

Operator's Guide

HP 9000 V2500/V2600 SCA Server

First Edition



A5845-96001

Customer Order Number: A5845-90001

July 1999

Printed in: USA

Revision History

Edition: First

Document Number: A5845-96001

Notice

© Copyright Hewlett-Packard Company 1999. All Rights Reserved.
Reproduction, adaptation, or translation without prior written permission is prohibited, except as allowed under the copyright laws.

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.

Contents

Preface	xiii
Notational conventions	xiv
Safety and regulatory information	xvi
Safety in material handling	xvi
USA radio frequency interference FCC Notice	xvi
Japanese radio frequency interference VCCI	xvii
EMI statement (European Union only)	xvii
Digital apparatus statement (Canada)	xvii
BCIQ (Taiwan)	xvii
Acoustics (Germany)	xviii
IT power system	xviii
High leakage current	xviii
Installation conditions (U.S.)	xix
Fuse cautions	xix
Associated documents	xx
Technical assistance	xxi
Reader feedback	xxii
1 Overview	1
V-Class System Components	2
The Service Support Processor	3
Server Console and Diagnostic Connections	4
V-Class Server Architecture	6
V2500/V2600 Crossbar Interconnection	6
V2500/V2600 Cabinet Components	8
Core Utilities Board	9
Processors	9
Memory	9
Input/Output	12
Multiple-Cabinet Server Connections	15
V2500/V2600 Cabinet Configurations	18
2 Indicators, switches, and displays	21
Operator panel	22
Key switch panel	23
Key switch	23
DC ON LED	23
TOC	24

DVD-ROM drive	24
Disk loading slot	24
Busy indicator	25
Eject button	25
Optional DAT drive	25
LEDs	25
Eject button	26
System Displays	27
LCD (Liquid Crystal Display)	28
Node status line	28
Processor status line	28
Message display line	30
Attention light bar	31
Environmental errors	32
3 SSP operation	35
SSP and the V-Class system	36
SSP log-on	37
SSP sppuser windows	37
Message window	40
Console window (sppconsole - complex console)	40
Console window (sppconsole - Node <i>X</i> console)	40
Console bar	40
ksh shell windows	40
Using the CDE (Common Desktop Environment) Workspace menu	41
CDE Workspace menu	41
Using the console	45
Creating new console windows	45
Starting the console	45
Starting the console from the Workspace menu	46
Starting the console using the sppconsole command	46
Starting the console using ts_config	47
Starting the console using the consolebar	48
Starting the console by logging back on	49
Console commands	49
Watching the console	50
Assuming control of the console	51
Changing a console connection	52
Accessing system logs	52
The set_complex command	52
Targeting commands to nodes	53
SSP file system	54
/spp/etc	54

/spp/bin	55
/spp/scripts	55
/spp/data/ <i>complex_name</i>	55
/spp/firmware	56
/spp/est	56
/spp/man	56
Device files	57
System log pathnames.....	58
4 Firmware (OBP and PDC)	59
Boot sequence.....	60
Boot process output	62
HP mode boot menu.....	64
Enabling Autoboot	67
Syntax	67
Examples.....	68
HElp command	69
Syntax.....	69
Examples.....	69
5 Configuration utilities	71
ts_config.....	72
Starting ts_config	72
ts_config operation	73
Configuration procedures.....	75
Upgrade JTAG firmware	75
Configure a Node.....	77
Configure the scub_ip address	81
Reset the Node	82
Deconfigure a Node.....	84
Add/Configure the Terminal Mux	84
Remove terminal mux.....	85
Console sessions	85
V2500/V2600 SCA (multinode) configuration.....	87
V2500/V2600 split SCA configuration.....	92
ts_config files	95
SSP-to-system communications	97
LAN communications.....	98
SSP host name and IP addresses	98
Serial communications.....	99
ccmd	100
xconfig	102
Menu bar.....	105

Node configuration map	106
Node control panel	107
Configuration utilities	110
autoreset	110
est_config	110
report_cfg	111
Effects of hardware and software deconfiguration	112
report_cfg summary report	112
report_cfg ASIC report	113
report_cfg I/O report	113
report_cfg memory report	114
report_cfg processor report	115
xsecure	115
6 HP-UX Operating System	117
HP-UX on the V2500/V2600	118
Displaying System Information	118
Listing the Server Hardware Configuration	118
Configuring HP-UX for V-Class Servers	120
HP-UX parameter sets	120
Multiple-cabinet kernel configurations	121
Process and Thread “Gang Scheduling”.	122
HP-UX 11.10 SCA Enhancements	122
HP-UX SCA Features	123
Starting HP-UX	125
Power-On Sequence.	126
Boot variables	127
Reviewing the state of the file system	128
Stopping HP-UX.	130
Shutdown considerations	130
Rebooting the system	132
Shutting down the system	133
Resetting the V2500/V2600 server hardware	134
7 Recovering from failures	137
Collecting information	138
Performance problems	139
System hangs	140
System panics.	141
Peripheral problem	142
Interface card and system problem	143
File system problem	144
LAN communication problem.	144

Logical Volume Manager (LVM) related problem	145
Recovery from other situations	145
Rebooting the system	146
Monitoring the system after a system panic.	146
Abnormal system shutdowns	147
Fast dump	147
Overview of the dump and save cycle	148
Crash dump destination and contents	148
New SCA-Extended Crash Dump Format	149
Memory Dumped on V2500/V2600 SCA Servers.	149
Configuration criteria	150
System recovery time	150
Crash information integrity	152
Disk space needs	154
Defining dump devices	155
Kernel dump device definitions	156
Runtime dump device definitions.	158
Dump order	160
What happens when the system crashes?.	160
Operator override options.	161
The dump.	161
The reboot	162
What to do after the system has rebooted?.	162
Using crashutil to complete the saving of a dump	163
Crash dump format conversion	163
Analyzing crash dumps.	164

Appendix A: LED codes 165

Power on detected errors	166
CUB detected memory power fail	171
CUB detected processor error	172
CUB detected I/O error	173
CUB detected fan error	174
CUB detected ambient air errors	175
CUB detected hard error.	176
CUB detected intake ambient air error	177
CUB detected dc error	178

Figures

Figure 1	Japanese radio frequency notice.	xvii
Figure 2	BCIQ (Taiwan).	xviii
Figure 3	V-Class Server Components: Cabinet and Service Support Processor.	2
Figure 4	Four-Cabinet V2500/V2600 Server Components.	3
Figure 5	Console and Diagnostic Connections for a Four-Cabinet V2500/V2600 Server.	5
Figure 6	Functional Diagram of a Single-Cabinet V2500/V2600 Server	7
Figure 7	V2500/V2600 HyperPlane Crossbar Connections	8
Figure 8	Conceptual Overview of V2500/V2600 Memory Board	10
Figure 9	Numbering and Locations of Single-Cabinet V2500/V2600 PCI I/O	13
Figure 10	Numbering and Locations of Multiple-Cabinet V2500/V2600 PCI I/O	14
Figure 11	Four-Cabinet V2500/V2600 Server CTI Cable Connections	16
Figure 12	Sample V2500/V2600 Cabinet Configurations	19
Figure 13	Operator panel.	22
Figure 14	Key switch panel	23
Figure 15	DVD-ROM drive	24
Figure 16	DDS-3 DAT drive front panel.	25
Figure 17	System displays.	27
Figure 18	Front panel LCD	28
Figure 19	SSP user windows for V2500/V2600 servers with one node	38
Figure 20	SSP user windows for V2500/V2600 servers with more than two nodes	39
Figure 21	SSP Workspace submenus for V2500/V2600	42
Figure 22	SSP Workspace submenus for V2500/V2600	42
Figure 23	SSP file system for V2500/V2600 servers	54
Figure 24	Boot process	61
Figure 25	ts_config sample display.	73
Figure 26	ts_config showing node 0 highlighted	76
Figure 27	ts_config “Upgrade JTAG firmware” selection.	76
Figure 28	Upgrade JTAG firmware confirmation panel	77
Figure 29	ts_config power-cycle panel	77
Figure 30	ts_config indicating Node 0 as not configured.	78
Figure 31	ts_config “Configure Node” selection.	78
Figure 32	ts_config node configuration panel	79
Figure 33	ts_config restart workspace manager panel.	80
Figure 34	ts_config indicating Node 0 is configured	80
Figure 35	ts_config “Configure ‘scub_ip’ address” selection	81
Figure 36	ts_config “SCUB OK” panel	81
Figure 37	ts_config scub_ip address configuration confirmation	82
Figure 38	ts_config scub_ip address set confirmation panel.	82
Figure 39	ts_config “Reset Node” selection	83
Figure 40	ts_config node reset panel	83

Figure 41	ts_config “Add/Configure Terminal Mux” selection	84
Figure 42	Terminal mux IP address panel	85
Figure 43	“Start Console Session” selection	86
Figure 44	Started console sessions	86
Figure 45	SSP supporting two single-node complexes	87
Figure 46	ts_config Configure Multinode complex selection	88
Figure 47	Configure Multinode Complex dialog window	88
Figure 48	Configure Multinode Complex dialog window with appropriate values.	90
Figure 49	Configuration started information box.	90
Figure 50	ts_config showing newly configured complexes.	92
Figure 51	ts_config Split Multinode complex operation.	93
Figure 52	ts_config Split Multinode complex panel	93
Figure 53	ts_config Split Multinode complex panel filled in	94
Figure 54	Split Multinode confirmation panel	94
Figure 55	ts_config Split Multinode operation complete	94
Figure 56	SSP-to-system communications	97
Figure 57	xconfig window—physical location names	103
Figure 58	xconfig window—logical names	104
Figure 59	xconfig window menu bar.	105
Figure 60	xconfig window node configuration map	106
Figure 61	xconfig window node control panel	108

Tables

Table 1	Valid CTI cache sizes	.12
Table 2	Indicator LED operation	.26
Table 3	Processor initialization steps	.29
Table 4	Processor run-time status codes	.29
Table 5	Message display line	.30
Table 6	Commands for creating console windows	.45
Table 7	sppconsole commands	.49
Table 8	Device file differences	.57
Table 9	System log pathnames	.58
Table 10	Boot menu commands	.65
Table 11	ts_config status values	.74
Table 12	report_cfg options	.111
Table 13	Hardware Path Numbering for V2500/V2600 Cabinets	.119
Table 14	Boot variables	.128
Table 15	CUB detects power on error	.166
Table 16	CUB detects memory power fail	.171
Table 17	CUB detects processor power fail	.172
Table 18	CUB detects I/O (IOB) power fail	.173
Table 19	CUB detects fan power fail	.174
Table 20	CUB detects ambient air error	.175
Table 21	Hard error	.176
Table 22	Ambient air (intake) error	.177
Table 23	dc error	.178

Preface

The *Operator's Guide HP 9000 V2500/V2600 Server* documents the information necessary to operate and monitor HP V-Class servers. This book is intended to be a reference for system administrators, system operators, and 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, <code>monospace</code> identifies command names, system calls, and data structures and types. In command examples, <code>monospace</code> identifies command output, including error messages.
<i>italic</i>	In paragraph text, <i>italic</i> identifies titles of documents. In command syntax diagrams, <i>italic</i> identifies variables that you must provide. The following command example uses brackets to indicate that the variable <i>output_file</i> is optional: command <i>input_file</i> [<i>output_file</i>]
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 **Cautions highlight procedures or information necessary to avoid injury to personnel. The caution should tell the reader exactly what will result from what actions and how to avoid them.**

WARNING **A warning highlights procedures or information necessary to avoid damage to equipment, damage to software, loss of data, or invalid test results.**

Safety and regulatory information

For your protection, this product has been tested to various national and international regulations and standards. The scope of this regulatory testing includes electrical/mechanical safety, radio frequency interference, ergonomics, acoustics, and hazardous materials. Where required, approvals obtained from third-party test agencies are shown on the product label.

Safety in material handling

CAUTION

Do not lift the node manually. To avoid physical injury you must use a mechanical lifting device.

USA radio frequency interference FCC Notice

The Federal Communications Commission (in CFR Part 15) has specified that the following notice be brought to the attention of the users of this product.

NOTE

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 Hewlett-Packard could result in the equipment being noncompliant with FCC Class A requirements and void the user's authority to operate the equipment.

Japanese radio frequency interference VCCI

Figure 1

Japanese radio frequency notice

この装置は、クラスA情報技術装置です。この装置を家庭環境で使用すると電波妨害を引き起こすことがあります。この場合には使用者が適切な対策を講ずるよう要求されることがあります。

VCCI-A

This equipment is a Class A category (Information Technology Equipment to be used in commercial and /or industrial areas) and conforms to the standards set by the Voluntary Control Council for Interference by Information Technology Equipment aimed at preventing radio interference in commercial and/or industrial areas.

Consequently, when used in a residential area or in an adjacent area thereto, radio interference may be caused to radios and TV receivers, etc. Read the instructions for correct handling.

EMI statement (European Union only)

This is a Class A product. In a domestic environment this product may cause radio interference in which case the user may be required to take adequate measures.

Digital apparatus statement (Canada)

This Class A digital apparatus meets all requirements of the Canadian Interference-Causing Equipment Regulations.

Cet appareil numérique de la classe A respecte toutes les exigences du Règlement sur le matériel brouilleur du Canada.

BCIQ (Taiwan)

This product has been reviewed, evaluated by GesTek Taiwan and is fully compliant to CNS 13438 (CISPR 22: 1993) Class A.

Figure 2

BCIQ (Taiwan)

檢磁 3862H354

Acoustics (Germany)

Laermangabe (Schalldruckpegel LpA) gemessen am fiktiver Arbeitsplatz bei normalem Betrieb nach DIN 45635, Teil 19: LpA =65.3 dB.

Acoustic Noise (A-weighted Sound Pressure Level LpA) measured at the bystander position, normal operation, to ISO 7779: LpA = 65.3 dB.

IT power system

This product has not been evaluated for connection to an IT power system (an AC distribution system having no direct connection to earth according to IEC 950).

High leakage current

CAUTION

High leakage current. Ground (earth) connection essential before connecting the supply.

Attention

Forts courants de peretes. Connection a une borne de terre est essentielle avant tout raccord electrique.

Achtung

Hoher ableitstrom. Vor inbetriebnahme schutzleiterverbindung herstellen.

Installation conditions (U.S.)

See installation instructions before connecting to the supply.

Voir la notice d'installation avant de raccorder au réseau.

CAUTION

Please note the following conditions of installation:

An insulated earthing conductor that is identical in size, insulation material, and thickness to the earthed and unearthed branch-circuit supply conductors except that it is green with or without one or more yellow stripes is to be installed as part of the branch circuit that supplies the unit or system. The earthing conductor described is to be connected to earth that the service equipment or, if supplied by a separately derived system, at the supply transformer or motor-generator set.

The attachment-plug receptacles in the vicinity of the unit or system are all to be of an earthing type, and the earthing conductors serving these receptacles are to be connected to earth at the service equipment.

CAUTION

For supply connections, use wires suitable for at least 60 °C.

Utiliser des fils convenant à une température de 60 °C pour les connexions d'alimentation.

Fuse cautions

CAUTION

Disconnect power before changing fuse.

Attention

Couper le courant avant de remplacer le fusible.

CAUTION

For continued protection against risk of fire, replace fuses only with same type and rating.

Attention

Pour ne pas compromettre la protection contre les risques d'incendie, remplacer par un fusible de même type et de mêmes caractéristiques nominales.

Associated documents

Associated documents include:

- *HP Diagnostic Guide: V2500/V2600 Servers*, (A5824-96002)
- *HP-UX SCA Programming and Process Management White Paper*
 - Available in `/usr/share/doc` for HP-UX 11.10
- *HP-UX 11.0 Configurable Kernel Parameters*
 - Available online at: <http://docs.hp.com/hpux/os>
- *HP-UX 11.10 Installation and Configuration Notes HP V2500 Servers*, (A5532-90005)
- *HP V-Class Server HP-UX Configuration Notes (for 11.0)*, (A4801-90001)
- *Managing Systems and Workgroups*, (B2355-90157)
- *PA-RISC 2.0 Architecture Reference Manual*, (ISBN 0-13-182734-0)
- *V2500 SCA HP-UX System Guide*, (A5532-90003)

Technical assistance

If you have questions that are not answered in this book, contact the Hewlett-Packard Response Center at the following locations:

- Within the continental U.S., call 1 (800) 633-3600.
- All others, contact your local Hewlett-Packard Response Center or sales office for assistance.

Reader feedback

This document was produced by the System Supportability Lab Field Engineering Support organization (SSL/FES). If you have editorial suggestions or recommended improvements for this document, please write to us.

Please report any technical inaccuracies immediately.

You can reach us through email at:

`fes_feedback@rsn.hp.com`

Please include the following information with your email:

- Title and part number of the document
- Edition number

This chapter introduces Hewlett-Packard V-Class system components and includes a brief overview of V2500/V2600 server hardware resources. Some basic details about HP-UX use also are provided. For details on the external cabinet controls and displays, see Chapter 2.

The V2500/V2600 model of V-Class server can have up to 128 processors, 128 Gbytes of memory, and 112 PCI I/O cards.

One new feature of the HP V2500/V2600 server is its Scalable Computing Architecture (SCA) design, which allows multiple V2500/V2600 cabinets to be connected to form a single HP-UX system. These SCA features are made available through HP's Coherent Toroidal Interconnect (CTI) technology.

