.

5.36 Hit the **Return** key when the prompt *Disks will all be set NOTRANSPORTABLE. Continue (y/n) [y]?* appears.

5.37 When all the disks have been moved into the *Config'd PTLs* column, hit the **Return** key to return to the Main Menu.

5.38 If there are no raidsets or sparesets, skip to step 5.54. Otherwise, enter **4** (*Add/delete raidsets/sparesets/failedsets*)

5.39 Enter **1** (*Create a raidset (elgible devices marked by a* ^))

5.40 Enter the number of disks for your raidset (based on what was decided in Section 4 or on the diagram for this RAID) at the prompt *Enter number of members for the raidset (3...6, or 2(only if reduced raidset)) []?* As an example, a raidset consisting of 6 disks (disk10000, disk20000, disk30000, disk40000, disk50000, and disk60000) will be used.

5.41 Enter **y** at the prompt *Include DISK10000 as a member (y/n/q) [n]?* to add this disk to the raidset.

5.42 Enter **n** at the prompt *Include DISK10100 as a member (y/n/q) [n]?* to skip this disk. Repeat this step until *Include DISK20000 as a member (y/n/q) [n]?* appears, then enter **y** to select this disk.

5.43 Repeat steps 5.41 and 5.42 to select DISK30000, DISK40000, DISK50000, and DISK60000.

5.44 Hit the **Enter** key at the prompt *Enter policy for auto-replacement of failed member (0 = BEST_PERFORMANCE, 1 = BEST_FIT, 2 = NOPOLICY (no auto-replacement)) [BEST_PERFORMANCE]?*

5.45 Hit the **Enter** key at the prompt *Enter priority for reconstruct (0 = NORMAL, 1 = FAST) [NORMAL]?*

5.46 Hit the **Enter** key at the prompt *Is this a previously configured REDUCED raidset (y/n) [n]?*

5.47 The following message will appear: *add raidset R1 POLICY=BEST_PERFORMANCE RECONSTRUCT=NORMAL DISK10000 DISK20000 DISK30000 DISK40000 DISK50000 DISK60000*. A warning message will appear if more than one disk from a channel is used for a raidset (Warning 3000). This warning can be ignored (press the **Enter** key to clear it).

5.48 Repeat steps 5.40 through 5.47 for any remaining raidsets which need to be configured.

5.49 If there is no hot spare being configured on this system, when all of the raidsets have been configured hit the **Enter** key to return to the MAIN MENU and skip to step 5.54. If there is a hot spare, proceed to step 5.50.

5.50 Enter **4** *(Add/delete device in SPARESET or FAILEDSET (submenu)).*

5.51 Enter **1** *(Add a device to the SPARESET (eligible devices marked by ^)).*

5.52 Enter **y** at the prompt *Add device DISK60500 to SPARESET (y/n/q)[n]?* (as an example, DISK60500 is used here).

5.53 Hit the **Enter** key twice to return to the MAIN MENU.

5.54 If there are no mirrorsets to be configured on this system, skip to step 5.60. Otherwise, enter **2** *(Add/delete mirrorsets)*.

5.55 Enter **1** *(Create a mirrorset (eligible devices marked by ^))*

5.56 Enter **2** when prompted *Enter number of members for the mirroset (1...) []?* As an example, a mirrorset consisting of disk50200 and disk60000 would be configured as shown in steps 5.57 and 5.58.

5.57 Enter **y** at the prompt *Include DISK50200 as a member (y/n/q) [n]?* to add this disk to the mirrorset.

5.58 Enter **y** at the prompt *Include DISK60000 as a member (y/n/q) [n]?* to add this disk to the mirrorset.

5.59 Hit the **Enter** key to return to the MAIN MENU.

5.60 Enter **6** (*Initialize devices and/or storagesets*).

5.61 Enter **1** (*initialize a device or storageset (eligible entities marked with ^)).*

5.62 Enter **y** at the prompt *Initialize raidset R1 (y/n/q) [n]?*

5.63 Enter **64** at the prompt *Enter chunksize (16...512) [DEFAULT = optimal based on characteristics of members]?* (skip this step for mirrorsets).

5.64 Enter **y** at the prompt *Initialize this container with metadata space reserved on each device for SAVE_CONFIGURATION info (y,n) [n]?*

5.65 Repeat steps 5.55 through 5.64 for any remaining mirrorsets (M2, M3, M4, etc.)

5.66 Hit the **Enter** key to return to the MAIN MENU.

5.67 Enter **7** (*Add/delete units).*

5.68 Enter **1** (*Create a unit (eligible entities marked by ^)).*

5.69 Enter **y** at the prompt *Create unit for raidset R1 (y/n/q) [n]?* (or M1).

5.70 Enter **0** at the prompt *Enter unit number for new unit:*

5.71 Hit the **Enter** key at the prompt *Set WRITE_PROTECT on unit (y/n) [n]?*

5.72 Hit the **Enter** key at the prompt *Set RUN on unit (answering NO disables the drives' ability to be spun up) (y/n) [y]?*

5.73 Enter **y** at the prompt *Enable WRITEBACK cache for this unit (y/n) [n]?*

5.74 If the SCSI IDs for this controller were not properly setup in step 5.7 through 5.15, an error message similar to the following appears: *Error 1140: Invalid unit number. Valid number ranges are 100 to 107. Press return to continue.* If this does happen, you must exit the cfmenu and carefully repeat step 5.7 through 5.15 prior to executing step 5.67.

