Operator's Guide

HP 9000 V-Class Server

First Edition

A3725-96005

Customer Order Number: A3725-90005

October 1997

Printed in: USA

Revision History

Edition: First

Document Number: A3725-96005

Notice

The information contained in this document is subject to change without notice.

Hewlett-Packard makes no warranty of any kind with regard to this material, including, but not limited to, the implied warranties of merchantability and fitness for a particular purpose. Hewlett-Packard shall not be liable for errors contained herein or for incidental or consequential damages in connection with the furnishing, performance or use of this material.

FCC Notice

This equipment has been tested and found to comply with the limits for a Class A digital device, pursuant to Part 15 of the FCC Rules. These limits are designed to provide reasonable protection against harmful interference when the equipment is operated in a commercial environment.

This equipment generates, uses, and can radiate radio frequency energy and, if not installed and used in accordance with the instruction manual, may cause harmful interference to radio communications. Operation of this equipment in a residential area is likely to cause harmful interference, in which case the user will be required to correct the interference at his own expense.

The user is cautioned that changes or modifications not expressly approved by the party responsible for compliance (Hewlett-Packard) could result in the equipment being noncompliant with FCC Class A requirements and void the user’s authority to operate the equipment.

Do not connect external equipment to the utility outlets in the Hewlett-Packard V-Class server. Unauthorized connection voids all agencies’ emissions certification.

Contents

Tables

Figures

Preface

The *Operator's Guide HP 9000 V-Class Servers* documents the information necessary to operate and monitor HP V-Class servers. This book is not intended to be a tutorial but a reference for system administrators, system operators, or system managers.

Notational conventions

This section describes notational conventions used in this book.

`bold monospace`	In command examples, **`bold monospace`** identifies input that must be typed exactly as shown.
`monospace`	In paragraph text, `monospace` identifies command names, system calls, and data structures and types. In command examples, `monospace` identifies command output, including error messages.
italic	In paragraph text, *italic* identifies titles of documents. In command syntax diagrams, *italic* identifies variables that you must provide. The following command example uses brackets to indicate that the variable *output_file* is optional: `command` *input_file* [*output_file*]
Brackets ([ ])	In command examples, square brackets designate optional entries.
Curly brackets ({}), Pipe (|)	In command syntax diagrams, text surrounded by curly brackets indicates a choice. The choices available are shown inside the curly brackets and separated by the pipe sign (|). The following command example indicates that you can enter either `a` or `b`: `command {a	b}`

Horizontal ellipses (...)	In command examples, horizontal ellipses show repetition of the preceding items.
Vertical ellipses	Vertical ellipses show that lines of code have been left out of an example.
Keycap	**Keycap** indicates the keyboard keys you must press to execute the command example.

NOTE A note highlights important supplemental information.

CAUTION A caution highlights procedures or information necessary to avoid damage to equipment, damage to software, loss of data, or invalid test results.

1 Overview

The V-Class server can contain four to 16 HP PA-8200, Precision Architecture - Reduced Instruction Set Computer (PA-RISC) processors and can support memory configurations of 256 Mbytes to 16 Gbytes.

This chapter provides a block diagram of the V-Class server and discusses its features and configurations.

The PA-8200 processor

The V-Class server uses the Hewlett-Packard PA-8200 processor, based on the concept of Reduced Instruction Set Computers (RISC). The PA-8200 was designed according to Hewlett-Packard's PA-RISC Architecture version 2.0 specifications.

The PA-RISC architecture is presented in the *PA-RISC 2.0 Architecture* reference manual. Please refer to that document for detailed information about the features of the PA-8200. This Operator's Guide does not attempt to duplicate information in that manual. Instead, it presents only V-Class server-specific information.

The processors of the system are supported by several Application Specific Integrated Circuits (ASIC) hardware controllers, an enhanced memory system, and a high-bandwidth I/O subsystem. Special hardware and software allow these processors to perform both as conventional single CPUs or together, in parallel, to solve more complex problems requiring parallel processing.

The node

The V-Class server can contain four to 16 processors. The terms node and system are used interchangeably in this book. The node uses a symmetric multiprocessor (SMP) design that can exploit fine-grain parallelism.

A conceptual block diagram of the node is shown in Figure 1. Centrally located in the diagram is the HP Hyperplane crossbar that is comprised of four Exemplar Routing Attachment Controllers (ERAC). The Hyperplane crossbar allows all of the processors to access all of the available memory. Processors are installed on Exemplar Processor Agent Controllers (EPACs). An EPAC allows the processor and the I/O subsystem (the Exemplar PCI-bus Interface controller—EPIC) access to the Hyperplane crossbar.

Also connected to the Hyperplane crossbar are the Exemplar Memory Access controllers (EMAC). Up to two processors are located on each EPAC. Memory is controlled by the EMAC. Input and output devices connect to the node through EPIC which is connected to the processor agents.

The Exemplar Core Utilities board (ECUB—commonly called the Utilities board) in the node contains a section of hardware called the core logic. It provides interrupts to all of the processors in the system through the core logic bus which connects to each processor agent. The ECUB attaches to the Exemplar Node Routing board (ENRB) centrally located in the node.

Figure 1 Functional block diagram of a V-Class system

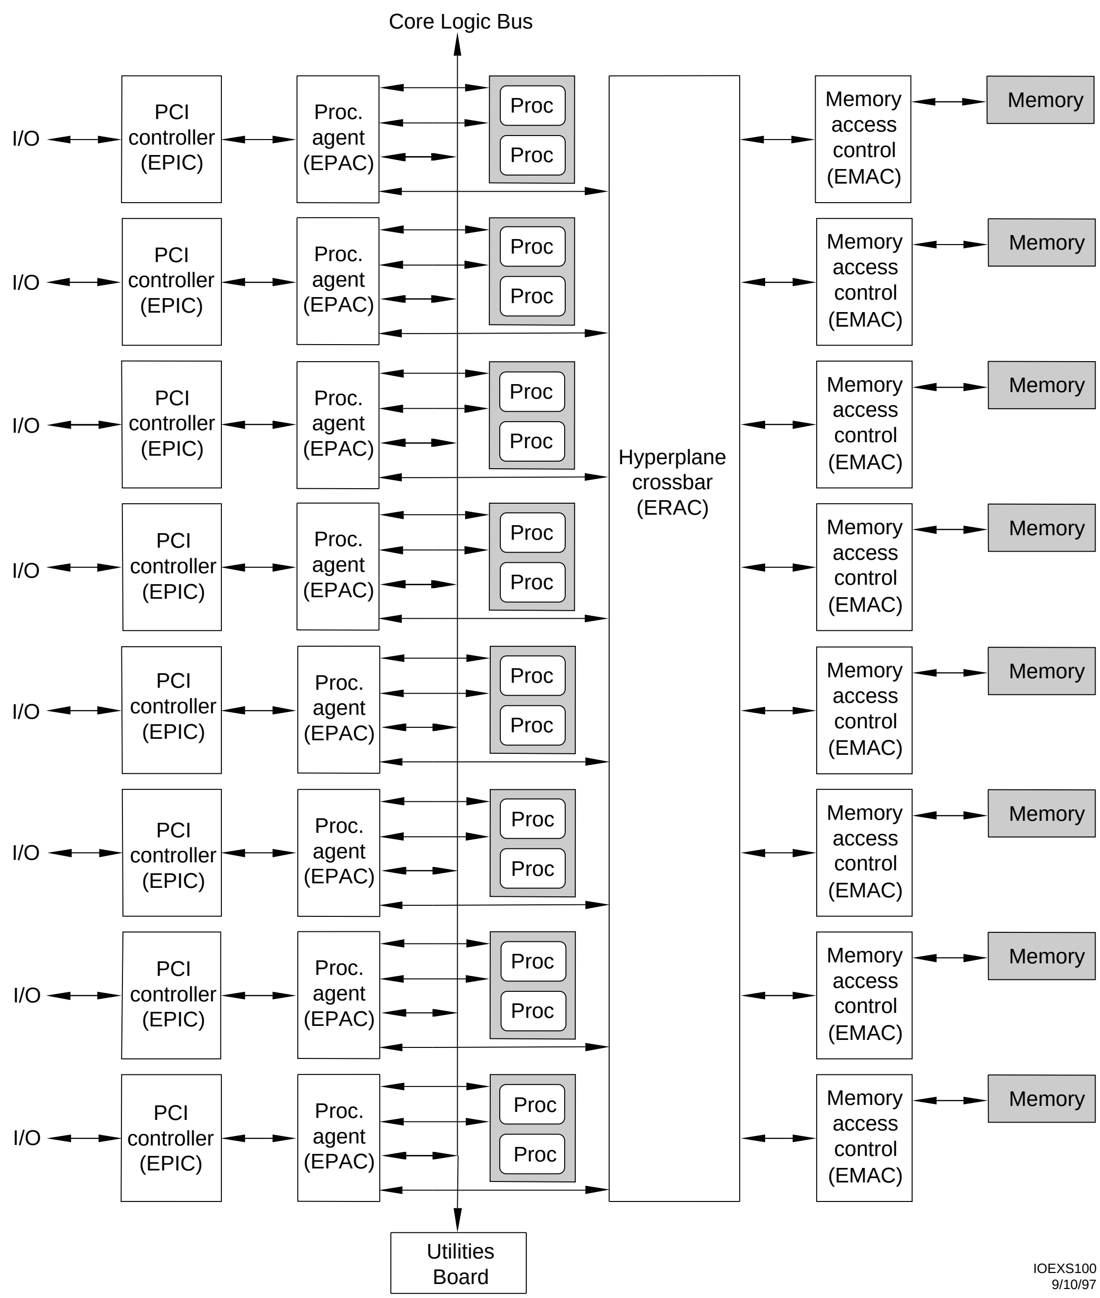

Control and status registers (CSRs)

System hardware is manipulated by control and status registers, CSRs, located in the processors and controllers.

CSRs provide control, status, or both to the processors and other hardware in the node. Each CSR is memory mapped and is available to all processors in the system. Many of the registers are described in detail throughout this book by functional groups, such as system configuration, messaging and data copy, error recovery, and so on.

Description of functional blocks

Each block in Figure 1 is described in the following sections.

Processor agent controller

The EPAC can connect to zero, one, or two PA-8200 processors. It can also connect to zero or one EPIC (the I/O controller). With no processors, the EPAC serves as an I/O-only interface. The EPAC has the following buses:

Runway bus (0, 1)—Two each, 64-bit, bidirectional buses for processor 0 and processor 1, respectively. These buses have a raw bandwidth of 960 Mbytes/sec.

Hyperplane crossbar port bus (0, 1)—Four 32-bit, unidirectional buses connected to two Hyperplane crossbar ERACs, two in each direction. These buses have a total raw bandwidth of 1.9 Gbytes/sec.

I/O port—Two 16-bit or 32-bit, unidirectional interfaces to an I/O device, one for reading data and one for writing data. The width of the bus depends on the width of the I/O device connected. Each bus has a bandwidth of 120 Mbytes/sec or 240 Mbytes/sec, depending on the width of the interface.

Core Logic Bus interface—Bidirectional bus that supports boot and support services.

The EPAC sends and receives transactions from the ERACs using four unidirectional data paths. There are four ERACs in the Hyperplane crossbar. Each processor agent, however, communicates with only two of the four ERACs.

The EPAC includes special hardware called the *data mover* for rapid message and data movement between memory within a node. This dedicated hardware greatly improves file I/O and networking over software versions.

Hyperplane crossbar—routing attachment controller

The Hyperplane crossbar is comprised of four ERACs that provide an interconnect for each processor and I/O device to memory.

Each of the four ERACs has the following buses:

EPAC Port (A, B, C, D)—Eight 32-bit, unidirectional interfaces to four EPAC ports, four in each direction. Each port has simultaneous (input and output) bandwidth of 960 Mbytes/sec.

EMAC Port (A, B, C, D)—Eight 32-bit, unidirectional interfaces to four EMAC ports, four in each direction. Each port has simultaneous (input and output) bandwidth of 960 Mbytes/sec.

Figure 2 shows how the ERACs connect to each EPAC and EMAC.