A V2500/V2600 server can include from one to four cabinets that contain the server resources, with each V2500/V2600 cabinet containing from two to 32 processors, from 512 Mbytes to 32 Gbytes of memory, and up to 28 PCI I/O cards.

Each V-Class system also includes a dedicated workstation connected to the server: the Service Support Processor (SSP workstation). The Service Support Processor is used for server booting, monitoring, and other operations. Details on using the Service Support Processor are provided in Chapter 3.


This book covers both single-cabinet and multiple-cabinet server configurations, support, and operations.

V-Class System Components

Each V-Class system includes two main components: a V-Class server and a Service Support Processor (SSP workstation) dedicated to supporting the server, as shown below in Figure 3.

Figure 3

V-Class Server Components: Cabinet and Service Support Processor




The V-Class cabinet contains all V-Class server resources, such as processors, memory, disks, power, and so forth. The Service Support Processor has software that allows you to monitor the resources in a V-Class cabinet. The V-Class server and the Service Support Processor run separate instances of the HP-UX operating system.

Multiple-cabinet servers may contain up to four V2500/V2600 cabinets, which are booted as a single HP-UX system. Each cabinet has its own cabinet ID (0, 2, 4, or 6) and contains processors, memory, and I/O resources that are available to HP-UX and the applications that run on the server. Cabinets are numbered based on their location in the server. Cabinet ID 0 is the “monarch” or “root” cabinet, which contains the I/O device used for booting and volume group 0. The other cabinets (IDs 2, 4, and 6) are “serf” cabinets, located as shown in Figure 5 on page 5.

Figure 4 shows a four-cabinet V2500/V2600 server and the Service Support Processor that is used for console, diagnostic, and other support work. The V2500/V2600 cabinets are tightly interconnected by Coherent Toroidal Interconnect (CTI) cables, as described in “Multiple-Cabinet Server Connections” on page 15. Connections among the Service Support Processor and V2500/V2600 cabinets are covered in “Server Console and Diagnostic Connections” on page 4.

Figure 4 **Four-Cabinet V2500/V2600 Server Components**



V25U074
10/6/98

The Service Support Processor

The Service Support Processor (SSP workstation) is an HP 712 or B180 workstation connected to the V-Class server. Key operations supported by the Service Support Processor include booting, configuring, and

monitoring the server hardware, as well as diagnostics operations. You also must use the Service Support Processor when installing or upgrading V-Class firmware.

The Service Support Processor runs HP-UX V10.20. In addition to HP-UX software, the Service Support Processor includes files and utility software for managing and monitoring the V2500/V2600 server. These items and all other V2500/V2600-related files, including log files, that are stored on the Service Support Processor can be found in the directory `/spp`.

The default user account for Service Support Processor operations is `sppuser`, with a home directory of `/users/sppuser`.

NOTE

The abbreviation “spp” stands for “scalable parallel processor” and is not to be confused with “SSP”.

See Chapter 3 for more detailed information on the Service Support Processor.

Server Console and Diagnostic Connections

The V2500/V2600 server's utilities board provides connections from the Service Support Processor to a V2500/V2600 server's cabinet or cabinets. Both the console port and diagnostic LAN on each cabinet are connected to the Service Support Processor for system monitoring, booting, and other operations.

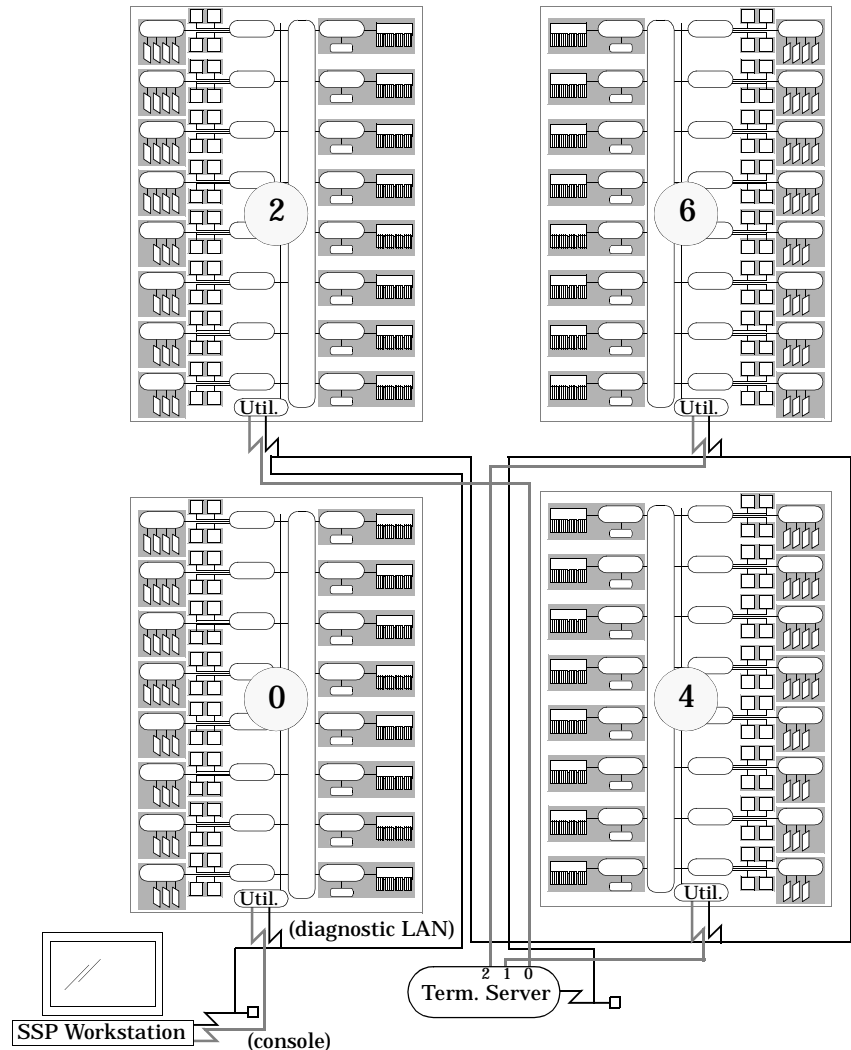
The Service Support Processor connections to a V2500/V2600 server provide only console, diagnostics, and preliminary booting support. For multiple-cabinet servers, the CTI cables between cabinets provide the multiple-cabinet interconnections that create a single, unified V2500/V2600 HP-UX system. Cross-cabinet connections are covered in the section “Multiple-Cabinet Server Connections” on page 15.

A single-cabinet V2500/V2600 server is connected directly to the Service Support Processor, as shown in Figure 6 on page 7.

As shown in Figure 5, multiple-cabinet V2500/V2600 servers have connections from each cabinet's utilities board to either the Service Support Processor or a terminal server.

While a four-cabinet V2500/V2600 server configuration is shown in Figure 5, a two-cabinet or three-cabinet configuration involves the same type of set up among the Service Support Processor, V2500/V2600 cabinets, and the terminal server.

Figure 5 Console and Diagnostic Connections for a Four-Cabinet V2500/V2600 Server



The console port on cabinet ID 0's utilities board connects to the Service Support Processor, and console ports on cabinet IDs 2, 4, and 6 connect to the terminal server (port numbers 2, 3, and 4, respectively).

The diagnostic LAN connects between, and is terminated at, the Service Support Processor and the terminal server. Between these two points, the diagnostic LAN runs in sequence to cabinet IDs 0, 2, 4, and 6.

V-Class Server Architecture

The V2500/V2600 server has a powerful set of interconnecting hardware components that allow the server's processors, memory, and I/O components to operate with minimal interruptions or contentions for resources.

The processor agents serve as a bus connection for a subset of the system's processors. Memory controllers provide cache-coherent access to a large, shared memory. PCI controllers are the connections for PCI I/O cards.

CTI controllers are an SCA feature used only in multiple-cabinet servers. The CTI controllers are connected to memory controllers and provide high-bandwidth connections to other cabinets that comprise the server. See Figure 11 on page 16 for an overview of cross-cabinet CTI connections.

V2500/V2600 Crossbar Interconnection

The primary interconnecting component of each V2500/V2600 server cabinet is the HyperPlane Crossbar, which provides connections from processors and I/O to memory.

The V2500/V2600 crossbar is a non-blocking 8x8 crossbar, which supports eight send messages and eight receive messages simultaneously. This crossbar provides a central connection among the processor agents, memory controllers, and PCI controllers within a V2500/V2600 cabinet. On multiple-cabinet V2500/V2600 servers high-speed CTI interconnections provide access to "remote" memory or other resources on remote cabinets. See "Multiple-Cabinet Server Connections" on page 15 for multiple-cabinet information.

As Figure 7 on page 8 shows, the crossbar has four Exemplar Routing Access Controllers (ERACs), each of which connects to 4 processor agents and four memory controllers. All memory controllers and processor agents connect to two separate ERACs, thus making the entire system's memory addressable by all processors and I/O devices in the system.

Figure 6 Functional Diagram of a Single-Cabinet V2500/V2600 Server

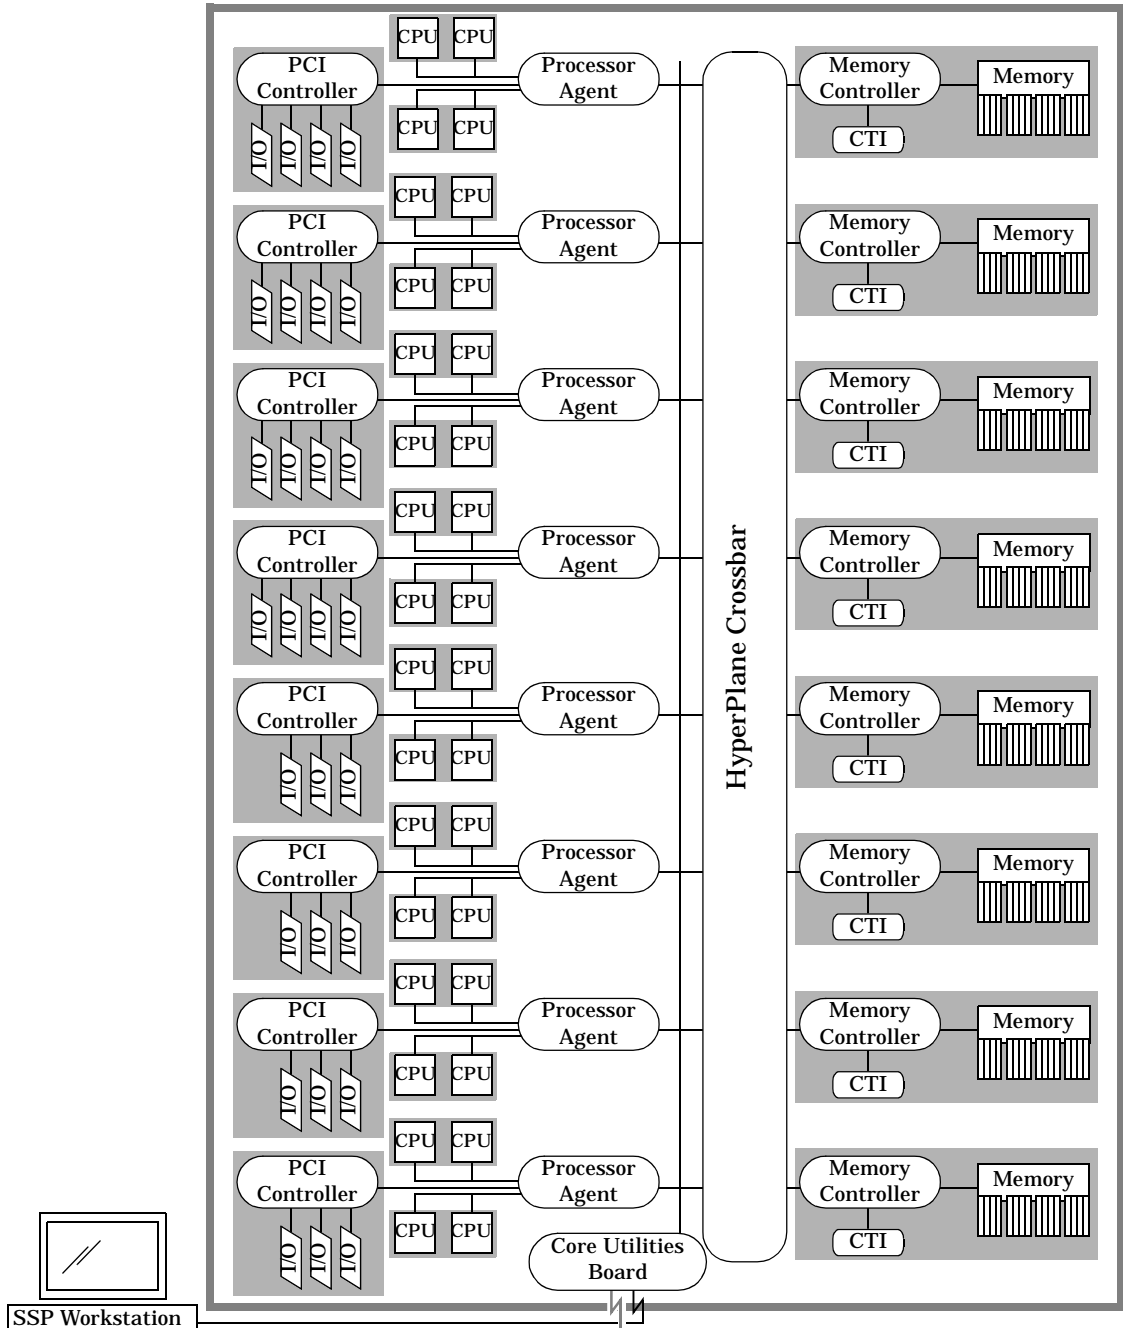



Figure 7

V2500/V2600 HyperPlane Crossbar Connections



Each ERAC has 16 ports, 4 send and 4 receive on each side, which may operate simultaneously.

V2500/V2600 Cabinet Components

The key components within a V2500/V2600 server cabinet include:

- “Core Utilities Board” on page 9
- “Processors” on page 9
- “Memory” on page 9
- “Input/Output” on page 12

In Chapter 2, you can find details on the V2500/V2600 cabinet external controls, such as the on/off key switch panel, and cabinet displays, including the LCD and attention light.

Core Utilities Board

The utilities board provides boot, diagnostics, and console connections from the V-Class cabinet to the Service Support Processor, as well as system clock, system LCD, and other functionality. It also stores the boot firmware and boot-time variable settings in non-volatile memory. For details on firmware use and configuration, refer to Chapter 4.

On multiple-cabinet servers, the Service Support Processor and a terminal server are connected to the console port and diagnostic LAN of each cabinet's utilities board. Details of these connections are shown in Figure 5 and in the section "Server Console and Diagnostic Connections" on page 4.

Processors

A V2500/V2600 server can include up to 128 processors. Each V2500/V2600 cabinet may contain from two to 32 64-bit processors. The V2500 uses the 440 MHz HP PA-8500 processor. The V2600 uses the 552 MHz PA-8600 processor. Each processor board contains one or two processors, with up to two processor boards connecting to each of the eight processor agents per cabinet.

The PA-8500 and PA-8600 processors are based on version 2.0 of Hewlett-Packard's Reduced Instruction Set Computer (RISC) processor architecture.

NOTE

The PA-RISC architecture is presented in the *PA-RISC 2.0 Architecture* reference manual. Please refer to that document for detailed information about processor features. This Operator's Guide does not duplicate information in that manual.

Other models of V-Class servers, including V2200 and V2250 servers, use Hewlett-Packard's PA-8200 processor. Upgrading these other models of servers to a V2500/V2600 involves upgrading the processor, among other components.

Memory

A four-cabinet V2500/V2600 server can contain up to 128 Gbytes of memory when all slots on all memory boards are populated with 256 MByte DIMMs. A maximum of 32 Gbytes of memory may be installed per V2500/V2600 cabinet. For both single-cabinet and multiple-cabinet servers, all memory boards within a cabinet must be configured identically.

Three DIMM sizes are supported for use in V2500/V2600 servers: 32 MByte, 128 MByte, and 256 MByte. Only specified mixed DIMM size configurations are supported.

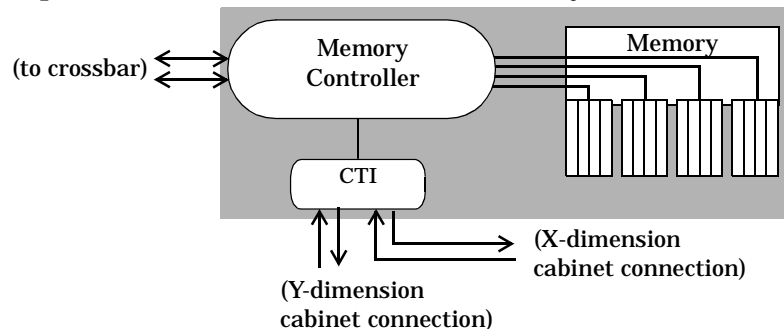
If planning for a multiple-cabinet server configuration, you must use 88-bit DIMMs and configure your V2500/V2600 server to be one-fourth, one-half, or fully populated with DIMMs.

Single-cabinet servers can instead use 80-bit DIMMs and may also be filled to three-fourths memory capacity (3 DIMMs in every quadrant).