5.75 Hit the **Enter** key when a warning message appears about WRITEBACK_CACHE. This message can be ignored. If an error about this being an *Invalid unit number* appears, the SCSI ID for the controller needs to be set per step 5.5.

5.76 The drives in raidset R1 will now be formatted. This process will take several hours to complete, during which time the lights on the drives will flicker constantly.

5.77 Repeat steps 5.68 through 5.76 for the remaining raidsets/mirrorsets (R2, R3, R4, M2, M3, etc.) NOTE: The unit number for R2 will be "1", for R3 will be "2", for R4 will be "3", etc.

5.78 Hit the **Enter** key **twice** to exit *cfmenu*.

5.79 The level of completion of the rebuild can be checked at any time by entering the command **sho r1** (or r2, r3, r4, etc.) The value listed after *RECONSTRUCT* will show how far the formatting has progressed. **DO NOT POWER DOWN THE**

RAID AT ANY TIME DURING THIS PROCESS OR IT WILL HAVE TO BE COMPLETELY REFORMATTED.

5.80 Once the disk lights stop flickering, the status of each raidset and mirrorset can be determined by entering the command **sho r1** for raidset R1 (repeat for other raidsets) and **sho m1** for mirrorset M1 (repeat for remaining mirrorsets). If a screen similar to the one below appears, raidset/mirrorset formatting is complete:

```
Name              Storageset          Uses               Used by
------------------------------------------------------------------------------------------
R1                raidset             DISK10000          D0
                                      DISK20100
                                      DISK30000
                                      DISK40100
                                      DISK50000
                                      DISK60000
        Switches:
         POLICY (for replacement) = BEST_PERFORMANCE
         RECONSTRUCT (priority) = NORMAL
         CHUNKSIZE = 64 blocks
        State:
         NORMAL
         DISK10000    (member  0)  is NORMAL
         DISK20100    (member  1)  is NORMAL
         DISK30000    (member  2)  is NORMAL
         DISK40100    (member  3)  is NORMAL
         DISK50000    (member  4)  is NORMAL
         DISK60000    (member  5)  is NORMAL
        Size: 177736020 blocks
```

5.81 To check on the full configuration, enter the command **run cfmenu** and check the configuration of the RAID. This should be similar to the following:

CFMENU Configuration Menu Utility

MAIN MENU:	Unconfig'd Dev. PTLs	Config'd PTLs	Device Name	Product ID	Stor. set Name	Stor. set Typ/Sz	Chnk size	Trn sp.	In-it'd	Re-duc	Unit	P T	W P	W B
1. Add/delete devices														
2. Add/delete mirrorsets														
3. Add/delete stripesets		miror: 10200	DISK10200	BB01811C9C	R2	RAD/6	64		Y	N	D1		N	Y
4. Add/delete raidsets/		20300	DISK20300	BB01811C9C	"	"	"		"	"	"		"	"
sparesets/failedsets		raid5: 30200	DISK30200	BB01811C9C	"	"	"		"	"	"		"	"
5. Partition Processing		40300	DISK40300	BB01811C9C	"	"	"		"	"	"		"	"
6. Initialize device		50100	DISK50100	BB01811C9C	"	"	"		"	"	"		"	"
and/or storagesets		60100	DISK60100	BB01811C9C	"	"	"		"	"	"		"	"
7. Add/delete units		10400	DISK10400	BB01811C9C	R3	RAD/6	64		Y	N	D2		N	Y
8. Setup terminal		20500	DISK20500	BB01811C9C	"	"	"		"	"	"		"	"
9. Exit CFMENU		30400	DISK30400	BB01811C9C	"	"	"		"	"	"		"	"
		40500	DISK40500	BB01811C9C	"	"	"		"	"	"		"	"
		50200	DISK50200	BB01811C9C	"	"	"		"	"	"		"	"
		60200	DISK60200	BB01811C9C	"	"	"		"	"	"		"	"
D=Scroll down U=Scroll up		spare: 60300	DISK60300	BB01811C9C										

5.82 Disk formatting is now complete. Refer to Engineering procedures for RAID disk partitioning, setting up RAID volumes, and Sybase installation.

6.0 INTERPRETING CONTROLLER LED CODES.

If the HSZ70 controller experiences any hardware failures, these are displayed on the six amber LEDs along with the green pushbutton which make up the operator control panel (OCP) LEDs on the front of the HSZ70. The following section details how to interpret these LEDs.

- During normal operation, the green reset LED on each controller flashes once per second, and the device bus LEDs are not lit.
- The amber LED for a device bus lights continuously when the installed devices do not match the controller configuration, or when a device fault occurs.
- The green reset LED lights continuously and the amber LEDs display a code when a controller problem occurs. Solid LED codes indicate a fault detected by internal diagnostic and initialization routines. Flashing LED codes indicate a fault that occurred during core diagnostics.