Figure 2 **ERAC interconnection**

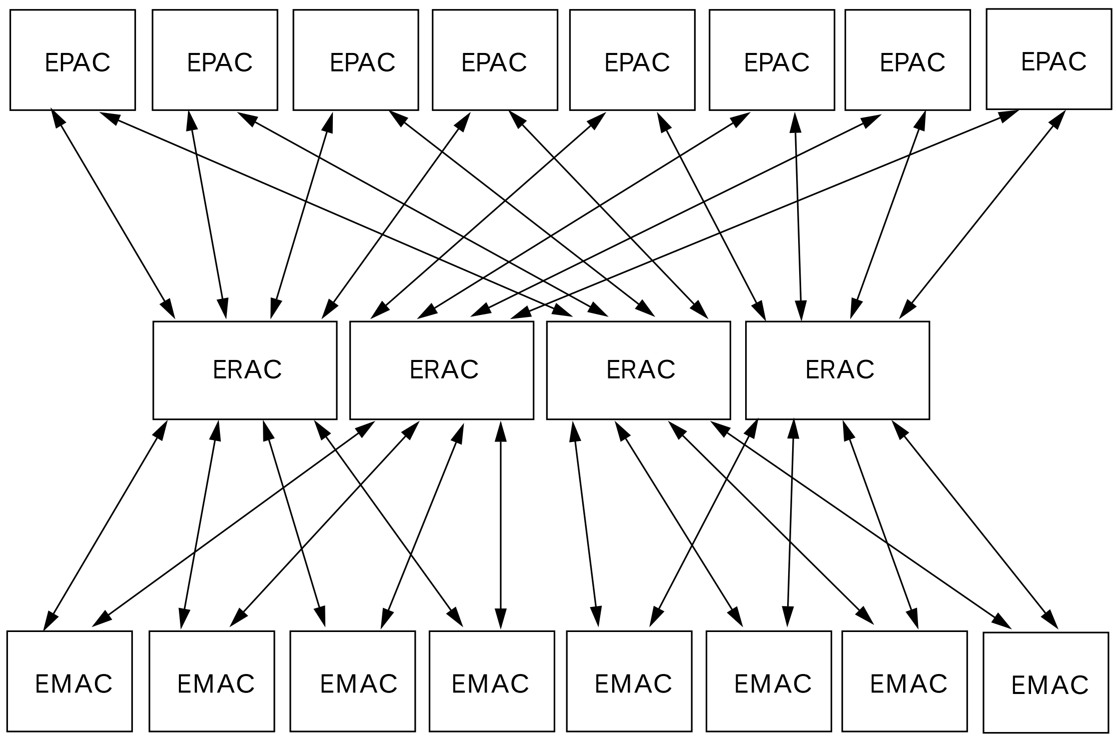

Memory access controller

The EMAC controls all accesses to memory. Each EMAC controls four banks of memory, allowing up to 32 banks in an eight-EMAC node. Memory banks consist of Single Inline Memory Modules (SIMMs) of Synchronous Dynamic Random Access Memory (SDRAM).

The EMAC has the following buses:

ERAC Port (A, B)—Four unidirectional, 32-bit interfaces, two in each direction. This interface supports a total simultaneous read-write bandwidth of 1.9 Gbytes/sec.

Even Memory—Bidirectional, 88-bit interface to the even memory banks associated with the EMAC.

Odd Memory—Bidirectional, 88-bit interface to the odd memory banks associated with the EMAC.

A processor accesses memory by sending a request, in the form of packets, to an ERAC. The request is then forwarded to one of the EMACs. The EMAC routes requests into even and odd pending queues. Some packets not destined for memory are routed from processor to processor through the EMAC. These packets are routed directly to the output ports.

The EMAC accesses one of four available memory banks, checking the Error Correction Code (ECC). Provided no additional coherency operations are required, the data accessed from memory is returned to the processor by sending a response back to the ERAC, which forwards the response to the EPAC.

Utilities board and core logic bus

The ECUB, or Utilities board, connects to the core logic bus and contains two field-programmable gate arrays (FPGAs): the EPUC and EMUC. The EPUC allows processors access to the system core logic and booting firmware, and the EMUC processes the environmental state of the node and interrupts the processors when appropriate. V-Class servers use the core logic bus primarily to boot the system and to issue node-local environmental interrupts.

The core logic bus is a low-bandwidth, multidrop bus that connects each processor to the control and interface logic (both RS232 and ethernet). A processor can write to control and status registers (CSRs) accessed using the core logic bus to initialize and configure the ERAC chips and Utilities board logic.

Node configurations

Table 1 shows the available configurations.

Table 1 **Node configurations**

Processors	Processor agents	Memory boards	Total memory (EMB)	I/O chassis
4	2	2	256	1
4	2	2	512	1
4	2	2	1024	1
4	2	4	2048	1
8	4	4	512	1 or 2
8	4	4	1024	1 or 2
8	4	4	2048	1 or 2
8	4	8	4096	1 or 2
12	6	8	1024	1, 2, or 3
12	6	8	2048	1, 2, or 3
12	6	8	3072	1, 2, or 3
12	6	8	4096	1, 2, or 3
16	8	8	1024	1, 2, 3, or 4
16	8	8	2048	1, 2, 3, or 4
16	8	8	3072	1, 2, 3, or 4
16	8	8	4096	1, 2, 3, or 4

Shared memory

V-Class servers provide exceptionally high performance using a shared-memory architecture; it allows the developer, compilers, and applications to view the system as a number of processors sharing a large physical memory and a number of high-bandwidth I/O ports.

Message passing (Data Mover) hardware provides high performance for applications developed using explicit MPI messaging.

Compilers use shared memory to provide automatic, efficient parallelization while viewing memory as a single contiguous virtual address space.

The Hyperplane crossbar provides high-bandwidth, low-latency nonblocking access from processors and I/O channels to the system memory. It prevents the performance drop-off associated with systems that employ a system-wide bus for processor and I/O memory traffic.

Sequential memory references (linearly ascending physical address) to shared memory are interleaved across up to eight memory boards on a 32-byte basis. See *V-Class Architecture* for information regarding memory interleaving.

With all processor references to memory, copies of the accessed data are encached into either the instruction or data caches of each processor. If the processor making the memory reference modifies the data and if another processor references that same data while a copy is still in the first processor cache, a condition exists whereby the data has become stale. The V-Class hardware continually works to ensure that the second processor does not use an outdated copy of the data from memory. The state that is achieved when both processors' caches always have the latest value for the data is called *cache coherence*.

To maintain updated coherent copies, V-Class servers operate under the following rules:

Any number of read encachements of a cache line can be made at a single time. The cache line can be read-shared in multiple caches.

To write data (store) into a cache line, the cache line must be “owned” exclusively by the processor. This implies that any other copies must be invalidated.

Modified cache lines must be written back to memory from the cache before being overwritten.

2 Indicators, switches, and displays

This section describes indicators, switches, and displays of the HP 9000 V-Class server.

Operator panel

The operator panel is located on the top left side of the server and contains the key switch panel, CDROM drive, optional DAT tape drive, and the LCD display. Figure 3 shows the location of the operator panel and its components.

Figure 3 Operator panel

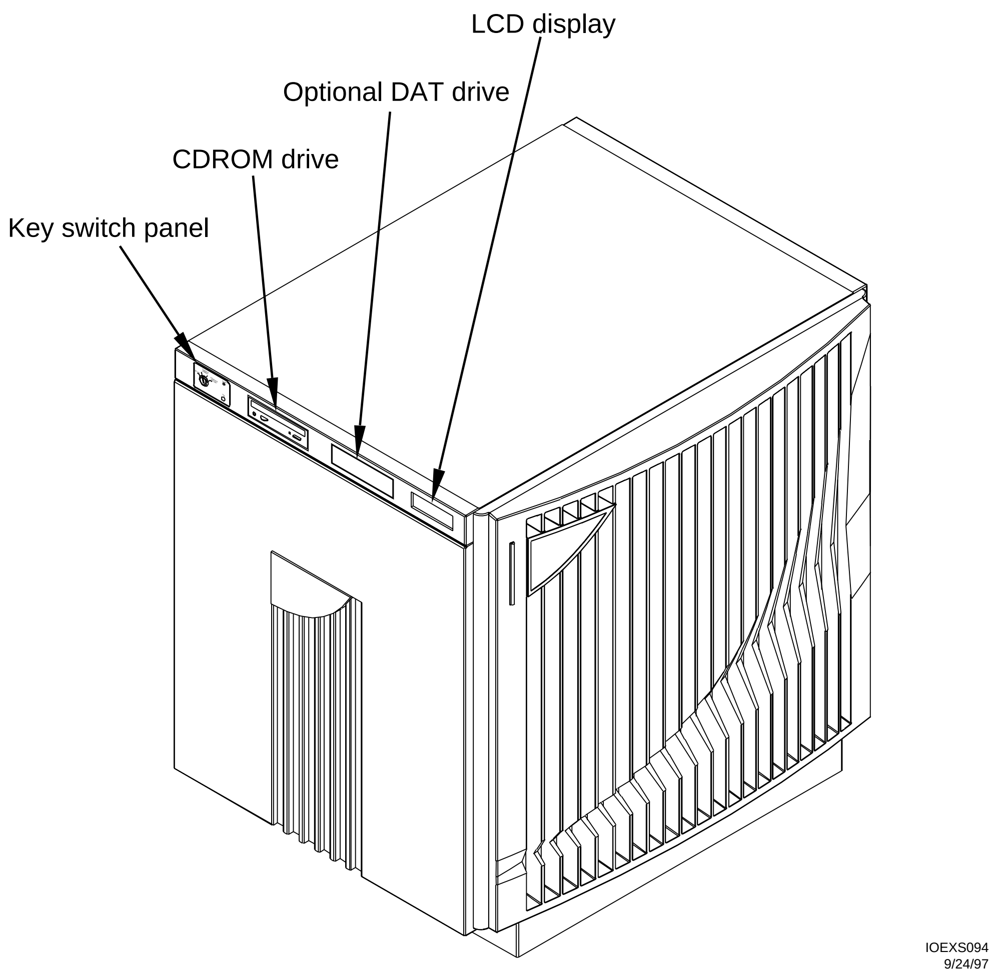

Key switch panel

The key switch panel is located on the left of the operator panel, as shown in Figure 3 on page 14. The key switch panel contains a two position key switch, a DC ON LED, and a TOC (Transfer Of Control) button, as shown in Figure 4.

Figure 4 **Key switch panel**

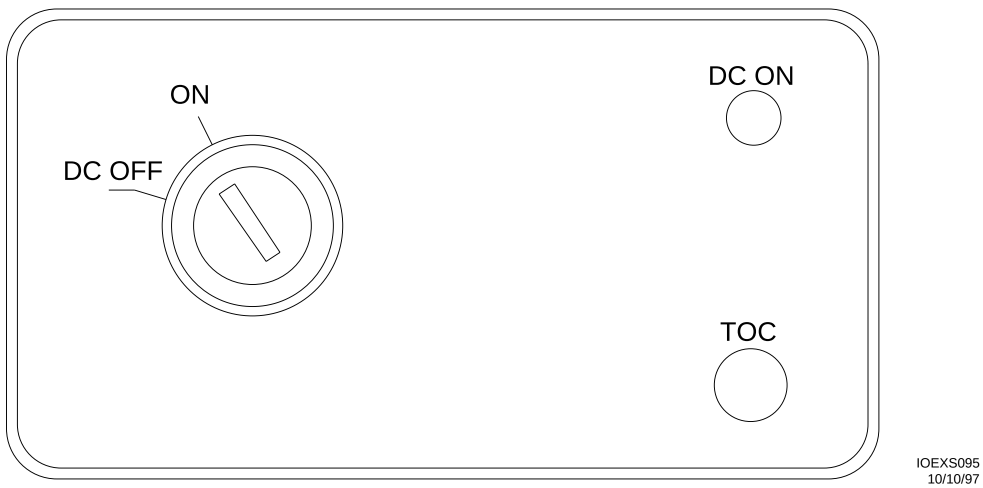

Key switch

The key switch has four positions. Placing the key switch in any position other than DC OFF powers up the V-Class server.

- DC OFF

  DC power is not applied to the system. Placing the keyswitch in this position is the normal method for turning off power to the system.

- ON

  DC power is applied to the system. POST (Power On Self Test) begins executing and brings up the system from an indeterminate state and then calls OBP.

DC ON LED

This LED indicates that DC power has been applied to the system.

TOC

The TOC (Transfer Of Control) button is a recessed switch that resets the system.

CDROM drive

The CDROM drive is located on the left of the operator panel, as shown in Figure 3 on page 14. Figure 5 shows the CDROM drive front panel in detail.

Figure 5 **CDROM drive**

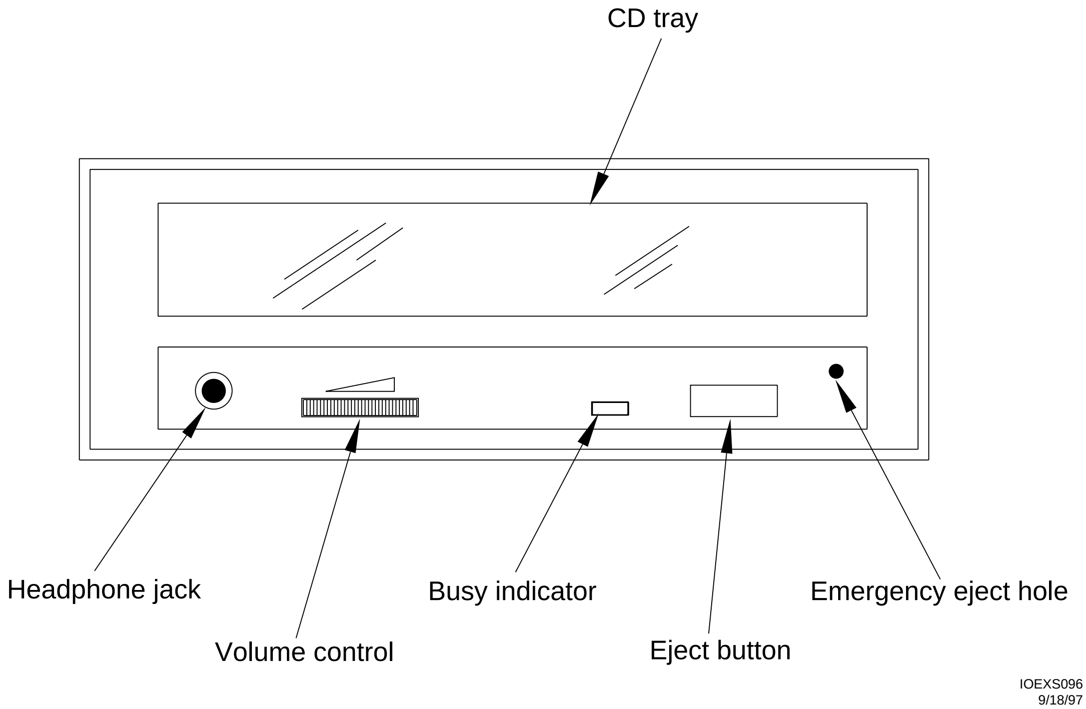

Busy indicator

The busy indicator LED flashes to indicate that a read operation is occurring.