As Figure 8 shows, each memory board includes a memory access controller, memory DIMMs, and a CTI controller. The CTI controller is not used in single-cabinet servers, but is present for connecting cabinets in multiple-cabinet servers. Details on multiple-cabinet CTI connections are available in “Multiple-Cabinet Server Connections” on page 15.

Figure 8

Conceptual Overview of V2500/V2600 Memory Board



The V-Class cabinet has a Symmetric Multi-Processing (SMP) design, which gives all processors equal access to all memory and a uniform latency for memory accesses.

Multiple-cabinet V2500/V2600 servers include memory on all cabinets. This provides SMP-like access to memory that resides on the same cabinet as the requesting processor and gives the whole system a cache-coherent non-uniform memory (ccNUMA) architecture. By dedicating some memory as “CTI cache” memory, you can enhance performance by allowing frequently used data that resides in memory on remote cabinets to be encached on the local cabinet. See “CTI Cache Memory” on page 11.

Memory Interleaving

Through the memory access controllers, each memory board provides separate read and write access to the memory DIMMs. Up to 16 DIMMs may be installed per board, providing up to 256-way memory interleaving per cabinet when all memory boards are fully populated.

Slots for DIMMs on each memory board are conceptually grouped in four quadrants. Each quadrant, as Figure 8 on page 10 shows, has a separate connection to the memory controller. All quadrants should have the same memory configuration; this also provides good performance by interleaving memory within a memory board.

Memory also is interleaved across memory controllers, allowing separate controllers and separate parts of the V2500/V2600 crossbar to simultaneously access memory on different controllers. See Figure 7 on page 8 for details.

CTI Cache Memory

A portion of the memory on each cabinet of a V2500/V2600 SCA system is used to implement the CTI cache memory. The CTI cache (also called “network cache”) is memory that is used to minimize the latency of remote memory accesses.

The CTI cache is directly mapped and physically indexed, unlike the processor data caches. The size of the CTI cache is tunable and may range from 8MB to 16GB, depending on the configuration of the memory boards.

To set the CTI cache size, use the `xconfig` utility. The `xconfig` program has a graphical interface and resides on (and is run from) the Service Support Processor. The `xconfig` utility’s **Memory** menu allows the CTI cache to be set through the **Configure Network Cache** menu item.

You also can configure CTI cache using the `ts_config` utility, which is on the Service Support Processor. For details, see Chapter 5, “Configuration utilities” or the `ts_config` man page.

The amount of memory dedicated as CTI cache must be set to be the same for each cabinet in the V2500/V2600 server. So, for example, a 128 MByte per-cabinet CTI cache setting would consume a total of 256 Mbytes of memory for a two-cabinet V2500/V2600 server. Single-cabinet servers should not have any memory dedicated as CTI cache, as it is unnecessary.

With small CTI cache sizes, additional aliasing between memory locations may occur, reducing the cache hit rate and increasing the latency for remote accesses. The bold entries in Table 1 show the minimal CTI cache sizes needed to avoid excessive aliasing.

Table 1 Valid CTI cache sizes

Node Memory size	Node CTI Cache size in MB
4 GB	32, 64, 128, 256, 512, 1024, 2048
8 GB	32, 64, 128, 256, 512, 1024, 2048, 4096
16 GB	32, 64, 128, 256, 512, 1024, 2048, 4096, 8192
24 GB	64, 128, 256, 512, 1024, 2048, 4096, 8192
32 GB	64, 128, 256, 512, 1024, 2048, 4096, 8192, 16384

CTI cache memory is configured at system boot time. The CTI cache memory is dedicated only to be used for encaching remote accesses, and is not available for any other use, such as use by HP-UX or applications.

Input/Output

A multiple-cabinet V2500/V2600 server can contain up to 112 PCI I/O cards, with each cabinet containing up to 28 PCI I/O cards. Each V2500/V2600 cabinet includes 64-bit PCI chassis, eight PCI buses, and connections for either three or four PCI cards per PCI bus.

The following I/O cards are supported on HP V2500/V2600 servers:

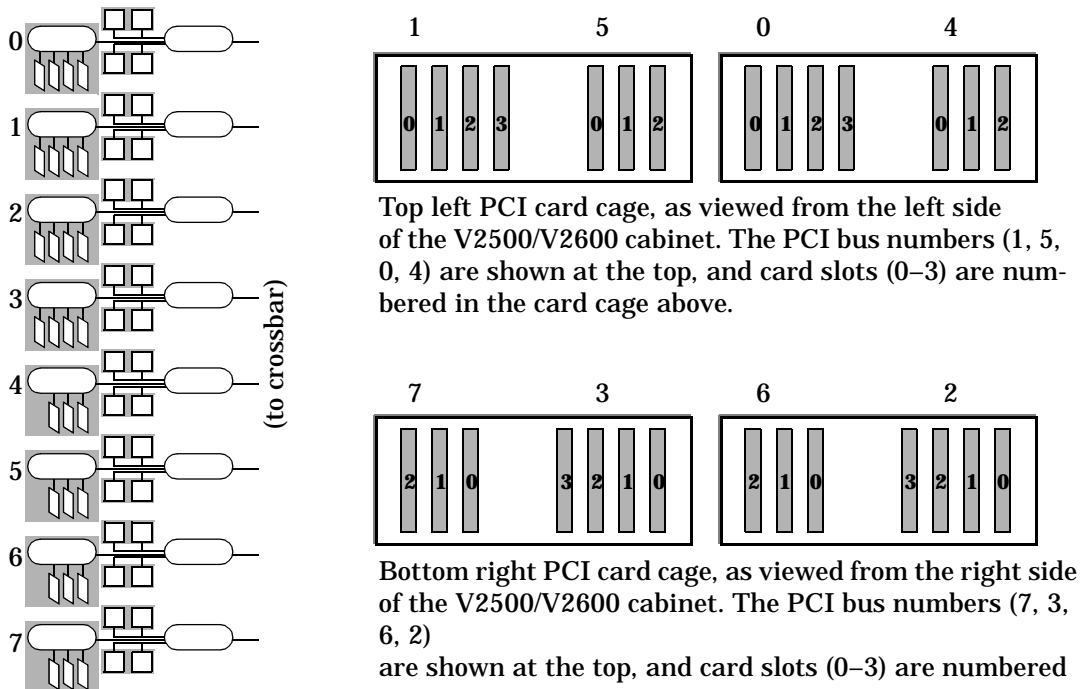
- Tachyon Fibre Channel
- HVD FWD SCSI
- LVD Ultra2 SCSI
- 10/100 Base T Ethernet
- 1000 Base SX Gigabit Ethernet
- FDDI

Each V2500/V2600 I/O port is capable of direct memory access (DMA), which eliminates processor involvement during data transfers and streamlines data transfer for large disk blocks and high-speed network connections.

The PCI bus controllers are numbered based on the V2500/V2600 cabinet in which they reside. The first component of the hardware path (such as reported by the HP-UX `ioscan` utility) indicates which cabinet a hardware component resides upon.

Figure 9 on page 13 shows the PCI bus numbers and card cage locations for a single-cabinet server. The I/O card cages are accessible from either the top-left or the bottom-right sides of the V2500/V2600 cabinet.

Figure 9 **Numbering and Locations of Single-Cabinet V2500/V2600 PCI I/O**



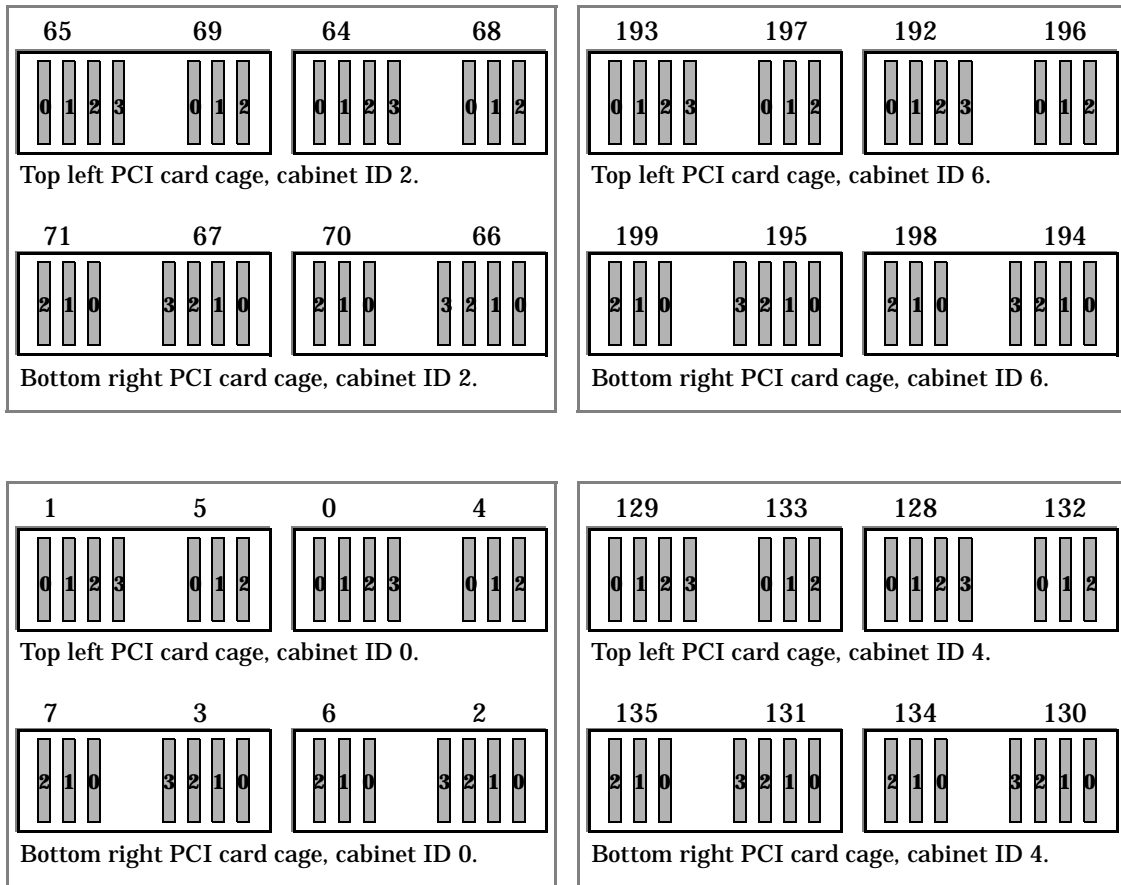
The PCI busses in a single-cabinet server are numbered from 0 to 7, as shown above. This numbering also is used for the PCI busses in cabinet 0 of a multiple-cabinet server.

For multiple-cabinet servers, the PCI bus numbering is as shown in Figure 10. The PCI bus number also serves as the first field of the associated devices' hardware path, so I/O devices on cabinet ID 0 are numbered with the first field of the hardware path of 0 to 7.

For cabinet ID 2, the PCI bus numbers are from 64 to 71. PCI busses on cabinet ID 4 are from 128 to 135, and cabinet ID 6 devices are numbered from 192 to 199.

The PCI bus and card slot numbering for multiple-cabinet servers is illustrated in Figure 10 on page 14.

Figure 10 **Numbering and Locations of Multiple-Cabinet V2500/V2600 PCI I/O**



For an example of listing I/O devices on various cabinets and details of listing other V-Class server hardware configuration details, see “Listing the Server Hardware Configuration” on page 118.

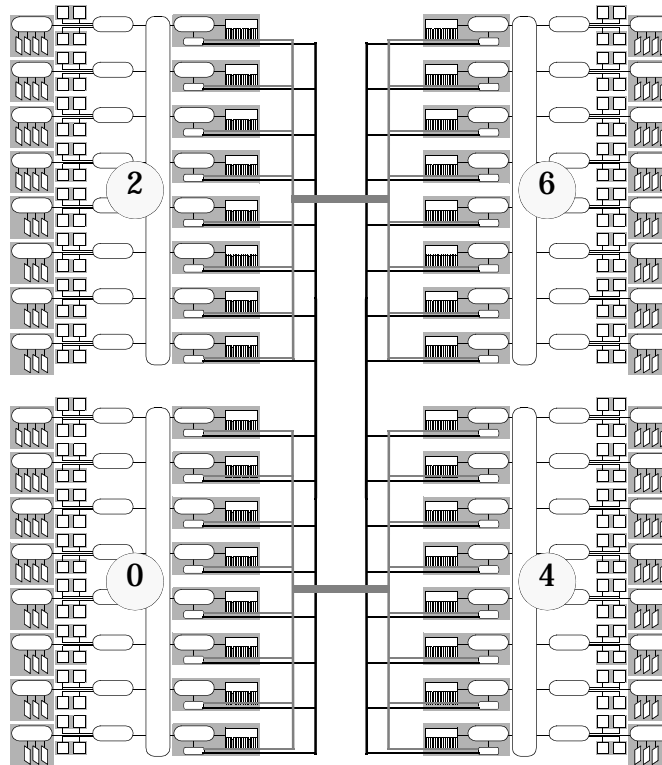
Multiple-Cabinet Server Connections

All cabinets in a multiple-cabinet V2500/V2600 server are tightly connected using HP’s Coherent Toroidal Interconnect (CTI) technology. CTI is an extension of the Scalable Coherent Interface standard defined by the IEEE.

CTI cables connect among the CTI controllers on the various cabinets. An overview of the CTI connections for a four-cabinet V2500/V2600 server is shown in Figure 11.

Each CTI controller connects to a corresponding CTI controller on a remote cabinet by cables that provide both send (local-to-remote) and receive (remote-to-local) connections among the cabinets.

Figure 11 Four-Cabinet V2500/V2600 Server CTI Cable Connections



Two dimensions of CTI connections are possible. Y-dimension cables connect between cabinets 0 and 2, and between cabinets 6 and 4. X-dimension cables connect cabinets 0 and 4, and cabinets 6 and 2.

Send and receive connections are provided in two dimensions on each controller, for a total of four connections per controller possible. In a two-cabinet server, cabinets 0 and 2 are connected via Y-dimension CTI cables only. For a three-cabinet server, cabinet 0 has Y-dimension CTI connections to cabinet 2 and X-dimension CTI connections to cabinet 4.

For Y-dimension connections, CTI cables connect to their counterparts on the remote cabinet. For example, the CTI cable connects from “Y-send” on memory board 0 to the remote cabinet’s “Y-receive” on memory board 0.

For X-dimension connections, CTI cables connect to the opposite controller on the remote cabinet. This means—for X-dimension CTI connections—memory boards connect in the following pairs: 0 and 2, 1 and 3, 4 and 6, and 5 and 7.

For details on CTI cable connections refer to qualified HP service personnel.

V2500/V2600 Cabinet Configurations

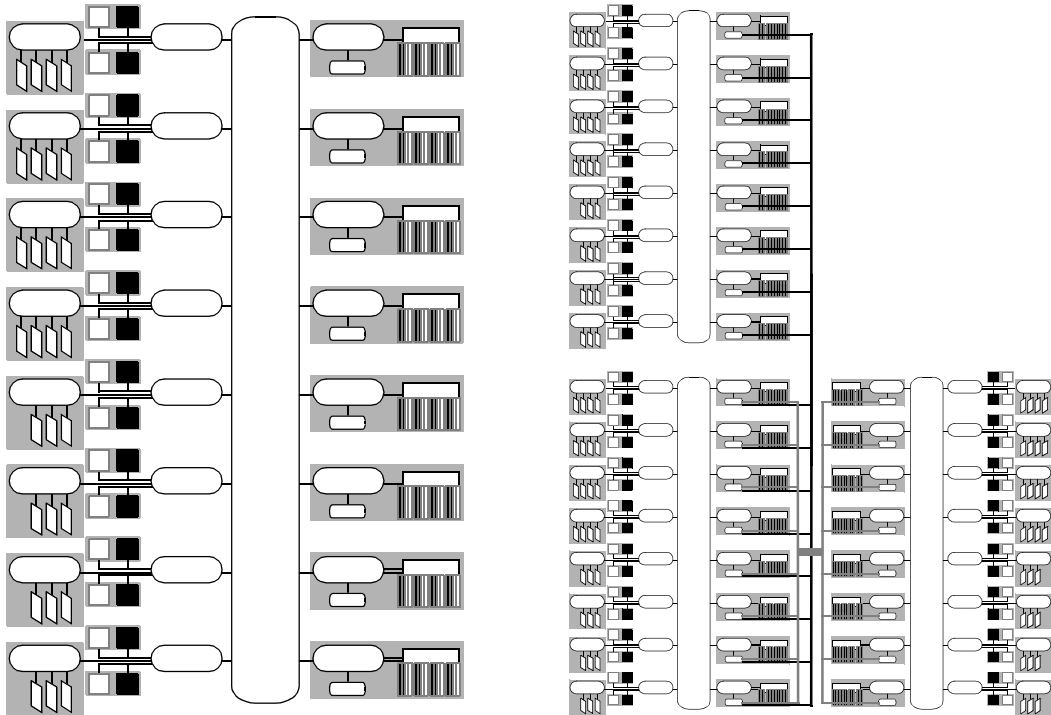
This section shows two sample V2500/V2600 server configurations: a single-cabinet system and a three-cabinet system, filled to one-half processor capacity and to one-half and full memory capacity, respectively.

Each V2500/V2600 cabinet can contain up to 32 processors, 32 Gbytes of memory, and 28 PCI cards, with up to four cabinets (up to 128 processors, 128 Gbytes of memory, and 112 I/O cards) comprising a V2500/V2600 server.