The symbols used in the tables below conform to the following standard:

0 = LED Off
1 = LED On
* = LED Flashing

Code	Description of Error	Corrective Action
1 * * * 0 * 0	There was an unexpected interrupt from a read cache or the present and lock bits are not working properly.	Replace controller module.
1 * * * 0 * *	There is an interrupt pending on the controller's policy processor when there should be none.	Replace controller module.
1 * * * * 0 0	There was an unexpected fault during initialization.	Replace controller module.
1 * * * * 0 *	There was an unexpected maskable interrupt received during initialization.	Replace controller module.
1 * * * * * 0	There was an unexpected nonmaskable interrupt received during initialization.	Replace controller module.
1 * * * * * *	An illegal process was activated during initialization	Replace controller module.
Code	**Description of Error**	**Corrective Action**

1 0 * 0 0 * *	The controller DRAB chip failed to detect forced parity, or detected parity when not forced.	Replace controller module.
1 0 * 0 * 0 0	The controller DRAB chip failed to verify the EDC correctly.	Replace controller module.
1 0 * 0 * 0 *	The controller DRAB chip failed to report forced ECC.	Replace controller module.
1 0 * 0 * * 0	The controller DRAB chip failed some operation in the reporting, validating, and testing of the multibit ECC memory error.	Replace controller module.
1 0 * 0 * * *	The controller DRAB chip failed some operation in the reporting, validating, and testing of the multiple single-bit ECC memory error.	Replace controller module.
1 0 * * 0 0 0	The controller main memory did not write correctly in one or more sized memory transfers.	Replace controller module.
1 0 * * 0 0 *	The controller did not cause an I-to-N bus timeout when accessing a "reset" host port chip.	Replace controller module.
1 0 * * 0 * 0	The controller DRAB chip did not report an I-to-N bus timeout when accessing a "reset" host port chip.	Replace controller module.
1 0 * * 0 * *	The controller DRAB chip did not interrupt the controller processor when expected.	Replace controller module.
1 0 * * * 0 0	The controller DRAB chip did not report an NXM error when nonexistent memory was accessed.	Replace controller module.
1 0 0 0 0 0 *	Program card EDC error.	Replace program card.
1 0 0 0 * 0 0	Timer zero in the timer chip will run when disabled.	Replace controller module.

Code	Description of Error	Corrective Action
1 0 0 0 * 0 *	Timer zero in the timer chip	Replace controller module.

	decrements incorrectly.	
1 0 0 0 * * 0	Timer zero in the timer chip did not interrupt the processor when requested.	Replace controller module.
1 0 0 0 * * *	Timer one in the timer chip decrements incorrectly.	Replace controller module.
1 0 0 * 0 0 0	Timer one in the timer chip did not interrupt the processor when requested.	Replace controller module.
1 0 0 * 0 0 *	Timer two in the timer chip decrements incorrectly.	Replace controller module.
1 0 0 * 0 * 0	Timer two in the timer chip did not interrupt the processor when requested.	Replace controller module.
1 0 0 * 0 * *	Memory failure in the I/D cache.	Replace controller module.
1 0 0 * * 0 0	No hit or miss to the I/D cache when expected.	Replace controller module.
1 0 0 * * 0 *	One or more bits in the diagnostic registers did not match the expected reset value.	Replace controller module.
1 0 0 * * * 0	Memory error in the nonvolatile journal SRAM.	Replace controller module.
1 0 0 * * * *	Wrong image seen on program card.	Replace program card.
1 0 * 0 0 0 0	At least one register in the controller DRAB does not read as written.	Replace controller module.
1 0 * 0 0 0 *	Main memory is fragmented into too many sections for the number of entries in the good memory list.	Replace controller module.
1 0 * 0 0 * 0	The controller DRAB chip does not arbitrate correctly.	Replace controller module.
1 1 1 1 1 1 1	DAEMON hard error.	Replace controller module.
1 1 1 1 1 1 0	Repeated firmware bugcheck.	Replace controller module.
1 1 1 1 1 0 1	NVMEM version mismatch.	Replace program card with later version of firmware.
1 1 1 1 1 0 0	NVMEM write error.	Replace controller module.
1 1 1 1 0 1 1	NVMEM read error.	Replace controller module.
Code	**Description of Error**	**Corrective Action**
1 1 1 1 0 1 0	NMI error within firmware bugcheck.	Reset the controller.
1 1 1 1 0 0 1	Inconsistent NVMEM	Reset the controller.

	structures repaired.	
1 1 1 1 0 0 0	Bugcheck with no restart.	Reset controller module.
1 1 1 0 1 1 1	Firmware induced restart following bugcheck failed to occur.	Replace controller module.
1 1 1 0 1 1 0	Hardware induced restart following bugcheck failed to occur.	Replace controller module.
1 1 1 0 1 0 1	Bugcheck within bugcheck controller.	Reset controller module.
1 1 1 0 0 1 1	NVMEM version is too low.	Verify the card is the lateset revision. If the problem still exists, replace the controller module.
1 1 1 0 0 1 0	Program card write fail.	Replace the card.
1 1 1 0 0 0 1	ILF, INIT unable to allocate memory.	Reset the controller.
1 1 1 0 0 0 0	Bugcheck before subsystem initialization completed.	Reset the controller.
1 0 0 0 0 0 0	No program card seen.	Try the card in another controller module. If the problem follows the card, replace the card. Otherwise, replace the controller module.
1 0 * * * 0 *	The controller DRAB chip did not report an address parity error when one was forced.	Replace controller module.
1 0 * * * * 0	There was an unexpected nonmaskable interrupt from the controller DRAB chip during the DRAB memory test.	Replace controller module.
1 0 * * * * *	Diagnostic register indicates there is no cache module, but an interrupt exists from the non-existent cache module.	Replace controller module.