CAUTION

Do not push the eject button while this LED is flashing. If you do, the operation in progress is aborted, and the CDROM is ejected, possibly causing a loss of data.

Eject button

Push the eject button to open the CD tray to remove CDROMs from the drive.

Emergency eject

Insert a paper clip into the emergency eject hole to open a jammed CD tray.

Optional DAT drive

The DAT drive is located on the right of the operator panel, as shown in Figure 3 on page 14. The DAT drive front panel contains two indicator LEDs and an eject button, as shown in Figure 6.

Figure 6 **DDS-3 DAT drive front panel**

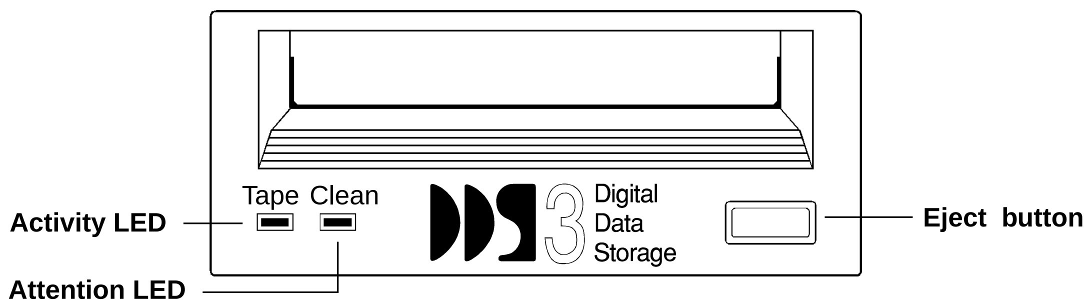

LEDs

The two LEDs provide operating information for normal as well as error conditions. Table 2 shows the meaning of the different LED patterns.

Table 2 **Indicator LED operation**

Tape (Activity) LED (green)	Clean (Attention) LED (amber)	Meaning
Flashing slowly	Off	A load or unload of a cartridge is in progress.
Flashing rapidly	Off	A cartridge is loaded and a read or write is in progress.
On	Off	A cartridge is loaded.

Tape (Activity) LED (green)	Clean (Attention) LED (amber)	Meaning
Any	Flashing slowly	Media caution signal. Indicates that a cartridge is near the end of its life or that the heads need cleaning.
Any	On	Fault
Flashing slowly	Off	Power-on (starts with two steady lights)

Eject button

Push the eject button to remove cartridges from the tape drive. The drive performs the following Unload sequence:

1. The tape is rewound to Beginning of Partition (BOP) for Partition 0.
2. If the tape is write-enable, the copy of the Tape log is written back to tape.
3. The tape is then rewound to Beginning of Media (BOM), unthreaded from the mechanism, and ejected.

CAUTION Do not push the eject button while the LED is flashing. If you do, the operation in progress is aborted and the cartridge is ejected, possibly causing a loss of data.

System Displays

The V-Class servers provide two means of displaying status and error reporting: an LCD and an Attention light bar.

Figure 7 **System displays**

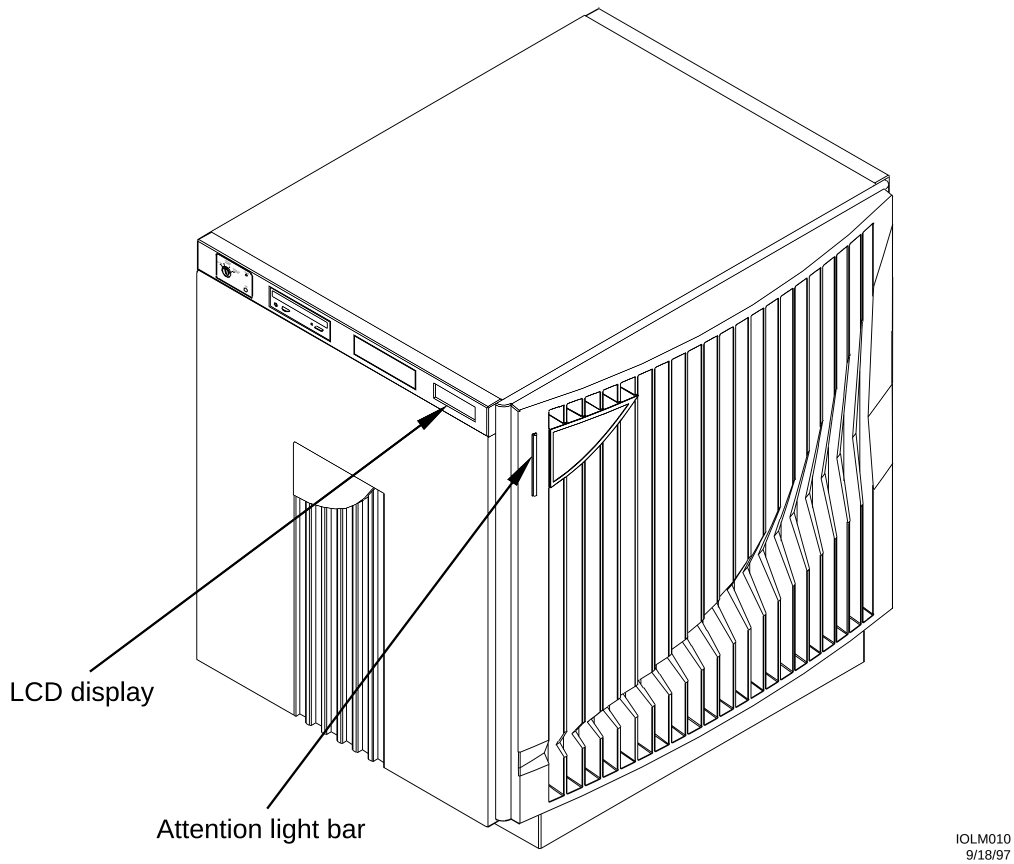

LCD (liquid crystal display)

The LCD panel is located on the right of the operator panel, as shown in Figure 7 on page 21. The LCD is a 20-character by 4-line liquid crystal display as shown in Figure 8.

Figure 8 **Front panel LCD**

```
0 (0,0) SN 20ffffff
IIII FIII IIII IIII
POST Rev 0.10.0.0
Monarch: PB2L
```

When the operator panel key switch is turned on, the LCD powers up, but is initially blank.

Power-On Self Test (POST) takes about 20 seconds to start displaying output to the LCD. POST is described in the *HP Diagnostics Guide: S-Class, X-Class, and V-Class Servers*. The following explains the output shown in Figure 8:

- The first row displays the serial number of the system
- The second row displays processor status
- The third row displays the revision of POST and the POST status codes
- The forth row displays which processor is in control

While HP-UX is running, the LCD displays the status of each processor.

- The first row displays HP-UX RUN status
- The second row displays processor status
- The third row is normally blank
- The forth row is blank

Table 3 shows the meaning of the processor status codes.

Table 3 **Processor status codes**

Code	Meaning
K	Kernel
I	Idle
F	FORTH (OBP)
S	Server
U	User
E	Emulator
M	Monarch
D	Installed but not responding during POST
d	User disabled
-	Processor down

Table 3 describes the POST status and progress code definitions.

Table 4 **POST status codes**

Code	Meaning
a	Monarch CPU is selected
b	POST code checksum
c	Test controller code checksum
e	Configuration parameter checksum
?	Rebuilding configuration parameter
f	Starting core logic hardware initialization
g	EPUC/ EMUC CSR initialization
h	Core logic SRAM pattern test, “A”s pattern
i	Core logic SRAM pattern test, “5”s pattern

Code	Meaning
j	Core logic SRAM address uniqueness test
k	Zero filling core logic SRAM
l	Starting C-runtime environment

Table 5 describes the processor status codes displayed while HP-UX is running.

Table 5 **HP-UX processor status codes**

Code	Meaning
F01F	The processor is least busy
F05F	The processor is more busy
F0AF	The processor is more busy
F0FF	The processor is most busy

Attention light bar

The Attention light bar is located at the top left corner on the front of the HP 9000 V-Class server as shown in Figure 7 on page 21. This light bar displays system status in three ways:

- Off—system powered down
- Steady on—system powered up
- Flashing—error condition

3 Teststation

This chapter provides information about the Hewlett-Packard workstation that serves as a teststation for the V-Class server. An ethernet bus called the test bus connects the teststation to the utilities board located in the V-Class server chassis. A teststation is required to enable service personnel to verify and troubleshoot the V-Class system.

Teststation features

The V2200 requires an HP workstation to function as the system console. There are two options:

- Product number A4082A

  The standard teststation consists of the basic PA9000 712/60 workstation, the special hardware to connect the workstation as a console and the teststation software.

- Product number B6044AA

  This option provides the special hardware and teststation software to connect an existing workstation as a console.

There are two ethernet LAN cards installed within the teststation: one is on the core I/O card and the other is a plug-in LAN card.

Additional software which provides all the necessary functions of the teststation has been added to the basic HP-UX 10.20. This software is located in the directory /spp.

Teststation functionality

Several functions of the teststation include:

- Running diagnostics
- Updating of Utility Board firmware
- Logging of environmental and system level events
- Configuration of hardware and boot parameters
- Booting the Operating System

The teststation is connected directly to the V-Class server via a dedicated LAN connection as well as a special RS-232 serial connection. Together the LAN and the serial connection provide control, status, and HPMC error information to the teststation where operators and CEs may view status.

Teststation logons

Two UNIX user accounts are created on the teststation during the HP-UX 10.20 operating system installation process.

sppuser — This user is the normal logon for the teststation during system operation, verification, and troubleshooting.

Default password: spp user
Please note the space between spp and user.

root — This user has the ability to modify and configure every parameter on the teststation.

NOTE If the passwords to these accounts are changed by the customer, the new passwords must be supplied to the Hewlett-Packard Customer Engineer (CE) upon request.

Teststation sppuser windows

When the sppuser is logged on to the teststation, the windows appear in the configuration shown in Figure 9.

Figure 9 Teststation sppuser windows

test station console - message output	sppconsole - complex console
Message window ● ccmd daemon status approx. 60 seconds after power on ● Failures and hard errors ● Prompt = *hostname:/path*	**Console window** ● Power-on self-test status ● HP mode (Boot menu) ● Prompt = Command: ● Prompt = Console login:
tcsh	**tcsh**
sppuser - tsh shell ● Commands and scripts executing on the teststation ● Prompt = *hostname*	**sppuser - tsh shell** ● Commands and scripts executing on the teststation ● Prompt = *hostname*

Message window (test station console - message output)

The message window (TOP LEFT) displays status messages about the V-Class server. Do not enter commands in this window.

This window displays status from the `ccmd` daemon running on the teststation. The `ccmd` daemon builds the node configuration file `/ssp/data/nod_0.cfg`.

Hard error logger also displays status in this window.

Console window (sppconsole - complex console)

This is the main console window (TOP RIGHT) for the V-Class server. All POST (Power-On Self-Test) status is displayed here. The user can type commands while in HP mode (boot menu) to boot and configure the node. See Chapter 4, “Firmware (OBP and PDC)” for more information about HP mode (boot menu). HP service personnel can also enter a special mode called forth mode (OBP) to perform special configuration commands.

tcsh shell windows (tcsh)

The tsh windows (LOWER LEFT and LOWER RIGHT) are local shell windows on the teststation. The user may enter commands into these windows which invoke scripts or functions on the V-Class server.

Some commands like the `do_reset` command are scripts that begin execution on the teststation and then control the V-Class server.

Using the mwm Root Menu

The teststation uses the Motif Window Manager (mwm) to control the X windows on your screen. The Root Menu is mwm's main menu. The options on the Root Menu affect your entire display. From the Root Menu, you can select menu items to create new windows, initiate diagnostic tools, and perform other tasks.

To select a command from the Root Menu:

Step 1. Press and hold down any mouse button over the root window. The root window is any part of your display that is not covered by a window. The Root Menu appears.

Step 2. Move the mouse pointer over an option. The option becomes boxed.

Step 3. Release the mouse button to select the option.

The Root Menu options include:

- **sspdsh window**—Creates a spp diagnostic shell window on your screen. This should be used by service personnel only.
- **ksh window**—Creates a new ksh window on your screen.
- **test station console**—Creates a new teststation window on your screen.
- **sppconsole menu**—Creates a new sppconsole window by selecting node 0/complex from this menu.
- **shell menu**—Lists several shell options for new windows including the csh.
- **diag tool menu**—Lists several diag tools. This should be used by service personnel only.

CAUTION Do not select `do_reset` (the first menu item) unless you are prepared to reset the V-Class server.

- **X tool menu**—Lists several X tools including the load average display and xlock (screen lock/saver).
- **shuffle up**—Moves the bottom window in a stack of windows to the top.