Additional server configuration and ordering information is available from the following Web site.

<http://eproducts.hp.com/>

Figure 12 **Sample V2500/V2600 Cabinet Configurations**



A single-cabinet V2500/V2600 server with 16 processors and 16 Gbytes memory, using 256 MByte

A three-cabinet V2500/V2600 server with 48 processors and 96 Gbytes memory, using 256 MByte

Overview
V2500/V2600 Cabinet Configurations

2

Indicators, switches, and displays


This section describes indicators, switches, and displays of the HP 9000 V2500 server.

Operator panel

The operator panel is located on the top left side of the server and contains the key switch panel, DVD-ROM drive, optional DAT tape drive, and the LCD display. Figure 13 shows the location of the operator panel and its components.

Figure 13

Operator panel




V25U102
3/24/99

Key switch panel

The key switch panel is located on the left of the operator panel, as shown in Figure 13 on page 22. 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 14.

Figure 14 Key switch panel



IOEXS095
10/10/97

Key switch

The key switch has two positions:

- **DC OFF**

DC power is not applied to the system. Placing the key switch 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.

DVD-ROM drive

The DVD-ROM drive is located on the left of the operator panel, as shown in Figure 13 on page 22. Figure 15 shows the DVD-ROM drive front panel in detail.

Figure 15

DVD-ROM drive



V25U101
3/17/99

Disk loading slot

Place the disk into the slot with the label side up. Gently push the front edge of the disk to load it into the drive. When an 8-cm disk is used, it must be set into an adapter prior to loading.

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 DVD-ROM is ejected, possibly causing a loss of data.

Eject button


Push the eject button to eject DVD-ROMs from the drive.

Optional DAT drive

The DAT drive is located on the right of the operator panel, as shown in Figure 13 on page 22. The DAT drive front panel contains two indicator LEDs and an eject button, as shown in Figure 16.

Figure 16

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.
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-enabled, 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.

WARNING


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 17

System displays




IOLM010
9/18/97

LCD (Liquid Crystal Display)

The LCD display is located on the right of the operator panel, as shown in Figure 17 on page 27. The LCD is a 20-character by 4-line liquid crystal display. Figure 18 shows the display and indicates what each line on the display means.

Figure 18 **Front panel LCD**



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: V2500/V2600 Servers*. The following explains the output shown in Figure 18:

Node status line

The Node Status Line shows the node ID in both decimal and X, Y topology formats.

Processor status line

The processor status line shows the current run state for each processor in the node. Table 3 shows the initialization step code definitions and Table 4 shows the run-time status codes. The M in the first processor status line stands for the monarch processor.

Table 3 Processor initialization steps

Step	Description
0	Processor internal diagnostic register initialization
1	Processor early data cache initialization.
2	Processor stack SRAM test.(optional)
3	Processor stack SRAM initialization.
4	Processor BIST-based instruction cache initialization.
5	Processor BIST-based data cache initialization
6	Processor internal register final initialization.
7	Processor basic instruction set testing. (optional)
8	Processor basic instruction cache testing. (optional)
9	Processor basic data cache testing. (optional)
a	Processor basic TLB testing (optional)
b	Processor post-selftest internal register cleanup. (optional)

Table 4 Processor run-time status codes

Status	Description
R	RUN: Performing system initialization operations.
I	IDLE: Processor is in an idle loop, awaiting a command.
M	MONARCH: The main POST initialization processor.
H	HPMC: processor has detected a high priority machine check (HPMC).
T	TOC: processor has detected a transfer of control (TOC).
S	SOFT_RESET: processor has detected a soft RESET.
D	DEAD: processor has failed initialization or selftest.

Status	Description
d	DECONFIG: processor has been deconfigured by POST or the user.
-	EMPTY: Empty processor slot.
?	UNKNOWN: processor slot status in unknown.

Message display line

The message display line shows the POST initialization progress. This is updated by the monarch processor. The system console also shows detail for some of these steps. Table 5 shows the code definitions.

Table 5 Message display line

Message display code	Description
a	Utilities board (SCUB) hardware initialization.
b	Processor initialization/selftest rendezvous.
c	Utilities board (SCUB) SRAM test. (optional)
d	Utilities board (SCUB) SRAM initialization.
e	Reading Node ID and serial number.
f	Verifying non-volatile RAM (NVRAM) data structures.
g	Probing system hardware (ASICs).
h	Initializing system hardware (ASICs).
i	Probing processors.
j	Initialing, and optionally testing, remaining SCUB SRAM.
k	Probing main memory.
l	Initializing main memory.
m	Verifying multi-node hardware configuration
n	Multi-node initialization starting synchronization
o	Multi-node hardware initialization

Message display code	Description
p	Multi-node hardware verification
q	Multi-node initialization ending synchronization
r	Enabling system error hardware.

Attention light bar

The Attention light bar is located at the top left corner on the front of the V2500/V2600 server as shown in Figure 17 on page 27. The light bar displays system status in three ways:

- OFF—dc power is turned off. Either the key switch or the side circuit breaker is in the off position.
- ON—Both the side circuit breaker and the keyswitch are in the on position and no environmental warning, error, or hard error exists.
- Flashing—There is an environmental error, warning, or hard error condition. Also indicates scanning during diagnostic execution.

NOTE

The light bar flashing during initial start up does not indicate a fault.

The types of environmental conditions that are monitored include:

- ASIC installation error sensing
- ASIC configuration or status
- 48V failure

NOTE

48V failures are cleared only after a power cycle.

- Power failure sensing
- Fan sensing
- Thermal sensing

Types of environmental control functions monitored include:

- Power-on
- Voltage margining (SSP interface)

Environmental errors

Environmental errors are detected by two basic systems in the V2500/V2600 server: Power-On and Environmental Monitor Utility Chip (MUC).

Power-On detected errors such as ASIC install or ASIC not OK are detected immediately and will not allow dc power to turn on until that condition is resolved.

MUC detected errors such as Ambient Air Hot allows the dc power to turn on for approximately 1.2 seconds before the dc power is turned off. If two or more fans fail simultaneously, the MUC will shut off dc power. Other MUC detected errors such as Ambient Air Warm will flash the LED and not turn off dc power.

Error codes may be viewed by using the SSP utility command `pce` to read the status of the CUB. However, this feature will only work after database generation is complete, not before.

Using the SSP utility `man leds` to decode the CUB status nibbles.

The current environmental temperature set-points are:

- Warm = 32 degrees Celsius (89.6 degrees Fahrenheit)
- Hot = 37 degrees Celsius (98.6 degrees Fahrenheit)

Displaying the CUB LED values using `pce`

Use the `sppdsh` command `pce` to display the value of the LEDs on the CUB.

Step 1. Bring up the `sppdsh` prompt at a `sppuser` window by entering:

```
$ sppdsh
```

Step 2. Use the `pce` command to display the LED values for all nodes, enter:

```
sppdsh: pce all
```

Node	IP address	Clocks	LEDS	@C	U	SHPT	Supply1	Supply2	Supply3	Supply4
0	15.99.111.116	Normal	0x00	25	1	0000	Nominal	Nominal	Nominal	Nominal
2	15.99.111.117	Normal	0x00	25	1	0000	Nominal	Nominal	Nominal	Nominal

For more information about the `pce` command see the `sppdsh` man page.

Step 3. Decode the LED values using Appendix A, "LED codes" .

Identifying a node with the blink command

The `blink` command is used to physically identify a node. This command forces the node attention light bar to blink or turns off blinking, provided an error does not exist on the node.

Step 1. Bring up the `sppdsh` prompt at a `sppuser` window by entering:

```
$ sppdsh
```

Step 2. Use the `blink` command to cause the attention light bar to blink on a specific node by entering the `blink` command followed by the node number. For example:

```
sppdsh: blink 0
```

For more information about the `blink` command see the `sppdsh` man page.

Step 3. After you have physically identified the node cause the attention light bar to return to a steady state by entering:

```
sppdsh: blink 0
```

Indicators, switches, and displays

System Displays

3

SSP operation

This chapter describes the operation the SSP in conjunction with a V-Class server and includes:

- SSP log-on
- Using the CDE (Common Desktop Environment) Workspace menu
- Using the console
- SSP file system
- System log pathnames

SSP and the V-Class system

The Service Support Processor (SSP) is either a Hewlett-Packard B180L or 712 workstation that performs the following functions for the V-Class system:

- Running diagnostics
- Updating of CUB (Core Utility Board) firmware
- Logging environmental and system level events
- Configuring of hardware and boot parameters
- Booting the operating system
- Accessing V-class console

The SSP is closely interfaced with the Core Utility Board (CUB), located on the Mid-plane Interconnect Board (MIB). It has the ability to access each section of the CUB, allowing for control, verification, testing, and normal management of the V-Class server.

A private ethernet bus called the “test bus” connects the SSP to the CUB located within the V-Class server.

The SSP has HP-UX installed and operates independently from the main server.

IMPORTANT

HP-UX 10.20 is required to allow the workstation to function as the SSP.

Additional software has been added to the basic HP-UX 10.20 to provide all the necessary functions of the SSP.

SSP log-on

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

sppuser	This user is the normal log-on for the SSP 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 SSP. Default password: serialbus

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.

SSP sppuser windows

When the user is logged on to the SSP on a V2500/V2600 server that consists of less than two nodes, the windows appear in the configuration shown in Figure 19. The Workspace, however, does not appear until a mouse button is clicked anywhere in the CDE backdrop.

If the V2500/V2600 server consists of more than two nodes, the console windows are replaced by a consolebar as shown in Figure 20.

SSP operation
 SSP log-on

Figure 19 SSP user windows for V2500/V2600 servers with one node

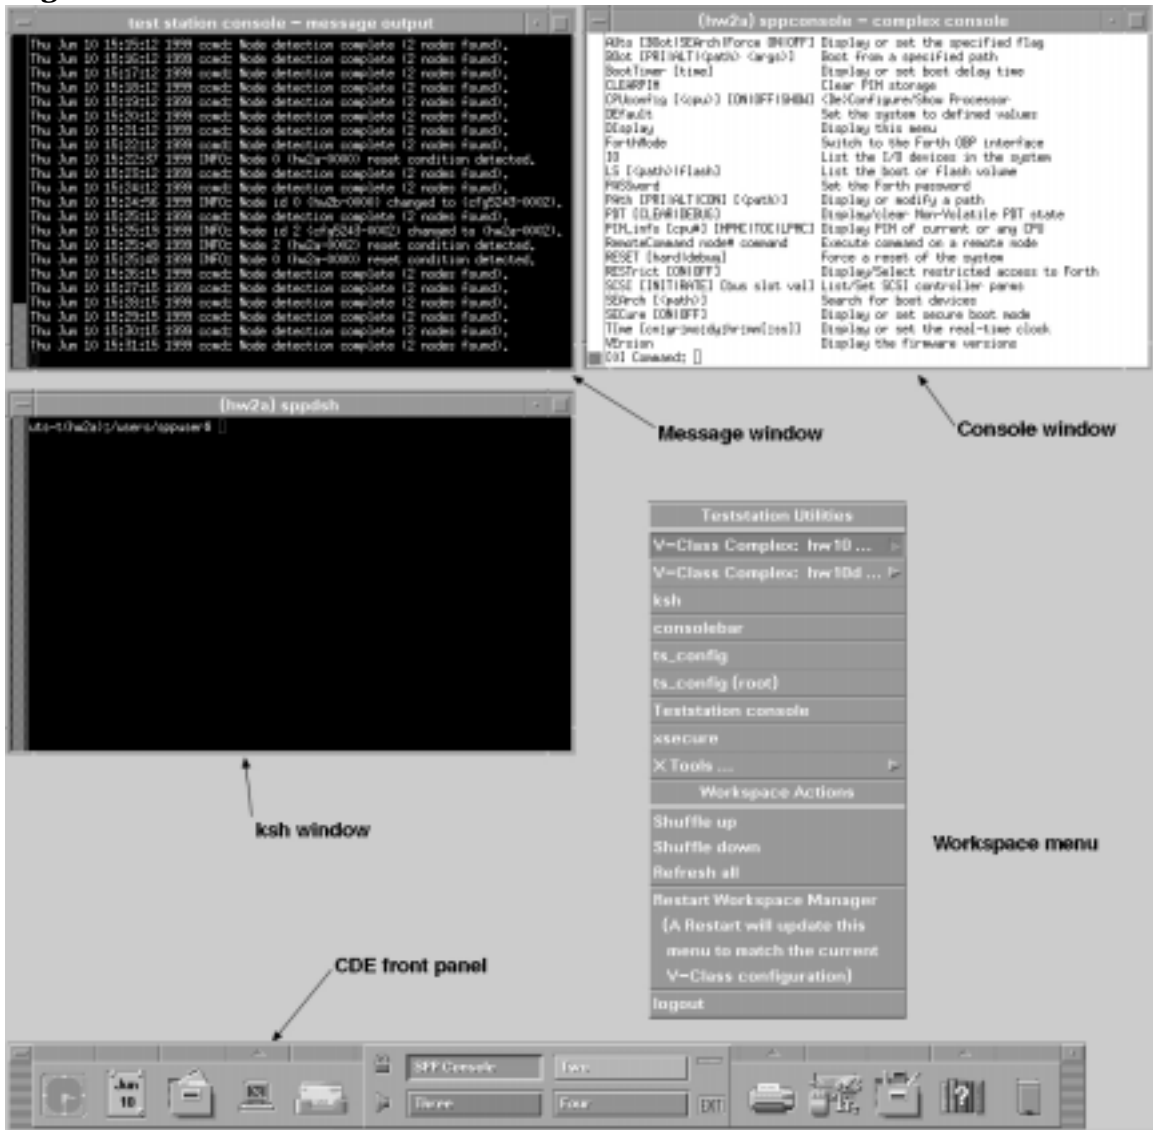
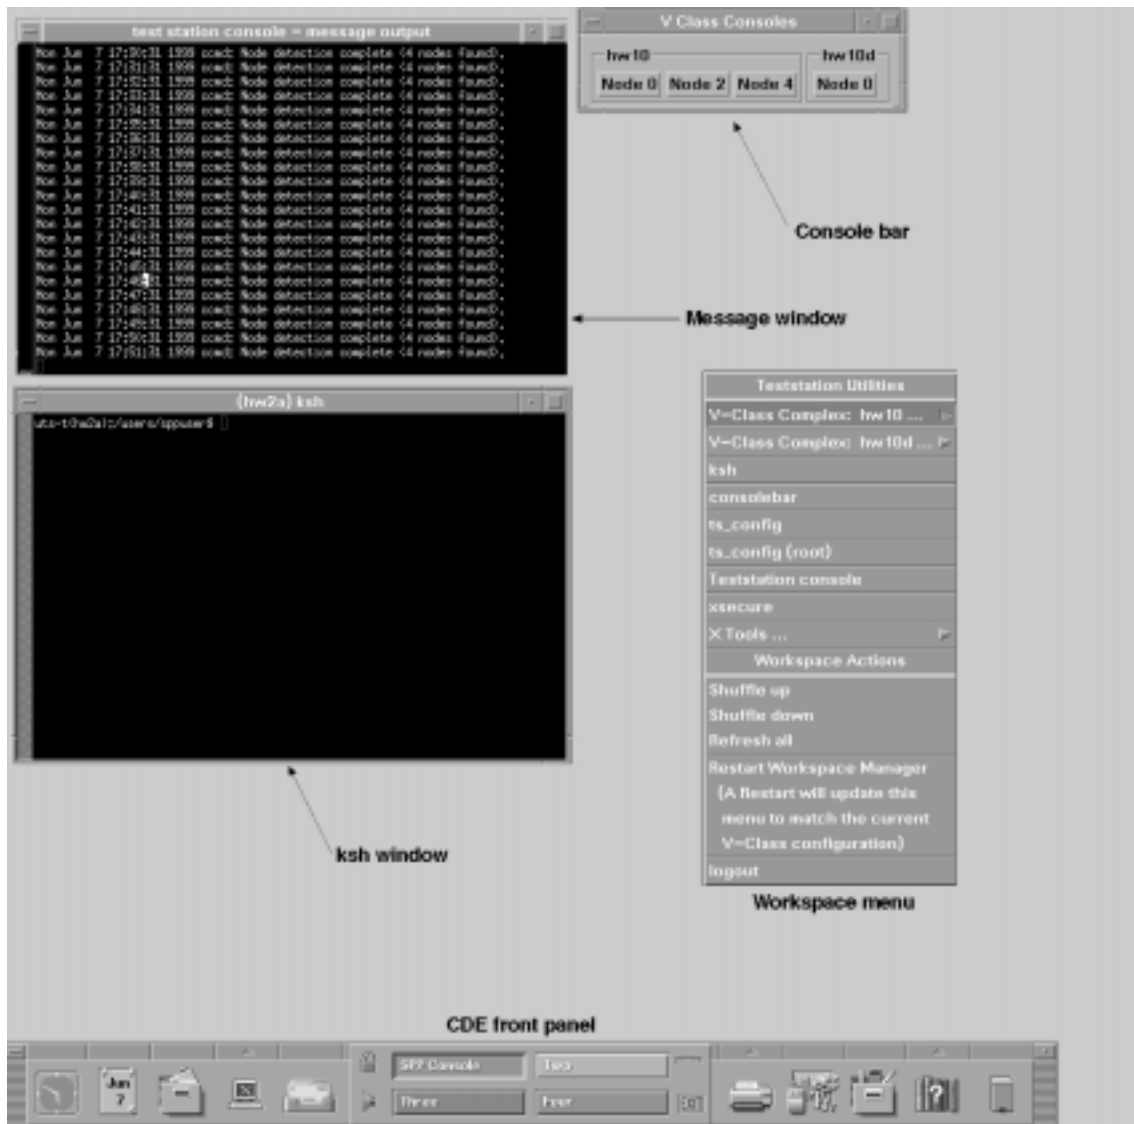


Figure 20 SSP user windows for V2500/V2600 servers with more than two nodes



SSP operation
SSP log-on

Message window

The message window displays status from the `ccmd` daemon running on the SSP approximately 60 seconds after power on. The hard error logger also displays status in this window. This is a display window only and does not accept input.