Code	Description of Error	Corrective Action
1 * 0 0 0 0 0	The required amount of memory available for the code image to be loaded from the program card is	Replace controller module.

	insufficient.	
1 * 0 0 0 0 *	The required amount of memory available in the pool area is insufficient for the controller to run.	Replace controller module.
1 * 0 0 0 * *	The required amount of memory available in the buffer area is insufficient for the controller to run.	Replace controller module.
1 * 0 0 * 0 0	The code image was not the same as the image on the card after the contents were copied to memory.	Replace controller module.
1 * 0 0 * 0 *	Diagnostic register indicates that the cache module does not exist, but access to that cache module caused an error.	Replace controller shelf backplane.
1 * 0 0 * * 0	Diagnostic register indicates that the cache module does not exist, but access to that cache module did not cause an error.	Replace controller shelf backplane.
1 * * 0 0 0 0	The journal SRAM battery is bad.	Replace controller module.
1 * 0 0 0 * 0	The required amount of memory in the buffer area is insufficient for the internal debugger to run.	Replace the controller.
1 * 0 * 0 0 *	Non-existent memory regions were not trapped as required.	Replace the controller.
1 1 0 1 1 1 1	Illegal SIMM configuration or mirrored cache with less than 64 MBtyes each or different amounts of memory.	Verify that SIMMs are installed as shown in step 1.4. If cache is mirrored, ensure that both cache modules contain the same amount of memory; either 64 MBytes or 128 MBytes.

Code	Description of Error	Corrective Action
1 1 0 1 1 1 0	Two or more cabinets are using the same PVA ID in an extended subsystem.	Reconfigure PVA ID to uniquely identify each cabinet in the subsystem. The cabinet with the controllers must be set to

		PVA ID 0. The additional cabinets must use PVA ID 2 and 3
1 1 0 1 1 0 1	The cabinet with the controllers is not configured to PVA ID 0.	Set PVA ID to 0 for the cabinet with the controllers. If problem persists, replace PVA module. If problem persists, replace EMU. If problem persists, replace the cabinet.
1 1 0 1 1 0 0	Termination power for one or more of the cabinet's device SCSI buses is absent or unacceptable.	Ensure that all of the cabinet's device SCSI buses have an I/O module. If problem persists, replace the failed I/O module.
1 1 0 1 0 1 0	Illegal I/O module configuration in an extended subsystem.	All I/O modules in an extended subsystem must be of the same kind: either single-ended or differential.
1 1 0 1 0 0 1	EMU microcode is incompatible with controller's software version.	Upgrade either the EMU microcode or the software so that they're compatible. See the release notes that accompanied the controller's software release.
1 0 0 0 0 0 0	No program card detected.	Ensure that the program card is properly seated, then reset the controller. If the error persists, try the card in another controller. If the error persists, replace the card, otherwise, replace the controller that reported the error.

7.0 CLI COMMANDS.

The cfmenu utility does not provide all of the features needed to work on HSZ70 RAID controllers. The following list provides additional CLI (Command List Interface) commands which often prove useful in maintaining and troubleshooting RAIDs. The full list of CLI commands can be found in the CLI Reference Manua included with the RAID accessoriesl.

sho this - Displays an abbreviated version of the RAID controller status. Useful for checking cache battery charge status and determining whether there is any unflushed data in the cache (something which is very important when trying to revive a corrupted RAID). When invoked, it displays a screen similar to the following:

```
Controller:
        HSZ70  2G93904673 Firmware V73Z-0, Hardware H04
        Not configured for dual-redundancy
        Device Port SCSI address 7
        Time: NOT SET
Host port:
        SCSI target (s) (0, 1, 2, 3)
        Preferred target (s)  (0, 1, 2, 3)
        TRANFSER_RATE_REQUESTED = 20 MHZ
        Host Functionality Mode = A
        Allocation Class                0
        Command Console LUN is disabled
Cache:
        64 megabyte write cache, version 4
        Cache is GOOD
        Battery is GOOD
        No unflushed data in cache
        CACHE_FLUSH_TIMER = DEFAULT (10 seconds)
        NOCACHE_UPS
Mirrored Cache:
        Not enabled
```

sho this full - This is an extended version of the sho this command. When executed, the following information is displayed:

```
Controller:
        HSZ70  2G93904673 Firmware V73Z-0, Hardware H04
        Not configured for dual-redundancy
        Device Port SCSI address 7
        Time: NOT SET
Host port:
        SCSI target (s) (0, 1, 2, 3)
        Preferred target (s)  (0, 1, 2, 3)
        TRANFSER_RATE_REQUESTED = 20 MHZ
        Host Functionality Mode = A
        Allocation Class                0
        Command Console LUN is disabled
Cache:
        64 megabyte write cache, version 4
        Cache is GOOD
        Battery is GOOD
        No unflushed data in cache
        CACHE_FLUSH_TIMER = DEFAULT (10 seconds)
        NOCACHE_UPS
Mirrored Cache:
        Not enabled
Extended information:
        Terminal speed 9600 baud, eight bit, no parity, 1 stop bit
        Operation control: 00000004 Security code: 70666
        Configuration backup enabled on 10 devices
```