- **shuffle down**—Moves the top window in a stack of windows to the bottom.
- **refresh all**—Refreshes your entire X display.
- **restart**—Stops and restarts mwm.
- **logout**—Closes all open windows and stops mwm.

Using the console

The console serves as a communication device between you and your V-Class server. Virtual consoles are also used to monitor specific operations, like a system software crash dump.

This section familiarizes you with the console and shows you how to:

- Start the console
- Access the console remotely
- Access system logs

Starting the console

The console server program automatically starts the console on the teststation when you log on as sppuser. If the console dies or stops running, restart it from the teststation using one of the following methods:

Using the Root Menu

To restart the console using the Root Menu complete the following steps:

Step 1. Press and hold down any mouse button over the root window. The root window is any part of your display that is not covered by a window. The Root Menu appears.

Step 2. Select sppconsole menu. The menu displays.

Step 3. Select node 0/complex. The new sppconsole window displays.

Using the `sppconsole` command

To restart the console using the `sppconsole` command complete the following steps:

Step 1. Select a shell window by placing the mouse pointer in the window.

Step 2. Start the console. Enter:

```
sppconsole
```

The new sppconsole window displays.

Logging out

To restart the console by logging out of the teststation and logging back on again complete the following steps:

Step 1. Press and hold down any mouse button over the root window. The root window is any part of your display that is not covered by a window. The Root Menu appears.

Step 2. Select logout. The teststation closes all open windows, stops mwm, and returns a HP-UX login prompt.

Step 3. Log into the teststation as sppuser. The new sppconsole window displays.

Accessing the console remotely

Use the sppconsole command to access the console from a system other than the teststation. Using control sequences with this command allows you to watch or to assume control of the console window.

Table 6 **sppconsole commands**

Command	Description
`^Ecw`	Display a list of other users connected to the console.
`^Ecf`	Force control of the console interface.
`^Ecs`	Relinquish control of the interface and return to spy mode.
`^Ec.`	Exit the console program.
`^Ec?`	List the console escape command sequences.

NOTE **^E** is the **Ctrl** and **e** keys pressed simultaneously. The **e** does not have to be an uppercase **E**.

Example: Performing a ^E command

To execute the `^Ecf` command complete the following steps:

1. Press the **Cntrl** key and the **e** key simultaneously.
2. Release the **Cntrl** key and the **e** key.
3. Press the **c** key.
4. Press the **f** key.

Watching the console remotely

Any root user can display the console via a remote login to the teststation, so it is possible to have many different processes watching the console at the same time. This is sometimes referred to as “spy mode”. Only one window can actually control the console; see “Controlling the console remotely” on page 36 for more information.

To monitor the console from a system other than the teststation, complete the following steps:

Step 1. Remotely log in to the teststation as sppuser (default password: spp user) with the following command:

`rlogin` *hostname*

```
login: sppuser
Password: spp user
```

Step 2. Access the system console with the following command:

`/spp/scripts/sppconsole`

At this point you are in “spy mode”, meaning that you can only monitor what is going on at the system console. If you try to enter a command the following message is displayed:

```
[read-only -- use ‘^Ecf’ to attach, ‘^Ec?’ for help]
```

Display a list of other console users

Step 1. Display a list of other users connected to the console with the following command:

CTRL-**Ecw**

Step 2. Exit the session with the following command:

CTRL-**Ec.**

The period is part of the command.

Controlling the console remotely

You can perform system maintenance or diagnostics remotely by assuming control of the console from a remote terminal. Upon gaining control of the console, you have write access to that window.

Only one window can be active at a time.

To assume control of the console, complete the following steps:

Step 1. Remotely log in to the teststation as sppuser (default password: spp user) with the following command:

`rlogin` *hostname*

```
login: sppuser
Password: spp user
```

Step 2. Access the system console with the following command:

`/spp/scripts/sppconsole`

At this point you are in spy mode, meaning that you can only monitor what is going on at the system console. If you try to enter a command the following message is displayed:

```
[read-only -- use '^Ecf' to attach, '^Ec?' for help]
```

Step 3. Assume control of the console by attaching to it with the following command:

CTRL-**`Ecf`**

Step 4. You can relinquish control of the console and return to spy mode with the following command:

CTRL-**`Ecs`**

Step 5. Exit the session with the following command:

CTRL-**`Ec.`**

Changing a console's connection

Once you have started a console as a watch or a control connection, you can change the connection type with escape characters.

To change a watch window to an active console window, enter:

CTRL-`Ecf`

To change an active console window to a watch window, enter:

CTRL-`Ecs`

Accessing system logs

You can monitor system status via two logs, event_log and consolelog, located in /spp/data on the teststation.

The event_log file periodically logs system status. Once the file reaches 4MB, the system renames it to event_log.old and creates a new event_log file.

The consolelog file logs data from the sppconsole window. When the file reaches 2MB, the system renames it to consolelog.old and creates a new consolelog file.

Teststation file system

The /spp directory located on the workstation's local disk (HP-UX version 10.20) contains all the necessary files necessary for the workstation to function as the V-Class server's teststation.

Figure 10 **Teststation file system**

/spp

etc	(spp specific daemons)
bin	(compiled executables)
scripts	(executable scripts)
data	(log and config files)
firmware	(firmware directory)
est	(support for scan testing)
man	(Man pages for /spp commands)
unsupported	(unsupported utilities)

/spp/etc

The /spp/etc directory contains many of the unique daemons which run on the teststation. These daemons provide services in the management of the V-Class node. Two daemons that are always running on the teststation are:

`cmd` The complex configuration management daemon that builds the configuration file /spp/data/node_0.cfg.

`conserver` The console-server that directs RS-232 console traffic from the Utility Board to the upper right console window on the teststation.

/spp/bin

In the /spp/bin directory you will find specific commands and daemons used to manage a V-Class node. The most frequently used commands are:

`ccmu`	A text-based tool that can modify the reconfiguration parameters initialized by POST.
`est`	The command (Exemplar Scan Testing) to initiate scan testing.
`do_reset`	The command executed on the teststation to reset the V-Class node remotely.
`pce_util`	The command to read the environmental status remotely from the Utility Board.
`event_logger`	Responsible for logging events such as HPMC (High Priority Machine Checks).

/spp/scripts

The /spp/scripts directory contains scripts that perform a variety of functions.

`dcm`	Dump Configuration Manager.
`hard_logger`	The hard error logger, HP Machine Check (HPMC) data collection script.
`sppconsole`	The console utility.

/spp/data

The /spp/data directory contains:

node_0.cfg	Configuration file with scan rings and configured hardware. Built by the `ccmd` daemon approximately every sixty seconds. This file describes all the ASIC chips populated in a V-Class node and also defines the scan rings which are used by the est (Exemplar Scan Test) utility. This file is a very useful troubleshooting tool for tracking scan ring failures to devices.
consolelog	A file containing all the console activity on the system.
est	The scan testing log.

hard_hist	Log of all hard failure information. Logs the output of all suspected ACIS (Application Specific Integrated Circuits). This file may be useful in troubleshooting intermittent ACIS failures.
event_log	Log of all event information. A read only file which captures information generated by the `ccmd` daemon.

/spp/firmware

The /spp/firmware directory is where firmware files are loaded temporarily before they are moved to the Utility Board (ECUB).

/spp/est

The est directory contains files used during scan testing.

/spp/man

The /spp/man directory contains the man (manual) pages and many of the teststation specific commands.

4 Firmware (OBP and PDC)

OpenBoot PROM (OBP) and SPP Processor Dependent Code (SPP_PDC) make up the firmware on HP V-Class servers that makes it possible to boot HP-UX.

This chapter discusses the boot sequence and the commands available from the boot menu.

Boot sequence

Once a machine powers on, the firmware controls the system until the operating system (OS) executes. If the system encounters an error any time during the boot process, it stops processing and goes to HP mode boot menu. See “HP mode boot menu” on page 46 for more information.

When the operator powers on or resets the machine, the following process occurs:

1. Power-On Self Test (POST) runs. POST is described in the *HP Diagnostics Guide: S-Class, X-Class, and V-Class Servers*.
2. OBP probes all the devices.
3. OBP loads SPP_PDC in RAM.
4. OBP starts the HP-UX loader, which in turns calls SPP_PDC to set up CPU’s, memory, and I/O devices in a way that HP-UX understands.
5. The next action depends on whether Autoboot is enabled.
    1. If autoboot is enabled, the operating system boots unless the user presses a key within 10 seconds.
    2. If autoboot is disabled or if the user presses a key within 10 seconds, the boot menu is displayed.

Figure 11 on page 43 illustrates the initialization and start-up process.

Figure 11 Boot process

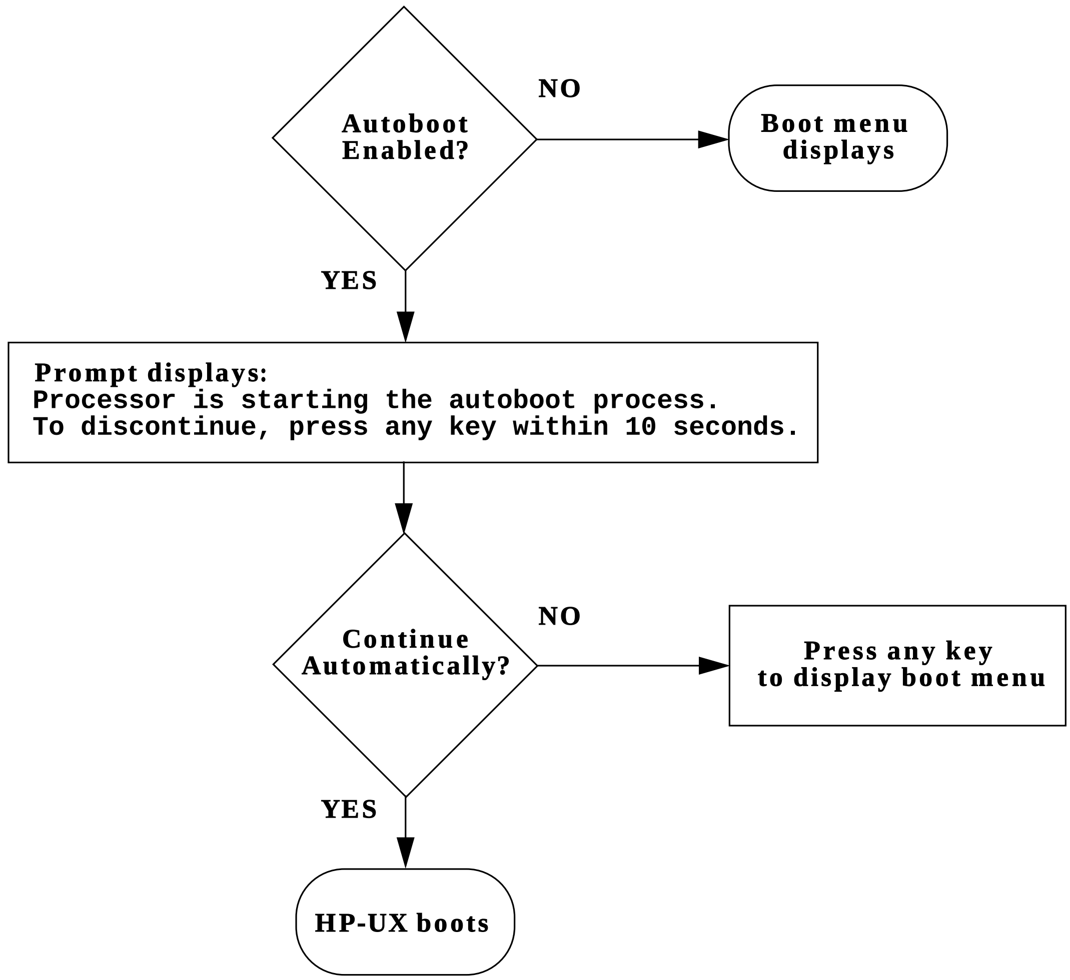

Boot process output

The following output illustrates what typically displays on the console as the system starts up:

```
SPP2000, POST version 4.0.0.1, compiled 1997/07/31 13:20:50

Node Id: 00000000

Monarch: PB2L

Probing CPUs.

Completing core SRAM initialization.

Initializing main memory.

  Probing memory: MB0L, MB1L, MB2R, MB3R, MB4L, MB5L, MB6R, MB7R

  Parallel memory initialization in progress

  PB3R MB0L ....----

  PB0R MB1L ....----

  PB1R MB2R ....----

  PB1L MB3R ....----

  PB2L MB4L ....----

  PB2R MB5L ....----

  PB0L MB6R ....----

  PB3L MB7R ....----

Nodemask=00000001

Booting OBP.

OBP Power-On Boot on [0:4]

----------------------------------------------------------------------

                       PDC Firmware Version Information

                            PDC_ENTRY version 3.1.0.29

                             POST Revision: 3.1.0.1

OBP Fieldtest Release 3.1.0.29, compiled 97/06/19 16:33:50 (2)

                            SPP_PDC version 1.1.7.3

----------------------------------------------------------------------
```

```
Proc type   Proc#   Proc Rev  Speed     State   Dcache   Icache
----------  -----   --------  -------  -------  -------  -------
HP,PA82000  0       4.0       200 MHz  Active   2048 KB  2048 KB
HP,PA82000  1       4.0       200 MHz  Active   2048 KB  2048 KB
HP,PA82000  2       4.0       200 MHz  Active   2048 KB  2048 KB
HP,PA82000  3       4.0       200 MHz  Active   2048 KB  2048 KB
HP,PA82000  4       4.0       200 MHz  Active   2048 KB  2048 KB
HP,PA82000  5       4.0       200 MHz  Active   2048 KB  2048 KB
HP,PA82000  6       4.0       200 MHz  Active   2048 KB  2048 KB
HP,PA82000  7       4.0       200 MHz  Active   2048 KB  2048 KB
HP,PA82000  8       4.0       200 MHz  Active   2048 KB  2048 KB
HP,PA82000  9       4.0       200 MHz  Active   2048 KB  2048 KB
HP,PA82000  10      4.0       200 MHz  Active   2048 KB  2048 KB
HP,PA82000  11      4.0       200 MHz  Active   2048 KB  2048 KB
HP,PA82000  12      4.0       200 MHz  Active   2048 KB  2048 KB
HP,PA82000  13      4.0       200 MHz  Active   2048 KB  2048 KB
HP,PA82000  14      4.0       200 MHz  Active   2048 KB  2048 KB
HP,PA82000  15      4.0       200 MHz  Active   2048 KB  2048 KB

16384 MB memory installed

Primary boot path   = 1/0:0.12.0
Alternate boot path = 15/3
Console path        = 15/1
Keyboard path       = 15/1

System is HP9000/800/V2200 series

Autoboot and Autosearch flags are both OFF or we are in HP core mode.
Processor is entering manual boot mode.
```