Console window (sppconsole - complex console)

The complex console window is the main console window for the V-Class server complex. It displays all POST (Power-On Self -Test) status for node 0. The user can boot and configure the server from this window using the boot menu (`Command:` prompt). The user can also enter a special mode called “forth mode” (OBP) to perform special configuration commands.

A second `sppconsole` window is spawned when two nodes exist.

Console window (sppconsole - Node X console)

`sppconsole` windows for each node in the complex are spawned using the “consolebar” which is available from the desktop Workspace menu. All POST (Power-On Self -Test) status for node *X* is displayed here. The user can boot and configure the server from this window using the boot menu. The user can also enter the special forth mode to perform special configuration commands.

Console bar

The console bar appears on the SSP workspace for systems that have more than two nodes (or complexes in the case of SCA systems). To see the complex console, click on the appropriate button in the console bar.

ksh shell windows

The `ksh` windows are local shell windows on the SSP. The user can enter commands into these windows to invoke scripts or functions on the SSP.

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

Using the CDE (Common Desktop Environment) Workspace menu

The SSP uses the CDE Workspace Manager to control the windows on the screen. The Workspace menu is Workspace Manager main menu. The Workspace menu selects create new windows, initiate diagnostic tools, and perform other tasks.

CDE Workspace menu

The following section describes how to use the CDE Workspace menu on V2500/V2600 servers:

- Step 1.** Move the pointer over the CDE workspace backdrop.
- Step 2.** Press and hold down any mouse button. The Workspace (root) menu appears.
- Step 3.** Drag the mouse pointer to an option.
- Step 4.** Release the mouse button to select the option.

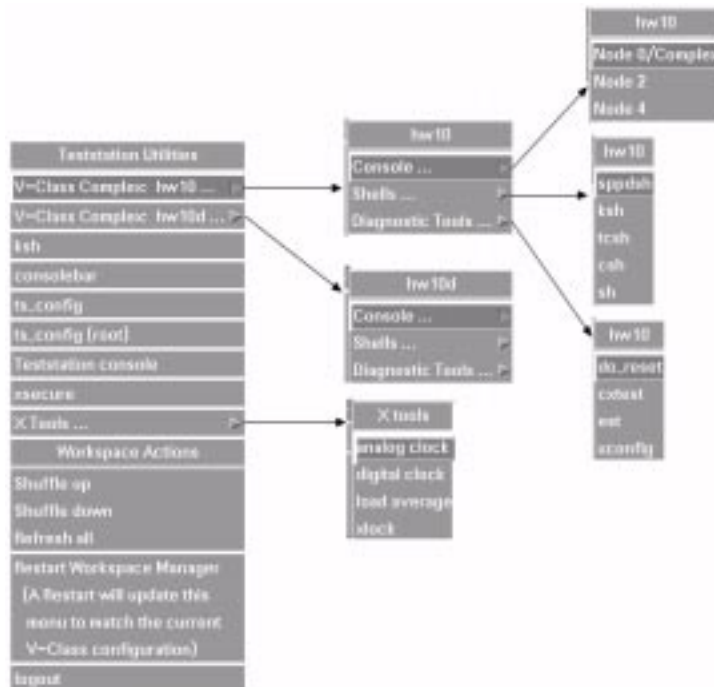
Figure 21 shows the V2500/V2600 SSP Workspace menu and Figure 22 shows the submenus.

SSP operation
 Using the CDE (Common Desktop Environment) Workspace menu

Figure 21 SSP Workspace submenus for V2500/V2600



Figure 22 SSP Workspace submenus for V2500/V2600



V2500/V2600 Workspace menu options include:

- **V-Class Complex:** *name*—Opens this submenu for the node/complex. If more than one node/complex has been configured, multiple V-Class Complexes are available by name.
 - **Console**—Creates a new console window for a list of available node/complexes.
 - **Shells**—Selects a shell: `sppdsh`, `ksh`, `tcsh`, `csch`, and `sh` shells.
 - **Diagnostic Tools**—Performs a `do_reset` or invokes `cctest`, `est`, or `xconfig`.
- **ksh**—Creates a new `ksh` window on the screen.
- **consolebar**—Creates GUI console select bar on the screen.
- **ts_config**—System monitoring, and diagnostic capabilities. Runs as `sppuser` allowing the resetting of nodes/complexes and console session generation.
- **ts_config (root)**—Configures the SSP console, system monitoring, and diagnostic capabilities. Runs as `root`, requires the password, allows reconfiguration of nodes multinode complexes, and configuration of the terminal mux.
- **teststation console**—Creates a new console window on the screen.

NOTE

Only one SSP console per node can be active at any time. If a new SSP console is started, any existing SSP console sessions for that particular node are disabled. The old console windows remain on the screen, but no new console messages are sent to the old sessions. Only the most recent SSP console will receive SSP console output.

-
- **xsecure**—Use this tool to disable modem and LAN activity.
 - **X tools**—Lists several X tools including the load average display and `xlock` (screen lock/saver).

Workspace actions include:

- **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 the entire X display.

SSP operation

Using the CDE (Common Desktop Environment) Workspace menu

- **Restart Workspace Manager**—Stops and restarts the Workspace Manager.
- **logout**—Closes all open windows and stops Workspace Manager.

Using the console

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

Creating new console windows

Console windows can also be created using the `sppconsole` and `xterm` commands from the SSP; see Table 6 for details.

Table 6

Commands for creating console windows

SSP command	Description
<code>/spp/scripts/sppconsole</code>	The <code>sppconsole</code> script provides an HP-UX console interface in the current SSP window.
<code>/usr/bin/X11/xterm</code>	Creates a SSP login window.
<code>/usr/bin/X11/xterm -C</code>	Creates a SSP console window.

Starting the console

The console server program automatically starts the console on the SSP when you log on as `sppuser`. If the console stops running, restart it from the SSP using one of the following methods: the Workspace menu, the `sppconsole` command, `ts_config`, `consolebar`, or logging back on.

Methods for starting the console V2500/V2600 servers are:

- Workspace menu
- `sppconsole` command
- `ts_config`
- `consolebar`
- logging back on

Starting the console from the Workspace menu

To start the console using the Workspace menu, complete the following steps:

- Step 1.** Move the pointer over the CDE workspace backdrop.
- Step 2.** Press and hold down any mouse button. The Workspace (root) menu appears.
- Step 3.** Drag the mouse pointer to the “V-Class Complex: *complex_name*” option.
- Step 4.** Release the mouse button to select the option.
- Step 5.** Drag the mouse pointer to the Console option.
- Step 6.** Release the mouse button to select the option.
- Step 7.** Select the desired V2500/V2600 complex.

NOTE

If the desired V2500/V2600 Complex name is not listed in the menu, restart the Workspace manager. The desktop Workspace menus are updated whenever a node is configured or deconfigured, but the new menus are not activated until the Workspace Manager is restarted.

- Step 8.** Select Console menu.
- Step 9.** Select Node *X*/complex. The new console window appears.

Starting the console using the `sppconsole` command

This method of starting the console works from the SSP or after logging on to the SSP from another system.

To start 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.** If more than one V-Class complex is connected to the SSP, use the `set_complex` command to select the desired complex by entering:
`set_complex.`

An output similar to the following will be displayed. Enter the desired complex from the list provided.

For example:

```
COMPLEX_NAME = [Select from colossus, guardian] colossus
```

Step 3. Start the console. Enter:

```
sppconsole
```

NOTE

Running `sppconsole` without any additional parameters defaults to Node 0 in the current complex. `sppconsole 2` would start a console on Node 2.

The new `sppconsole` window appears.

In the example above, even if the user's "default complex" is set to `colossus`, the user could start a console on guardian Node 0 by entering the Node Name in the following format:

Complex Name-nnnn

where `nnnn` is the Node ID extended to four digits and zero-filled on the left. These names can be viewed using `jf-ccmd-info` or `ts_config`.

For example:

```
sppconsole guardian-0000
```

To start a console on Node 2 of the complex named `guardian` enter:

```
sppconsole guardian-0002
```

Starting the console using `ts_config`

This method of starting the console works from the SSP or after logging on from another system.

To start a console session from within `ts_config`, complete the following steps:

- Step 1.** Move the pointer over the CDE workspace backdrop.
- Step 2.** Press and hold down any mouse button. The Workspace (root) menu appears.
- Step 3.** Drag the mouse pointer to the `ts_config` (root) option.
- Step 4.** Release the mouse button to select the option.

Step 5. Enter the root password.

Refer to “Starting ts_config” on page 92 for information on starting ts_config from a local or remote shell.

Step 6. Select the desired node(s) from the list in the display panel. For example, clicking on node 0 in the list highlights that line in the window.

Step 7. Start the console session by doing one of the following:

- Select “Actions” to drop the pop-down menu and then click “Start Console Session.”
- Click the right-mouse button and select “Start Console Session.”

The console window(s) appears.

Starting the console using the consolebar

This method of starting the console works from the SSP or after logging on from another system.

The `consolebar` utility is an GUI that shows the configured nodes, grouped by complex. Each node is a push-button that, when pushed, activates a console session for that node.

If more than two nodes are configured on a V2500/V2600 SSP, `consolebar` automatically starts when the `sppuser` logs in at the SSP display.

To start a console session using `consolebar` complete the following steps:

Step 1. Move the pointer over the CDE workspace backdrop.

Step 2. Press and hold down any mouse button. The Workspace (root) menu appears.

Step 3. Drag the mouse pointer to the consolebar option.

Step 4. Release the mouse button to select the option.

Step 5. Enter the root password. The console bar appears.

Step 6. Click on the button of the node to start

Starting the console by logging back on

This method of starting the console works from the SSP or after logging on from another system.

To start the console by logging out of the SSP and logging back on again, complete the following steps:

- Step 1.** Move the pointer over the CDE workspace backdrop.
- Step 2.** Press and hold down any mouse button. The Workspace (root) menu appears.
- Step 3.** Drag the mouse pointer to the logout menu option.
- Step 4.** Release the mouse button to select the option. The SSP closes all open windows and returns a HP-UX login prompt.
- Step 5.** Log into the SSP as sppuser. The new sppconsole window displays.

Console commands

Use the sppconsole commands to control the console. Using these commands allows the user to watch or to assume control of the console window.

Table 7 sppconsole commands

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

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

Any user can display the console via a remote login to the SSP, 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 “Assuming control of the console” on page 51, for more information.

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

- Step 1.** Remotely log in to the SSP 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:

```
sppconsole
```

At this point the console is in “spy mode,” meaning the user can only monitor what is going on at the system console. If commands are entered 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.

Assuming control of the console

System maintenance or diagnostics can be performed remotely by assuming control of the console from a remote terminal. Upon gaining control of the console, the user has 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 SSP 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:

```
sppconsole
```

At this point the console is in spy mode, meaning the user can only monitor what is going on at the system console. If commands are entered 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.** To relinquish control of the console and return to spy mode, enter the following command:

```
CTRL-ECS
```

- Step 5.** Exit the session with the following command:

```
CTRL-Ec .
```

Changing a console connection

Once the console is started as a watch or a control connection, the connection type can be changed with escape characters.

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

CTRL-Ecf

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

CTRL-Ecs

Accessing system logs

Monitor system status via two logs, `event_log` and `consolelogX` (where X is the `node_id`), located in `/spp/data/complex_name` on the SSP.

The `event_log` file periodically logs system status. Once the file reaches 1 Mbyte, the system compresses it to `event_log.old.Z` and creates a new `event_log` file.

The `consolelogX` files grow without bounds. These need to be periodically checked by the system administrator.

The `set_complex` command

From an existing shell on the SSP, the user can set the “default complex” for the shell by executing the `set_complex` command.

```
set_complex complex name
```

The `set_complex` command lists the configured complex names and prompts the user for a selection. After a complex has been selected, the user can issue diagnostic, scan, and console commands against a particular node ID (e.g. 0). The SSP software accesses the correct node by combining the user-specified node ID with the complex name selected by the last `set_complex` command run from the shell. The default complex name is always visible. It is enclosed in parenthesis as part of the shell prompt and included in the Window title if the shell is on the SSP desktop.

```
set_complex default
```

If there is only one complex configured on a SSP, there is no need to run the `set_complex` command since there is only one complex (in fact, `set_complex` automatically assigns the default complex name without

prompting the user if only one complex is configured). This utility accesses the desired node based on node ID. However, the single node must still be configured by `ts_config` and assigned a complex name before it can be accessed.

Targeting commands to nodes

Use the `jf-ccmd_info` command to determine what names or IP addresses the JTAG interfaces have been set to on an SPP. The command provides a list with the following information:

- Ethernet Address (MAC)
- IP Address
- Complex Serial Number
- Node Number
- Environmental LED's
- Power Status
- SCUB Status
- Diagnostic (JTAG) node names

The diagnostic node name or the JTAG IP address is required when using the `load_eprom` command. Since the SSP software allows changing the default JTAG hostnames and IP addresses, the user may need to run this command to view the unique names and active IP addresses for JTAG.

Example:

```
load_eprom -n COMPLEX_NAME-0000 -l /spp/firmware/diaglifhdr.fw
```

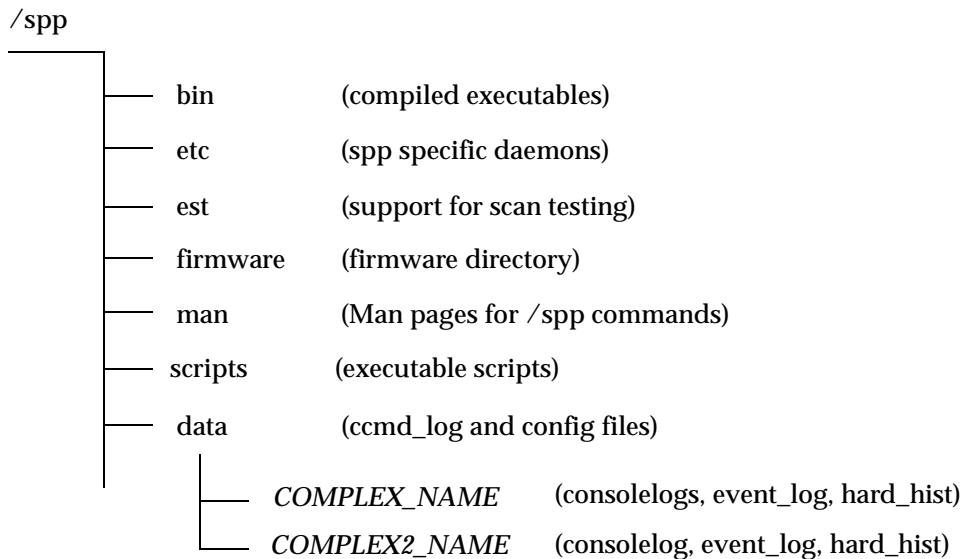
SSP file system

The `/spp` and `/users/sppuser` directories contain most of the SSP specific files. Other files in various directories are also modified. This section restricts, however, its discussion to the `/spp` directory.

Figure 23 shows the SSP file system structure for V2500/V2600 servers.

Figure 23

SSP file system for V2500/V2600 servers



`/spp/etc`

The `/spp/etc` directory contains many of the unique daemons that run on the SSP. These daemons manage of the V-Class node. Two daemons that are always running on the SSP are:

`ccmd` A daemon that maintains a database of information about the V-Class hardware. It also monitors the system and reports any significant changes in system status. For more information, see the `ccmd` man page.

`conserver` The console-server that directs RS-232 console traffic from the Utility Board to the various sppconsole sessions.

/spp/bin

In the `/spp/bin` directory are specific commands and daemons that manage a V-Class node. Some of these are:

`est` The command (Exemplar Scan Test) to initiate scan testing.

`do_reset` The command executed on the SSP to reset the V-Class node remotely.

`sppdsh` An enhanced version of the Korn Shell (ksh) with all of the functionality of ksh, as well as new commands that are suited to a diagnostic environment.

`event_logger` A daemon that receives messages from diagnostic utilities through rpc calls and writes them to the event log for later review or processing.

`dcm` Dump Configuration Map. `dcm` dumps the boot configuration map information for the specified node.

/spp/scripts

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

`sppconsole` The console utility.

`hard_logger` The hard error logger script run automatically by `ccmd`.

/spp/data/*complex_name*

The `/spp/data/complex_name` directory contains:

`node_0.cfg` Configuration file with scan rings and configured hardware. 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.