sho raid - To display any raidsets on this RAID, enter this command and observe a screen similar to the one depicted below (if there are none, it will display *No RAIDSETS*):

```
Name            Storageset          Uses                Used by
------------------------------------------------------------------------------
R1              raidset             DISK10000           D0
                                    DISK20100
                                    DISK30000
                                    DISK40100
                                    DISK50000
                                    DISK60000

R2              raidset             DISK10200           D1
                                    DISK20300
                                    DISK30200
                                    DISK40300
                                    DISK50100
                                    DISK60100

R3              raidset             DISK10400           D2
                                    DISK20500
                                    DISK30400
                                    DISK40500
                                    DISK50200
                                    DISK60200
```

show r1 - This is an abbreviated version of the *sho raid* command. It will show all the details about a specific raidset. An example is shown below:

```
Name            Storageset          Uses                Used by
------------------------------------------------------------------------------
R1              raidset             DISK10000           D0
                                    DISK20100
                                    DISK30000
                                    DISK40100
                                    DISK50000
                                    DISK60000
        Switches:
          POLICY (for replacement) = BEST_PERFORMANCE
         RECONSTRUCT (priority) = NORMAL
          CHUNKSIZE = 64 blocks
        State:
          NORMAL
          DISK10000     (member  0)  is NORMAL
          DISK20100     (member  1)  is NORMAL
          DISK30000     (member  2)  is NORMAL
          DISK40100     (member  3)  is NORMAL
          DISK50000     (member  4)  is NORMAL
          DISK60000     (member  5)  is NORMAL
        Size: 177736020 blocks
```

sho mirror - This command is the same as *sho raid*, except that it displays any mirrorsets currently configured on the system (if there are none, the message *No MIRRORSETS*) will appear:

```
Name            Storageset          Uses                Used by
------------------------------------------------------------------------------
M1              raidset             DISK10000           D0
                                    DISK20100
```

sho disk10000 - To display the characteristics of a particular disk, enter this command. Note especially the last line, which refers to the fact that the configuration is being backed up on this disk. This means that *any* changes made to *any* of the RAID disks, even if it doesn't affect the storage set (mirrorset or raidset) this disk belongs to, will be recorded on this disk. This is *vital* if a controller crashes, as it will configure the replacement controller automatically, identifying all of the disks in their proper storage sets. This parameter can be changed either when running ***run cfmenu*** or by using the ***initialize*** command. A display similar to the following will appear:

```
Name            Type            Port    Targ    Lun             Used by
------------------------------------------------------------------------------
DISK10000       disk             1       0       0               R1
        DEC             BB01811C9C      3B05
        Switches:
         NOTRANSPORTABLE
         TRANSFER_RATE_REQUESTED = 20 MHZ  (synchronous 10.00 MHZ negotiated)
        Size:  35555889  blocks
        Configuration being backed up on this container
```

sho disks full - This command displays all of the information on every disk on the RAID (the data scrolls past quickly, but, if using an Xterm window, the user can scroll back up the list to view the data at their leisure). The data takes the following form:

```
Name            Type            Port    Targ    Lun             Used by
------------------------------------------------------------------------------
DISK10000       disk             1       0       0               R1
        DEC             BB01811C9C      3B05
        Switches:
         NOTRANSPORTABLE
         TRANSFER_RATE_REQUESTED = 20 MHZ  (synchronous 10.00 MHZ negotiated)
        Size:  35555889  blocks
        Configuration being backed up on this container
DISK20100       disk             2       1       0               R2
        DEC             BB01811C9C      3B05
        Switches:
         NOTRANSPORTABLE
         TRANSFER_RATE_REQUESTED = 20 MHZ  (synchronous 10.00 MHZ negotiated)
        Size:  35555889  blocks
        Configuration being backed up on this container
DISK30300       disk             3       3       0               R3
        DEC             BB01811C9C      3B05
        Switches:
         NOTRANSPORTABLE
         TRANSFER_RATE_REQUESTED = 20 MHZ  (synchronous 10.00 MHZ negotiated)
        Size:  35555889  blocks
        Configuration being backed up on this container
DISK60300       disk             6       3       0          SPARESET
        DEC             BB01811C9C      3B05
        Switches:
         NOTRANSPORTABLE
         TRANSFER_RATE_REQUESTED = 20 MHZ  (synchronous 10.00 MHZ negotiated)
        Size:  35555889  blocks
        Configuration being backed up on this container
```

sho spareset - This command displays any sparesets assigned to this RAID. A display similar to the following will appear (assuming this RAID has sparesets):

Name	*Storageset*	*Uses*	*Used by*
SPARESET	*spareset*	*DISK60300*	*D0*

sho failedset - Whenever a disk fails, this command will display the details as follows:

Name	*Storageset*	*Uses*	*Used by*
FAILEDSET	*failedset*	*DISK50200*	
Switches:			
NOAUTOSPARE			

sho disks - To display all of the disks on the RAID, enter this command. The information will be shown as follows :

Name	*Type*	*Port*	*Targ*	*Lun*	*Used by*
DISK10000	*disk*	*1*	*0*	*0*	*R1*
DISK10200	*disk*	*1*	*2*	*0*	*R2*
DISK10400	*disk*	*1*	*4*	*0*	*R3*
DISK20100	*disk*	*2*	*1*	*0*	*R1*
DISK20300	*disk*	*2*	*3*	*0*	*R2*
DISK20500	*disk*	*2*	*5*	*0*	*R3*
DISK30000	*disk*	*3*	*0*	*0*	*R1*
DISK60300	*disk*	*6*	*3*	*0*	*SPARESET*

shutdown this - To perform an orderly shutdown, the RAID controller must be disabled prior to powering down the RAID controller and drives. This command will perform that task. When this command has been successfully executed, the flashing green button on the front of the HSZ70 will be on constantly (no pulsing) and the controller will not respond to any commands which are typed in. To boot up the controller, simply press the green lighted button. The six amber LEDs on the front of the controller will flash through a pattern and, when successfully completed, the green lighted button will begin pulsing again. The *HSZ>* prompt will reappear after a series of coded messages scroll past on the display and the controller will again respond to commands which are keyed in.