HP mode boot menu

In some instances, the boot menu displays; otherwise the operating system boots and the system is ready for use. The boot menu displays when one of the following occurs:

- The system encounters a problem while booting
- Autoboot is disabled
- The operator interrupts the boot process

```
Command                          Description

-------                          -----------

AUto [BOot|SEArch ON|OFF]        Display or set the specified flag
BOot [PRI|ALT|<path> <args>]     Boot from a specified path
BootTimer [time]                 Display or set boot delay time
CLEARPIM                         Clear PIM storage
CPUconfig [<proc>] [ON|OFF]      Configure/Deconfigure Processor
DEfault                          Set the system to defined values
DIsplay                          Display this menu
ForthMode                        Switch to the Forth OBP interface
IO                               List the I/O devices in the system
LS [<path>|flash]                List the boot or flash volume
OS [hpux|sppux]                  Display/Select Operating System
PASSword                         Set the Forth password
PAth [PRI|ALT|CON] [<path>]      Display or modify a path
PDT [CLEAR|DEBUG]                Display/clear Non-Volatile PDT state
PIM_info [cpu#] [HPMC|TOC|LPMC]  Display PIM of current or any CPU
RESET [hard|debug]               Force a reset of the system
RESTrict [ON|OFF]                Display/Select restricted access to Forth
SCSI [INIT|RATE] [bus slot val]  List/Set SCSI controller parms
SEArch [<path>]                  Search for boot devices
SECure [ON|OFF]                  Display or set secure boot mode
TIme [cn:yr:mo:dy:hr:mn[:ss]]    Display or set the real-time clock
VErsion                          Display the firmware versions
Command:
```

At this point, you can either enter any command on the menu or continue the boot process. To continue booting, enter the following command:

`Command:` **`boot`**

Commands and options are case-insensitive (auto = AUTO). Each command has a shortcut, which is the minimum letters that can be entered to execute the command. For example, to execute the search command, enter **SEA**, **SEAR**, **SEARC**, or **SEARCH**. This shortcut is indicated by capital letters in Table 7 and in the rest of this chapter.

Table 7 lists the commands available from the `Command:` prompt.

Table 7 Boot menu commands

Command	Description
AUto [BOot | SEArch ON | OFF]	Displays or sets the Autoboot or Search flag. If Autoboot is on, the system boots automatically after reset. If AutoSearch is on, the system searches for and displays all I/O devices that the system can boot from.
BOot [PRI | ALT | *path args*]	Initiates the boot sequence. A default or specified path to the boot device can be used.
BootTimer [time]	Displays or sets a delay time for the system to wait for external mass storage devices to come online.
CLEARPIM	Clears (zeros) Processor Internal Memory (PIM) storage after a system crash. CAUTION: this command can delete important troubleshooting information; do not enter the CLEARPIM command unless directed to.
CPUconfig [proc] [ON | OFF]	Displays or sets the configuration of processors.
DEfault	Sets the system environment variables to defined values and changes certain HP variables so that HP-UX can boot.
DIsplay	Displays this menu.
ForthMode	Switches to the Forth OBP interface. For use by service personnel only.
IO	Displays all I/O devices in the system whose SCSI controller cards are enabled.
LS [*path* | flash]	Displays the LIF contents (boot or flash volume) of a device.
OS [hpux | sppux]	Displays or sets which OS is going to boot—HP-UX or SPP-UX. For V-Class, this should be set to HP-UX.
PASSword	Defines the password used to control access to ForthMode. Same as UNIX `password` command.

Command	Description
PAth [PRI | ALT | CON] [*path*]	Displays or sets primary, alternate, console, and keyboard hardware paths. Keyboard path cannot be modified.
PDT [CLEAR | DEBUG]	Displays or clears Page Deallocation Table (PDT) information. For use by service personnel only.
PIM_info [cpu#] [HPMC | TOC | LPMC]	Displays Processor Internal Memory (PIM) information for current or any CPU.
RESET [hard | debug]	Resets the system state.
RESTrict [ON | OFF]	Displays or sets restricted access to Forth mode.
SCSI [INIT | RATE] [bus slot val]	Displays or sets SCSI controller initiator ID or transfer rate.
SEArch [*path*]	Displays pathnames for devices with bootable media in the system.
SECure [ON | OFF]	Displays or sets secure boot mode. If secure mode is set, the boot process cannot be interrupted. Only useful if autoboot is on; the system will autosearch and autoboot.
TIme [cn:yr:mo:dy:hr:mn[:ss]]	Displays or sets the realtime clock.
VErsion	Displays the internal firmware versions.

Enabling Autoboot

AUto displays or sets the Autoboot or Search flag, which sets the way a system will behave after powering on. If Autoboot is ON, the system boots automatically after reset. If AutoSearch is ON and Autoboot is OFF, the system searches for and displays all I/O devices that the system can boot from. Changes to a flag take effect after a system reset or power-on. The default value for both Autoboot and Autosearch is OFF.

Syntax

AUto [BOot | SEArch] [ON | OFF]

- Used alone, this command displays the current status of the Autoboot and Autosearch flags.
- BOot - If ON, the OS is automatically loaded from the primary boot path after a power-up or reset. Otherwise, the system displays the boot menu and waits for interactive boot commands. During an autoboot, the process pauses for 10 seconds to allow the operator to stop the boot process.
- SEArch - If ON, the system searches for all I/O devices that it can boot from and displays a list. Usually disabled because the search can be time-consuming.
- ON enables the indicated feature.
- OFF disables the indicated feature.

Examples

au

This command displays the status of the Autoboot and Autosearch flags.

```
Autoboot:ON
Autosearch:ON
```

au bo

This command displays the current setting of the Autoboot flag.

```
Autoboot:ON
```

au bo on

This command sets the Autoboot flag ON.

```
Autoboot:ON
```

HElp command

The `help` command displays help information for the specified command or redisplays the boot menu.

Syntax

`HElp` [*command*]

Used alone, `HElp` displays the boot menu. Specifying *command* displays the syntax and description of the named command.

Examples

The following example illustrate use of this command:

`help au` This command displays information for the `auto` command.

```
AUto[BOot|SEArch] [ON|OFF] Display or set the specified flag

AUto boot on                    Enable  auto boot on next boot.
AUto boot off                   Disable auto boot on next boot.
AUto search on                  Enable  auto search on next boot.
AUto search off                 Disable auto search on next boot.

Auto search enables the automatic search of a boot device.
Auto boot enables the autoboot sequence.
```

5 Starting and stopping HP-UX

This chapter provides information about starting and stopping HP-UX:

- Starting your HP 9000 V-Class Server
- Using the console
- Starting HP-UX
- Stopping HP-UX

Terms

Terms used in this chapter include:

Open Boot PROM (OBP)	Tests the server's hardware, provides a complete and accurate description of all available hardware, and boots the operating system. Resides on the V-Class server.
node	Set of processors and physical memory organized as a symmetric multiprocessor (SMP) running a single image of the operating system. A V-Class server consists of one node.
multiuser mode	Allows multiple users to access the system simultaneously. See also single-user mode.
run level	An HP-UX mode of operation. The /etc/inittab file defines which terminals and processes are active at each run level.
run-level s	Restricts user input to the console. Also known as single-user mode.
single-user mode	Restricts user input to the console. Also known as run-level s, system administrators use this mode to prevent users from accessing the system during system maintenance.

Starting your HP 9000 V-Class Server

Bringing your V-Class server to a usable state involves two systems and their hardware and software. This section provides a brief overview of the process; for complete instructions, see the *Managing Systems and Workgroups*.

A V-Class server consists of two systems:

- Teststation
  - Boots the V-Class server
  - Monitors the V-Class server for hardware errors
  - Debugs a hung system
  - Runs HP-UX
- V-Class node
  - Hosts OpenBoot PROM (OBP) software
  - Runs HP-UX

The boot procedure differs according to the value of the Autoboot flag. See “Enabling Autoboot” on page 49 for information on how to set Autoboot. After you power-up your V-Class server, if Autoboot is set to:

- `True`, OBP automatically starts HP-UX. (Press the **ESC** key within 10 seconds to interrupt the boot process and enter bootmenu commands.)
- `False`, you must:
  - Start OBP at the teststation’s HP-UX prompt by entering the following command:

    `do_reset`

  - Start the V-Class server using default values at OBP’s default prompt by entering the following command:

    `boot`

Starting HP-UX

You must start up, or boot, HP-UX after the operating system has been completely shut down or after you have partially shut down the operating system to perform system administration tasks. This section provides a brief overview of the process; for complete instructions, see the *Managing Systems and Workgroups*. This section describes:

- Starting HP-UX
- Reviewing the state of the file system
- Restarting HP-UX

HP-UX boot process

The process begins when you power up the V-Class server. OBP tests the hardware, provides a complete and accurate description of all available hardware, and boots the operating system.

OBP then passes control to the `/sbin/init` process, which sequentially executes the contents of `/etc/inittab`. The `inittab` file executes the `/sbin/bcheckrc` and `/sbin/rc` scripts to check your file system and initialize the system.

The operating system continues to boot and displays additional information about your system. After the boot process completes, an HP-UX login prompt appears in the sppconsole window.

Reviewing the state of the file system

During the start-up process, the `/sbin/bcheckrc` script executes `/usr/sbin/fsclean`. This command determines the shut down status of the system and returns three possibilities:

1. Proper file system shut down

   The startup process continues, and you see the following message:

```
/usr/sbin/fsclean:/dev/dsk/0s0(root device) ok file system is OK, not running
fsck
```

2. Improper file system shut down

   The start-up process is interrupted, and you see:

```
/usr/sbin/fsclean:/dev/dsk/0s0 not ok run fsck FILE SYSTEM(S) NOT PROPERLY
SHUTDOWN, BEGINNING FILE SYSTEM REPAIR.
```

At this point, the system runs `/usr/sbin/fsck` in a mode that corrects certain inconsistencies in the file systems without your intervention and without removing data. The `fsck` command does one of the following:

- Repairs and reboots the system, incorporating the changes
- Prompts you to run the `fsck` command manually. If you need to run `fsck` manually, see the fsck(1m) manpage

3. Other errors detected

   An error message displays (for example, unable to open a specified device file), the start-up process ends, and you need to solve the problem.

Restarting HP-UX

To restart the system after a reboot or hang, you must bring the system to single-user state. Characteristics of the single-user state include:

- The only access to the system is through the sppconsole
- The only processes running on the system are the
  - Shell on the console
  - Background daemon processes started by `/sbin/rc`
  - Processes that the root user invokes

Restarting HP-UX from single-user mode

Complete the following steps to restart your system:

Step 1. Select the sppconsole window on the teststation or assume control remotely. See “Accessing the console remotely” on page 34.

Step 2. Log in as root.

Step 3. Reboot the system with the `reboot` command. Enter:

`reboot`

Restarting HP-UX from multiuser mode

Complete the following steps to restart your system:

Step 1. Select the sppconsole window on the teststation or assume control remotely. See “Accessing the console remotely” on page 34.

Step 2. Log in as root.

Step 3. Change to the root directory. Enter:

`cd /`

Step 4. Bring the system to single-user state with the `shutdown` command. Enter:

`shutdown`

Step 5. Reboot the system with the `reboot` command. Enter:

`reboot`

Stopping HP-UX

This section provides a brief overview of the process; for complete instructions, see the *Managing Systems and Workgroups*.

Typically, you shut down the system to:

- Put it in single-user state so you can update the system or check files systems.
- Turn it off in order to perform a task such as installing a new disk drive.

CAUTION Never stop the system by turning off the power. Stopping the system improperly can corrupt your file systems. Use `shutdown`.

Shutdown considerations

Only the system administrator or a designated superuser can shut down the system.

The `/sbin/shutdown` command:

- Warns all users to log out of the system within a grace period you specify
- Halts daemons
- Kills unauthorized processes
- Unmounts file systems
- Puts the system in single-user mode
- Writes the contents of the I/O buffers to a disk

CAUTION Do not run `shutdown` from a remote system via `rlogin` if you use a network service. The shutdown process logs you out prematurely and returns control to the console. Run `shutdown` from the sppconsole window on the teststation.