SSP operation
SSP file system

<code>consolelogX</code>	A file containing all the console activity on the system, where <i>X</i> is the node ID.
<code>est.log</code>	The scan testing log.
<code>hard_hist</code>	Log of all hard failure information. Logs the output of all suspected ASIC (Application Specific Integrated Circuits). This file may be useful in troubleshooting intermittent ASIC failures.
<code>event_log</code>	Log of all event information. A read only file which captures information generated by the <code>ccmd</code> daemon.

/spp/firmware

The `/spp/firmware` directory is where firmware files are written when SSP software is installed. The firmware files are loaded from this directory into flash or SRAM.

/spp/est

The `est` directory contains files used during scan testing.

/spp/man

The `/spp/man` directory contains the man (manual) pages on many of the SSP specific commands.

Device files

Table 8 shows the differences in the device files between the HP B180L and HP 712 SSPs.

Table 8 **Device file differences**

Device	712		B180L	
	Location on workstation	Device file	Location on workstation	Device file
Private/diagnostic LAN	RS-232	/dev/lan1	LAN-AUI	/dev/lan0
Global/customer LAN	LAN-TP	/dev/lan0	Slot 2	/dev/lan1
Node 0 console port	RS-232	/dev/tty0p0	Serial 1	/dev/tty1p0
Terminal mux configuration port	9-PIN connector of the Y-Cable	/dev/ty1p0	Serial 2	/dev/tty0p0
Remote modem	9-PIN connector of the Y-Cable	/dev/ttyd1p0 or /dev/cua1p0 or /dev/cul1p0	Serial 2	/dev/ttyd0p0 or /dev/cua0p0 or /dev/cul0p0

System log pathnames

To separate the configuration and log files for each complex, several files have been moved to complex-specific directories. In Table 9, *complex* denotes specific complex names. These are assigned by the operator using `ts_config`.

The `/spp/data/complex` directories are created by `ts_config` during the “Configure Node” process. Configuration and log files are then created by the various daemon and utility programs as necessary.

The old/new pathname mappings are shown in Table 9.

Table 9

System log pathnames

Log name	V2500/V2600 pathname
CCMD log file	<code>/spp/data/ccmd_log</code>
CCMD old log file	<code>/spp/data/ccmd_log.old.Z</code>
Node 0 consolelog	<code>/spp/data/complex/consolelog0</code>
Node 2 consolelog	<code>/spp/data/complex/consolelog2</code>
Node CFG file	<code>/spp/data/complex/node_0.cfg</code>
Node PWR file	<code>/spp/data/complex/node_0.pwr</code>
Event log	<code>/spp/data/complex/event_log</code>
Event log archive	<code>/spp/data/complex/event_log.old.Z</code>
Hard Logger log	<code>/spp/data/complex/hard_hist</code>
Hard Logger temp	<code>/spp/data/complex/hl</code>
autoreset config	<code>/spp/data/complex/.ccmd_reboot</code>
EST log	<code>/spp/data/complex/est.log</code>
EST prior log	<code>/spp/data/complex/est.log.old</code>
cxttest log (GUI)	<code>/spp/data/complex/cxttest.log</code>

4

Firmware (OBP and PDC)

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

Boot sequence

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.

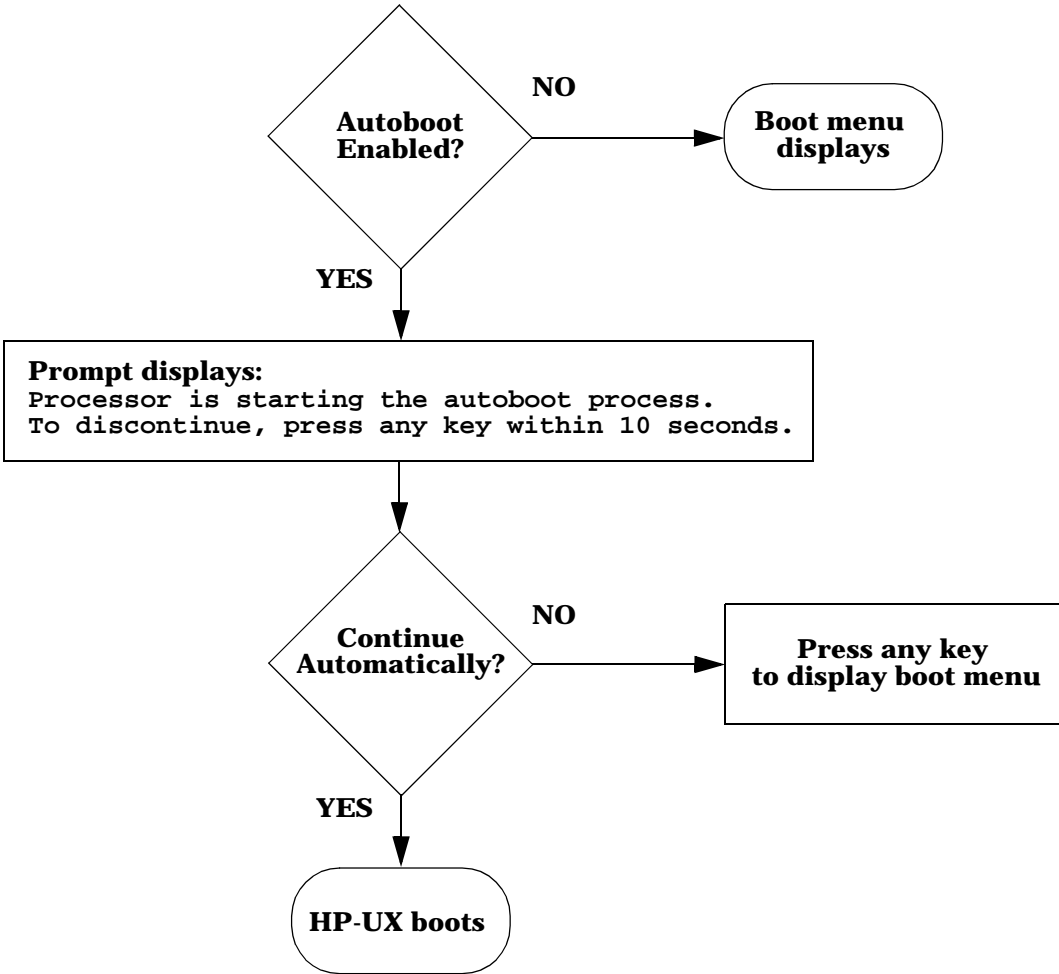
Once a machine powers on, the firmware controls the system until the operating system (OS) executes. If the system encounters an error any time during the boot process, it stops processing and goes to HP mode boot menu. See “HP mode boot menu” on page 64 for more information.

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

1. Power-On Self Test (POST) runs. POST is described in “LCD (Liquid Crystal Display)” on page 28 and in the *HP Diagnostics Guide: V2500/V2600 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. For further information see, “Enabling Autoboot” on page 67.

Figure 24 on page 61 illustrates the initialization and start-up process.

Figure 24 **Boot process**



Boot process output

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

```
POST Hard Boot on [0:PB4L_A]
HP9000/V2500 POST Revision 1.0.0.2, compiled 1999/04/12 11:51:10
Probing CPUs: PB4L_A
Completing core logic SRAM initialization.
Starting main memory initialization.
Probing memory: MB0L MB1L MB2R MB3R
Installed memory: 2048 MBs, available memory: 2048 MBs.
Initializing main memory.
      r0          r1          r2          r3
PB4L_A MB0L [:::: ____][:::: ____][____ ____][____ ____]
PB4L_A MB1L [:::: ____][:::: ____][____ ____][____ ____]
PB4L_A MB2R [:::: ____][:::: ____][____ ____][____ ____]
PB4L_A MB3R [:::: ____][:::: ____][____ ____][____ ____]
Building main memory map.
Main memory initialization complete.

Starting multinode initialization.
Collecting memory configuration from nodes: 0,2
Initializing ERI rings for node 0,2

Synchronizing nodes: 0,2
Initializing CTI cache.
      r0          r1          r2          r3
PB4L_A MB0L [LLLL ____][CCCC ____][____ ____][____ ____]
PB4L_A MB1L [LLLL ____][CCCC ____][____ ____][____ ____]
PB4L_A MB2R [LLLL ____][CCCC ____][____ ____][____ ____]
PB4L_A MB3R [LLLL ____][CCCC ____][____ ____][____ ____]
Synchronizing nodes: 0,2
Verifying remote memory access.
Enabling Time of Century synchronization routing.
Synchronizing nodes: 0,2
Multinode initialization complete.

Booting OBP
OBP Power-On Boot on [0:8]

Node 0 OBP attempting to synchronize with node 2.

Node(s) { 2 } now synchronized.
```

```
-----  
PDC Firmware Version Information  
PDC_ENTRY version 4.2.0.4  
POST Revision: 1.0.0.2  
OBP Release 4.2.0, compiled 99/01/06 14:00:18 (3)  
SPP_PDC 2.0.0 (03/08/99 11:47:51)  
-----
```

```
-----  
Multi-node Configuration Summary  
=====
```

```
NODE 0
```

```
UART? No  
CORE MAC address 0:a0:d9:0:bf:16 IP# 15.99.111.166 (0x0f636fa6)  
JTAG MAC address 0:a0:d9:0:bf:3 IP# 15.99.111.116 (0x0f636f74)  
MEMORY 2048 MB memory installed 1024 MB CTI cache configured  
CPUs 8  
PACs 0,1,2,3,4,5,6,7 MACs 0,1,2,3  
TACs 0,1,2,3 PCIs 0,3,4,7  
-----
```

```
NODE 2
```

```
UART? No  
CORE MAC address 0:a0:d9:0:be:eb IP# 15.99.111.167 (0x0f636fa7)  
JTAG MAC address 0:a0:d9:0:c3:a3 IP# 15.99.111.117 (0x0f636f75)  
MEMORY 2048 MB memory installed 1024 MB CTI cache configured  
CPUs 15  
PACs 0,1,2,3,4,5,6,7 MACs 0,1,2,3  
TACs 0,1,2,3 PCIs 2,6  
-----
```

```
4096 MB memory installed 2048 MB CTI cache configured (total, all nodes)
```

```
Primary boot path = 1/0/0.6.0  
Primary boot arguments =  
Alternate boot path = 15/3 NFS 15.99.111.99:/spp/os/uxinstlf  
Alternate boot arguments =  
Console path = 15/1  
Keyboard path = 15/1
```

```
[*** Manufacturing (or Debug) Permissions ON ***]
```

```
System is HP9000/800/V2500 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 Force 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 [<cpu>] [ON OFF SHOW]	(De)Configure/Show 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
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
RemoteCommand node# command	Execute command on a remote node
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
[0] Command:	

At this point, the user 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 10 and in the rest of this chapter.

Table 10 lists the commands available from the `Command: prompt`.

Table 10 **Boot menu commands**

Command	Description
AUto [BOot SEArch Force 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. If Autoforce is on, OBP allows HP-UX to boot even if one or more cabinets does not complete power on self test.
BObot [PRI ALT <i>path args</i>]	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 [<i>path</i> flash]	Displays the LIF contents (boot or flash volume) of a device.
PASSword	Defines the password used to control access to ForthMode. Same as UNIX <code>password</code> command.

Firmware (OBP and PDC)
HP mode boot menu

Command	Description
Path [PRI ALT CON] [<i>path</i>]	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.
RemoteCommand node# command	Executes the specified command on the remote node identified by node number.
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 [<i>path</i>]	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 from which the system can boot. 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 redisplay 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.
```

Firmware (OBP and PDC)

HE1p command

This chapter describes server configuration management and includes:

- `ts_config`
- `ccmd`
- `xconfig`
- Configuration utilities

Two utilities, `sppdsh` and `xconfig`, allow reading or writing configuration information. OBP can also be used to modify the configuration.

The SSP allows the user to configure the node using the `ts_config` utility. This is the preferred method for V2500/V2600 servers. `ts_config` configures the SSP to communicate with the node. The SSP daemon, `ccmd`, monitors the node and reports back configuration information, error information and general status. `ts_config` must be run before using `ccmd`.

ts_config

```
ts_config [-display display name]
```

Any V2500/V2600 nodes added to the SSP must be configured by `ts_config` to enable diagnostic and scan capabilities, environmental and hard-error monitoring, and console access.

Once the configuration for each node is set, it is retained when new SSP software is installed.

`ts_config` tasks include:

- Configuring a node—Adding and removing a node to the SSP configuration
- Configuring the terminal mux—Configuring and removing the terminal mux on the SSP
- Installing a node—Upgrading JTAG firmware, configuring a node `scub_ip` address, and resetting a node
- Configuration of Multiple-node complex—Configuring V2500/V2600 nodes into a single complex and splitting V2500/V2600 nodes out of a multiple-node complex.
- Operational support—Resetting a V2500/V2600 node or multiple-node complex and starting console sessions.

The user must have root privilege to configure a node or the terminal mux, because several HP-UX system files are modified during the configuration.

Starting ts_config

To start `ts_config` from the SSP desktop, click on an empty area of the background to obtain the Workspace menu and then select the `ts_config (root)` option. Enter the root password.

To start `ts_config` from a shell (local or remote), ensure that the `DISPLAY` environment variable is set appropriately before starting `ts_config`.

For example:

```
$ DISPLAY=myws:0; export DISPLAY      (sh/ksh/sppdsh)
% setenv DISPLAY myws:0              (csh/tcsh)
```

Also, the `-display` start-up option may be used as shown below:

For example:

```
# /spp/bin/ts_config -display myws:0
```

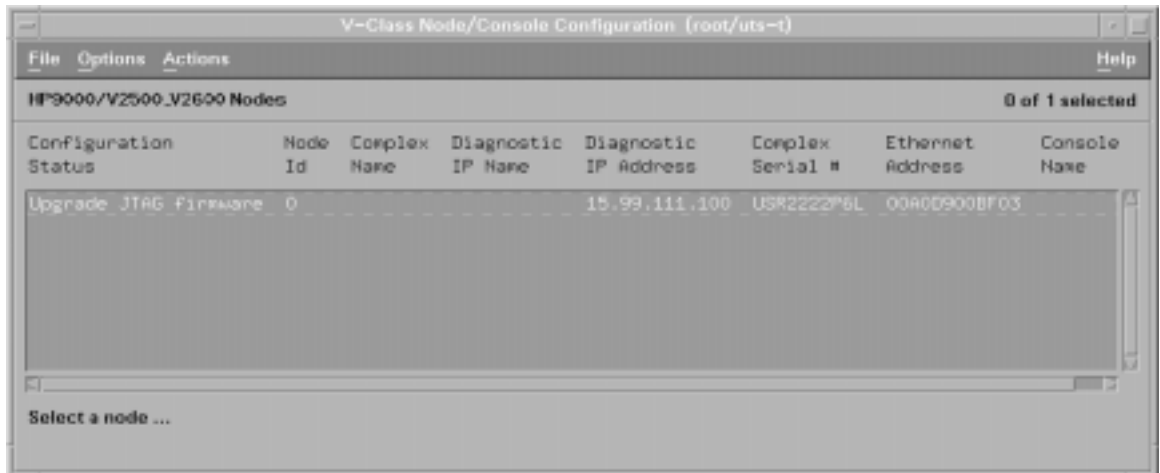
NOTE

For shells that are run from the SSP desktop, the `DISPLAY` variable is set (at the shell start-up) to the local SSP display.

ts_config operation

The `ts_config` utility displays an active list of nodes that are powered up and connected to the SSP diagnostic LAN. The operator selects a node and configures the selected node. A sample display is shown below.

Figure 25 `ts_config` sample display



The window has three main parts: the drop-down menu bar, the display panel, and the message panel. The display panel contains a list of nodes and their status. To select a node, click with the left-mouse button the line containing the desired node entry in the list. When a node is selected, information about that node is shown in the message panel at the bottom of the `ts_config` window. If an action needs to be performed to configure the node, specific instructions are included.

ts_config automatically updates the display when it detects either a change in the configuration status of any node or a newly detected node. The node display is not updated while an Action is being processed or while the user is entering information into an Action dialog.

The upper right corner of the ts_config window indicates whether a node has been selected.

The ts_config window title includes in parenthesis the name of the effective user ID running ts_config, either root or sppuser along with the host name of the SSP.

The ts_config display shows the configuration status of the nodes. Table 11 shows the possible status values.

Table 11 **ts_config status values**

Configuration Status	Description	Action Required
Upgrade JTAG firmware	The version of JTAG firmware running on the SCUB does not support the capabilities required to complete the node configuration process.	Select the node and follow the instructions given at the bottom of the ts_config window. ts_config guides the operator through the JTAG firmware upgrade procedure.
Not Configured	ts_config has detected the node on the Diagnostic LAN and the JTAG firmware is capable of supporting the node configuration activity and the node needs to be configured.	Select the node and follow the instructions given at the bottom of the ts_config window. ts_config guide the operator through the node configuration procedure described later in this document.

Configuration Status	Description	Action Required
Active	The node is configured and answering requests on the Diagnostic LAN.	None required. This is the desired status.
Inactive	The SSP node configuration file contains information about the specified node, but the node is not responding to requests on the Diagnostic LAN. This status is also shown if a node was configured and then removed from the SSP LAN without being deconfigured.	Power-up the node and/or check for a LAN connection problem. If the node information shown is for a node that has been removed, select the node then select "Actions," "Deconfigure Node," and click "Yes."
Node Id changed	The node is configured and answering requests on the diagnostic LAN, but the node ID currently reported by the node does not match the SSP configuration information.	Select the node to obtain additional information. If the node COP information was changed to a different node ID and the new node ID is correct, select "Actions," "Configure Node," then click "Configure." The SSP configuration information is updated using the new node ID.