restart this - As with *shutdown this*, this command will halt the controller, but then it will restart it without the user having to punch the green lighted button. This is useful when the RAID ends up in an unknown state (such as after a power outage).

clear_errors cli - This command clears out messages from the CLI error buffer after a problem has been fixed. This cures those annoying *Bad shelf* errors which appear when a fan has been pulled off a shelf.

sho hist - A history file of errors and reboots is maintained by the controller. This command clears that history file out, which is useful if a series of problems has been fixed and the user doesn't want to have to scroll through every problem the controller has experienced since day one.

clear_errors this invalid_cache destroy_unflushed_data - When attempting to restart a crashed RAID, there is often data stuck in the cache. The controller will work in a crippled manner with a reduced command list (effectively preventing the user from doing anything useful with the RAID). This command flushes the bad data out of the cache, allowing the user to fully access all of the commands again. This should probably be run whenever a RAID has crashed, especially if the error message *...Unflushed data in cache. Using reduced instruction set* appears.

clear_errors lost_data d0 (or d1, d2, etc) - When a RAID has suffered a power outage or other failure, there may be lost data still clogging up the storagesets, which prevents the RAID from working properly. This command clears out that data. However, once started, it may take several hours to complete, depending on the size of the database installed on the RAID.

set this initial_configuration - This command resets the HSZ70 controller to its factory defaults. This command can be used if a new controller has been installed on a RAID after the old controller failed. The new controller may have an old configuration which is different from the actual RAID configuration. In order to make the controller recognize the new configuration, this command must be run to purge the old configuration from the new controller's memory.

delete d0 - This command performs the same function as deleting units with *cfmenu*. It is used to break a unit so that the raidset or mirrorset can be broken up to use the disks elsewhere.

add unit d0 m1 writeback - This command assigns mirrorset #1 to unit d0 with the *writeback* parameter (meaning that there is a cache battery on the controller). This command can also be executed from the *cfmenu* utility by adding mirrorset #1 to unit d0 and toggling the unit to be *writeback*.

add unit d1 r1 writeback - This command assigns raidset #1 to unit d0 with the *writeback* parameter (meaning that there is a cache battery on the controller). This command can also be executed from the *cfmenu* utility by adding raidset #1 to unit d0 and toggling the unit to be *writeback*.

locate disk10000 - To confirm the location of a disk within a RAID, this command comes in handy. The amber LED on this particular disk will begin flashing and will continue to flash until the command ***locate cancel*** is entered. This can also be used for raidsets and mirrorsets with the command ***locate r1*** or ***locate m1***.

8.0 TROUBLESHOOTING.

When problems occur with the RAID drives or the controller, the following suggestions can be useful for determining the correct steps to take to resolve these problems.

Problem #1: RAID disks do not show up when doing a SCSI inquiry. SCSI I/O errors appear when attempting to mount RAID volumes. The HSZ70 controller displays the error message:
Invalid cache -- CLI command set reduced. Type SHOW THIS_CONTROLLER. Please see user guide to determine corrective action.
When the command ***sho this*** is entered, the *Cache:* area displays:
Cache is INVALID. Cache containing unflushed data has been removed from this controller.

Solution: Enter the command ***clear_errors this invalid_cache destroy_unflushed_data*** to remove (flush) the corrupt data from the cache. It only takes a moment or two to complete this operation. The *Invalid cache* error message will appear again when this command is entered, but when a *sho this* command is entered, the *Cache:* area will show that *Cache is GOOD*. To confirm that the RAID drives are accessible, enter the command ***restart this*** to reboot the controller. If there are any problems with corrupted data on the drives, an error message similar to the following will appear:

Unable to rundown the following unit on this controller: D0 -- cannot rundown unit

If this error does appear, proceed to *Problem #2* for resolution. If the controller restarts without errors, entering the command ***mount -a*** on the computer should remount the RAID volumes.

Problem #2: When an attempt is made to shutdown the controller, it displays the error message *Unable to rundown the following unit on this controller: D0 -- cannot rundown unit.* SCSI I/O errors appear when attempting to mount RAID volumes. This generally occurs when the RAID was improperly shutdown.

Solution: To correct this problem, first enter the command:

clear_errors this invalid_cache destroy_unflushed_data

to clear corrupt data out of the cache module, then enter the command:

clear_errors lost_data d0

Repeat this for any other units which show up in the error statement. The process of clearing out the data may take several hours to complete. NOTE:

The RAID should not be accessed for reads or writes to the volumes during this time. Once the lights stop flashing on the RAID it is safe to access the volumes again.

Problem #3: RAID is not functioning, even though power is applied, the green-lighted button on the front of the HSZ70 stays solid green (fails to blink even when pushed). There is no display on either a VT or the *tip cli* screen.

Solution: The controller has died. A new controller must be installed.

Problem #4: A disk which failed is ready to be replaced with a new disk.