See the shutdown(1M) man page for a complete description of the shutdown process and available options.

Rebooting or shutting down the system

To reboot or shut down your V-Class server, perform the following steps:

Step 1. Select the sppconsole window on the teststation. See “Teststation sppuser windows” on page 28.

Step 2. Log in as root.

Step 3. Change to the root directory. Enter:

```
cd /
```

Step 4. Shut down the system using the `shutdown` command. Enter:

```
shutdown
```

Progress messages detailing system shutdown activities print to your terminal. Upon reaching run-level 0, the system:

- Restarts in single-user mode
- Displays the root prompt

Step 5. Bring the system to a complete stop with the `reboot` command. Enter:

```
reboot -h
```

CAUTION Turn power off to the node only after the words `CPU halted` have been displayed in the sppconsole window.

6 Power

This chapter describes how to configure the UPS and what to do after a power failure. Tasks covered include:

- Configuring a PowerTrust UPS on V-Class systems
- Monitoring the PowerTrust UPS
- Handling ac power failures.

Configuring a PowerTrust UPS on V-Class systems

You can use SAM (Peripheral Devices) to modify the UPS configuration. You can specify:

- The serial port(s) to be used for the UPS.
- Whether to activate or deactivate the automatic shutdown feature entirely.
- The parameter `shutdown_delay_mins` (default = 1 minute). This parameter specifies the number of minutes to wait before initiating `shutdown -h` following notification that ac power to the UPS is lost. This interval allows the computer to continue operation through brief electrical glitches.

NOTE Consider increasing the value of `shutdown_delay_mins` if the site commonly experiences momentary power interruptions greater than one (1) minute for which recovery of power is expected.

- The parameter `shutdown_timeout_mins`. This parameter specifies the number of minutes to monitor the `shutdown -h` operation before initiating reboot with the halt option. In this way a reboot is executed even if the shutdown process changes.

NOTE Consider increasing the value of `shutdown_timeout_mins` if `shutdown`(1M) will take longer than 5 minutes (including the 60 second delay by `shutdown` (1M) after notification of users).

Consider decreasing the value of `shutdown_timeout_mins` if `shutdown` (1M) will take less than 5 minutes; small systems can take advantage of this.

For more information on UPS on HP-UX, see the man pages on `ups_mond`(1M) and `ups_conf`(4).

Monitoring the PowerTrust UPS

If the system has one or more PowerTrust Uninterruptible Power Systems (UPS), you can monitor their operation in three ways:

- Check the switches and indicators on the front panel of the PowerTrust UPS.
- Look at the teststation. Some messages from the UPS are displayed on the console.
- Look in a log file. UPS messages are recorded in a system log file. On HP-UX, the log files is /usr/adm/syslog, a text file which can be viewed with `vi` or another editor.

  For more information, see the *PowerTrust System Guide* (PN 3329-90002) that accompanied the PowerTrust UPS.

If ac power fails

If the automatic shutdown feature has been enabled, follow the procedure in “ac power failure with automatic shutdown enabled”. If the automatic shutdown feature is disabled, follow the procedure in “ac power fails without automatic shutdown enabled” on page 65.

ac power failure with automatic shutdown enabled

1. You will hear an alarm (a single beep every 10 seconds). If desired, you can turn off the auditory alarm by pressing the switch on the PowerTrust UPS labeled Silence Alarm.

CAUTION Do NOT press the switch on the PowerTrust UPS labeled Output On/ Output Off. Pressing this switch would shut off power to the node.

2. With the OS properly configured for UPS, you will see UPS messages on the console. See Appendix A in P*owerTrust System Guide* (PN 3329-90002) for explanations of the messages.
3. The indicators on the PowerTrust UPS front panel labeled AC Output and Battery Power will be lit.

If ac power does not return

If ac power does not return within the time specified by the parameter, shutdown_delay_mins (default = 1 minute), the `shutdown(1M)` command with the `-h` option is executed. Normal shutdown processing occurs. In case the shutdown process changes, the reboot command with the halt option is executed after an interval specified by shutdown_delay_mins (default = 5 minutes).

CAUTION Do not turn off power yourself, no matter what messages you see on the console. Power will be turned off automatically as part of the automatic shutdown. If you turn off power yourself, the system may be powered off automatically as it is rebooting, causing a hard crash.

If the unit nears full battery discharge

If the unit comes within approximately three (3) minutes of full battery discharge, the reboot command with the halt option will execute, as described in “If ac power does not return” on page 64. The PowerTrust UPS will issue three beeps every 10 seconds.

For more information, see the PowerTrust System Guide (PN 3329-90002) that accompanied the PowerTrust UPS.

The operating system should have been configured for the UPS as part of the system installation. See “Configuring a PowerTrust UPS on V-Class systems” on page 62.

ac power fails without automatic shutdown enabled

The PowerTrust Uninterruptible Power System (UPS) will supply power to the system from a battery for approximately 15 minutes if the battery is fully charged.

1. You will hear an auditory alarm (a single beep every 10 seconds). If desired, you can turn off the auditory alarm by pressing the switch on the PowerTrust UPS labeled Silence Alarm.

CAUTION Do NOT press the switch on the PowerTrust UPS labeled Output On/Output Off. Pressing this switch would shut off power to the node and other units connected to the UPS.

2. With the OS properly configured for UPS, you will see UPS messages on the console. See Appendix A in P*owerTrust System Guide* (PN 3329-90002) for explanations of the messages.
3. The indicators on the PowerTrust UPS front panel labeled “AC Output” and “Battery Power” will be lit.

If you anticipate a long power outage

If you anticipate a long period without ac power, you may want to shutdown your system during the 15 minutes of reserve power.

1. Follow the usual OS shutdown procedure.
2. Power off the protected equipment.
3. Set the Output On/Output Off switch to Output Off.

4. Set the UPS/BATTERY switch to Disable.
5. Set the MAIN and BYPASS Input breakers to the Off position.
6. Set the Output breakers to the Off position.
7. If the unit is to be stored or shipped, disconnect the battery cable.

If the unit nears full battery discharge

If the unit comes within approximately two (2) minutes of full battery discharge, the PowerTrust UPS will issue three beeps every 10 seconds. Prepare for an immediate loss of battery power.

For more information, see the PowerTrust System Guide (PN 3329-90002) that accompanied the PowerTrust UPS.

The operating system should have been configured for the UPSs as part of the system installation. See “Configuring a PowerTrust UPS on V-Class systems” on page 62.

7 Recovering from failures

This chapter provides detailed information on recovering from HP-UX system interruptions.

Usually, the first indication of a problem is that the system does not respond to user input. This lack of response indicates either a performance problem or system interruption.

Performance problems are generally characterized by:

- The system responds to one or more programs/users but not all, or sluggishly to others
- The system seems to be very slow

System interruptions usually result in a total loss of CPU resources for all users/programs due to a:

- System hang
- System panic
- HPMC

Collecting information

Providing the Response Center with a complete and accurate symptom description is important in solving any problem. The V-Class server's teststation automatically records information on environmental and system level events in several log files. See "Teststation file system" on page 38 for more information about these files.

Use the following procedure to collect troubleshooting information:

Step 1. If an error message is displayed on the system console, record it.

Step 2. Record the information displayed on the system LCD. See "LCD (liquid crystal display)" on page 22 for more information.

Step 3. Record any relevant information contained in the log files in the /spp/data directory on the teststation:

- event_log
  Main log file
- hard_hist
  Filtered output from the hard_logger, appended after each error

Step 4. Record any relevant information contained in the syslog.log file in the /var/adm/syslog directory on the system disk. Access to this log file may require rebooting the system if it has hung or crashed.

Performance problems

Performance problems are generally perceived as:

- Sluggish response at the operating system prompt
- Slow program execution
- Some users/programs unable to get a response

Use the following procedure to troubleshoot a performance problem:

Step 1. At the console window of the teststation, use one or both of the following commands to check for active processes making heavy use of system resources:

- `ps`
- `top`

See *Managing Systems and Workgroups* and the `ps` and `top` man pages for more information about options and usage.

Step 2. Enter a **Ctrl-C** from the terminal exhibiting the problem to abort an executing command.

Step 3. Check another terminal to verify that the problem is not just a console hang.

Step 4. Contact your CE or the HP Customer Response Center.

System hangs

System hangs are characterized by users unable to access the system, although the LCD display and attention light may not indicate a problem exists. The system console may or may not be hung.

Use the following procedure to troubleshoot a system hang:

Step 1. Press **Enter** at a terminal several times and wait for a response.

Step 2. Press **Ctrl-C** at a terminal to abort an executing command.

Step 3. Check another terminal to verify that the problem is not just a console hang.

Step 4. At the console window of the teststation, use one or both of the following utilities to communicate with the server:

- `ping`
- `telnet`

See the `ping` and `telnet` man pages for more information about options and usage.

Step 5. If possible, wait about 15 minutes to see if the computer is really hung or if it has a performance problem. With some performance problems, a computer may not respond to user input for 15 minutes or longer.

Step 6. If the computer is really hung, reset the server by issuing a `do_reset` command from the console window of the teststation.

`do_reset` *[node] [level]*

node	node to reset. Default is all nodes connected to the teststation.
level	level to reset. There are four:

1	Jtag controller—core utility board reset, hard reset, clear option bits and send messages to `ccmd`. (default)
2	Jtag controller—core utility board reset and system soft reset

3	Jtag controller—core utility board reset
4	TOC reset. Used to produce a crash dump

Step 7. Save the core dump file and contact the HP Response Center to have the core dump file analyzed. Refer to your service contract for the phone number of your HP Response Center. See “Fast dump” on page 79 for more information.

System panics

A system panic is the result of HP-UX encountering a condition that it is unable to respond to and halting execution.

System panics are rare and are not always the result of a catastrophe. They may occur on bootup, if the system was previously shut down improperly. Sometimes they occur as a result of hardware failure.

Recovering from a system panic can be as simple as rebooting the system. At worst, it may involve reinstalling HP-UX and restoring any files that were lost or corrupted. If the system panic was caused by a hardware failure such as a disk head crash, repairs have to be made before reinstalling HP-UX or restoring lost files.

NOTE It is important to maintain an up-to-date backup of the files on your system so that you can recover your data in the event of a disk head crash or similar situation. How frequently you update these backups depends on how much data you can afford to lose. For information on how to back up data, refer to *Managing Systems and Workgroups*.

After HP-UX experiences a system panic, the system:

- May display an HPMC tombstone on the console if panic was caused by an HPMC. A tombstone is a list of register values used for troubleshooting.
- May attempt to save a core file (an image of physical memory) to the dump device (by default this is the primary swap device).
- Attempts to reboot.
- Usually displays a panic message on the console. A panic message consists of several lines of text starting with the heading System Panic.
- May attempt to copy the core file to the file system (by default, to the directory /emp/syscore) if HP-UX can reboot.

Use the following procedure to troubleshoot a system panic:

Step 1. If an HPMC tombstone appears on the console, copy or print out the “Machine Check Parameters” field, and all information that follows them.

Step 2. Record the panic message displayed on the system console. Look for text on the console that contains terms like:

- System Panic
- HPMC
- Privilege Violation
- Data Segmentation Fault
- Instruction Segmentation Fault

Step 3. Categorize the panic message. The panic message describes why HP-UX panicked. Sometimes panic messages refer to internal structures of HP-UX (or its file systems) and the cause might not be obvious.

The wording of the panic message should allow you to classify the problem into one of the following areas:

- Peripheral problem
- Server or I/O card problem
- File system problem
- LAN communication problem
- Logical Volume Manager (LVM) related problem
- Other

Peripheral problem

Use the following procedure to troubleshoot an apparent peripheral hardware failure:

Step 1. Check to ensure the device is powered on and online.

CAUTION Do not connect or disconnect cables or power off or on SCSI devices while the V-Class server is powered on. Doing so could lead to corruption of disk data.

Step 2. Check the device’s error display. If an error is displayed:

1. Record the error message or display.

2. Take the device offline.

3. Power down the device.

4. If it is a disk drive, wait for the disk to stop spinning.

5. Power up the device.

6. Place the device back online.

Step 3. Check to ensure the device address or ID is correct.

Step 4. Check cable and terminator connections.

Step 5. If the system does not reboot by itself, reboot the computer by issuing the `reset` command in the console window or `do_reset` command at the tsh-shell window. For more information about rebooting the system see “Rebooting the system” on page 77.

Step 6. If the problem reappears on the device you may have an interface card or system problem. See “Interface card and system problem”.

If the problem reappears, it might be necessary to have the problem fixed by Hewlett-Packard service personnel.

Interface card and system problem

Use the following procedure if a hardware failure appears to be associated with an interface card or with the an internal component of the system:

Step 1. If an HPMC tombstone is displayed, record it.

Step 2. Record the information displayed on the LCD. See “LCD (liquid crystal display)” on page 22 for more information.

Step 3. Record any relevant information contained in the following log files in the /spp/data directory on the teststation:

- event_log
  Main log file

- hard_hist
  Filtered output from the hard_logger, appended after each error

- consolelog
  Complete log of all input/output from the sppconsole window

Step 4. If the system does not reboot by itself, reboot the computer by issuing the `reset` command in the console window or `do_reset` command at the tsh-shell window. For more information about rebooting the system see “Rebooting the system” on page 77.