Configuration procedures

The following procedures provide additional details about each configuration action and are intended as a reference. `ts_config` automatically guides the user through the appropriate procedure when a node is selected.

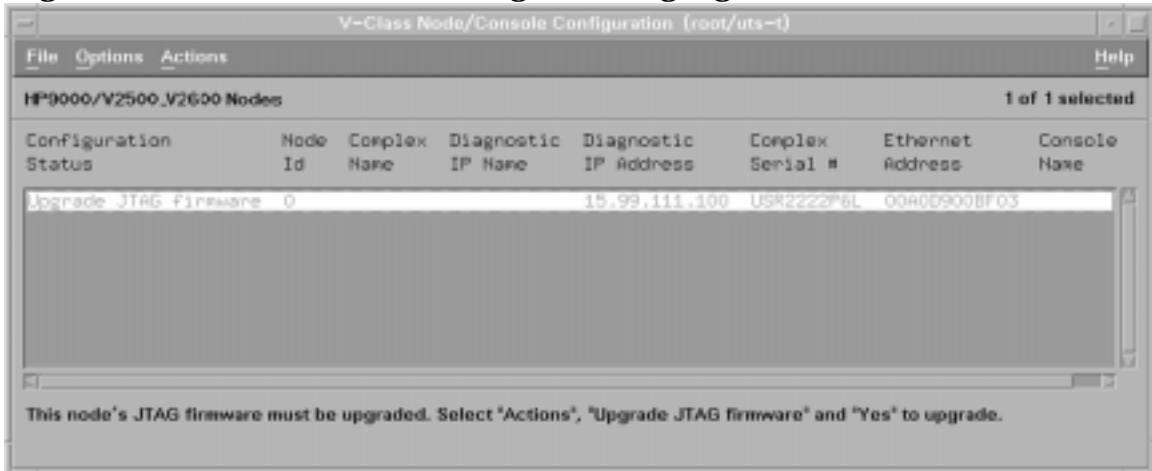
Upgrade JTAG firmware

NOTE

If the node shows "Not Configured," do not perform this procedure. Perform the following procedure only when the status shows "Upgrade JTAG firmware."

- Step 1.** Select the node from the list in the display panel. For example, clicking on node 0 in the list highlights that line as shown in Figure 26.

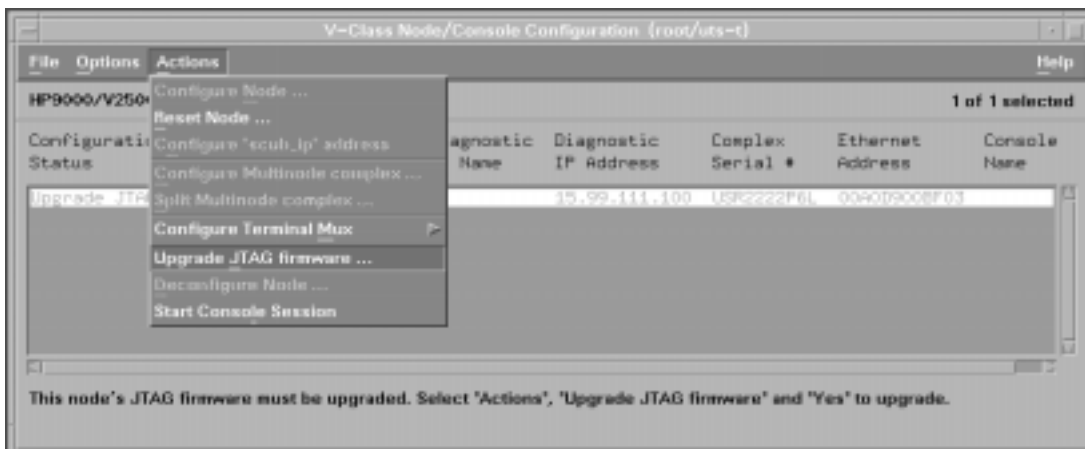
Figure 26 ts_config showing node 0 highlighted



Notice that after the node has been highlighted that `ts_config` displays information concerning the node. In this step, it tells the user what action to take next, “This node’s JTAG firmware must be upgraded. Select “Actions,” “Upgrade JTAG firmware” and “Yes” to upgrade.”


Step 2. Select “Actions” to drop the pop-down menu and then click “Upgrade JTAG firmware,” as shown in Figure 27.

Figure 27 ts_config “Upgrade JTAG firmware” selection.




Step 3. A message panel appears as the one shown in Figure 28. Read the message. If this is the desired action, click “Yes” to begin the upgrade.

Figure 28 Upgrade JTAG firmware confirmation panel



Step 4. After the firmware is loaded a panel appears as the one shown in Figure 29. Click “OK” and then power-cycle the node to activate the new firmware.

Figure 29 ts_config power-cycle panel




When the node is powered up, the “Configuration Status” should change to “Not Configured.”

Configure a Node

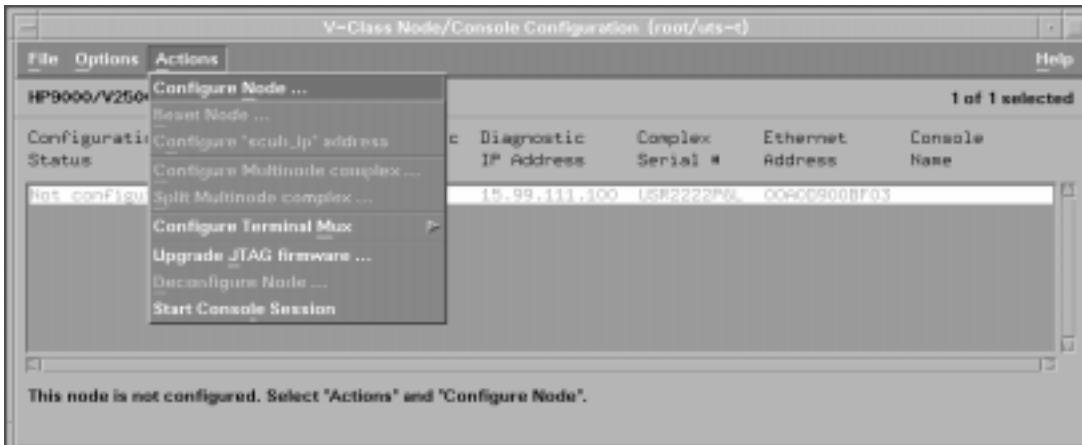
Step 1. Select the desired node from the list of available nodes. When the node is selected, the appropriate line is highlighted as shown in Figure 30. Notice the bottom of the display indicates the Node 0 is not configured and provides the steps necessary to configure the node.

Figure 30 ts_config indicating Node 0 as not configured




Step 2. Select “Actions” and then click “Configure Node,” as shown in Figure 31.

Figure 31 ts_config “Configure Node” selection.



After invoking `ts_config` to configure the node, a node configuration panel appears as the one in Figure 32.

Figure 32 ts_config node configuration panel



- Step 3.** Enter a name for the V2500/V2600 System. The SSP uses this name as the “Complex Name” and to generate the IP host names of the Diagnostic and OBP LAN interfaces. Select a short name that SSP users can easily relate to the associated system (for example: **hw2a**, **swtest**, etc.).
- Step 4.** Select an appropriate serial connection for the V2500/V2600 console from the pop-down option menu in the node configuration panel.

ts_config automatically assigns the first unused serial port. If the terminal mux has been configured, the terminal mux ports are included in the list of available serial connections.


The IP address information for the Diagnostic interface is provided. The ts_config utility automatically changes the IP address of the diagnostic LAN interface to prevent a duplicate when other nodes are added to this SSP configuration.

ts_config automatically updates the local /etc/hosts file with the names and addresses of the Diagnostic and OBP LAN interfaces.

- Step 5.** Click “Configure.”


This updates several SSP files. The node configuration confirmation panel appears as the one in Figure 33.

Figure 33 ts_config restart workspace manager panel.



Step 6. Read the panel and click “OK.” When the configuration process is complete, the “Configuration Status” of the node changes to “Active,” as shown in Figure 34.

Figure 34 ts_config indicating Node 0 is configured



Step 7. Restart the Workspace Manager: Click the right-mouse button on the desktop background to activate the root menu. Select the “Restart” or “Restart Workspace Manager” option, then “OK” to activate the new desktop menu.

NOTE

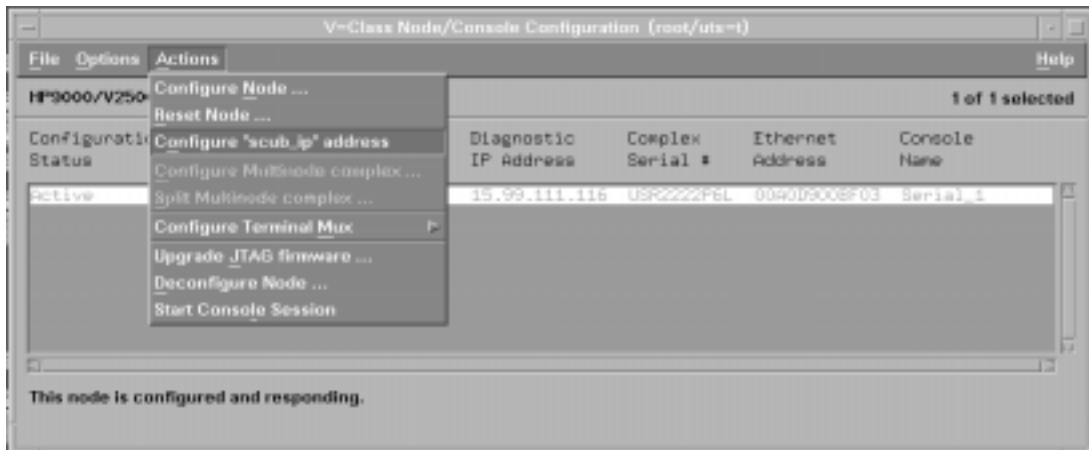
If adding multiple nodes to the SSP, wait until the final node is added before restarting the Workspace Manager.

Configure the scub_ip address

Step 1. Select the desired node from the list of available nodes.

Step 2. In the ts_config display panel, select “Actions” and then “Configure ‘scub_ip’ address,” as shown in Figure 35.


Figure 35 ts_config “Configure ‘scub_ip’ address” selection



ts_config checks the scub_ip address stored in NVRAM on the SCUB in the node. This would initially be the default address set at the factory.

If the scub_ip address is correct, the panel shown in Figure 36 is displayed and no action is required. If the node is not detected and scanned by ccmd, ts_config may ask you to try again later. The ccmd detection scan process should take less than a minute.

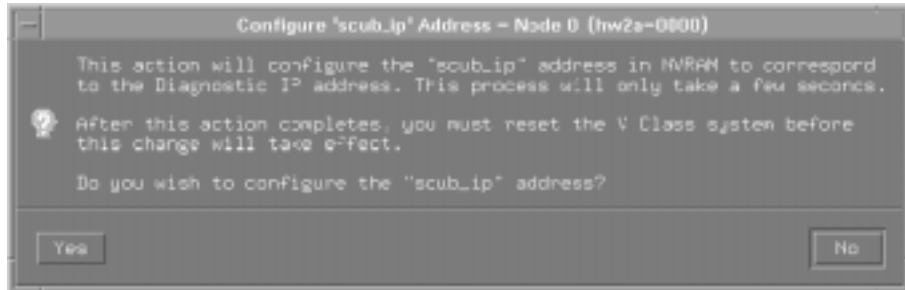
Figure 36 ts_config “SCUB OK” panel



Step 3. If prompted by ts_config (as indicated by the panel in Figure 37), click “Yes” to correctly set the scub_ip address.

Figure 37


ts_config scub_ip address configuration confirmation



Step 4. A panel as the shown in Figure 38 appears confirming that the scub_ip address is set. Click OK.

Figure 38

ts_config scub_ip address set confirmation panel



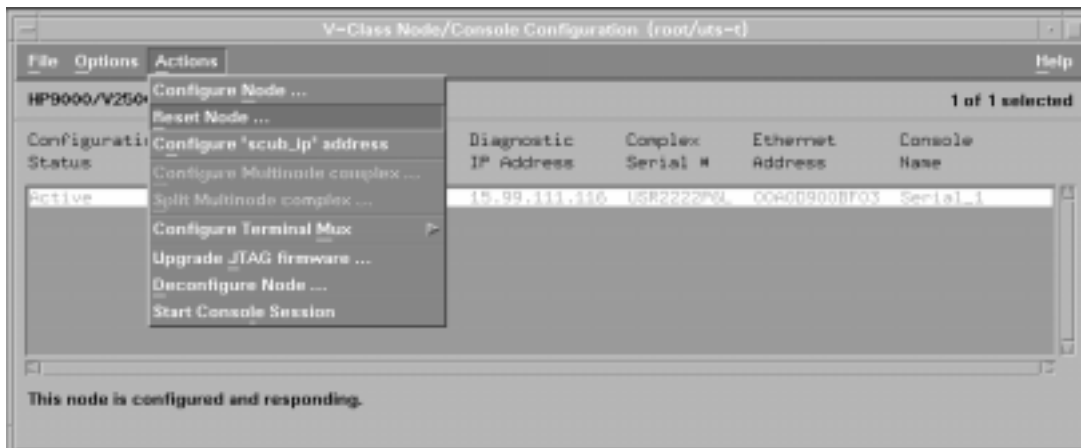
Initiate a node reset to activate the new scub_ip address.

Reset the Node

Step 1. Select the desired node from the list of available nodes.


Step 2. Select "Actions," then "Reset Node." This is indicated in Figure 39.

Figure 39 ts_config “Reset Node” selection



A panel as the one shown in Figure 40 appears.

Figure 40 ts_config node reset panel



Step 3. In the Node Reset panel, select the desired “Reset Level” and “Boot Options,” then click Reset.”

Configuration utilities
ts_config

Deconfigure a Node

Deconfiguring a node removes the selected node from the SSP configuration. The SSP will no longer monitor the environmental and hard-error status of this node. Console access to the node is also be disabled.

- Step 1.** Select the desired node from the list of available nodes.
- Step 2.** Select “Actions,” then “Deconfigure Node,” then click “Yes.”

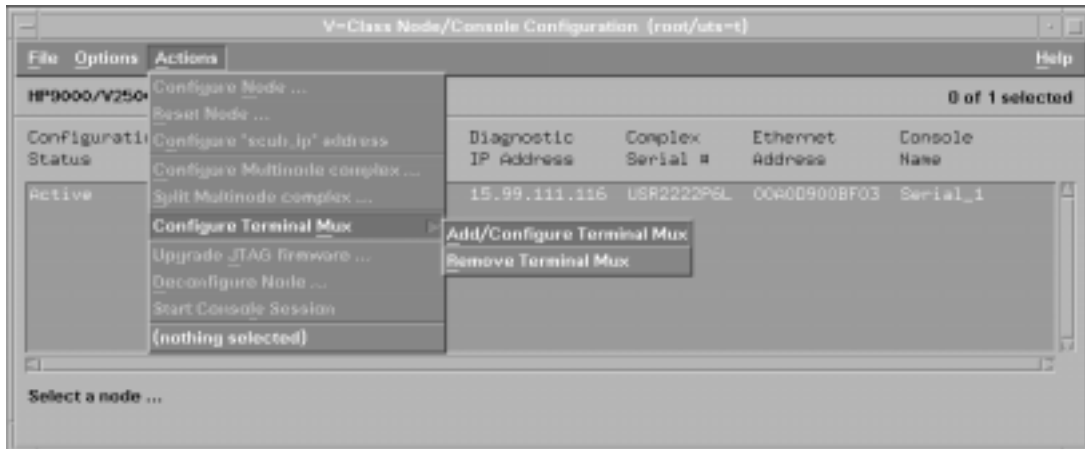
Add/Configure the Terminal Mux

To add or reconfigure the terminal mux, perform the following procedure.

- Step 1.** In the ts_config display, select “Actions,” then “Configure Terminal Mux.”

Select “Add/Configure Terminal Mux.” This is indicated in Figure 41.


Figure 41 ts_config “Add/Configure Terminal Mux” selection



- Step 2.** Connect a serial cable from serial port 2 on the SSP to port 1 on the terminal mux.
- Step 3.** A panel shown in Figure 42 display the IP address.

Figure 42

Terminal mux IP address panel



Remove terminal mux

ts_config does not remove the terminal mux if any node consoles are assigned to terminal mux ports.

- Step 1.** Select “Actions,” then “Configure Terminal Mux.”
- Step 2.** Select “Remove Terminal Mux,” then click “Yes.”

Console sessions

ts_config may also start console sessions by selecting the desired node(s) and then selecting the “Start Console Session” action as shown in Figure 43. Figure 44 shows the started console sessions.