Solution: A detailed procedure for restoring the RAID to its original configuration is given below. NOTE: For purposes of this example, DISK10000 is used as the failed disk and DISK40300 was the Hot Spare prior to DISK10000 failing. To put DISK10000 back into the raidset and DISK40300 back to the Hot Spare, execute the steps outlined below:

a) At the *HSZ*> prompt, enter the command ***set failedset autospare***
b) Enter the command ***delete failedset DISK10000***
c) Enter the command ***set raidset r1 nopolicy***
d) Remove DISK10000 from the RAID rack.
e) Push the button marked **1** in the row of six buttons located next to the green-lighted button on the front of the HSZ70 controller.
f) Place the new drive into the shelf where DISK10000 was located.
g) Enter the command ***add disk disk10000 1 0 0 0 0***
h) Enter the command ***set r1 remove=disk40300***
i) Enter the command ***set r1 replace=disk10000***
j) Enter the command ***delete failedset disk40300***
k) Enter the command ***add spareset disk40300***
l) Enter the command ***set r1 policy=best_performance***
Wait for the LEDs to stop flickering before accessing the RAID database.

Problem #5: The cache battery has failed (either the ***sho this*** command lists the battery as failed or the LED never stops flashing).

Solution: The cache battery must be replaced. If the RAID needs to run without the battery, enter the following command:

set this cache_policy=b

This command will allow the RAID drives to work for 10 hours without having a battery attached to the cache module. If a replacement battery cannot be obtained within 10 hours, a potentially risky option is to tell the

system that the cache module is being protected by a UPS (uninterruptable power supply). This can be done by issuing the following command:

set this cache_ups

This is dangerous to use, as any data in the cache module at the time of a power outage will not be saved. If a system absolutely has to run, however, this will perform that function for a prolonged period of time (more than the 10 hours which the previous command provides).

Problem #6: SCSI error messages appear when the Alpha computer attached to the RAID boots up. The flashing green button on the HSZ70 controller goes to steady green and one or more of the amber LEDs on the HSZ70 controller lights up. RAID volumes are missing during bootup.

Solution: SCSI cable(s) to the channel(s) with the amber LED(s) lighted is/are not properly connected. Confirm that the SCSI cable(s) is/are good or replace with known-good cable(s).

Problem #7: Can't exit from window created with ***tip cli*** command due to typing mistake.

Solution: Close the X-term or DEC-term window and open another to work in.

Problem #8: Controller has failed and replacement is ready to install. Disks were originally formatted with configuration saved. NOTE: If unsure about whether the disks have a saved configuration on them, at the *HSZ>* prompt enter the command:

sho disk10000 full.

If the configuration was saved, the response will contain the following line:

Configuration being backed up on this container

If this message does not appear, see *Problem #9* for the restore procedure for a RAID with an unsaved configuration.

Solution: As soon as the new controller is installed and the RAID powered up, the disks with the saved configuration will notify the new controller about what raidset/mirrorset/spareset each disk belongs to. The user does not need to do anything (except refrain from issuing commands while the restore is being done - a process which takes about five minutes to complete).

Problem #9: Controller has failed and replacement is ready to install. Disks were not originally formatted with configuration saved. NOTE: If unsure about whether the disks have a saved configuration on them, at the *HSZ>* prompt enter the command:

sho disk10000 full.

If the configuration was saved, the response will contain the following line:

Configuration being backed up on this container

If this message does not appear, then the configuration was not saved. Since there is no saved configuration on the RAID disks, there is a good chance that the configuration cannot be restored. However, if precise information about the RAID configuration is available (either as a drawing or as labels on the disks themselves), then there is a chance that the configuration can be restored and the data saved.

Solution: Follow the steps below to attempt this restore:

a) Collect as much information about which drives belong to which raidsets as possible. Without precise information on this, it would probably be best not to proceed further, but to simply reload from scratch.
b) Follow the steps outlined in Section 5 beginning with step 5.1. When step 5.18 is reached, proceed to step 5.35 (as though this RAID controller has never been configured).
c) Create raidsets and mirrorsets identical to the information you have on the old configuration for steps 5.38 through 5.59.
d) ***DO NOT DO STEPS 5.60 THROUGH 5.66***! Implementing these steps will erase all data from the disks. Proceed to step 5.67.
e) Execute steps 5.67 through 5.78 and exit cfmenu.
f) Restart the RAID controller with the command **restart this** (if this is a dual-redundant controller, you must also do a ***restart other*** command).
g) Halt the RAID computer.
h) Cycle power on the RAID computer.
i) Reboot the RAID computer.
j) If the configuration was correct and the raidsets and mirrorsets recovered, then the RAID should be back to its original configuration with the data intact and readily accessible. If not, then Section 5 will have to be repeated, this time with all steps implemented, including 5.60 through 5.66. The data will then have to be reloaded from tape.