If the problem reappears, it might be necessary to have the problem fixed by Hewlett-Packard service personnel.

File system problem

If the panic message indicates a problem with one of your file systems, reboot the system and run the file system checker `fsck` to check and correct the problem. Follow all directions that fsck gives you. Especially when your root file system (the one with the / directory) has problems, it is important to use the `-n` option to the `reboot` command, right after fsck completes. `fsck` is normally run automatically at boot time. See “Rebooting the system” on page 77.

LAN communication problem

Use the following procedure if the panic messages indicate a problem with LAN communication:

Step 1. Check LAN cable and media access unit (MAU) connections.

Step 2. Ensure that all vampire taps are tightly connected to their respective cables.

Step 3. Ensure that the LAN is properly terminated. Each end of the LAN cable must have a 50-ohm terminator attached. Do not connect the system directly to the end of a LAN cable.

If the problem reappears or if the hardware failure appears to be associated with a LAN card or an internal component of the V-Class server, it might be necessary to have the problem fixed by Hewlett-Packard service personnel.

Logical Volume Manager (LVM) related problem

If the size of a logical volume that contains a file system is reduced such that the logical volume is smaller than the file system within it, the file system will be corrupted. When an attempt is made to access a part of the truncated file system that is beyond the new boundary of the logical volume a system panic will often result.

The problem might not show up immediately. It will occur when the truncated part of the file system is overwritten by something else (such as a new logical volume or the extension of a logical volume in the same volume group as the truncated file system).

For more information on LVM, see the *Managing Systems and Workgroups*.

Recovery from other situations

When you suspect the problem was something other than previously discussed or you do not know where to classify it, proceed to “Rebooting the system” on page 77. Be sure to record the exact text of the panic message be recorded for future troubleshooting purposes. See “Collecting information” on page 68 for further information.

Rebooting the system

Once you have corrected any problem, you are ready to reset and reboot your system.

Step 1. Reset the V-Class node using one of four different methods:

- Power cycle the V-Class node by turning the keyswitch to the DC OFF position then back to the On position. See “Key switch panel” on page 15 for more information.
- Press the keyswitch panel TOC button.
- Type `reset` into console window (if in menu or fourth mode only).
- `do_reset` executed from one of the teststation tsh windows.

There may be differences in the boot up displays/activities as compared with your normal boot up sequence.

The system may have saved a system core file to disk. See “Abnormal system shutdowns” on page 79.

Step 2. If the system panicked due to a corrupted file system, fsck will report the errors and any corrections it makes. If fsck terminates and requests to be run manually, refer to the *Managing Systems and Workgroups* for further instructions. If the problems were associated with your root file system, fsck will ask you to reboot your system when it finishes. Use the command:

```
reboot -n
```

The `-n` option tells reboot not to sync the file system before rebooting. Since fsck has made all the corrections on disk, you do not want to undo the changes by writing over them with the still corrupt memory buffers.

Monitoring the system after a system panic

If the system successfully reboots, there is a good chance that you can resume normal operations. Many system panics are isolated events, unlikely to reoccur.

Check applications to be sure that they are running properly and monitor the system closely for the next 24 hours. For a short while, backups may be done more frequently than normal until confidence in the system has been restored.

Abnormal system shutdowns

Abnormal systems shutdowns (often referred to as system crashes) can occur for many reasons. In some cases, the cause of the crash can be easily determined. In some extreme cases, however, you may need to analyze a snapshot (called a core dump or simply dump) of the computer's memory in order to determine the cause of the crash or have Hewlett-Packard do it for you.

V-Class servers using HP-UX Release 11.0 or greater employ a more efficient dump mechanism than other HP servers using previous releases of HP-UX. This mechanism is called *fast dump*.

Fast dump

When a system crashes, you can now choose whether or not to dump, and if so, whether the dump should contain the relevant subset of memory or all memory (without operator interaction).

By default fast dump selectively dumps only the parts of memory that are expected to be useful in debugging. It improves system availability in terms of both the time and space needed to dump and analyze a large memory system.

The following commands allow you to configure, save, and manipulate the fast core dump:

- `crashconf`—Configures the destination and contents of a crash dump without rebooting. See the crashconf(1M) man page for more information.
- `savecrash`—Runs at boot time and saves any information that may be overwritten by normal system activity. See the savecrash(1M) man page for more information.
- `crashutil`—Saves or manipulates the crash dump (if desired). It can format the dump snapshot so that it can be read by the older commands. See the crashutil(1M) man page for more information.

Installations that used to call `savecore` in any way other than by the HP-supplied, unmodified /sbin/init.d/savecore script need to be updated to use `savecrash` and/or `crashutil`.

The on-disk and file system formats of a crash dump have changed with HP-UX 11.0.

`libcrash`(3) is a new library provided to allow programmatic access to a crash dump. supports all past and current crash dump formats. By using `libcrash`(3) under certain configurations, crash dumps no longer need to be copied into the file system before they can be debugged. See the `libcrash`(3) man page for more information.

Overview of the dump and save cycle

When the system crashes, HP-UX saves the image of physical memory or certain portions of it to predefined locations called dump devices. When you next reboot the system, a special utility copies the memory image from the dump devices to the HP-UX file system area. Once copied, the memory image can be analyzed with a debugger or saved to tape for later analysis.

Prior to HP-UX 11.0, dump devices had to be defined in the kernel configuration, and they still can be using Release 11.0. Beginning with Release 11.0, however, a new more-flexible method for defining dump devices is available using `crashconf`.

Beginning with HP-UX Release 11.0, there are three places where dump devices are configured:

1. In the kernel (same as releases prior to Release 11.0)
2. During system initialization when the initialization script for `crashconf` runs (and reads entries from the `/etc/fstab` file)
3. During runtime, by the operator or administrator manually running the `/sbin/crashconf` command.

Crash dump destination and contents

Defining the contents and destination of the crash dump are two important factors to consider when preparing for the dump. The destination and contents are configurable without rebooting, using the `crashconf` interface. See the crashconf(1M) man page for more information.

In order to capture the memory image of the system when a crash occurs, the image storage location(s) must be defined in advance. They can be on local disk devices or logical volumes.

It is important to have sufficient space to capture the part of memory that contains the instruction or data that caused the crash. You can define more than one dump device so that if the first one fills up, the next one continues dumping until the dump is complete or no more defined space is available. To ensure enough dump space, define a dump area that is at least as big as your computer's physical memory plus 1 Mbyte.

Setting the amount of memory dumped and the classes of the memory pages determines the size of the dump. The content can be configured while the system is running and changed without rebooting the system. The larger the size of the system's physical memory, the longer it takes to dump it to disk (and the more disk space it consumes).

Configuration criteria

There are three main criteria to consider when making decisions regarding how to configure system dumps. The criteria are:

- System recovery time—Get the system back up as soon as possible
- Crash information integrity—Capture the correct information
- Disk space needs—Conserve available disk space

System recovery time

To get the system back up and running as soon as possible, consider the following factors:

- Dump level
- Compressed save vs. noncompressed save
- Using a device for both paging and dumping
- Partial save

These factors are discussed in the following sections.

Dump level

With HP-UX 11.0 you can select three levels of core dumps: no dump, selective dump, or full dump.

Selective dump causes only the selected memory pages to get dumped (see the crashconf(1M) man page for more information).

NOTE In some specific cases, HP-UX will override the selective dump and request a full dump. The operator is given ten seconds to override HP-UX and continue with a selective dump.

The fewer pages dumped to disk (and on reboot, copy to the HP-UX file system area), the faster the system can be back up and running. Therefore, avoid using the full dump option.

When defining dump devices, whether in a kernel build or at run time, you can list which classes of memory *must always get dumped*, and which classes of memory *should not be dumped*. If you leave both of these “lists” empty, HP-UX decides for you which parts of memory should be dumped based on what type of error occurred. In nearly all cases, leaving the lists empty is preferred.

NOTE Even if you have defined (in the kernel or at run time) that you do not want a full dump to be performed, you can override the definitions (within a ten second window) and request a full dump after a system crash. Likewise, if you know what caused the crash, you can override those definitions and request no dump be performed.

Compressed save vs. noncompressed save

System dumps can be so large that they tax the HP-UX file system area.

The boot time utility, `savecrash`, can be configured (by editing the file `/etc/rc.config.d/savecrash`) to compress or not compress the data as it copies the memory image from the dump devices to the HP-UX file system area during the reboot process. This effects system recovery time in that data compression takes longer. Therefore, if you have enough disk space and require the fastest system recovery, configure `savecrash` to not compress the data. See the savecrash(1M) man pages for more information.

Using a device for both paging and dumping

It is possible to use a specific device for both paging (swapping) and as a dump device. If system recovery time is critical, do not configure the primary paging device as a dump device.

When the primary paging device is not used as one of the dump devices or after the crash image on the primary paging device has been saved, by default, `savecrash` runs in the background. This reduces system boot time by running the system with only the primary paging device.

Another advantage to keeping paging and dump devices separate is that paging does not overwrite the information stored on a dump device, no matter how long the system has been up or how much activity has taken place. Disabling `savecrash` processing at boot time (by editing the file `/etc/rc.config.d/savecrash`) reduces system recovery time. After the system recovery, you can run `savecrash` manually to copy the memory image from the dump area to the HP-UX file system area.

Partial save

If a memory dump resides partially on dedicated dump devices and partially on devices that are also used for paging, you can save only those pages that are endangered by paging activity.

Pages residing on the dedicated dump devices can remain there. If you know how to analyze memory dumps, it is even possible to analyze them directly from the dedicated dump devices using a debugger that supports this feature. If, however, you need to send your memory dump to someone else for analysis, you first have to finish moving the pages on the dedicated dump devices to the HP-UX file system area, so that you can use a utility such as `tar` to bundle them for shipment. To do that, use the command `/usr/sbin/crashutil` instead of `savecrash` to complete the copy.

Crash information integrity

This section discusses how to make sure you capture the part of memory that contains the instruction or piece of data that caused the crash. The factors you have to consider are:

- Full dump vs. selective dump
- Dump definitions built into the kernel vs. defined at runtime
- Using a device for both paging and as a dump device

Full dump vs. selective dump

The only way to guarantee capturing the specific instruction or data that caused the crash is to capture everything. This means selecting a full dump of memory.

Be aware, however, that this can be costly in terms of time and disk space. A large amount of time and disk space is needed to dump the entire contents of memory in a system with a large memory configuration or to copy a large memory image to the HP-UX file system area during the reboot process.

The amount of dump area should at least be equal to the amount of memory in your system; depending on a number of factors, additional disk space greater than the amount of physical memory in your system may be needed

Dump definitions built into the kernel vs. defined at runtime

There are three places that you can define which devices are to be used as dump devices:

1. During kernel configuration
2. At boot time (entries defined in the `/etc/fstab` file)
3. At run time (using the `/sbin/crashconf` command)

Definitions at each of these places add to or replace any previous definitions from other sources. However, consider the following situation:

Example

A system called appserver has 1 Gbyte of physical memory. If you were to define dump devices for this system with a total of 256 Mbytes of space in the kernel file and then define an additional 768 Mbytes of disk space in the `/etc/fstab` file, you would have enough dump space to hold the entire memory image (a full dump).

If the crash occurs, however, before `/etc/fstab` is processed, only the amount of dump space already configured is available at the time of the crash; in this example, it is 256 Mbytes of space.

Define enough dump space in the kernel configuration if it is critical to capture every byte of memory in all instances, including the early stages of the boot process.

NOTE This example is presented for completeness. The actual amount of time between the point where kernel dump devices are activated and the point where runtime dump devices are activated is very small (a few seconds), so the window of vulnerability for this situation is practically nonexistent.

Using a device for both paging and as a dump device

It is possible to use a specific device for both paging purposes and as a dump device. If, however, crash dump integrity is critical, this is not recommended.

If `savecrash` determines that a dump device is already enabled for paging and that paging activity has already taken place on that device, a warning message indicates that the dump may be invalid. If a dump device has not already been enabled for paging, `savecrash` prevents paging from being enabled to the device by creating the file `/etc/savecore.LCK`. `swapon` does not enable the device for paging if the device is locked in `/etc/savecore.LCK`.

Systems configured with small amounts of memory and using only the primary swap device as a dump device might not be able to preserve the dump (copy it to the HP-UX file system area) before paging activity destroys the data in the dump area. Larger memory systems are less likely to need paging (swap) space during start-up and are therefore less likely to destroy a memory dump on the primary paging device before it can be copied.

Disk space needs

This section discusses how to manage limited disk resources on your system for the post-crash dump and/or the post-reboot save of the memory image. The factors to consider are:

- Dump level
- Compressed save vs. noncompressed save
- Partial save (savecrash -p)

Dump level

There are three levels of core dumps: full dump. selective dump, and no dump. The fewer pages required to dump, the less space is required to hold them. Therefore, a full dump is not recommended. If disk space is really at a premium, one option is no dump at all.

A third option is called a selective dump. HP-UX 11.0 can determine which pages of memory are the most critical for a given type of crash, and save only those pages. Choosing this option can save a lot of disk space on the dump devices and again later on the HP-UX file system area. For instructions on how to do this see “Defining dump devices” on page 88.

Compressed save vs. noncompressed save

Regardless of whether you choose full or selective dump, whatever is saved on the dump devices needs to be copied to the HP-UX file system area before you can use it.