9.0 FAULT MANAGEMENT UTILITY.

HSZ70 controllers sometimes display cryptic error messages. The Fault Management Utility (FMU) can be used to decipher these messages. This section gives explanations and examples on using this utility. For more detailed information on failures, refer to the ***Service Manual*** *included with the* ***Digital StorageWorks HSZ70 Array Controller*** *manuals.*

9.1 To invoke the Fault Management Utility (FMU), at the *HSZ>* prompt, enter the command **run fmu**.

9.2 To display the most recent failure, enter the command **sho last most** (which is an abbreviation of the ***sho last_failure most_recent*** command). This will display the most recent error message. An example is given below:

```
Last Failure Entry: 4. Flags: 000FF380
Template: 1.(01) Description: Last Failure Event
Power On Time: 0. Years, 135. Days, 21. Hours, 25. Minutes, 49. Seconds
Controller Model: HSZ70-AX
Serial Number: ZG74838477 Hardware Version:  A03(03)
Firmware Version: V51Z(51)
Informational Report
Instance Code: 01010302 Description:
An unrecoverable hardware detected fault occurred.
Reporting Component: 1.(01) Description:
Executive Services
Reporting component's event number: 1.(01)
Event Threshold: 2.(02) Classification:
HARD. Failure of a component that affects controller performance or
precludes access to a device connected to the controller is indicated.
Last Failure Code: 018700A0 (No Last Failure Parameters)
Last Failure Code: 018700A0 Description:
A processor interrupt was generated with an indication that the (//) RESET
button on the controller module was depressed.
Reporting Component: 1.(01) Description:
Executive Services
Reporting component's event number: 135.(87)
Restart Type: 2.(02) Description: Automatic hardware restart
```

9.3 Entering the command **describe last most** (an abbreviation for ***describe last_failure most_recent***) yields the following result:

```
FMU> describe instance 01010302
 Instance Code: 01010302 Description:
  An unrecoverable hardware detected fault occurred.
 Reporting Component: 1.(01) Description:
  Executive Services
 Reporting component's event number: 1.(01)
 Event Threshold: 2.(02) Classification:
  HARD. Failure of a component that affects controller performance or
  precludes access to a device connected to the controller is indicated.
```

which translates into a controller reset due to a bad disk in the RAID. This disk had failed during configuration, displaying a blinking amber LED, effectively preventing the controller from finishing its system configuration.

9.4 The next example demonstrates the use of the **describe last** (short for ***describe last_failure***) command to translate the error code 018700A0, which in this case indicates that the reset button on the controller was pressed.

```
FMU> describe last_failure  018700A0
Last Failure Code: 018700A0 Description:
A processor interrupt was generated with an indication that the (//) RESET
button on the controller module was depressed.
Reporting Component: 1.(01) Description:
Executive Services
Reporting component's event number: 135.(87)
Restart Type: 2.(02) Description: Automatic hardware restart
```

9.5 If unsure about what parameters are available for a given command, enter the part of the command you know with a question mark (?) at the end. A listing of possible choices will be listed. As an example, if use of the *describe* command is desired but the user doesn't know what options exist, typing in the command **describe ?** will result in the following being displayed:

```
FMU> describe ?
Your options are:
  ASC_ASCQ_CODE
  COMPONENT_CODE
  CONTROLLER_UNIQUE_ASC_ASCQ_CODE
  DEVICE_TYPE_CODE
  EVENT_THRESHOLD_CODE
  INSTANCE_CODE
  LAST_FAILURE_CODE
  REPAIR_ACTION_CODE
  RESTART_TYPE
  SCSI_COMMAND_OPERATION_CODE
  SENSE_DATA_QUALIFIERS
  SENSE_KEY_CODE
  TEMPLATE_CODE
```

10.0 CFMENU HEADING DESCRIPTION.

The cfmenu display has several headings which aren't fully explained on screen. This reference has been included for completeness. A typical heading is reproduced below:

```
------------------------------------------------------------------ CFMENU Configuration Menu Utility ------------------------------------------------------------------
MAIN MENU:          |Unconfig'd |     Config'd  Device          Product          Stor. set  Stor. set  Chnk  Trn  In-  Re-          P W W
                    | Dev. PTLs |       PTLs   Name            ID               Name       Typ/Sz     size  sp.  it'd duc  Unit    T P B
```

Main Menu	Shows the tasks you can accomplish with CFMENU
Unconfig'd Dev. PTLs	Shows the PTL locations of devices that have not yet been added to the controller's configuration. Use these devices to create single-disk units and storagesets, such as mirrorsets and raidsets.
Config'd PTLs	Shows the PTL locations of all devices that are used in - or are eligible to be used in - a storageset or as a single-disk unit.
Device Name	Shows the names of all devices that are used in - or are eligible to be used in - a storageset or as a single-disk unit.
Product ID	Shows the model numbers of all devices that are used in - or eligible to be used in - a storageset or as a single-disk unit.
Stor.set Name	Shows the name of all storagesets in the controller's list of configured storagesets: by convention, S*n* for stripesets, M*n* for mirrorsets, and R*n* for raidsets, where the *n* is the number of the storageset, such as M1.
Stor.set Typ/Sz	Shows the types of storagesets and their number of members. For example, MIR/2 is a mirrorset that contains two disk drives.
Chunk Size	Shows the chunksize for stripesets and raidsets. This column is marked *unk* (unknown) until you initialize the storageset.
Trnsp.	Displays *Y* if you selected *Transportable* when initializing the disk(s).
Init'd	Displays *Y* if you initialized the disk(s).
Reduc	Displays *Y* if you selected *Reduced* when initializing the disk(s) or if one or more disks failed in this storageset.
Unit	Shows the unit numbers for all storagesets or devices.
PT	Displays *P* if the unit is partitioned.
WP	Displays *Y* if you selected *Write-Protect* when initializing the disk(s).
WB	Displays *Y* if you selected *Write-Back Cache* when initializing the disk(s).