NOTE With HP-UX 11.0, it is possible to analyze a crash dump directly from dump devices using a debugger that supports this feature. If, however, you need to save it to tape or send it to someone, copy the memory image to the HP-UX file system area first.

If there is a disk space shortage in the HP-UX file system area (as opposed to dump devices), you can elect to have `savecrash` (the boot time utility that does the copy) compress your data as it makes the copy.

Partial save (savecrash -p)

If the system has plenty of dump device space but is limited in HP-UX file system space, consider using the `-p` option for the `savecrash` command. This option copies only those pages on dump devices that are endangered by paging activity (i.e. pages on the devices used for both paging and as dump devices). Pages that are on dedicated dump devices remain there.

To configure this option into the boot process, edit the file `/etc/rc.config/savecrash` and comment out the line that sets the environment variable SAVE_PART=1.

Defining dump devices

When defining dump devices, it is important to accurately determine the amount of space needed to hold the dump without wasting disk space. To save a full dump, the amount of dump space needed is equal to the size of the system's physical memory.

For selective dumps, the size of dump space varies, depending on the classes of memory be saved. To determine amount of space needed, perform the following procedure:

Step 1. When the system is running with a typical workload, enter the following command:

```
/sbin/crashconf -v
```

The following typical output appears:

```
CLASS          PAGES  INCLUDED IN DUMP  DESCRIPTION
--------  ----------  ----------------  ------------------------------------

UNUSED          2036  no,  by default   unused pages
USERPG          6984  no,  by default   user process pages
BCACHE         15884  no,  by default   buffer cache pages
KCODE           1656  no,  by default   kernel code pages
USTACK           153  yes, by default   user process stacks
FSDATA           133  yes, by default   file system metadata
KDDATA          2860  yes, by default   kernel dynamic data
KSDATA          3062  yes, by default   kernel static data

Total pages on system:            32768
Total pages included in dump:      6208

DEVICE        OFFSET(kB)   SIZE (kB)  LOGICAL VOL.  NAME
------------  ----------  ----------  ------------  -------------------------
31:0x00d000        52064      262144  64:0x000002  /dev/vg00/lvol2
                          ----------
                              262144
```

Step 2. Multiply the number of pages listed in "`Total pages included in dump`" by the page size (4 Kbytes) and add 25% for a margin of safety. In the above example, the calculation would be:

`(6208 x 4 Kbytes) x 1.25 = approx. 30 Mbytes`

Kernel dump device definitions

Capturing dumps for crashes that occur during early stages of the boot process requires sufficient dump space in the kernel configuration.

Using SAM to configure dump devices into the kernel

The easiest way to configure dump devices is to use SAM. A screen for dump device definition is located in the Kernel Configuration area. After changing the dump device definitions, you *must* build a new kernel and reboot the system using the new kernel file to make the changes take effect. To configure dump devices into the kernel, perform the following procedure:

Step 1. Run SAM and select the Kernel Configuration Area

Step 2. From the Kernel Configuration Area, select the Dump Devices area

A list of dump devices configured into the next kernel built by SAM is displayed. This is the list of pending dump devices.

Step 3. Use the SAM action menu to add, remove, or modify devices or logical volumes.

NOTE The order of the devices in the list is important. Devices are used in reverse order from the way they appear in the list. The last device in the list is as the first dump device.

Step 4. Follow the SAM procedure for building a new kernel.

Step 5. Boot the system from the new kernel file to activate the new dump device definitions.

Using HP-UX commands to configure dump devices into the kernel

You can also edit your `system` file and use the `config` program to build the new kernel. For details see *Managing Systems and Workgroups*. Perform the following procedure to configure dump devices into the kernel using HP-UX commands:

Step 1. Edit `system` file (the file that `config` uses to build your new kernel). This is usually the file `/stand/system` but can be another file if you prefer.

Dump to Hardware Device—For each hardware dump device you want to configure into the kernel, add a dump statement in the area of the file designated "`* Kernel Device info`" immediately prior to any tunable parameter definitions. For example:

`dump 2/0/1.5.0`

`dump 56/52.3.0`

Dump to Logical Volume—For logical volumes, it is not necessary to define each volume used as a dump device. For dumping to logical volumes, the logical volumes must meet all of the following requirements:

- Each logical volume to be used as a dump device must be part of the root volume group (vg00). For details on configuring logical volumes as kernel dump devices, see the lvlnboot (1M) manpage.
- The logical volumes must be contiguous (no disk striping or bad-block reallocation is permitted for dump logical volumes).

- The logical volume cannot be used for file system storage, because the whole logical volume is used.

To use logical volumes for dump devices (no matter how many logical volumes you require), include the following dump statement in the system file:

```
dump lvol
```

Configuring No Dump Devices—To configure a kernel with no dump device, use the following dump statement in the system file:

```
dump none
```

To configured the kernel for no dump device, you must use the above statement (**`dump none`**).

NOTE Omitting dump statements altogether from the `system` file results in a kernel that uses the primary paging device (swap device) as the dump device.

Step 2. Once you have edited the `system` file, build a new kernel file using the `config` command.

Step 3. Save the existing kernel file (probably `/stand/vmunix`) to a safe place (such as `/stand/vmunix.safe`) in case the new kernel file can not be booted.

Step 4. Boot the system from the new kernel file to activate the new dump device definitions.

Runtime dump device definitions

If you are not concerned about capturing a dump that occurs during the earliest stages of the boot process, you can replace or supplement any kernel dump device definitions while the system is booting or running. There are two ways to do this:

1. Using `crashconf` to read dump entries in the `/etc/fstab` file (using `crashconf`'s `-a` option)
2. Using arguments to the `crashconf` command, directly specifying the devices to be configured

The /etc/fstab file

You can define entries in the `fstab` file to activate dump devices during the HP-UX initialization (boot) process or when `crashconf` reads the file. The format of a dump entry for `/etc/fstab` looks like the following:

devicefile_name `/ dump defaults 0 0`

Examples:

```
/dev/dsk/c0t3d0 / dump defaults 0 0
/dev/vg00/lvol2 / dump defaults 0 0
/dev/vg01/lvol1 / dump defaults 0 0
```

Define one entry for each device or logical volume you want to use as a dump device.

NOTE Unlike dump device definitions built into the kernel, with run time dump definitions you can use logical volumes from volume groups other than the root volume group.

The crashconf command

You can also use the `/sbin/crashconf` command to add to, remove, or redefine dump devices. The following are two ways to do this:

- Reread the `/etc/fstab` file using the `crashconf -a` option
- Use device arguments with `crashconf` to configure the devices

With either method, you can use the `crashconf -r` option to specify that new definitions replace, rather than add to, any previous dump device definitions.

Examples:

To have `crashconf` read the `/etc/fstab` file (thereby adding any listed dump devices to the currently active list of dump devices), enter the following command:

```
/sbin/crashconf -a
```

To have `crashconf` read the `/etc/fstab` file (thereby replacing the currently active list of dump devices with those defined in `fstab`), enter the following:

```
/sbin/crashconf -ar
```

To have `crashconf` add the devices represented by the block device files `/dev/dsk/c0t1d0` and `/dev/dsk/c1t4d0` to the dump device list, enter the following:

`/sbin/crashconf /dev/dsk/c0t1d0/dev/dsk/c1t4d0`

To have `crashconf` replace any existing dump device definitions with the logical volume `/dev/vg00/lvol3` and the device represented by block device file `/dev/dsk/c0t1d0`, enter the following:

`/sbin/crashconf -r /dev/vg00/lvol3 /dev/dsk/c0t1d0`

Dump order

The order that devices dump after a system crash is important when using the primary paging device along with other devices as a dump device.

Regardless of how the list of currently active dump devices was built (from a kernel build, from the `/etc/fstab` file, from use of the `crashconf` command, or any combination of the these) dump devices are used (dumped to) in the reverse order from which they were defined. The last dump device in the list is the first one used, and the first device in the list is the last one used.

Place devices that are used for both paging and dumping early in the list of dump devices so that other dump devices are used first and overwriting of dump information due to paging activity is minimized.

What happens when the system crashes?

This section discusses the unlikely event of a V-Class system crash. A system panic means that HP-UX encountered a condition that it could not to handle. Sometimes the cause of the crash is apparent, but many times an in-depth analysis is required. HP-UX is equipped with a dump procedure to capture the contents of memory at the time of the crash.

Operator override options

When the system crashes, the system console displays a panic message similar to the following:

```
*** A system crash has occurred.  (See the above messages for details.)
*** The system is now preparing to dump physical memory to disk, for use
*** in debugging the crash.


*** The dump will be a SELECTIVE dump:  21 of 128 megabytes.

*** To change this dump type, press any key within 10 seconds.


*** Select one of the following dump types, by pressing the corresponding key:
N) There will be NO DUMP performed.
S) The dump will be a SELECTIVE dump:  21 of 128 megabytes.
F) The dump will be a FULL dump of 128 megabytes.
O) The dump will be an OLD-FORMAT dump of 128 megabytes.
*** Enter your selection now.
```

The operator can override any dump device definitions by entering N (for no dump) at the system console within the 10-second override period.

If disk space is limited, but the operator feels that a dump is important, the operator can enter S (for selective dump) regardless of the currently defined dump level.

The dump

After the operator overrides the current dump level, or the 10-second override period expires, HP-UX writes the physical memory contents to the dump devices until one of the following conditions is true:

- The entire contents of memory are dumped (if a full dump was configured or requested by the operator).
- The entire contents of selected memory pages are dumped (if a selective dump was configured or requested by the operator).
- Configured dump device space is exhausted

Depending on the amount of memory being dumped, this process can take from a few seconds to hours.

NOTE

During the dump, status messages on the system console indicate the progress. You can interrupt the dump at any time by pressing the **ESC** key. However, if you interrupt a dump, all information is lost.

Following the dump, the system attempts to reboot.

The reboot

When dumping of physical memory pages is complete, the system attempts to reboot (if the Autoboot is set). For information on the Autoboot flag, see “Enabling Autoboot” on page 49.

savecrash processing

During the boot process, you can run a process called `savecrash` that copies (and optionally compresses) the memory image stored on the dump devices to the HP-UX file system area.

Dual-mode devices (dump / swap)

By default, `savecrash` performs its copy during the boot process. You can disable this operation by editing the file: `/etc/rc.config.d/savecrash` and setting the SAVECRASH environment variable to a value of zero. This is generally safe to do if the dump devices are not also being used as paging devices.

CAUTION If using devices for both paging and dumping, do not disable savecrash boot processing. Loss of the dumped memory image to subsequent system paging activity can occur.

What to do after the system has rebooted?

After the system reboots, make sure that the physical memory image dumped to the dump devices is copied to the HP-UX file system area then you can either package and send it in for analysis or analyze it yourself using a debugger.

NOTE With HP-UX 11.0, it is possible to analyze a crash dump directly from dump devices. If, however, you need to save it to tape or send it to someone, first copy the memory image to the HP-UX file system area.

Unless you specifically disable `savecrash` processing during reboot (see savecrash processing), the `savecrash` utility copies the memory image for you during the reboot process. The default HP-UX directory for the memory image in is `/var/adm/crash`. You can specify a different location by editing the file `/etc/rc.config.d/savecrash` and setting the environment variable called SAVECRASH_DIR to the name of the directory the dumps are to be located.

Using crashutil to complete the saving of a dump

If you are using devices for both paging (swapping) and dumping, it is very important to not disable savecrash processing at boot time. If you do, there is a chance that the memory image in the dump area will be overwritten by normal paging activity. If, however, there are separate dump and paging devices (no single device used for both purposes), you can delay copying the memory image to the HP-UX file system area in order to speed up the boot process. To do this, edit the file /etc/rc.config.d/savecrash and set the environment variable called SAVECRASH=0.

If you have delayed copying the physical memory image from the dump devices to the HP-UX file system area, run savecrash manually to do the copy when the system is running and when you have configured enough space to hold the copy in your HP-UX file system area.

If you chose to do a partial save, the only pages copied to your HP-UX file system area during the boot process are those that were on paging devices. Pages residing on dedicated dump devices are still there. A partial save can be selected by leaving the SAVECRASH environment set to 1 and setting the environment variable SAVE_PART=1 in /etc/rc.config.d/savecrash. To copy the remaining pages to the HP-UX file system area when your system is running again, use the command crashutil. See the crashutil(1M) manpage for details.

Example

```
/usr/sbin/crashutil -v CRASHDIR /var/adm/crash/
crash.0
```

Crash dump format conversion

If you are analyzing a crash dump on a computer running a different version of HP-UX than the V-Class server, or if you are using a debugging tool that does not recognize the specific format of the saved file, you may not be able to analyze the crash dump in its current format. Use `crashutil` to convert the file format.

The basic format of the `crashutil` command to do a conversion is:

`/usr/sbin/crashutil -v` *version source* [*destination*]

version — Designates the version of the destination format.

source — Designates the pathname of the crash dump to be converted.

destination — Designates the pathname where the converted file will be written. If no *destination* is specified the *source* will be overwritten.

See the crashutil(1M) manpage for more information.

Analyzing crash dumps

Analyzing crash dumps is not a trivial task. It requires intimate knowledge of HP-UX internals and the use of debuggers. It is beyond the scope of this document to cover the actual analysis process. If you need help analyzing a crash dump, contact your Hewlett-Packard representative.

Index