Figure 43 “Start Console Session” selection


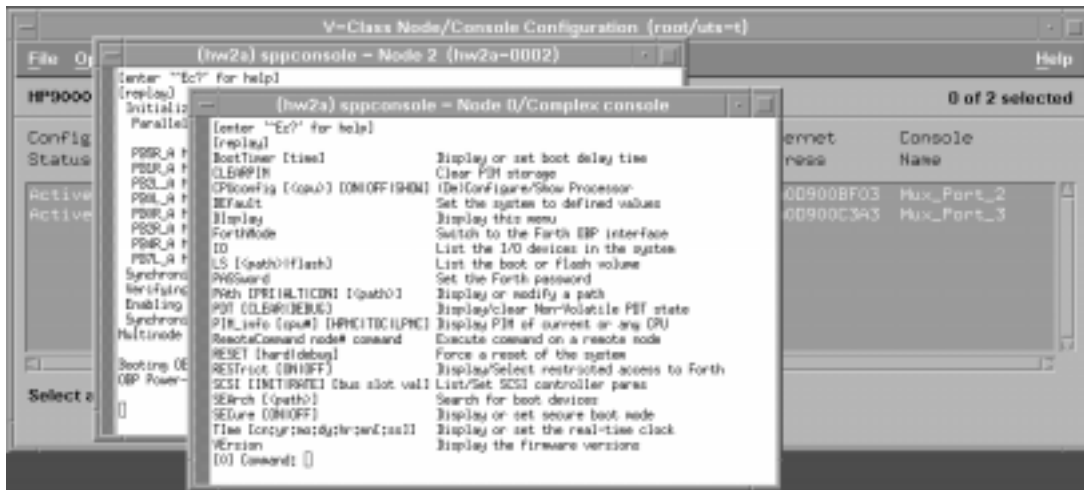


Figure 44 Started console sessions




V2500/V2600 SCA (multinode) configuration

ts_config can also configure a V2500/V2600 SCA system. An example to follow describes how. The example assumes that there are two active single-node complexes. After the system has rebooted to OBP, node 0 becomes the console for the SCA complex.

To configure the two-node system in the example, start ts_config as described in “Starting ts_config” on page 72. Once ts_config has started, a window like that shown in Figure 45 is displayed.

Figure 45 SSP supporting two single-node complexes

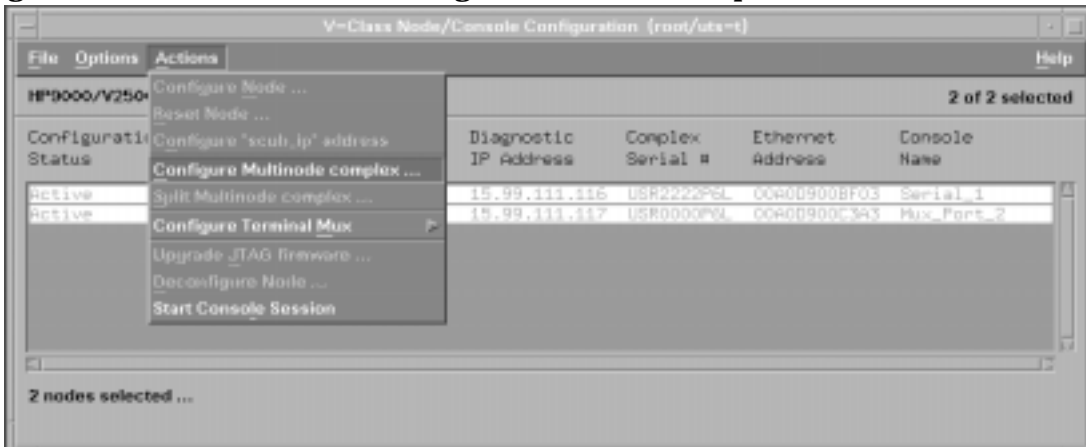


The following procedure configures the two-node SCA system in the example.

- Step 1.** Select the nodes by clicking anywhere in the information display for each node. As the nodes are selected, they become highlighted.
- Step 2.** Select “Action” and then “Configure Multinode complex” as shown in Figure 46.

Configuration utilities
ts_config

Figure 46 ts_config Configure Multinode complex selection



Step 3. When “Configure Multinode complex” is selected, a configuration dialog appears as shown in Figure 57.

Figure 47 Configure Multinode Complex dialog window



- Step 4.** Enter the required fields into the Configure Multinode Complex dialog window.
- V-Class Complex Name—Current complex name of either node or a new complex name.
 - Complex Serial Number—Unique serial number of the complex. This is not required if the nodes have the same serial numbers.
 - Complex Key—Number required to enter the Complex Serial Number.

NOTE

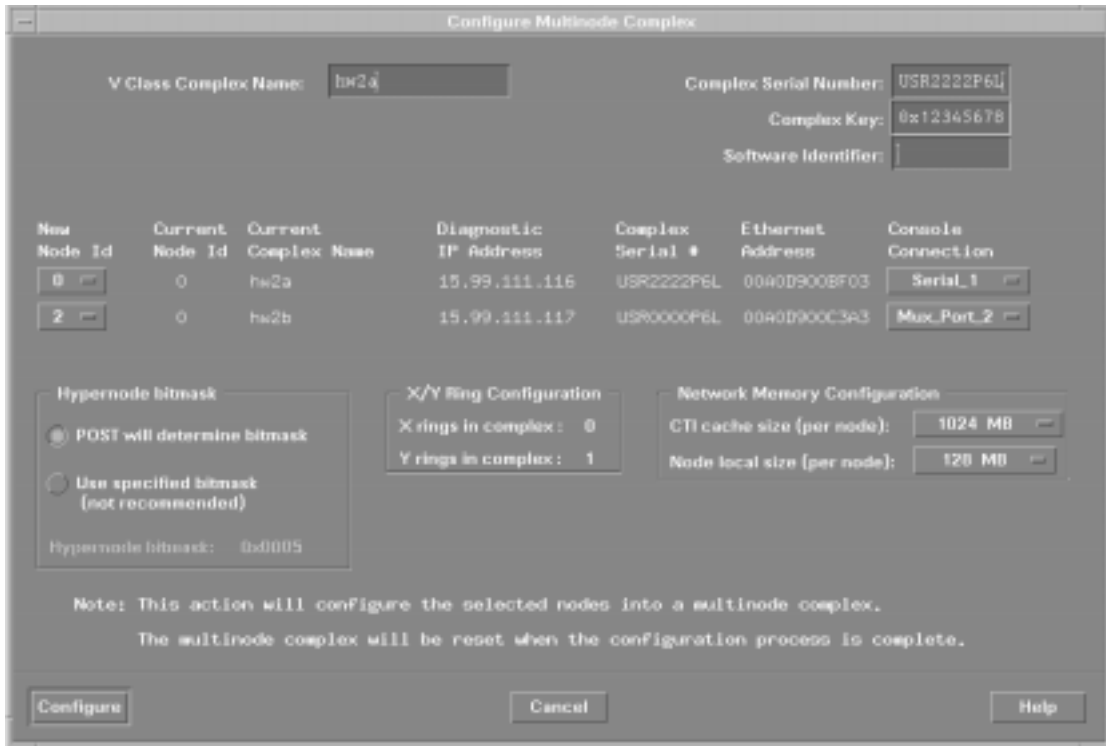
If all of the SCA system complex serial numbers are the same, no complex key is required.

- Step 5.** Select the desired node IDs from the “New Node ID” drop-down lists.
- Step 6.** If the console connection must be changed, select appropriate connection from the “Console Connection” drop-down list.
- Step 7.** In the “Hypernode bitmask” section select “POST will determine bitmask.”
- Step 8.** If necessary, select the desired CTI cache size from the “CTI cache size” drop-down list.
- Step 9.** If necessary, select the node-local memory size from the “Node local size” pull-down list.

NOTE


The default settings for CTI cache and node-local memory are recommended.

Figure 48 **Configure Multinode Complex dialog window with appropriate values**



Step 10. Click the “Configure” button to start the configuration. A message box appears indicating that the configuration has started.

Figure 49 **Configuration started information box**



The following activities occur during the configuration process:


- SSP files are updated based on the new complex and node names.
- Essential console server processes are started, and the now-obsolete server processes are halted.
- New node information is written to the COP chip in each node.

This information includes:

- Node ID
- Complex serial number (if it has been modified)
- Requested or auto-generated software identifier
- Configuration Manager Daemon, `ccmd`, is notified of the new configuration.
- The shared-memory database of node information is updated.
- Multinode configuration parameters are written to NVRAM in each node. These include:
 - Hypernode bitmask
 - X- and Y-ring information
 - Node count
 - CTI cache size
 - Node-local memory size
- The boot vector of each node is set to OBP and each node is reset.

When the configuration process is complete, `ts_config` shows the new multinode complex, as in Figure 50. The restart process activates the new SSP root menu which includes customized menus for each complex. The Workspace Manager must be restarted or else the root menu will be outdated (the rest of the configuration process is complete).

Figure 50 ts_config showing newly configured complexes



When remotely running `ts_config`, the Restart Workspace Manager step cannot be performed, because it is the SSP Workspace Manager that needs to be restarted. The Workspace Manager can be restarted at any time by clicking on the desktop background and selecting Restart Workspace Manager, then OK.

Any of the configurable parameters on the Multinode Configuration dialog may be changed by selecting each node and choosing the "Configure Multinode complex" action. Set the desired options and click Configure. During a reconfiguration, several of the required fields in the Multinode Configuration dialog are filled in by `ts_config`.

V2500/V2600 split SCA configuration

`ts_config` also provides a "Split Multinode complex" action that takes an SCA complex and logically splits it into single node systems. Each node becomes Node 0 in a new complex.

The following procedure allows the user to split the SCA system:

- Step 1.** In the `ts_config` window, select the desired nodes.
- Step 2.** Select every node in the desired complex, then "Actions," and then "Split Multinode complex." Figure 52 shows the `ts_config` Split Multinode complex panel.

Figure 51 ts_config Split Multinode complex operation

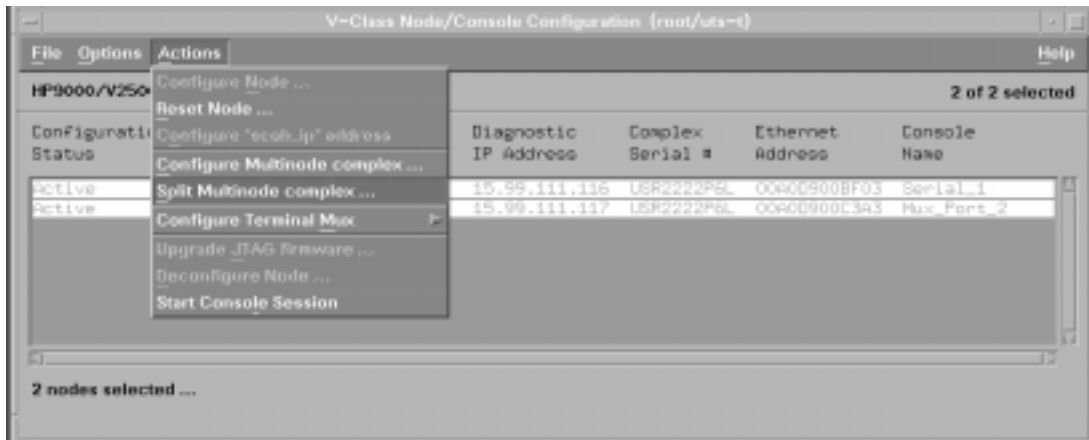
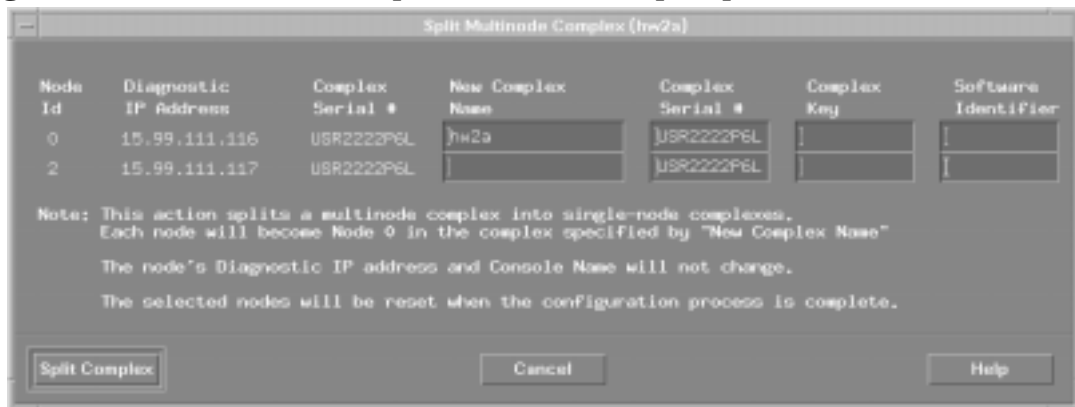
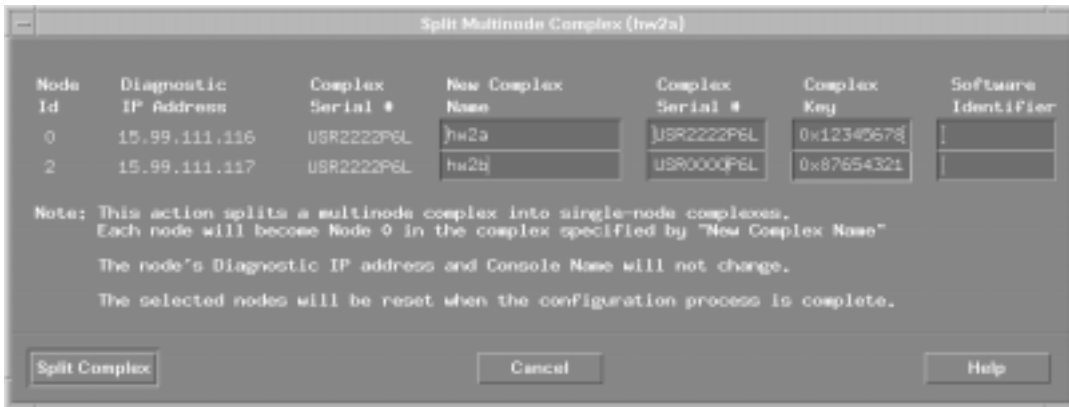


Figure 52 ts_config Split Multinode complex panel



Step 3. Enter the complex names for each node. New complex serial numbers may be assigned. Each node becomes node 0 in a new complex. Figure 53 shows the Split Multinode panel filled in. Click the Split Complex button to initiate the configuration process.

Figure 53 ts_config Split Multinode complex panel filled in



The message shown in Figure 54 appears indicating the configuration is taking place.

Figure 54 Split Multinode confirmation panel


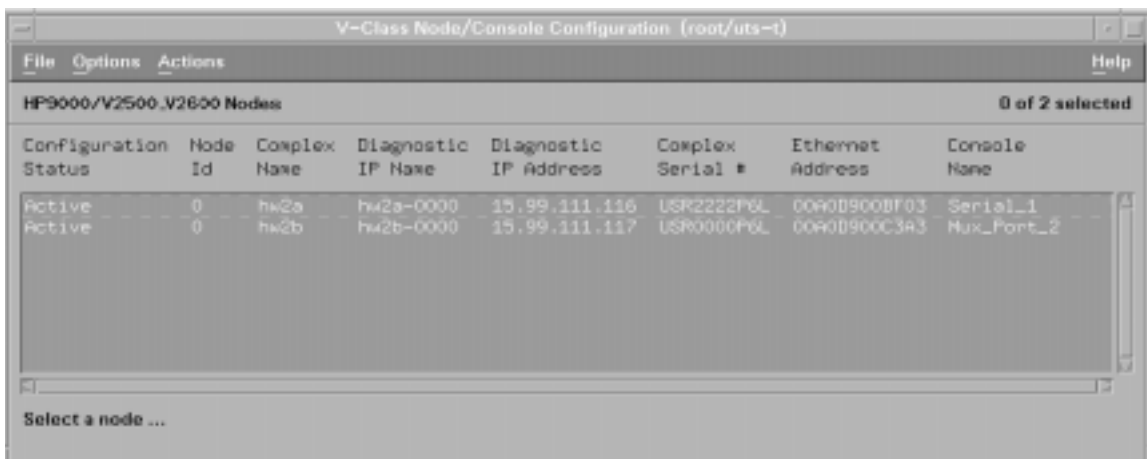


Figure 55 shows the main ts_config display after the split multinode operation has completed. It shows the resulting configuration: two single node complexes (two node 0s) with names assigned in the prior step.

Figure 55 ts_config Split Multinode operation complete



ts_config files

ts_config either reads or maintains the following SSP configuration files:

/etc/hosts The standard system hosts file, includes entries for the cabinet related IP addresses.

/etc/services Service definitions for the console interface.

/etc/
inetd.conf Contains entries for starting console related processes

/spp/data/
nodes.conf Contains entries which define the complexes (either single cabinet or multi-cabinet) managed by the SSP. This file is maintained by the Configure Node Action of ts_config, but other commands can also update this file: delete_node, configure_node, and split_multinode.

/spp/data/
conserver.cf Connection definitions for the console interface.

/spp/data/
consoles.conf Console name to ttylink number resolution. This file is maintained by the Configure Mux Action of ts_config.

/spp/data/
<complex_name> For each newly configured complex, there is a complex-specific directory that contains complex-specific files, such as event and console logs. ts_config generates each <complex_name> during configuration.

The nodes.conf file contains most of the configuration management information. It defines the relationship between complex names, cabinets (nodes) within that complex and the associated host names and console port connections.

Each node has an entry in the nodes.conf file as follows:

NODE Complex Node ID JTAG-hostname OBP-hostname SSP-hostname Console-port

The variables of the entry are defined as follows:

NODE—Keyword designating a cabinet (node) entry.

Complex—Name to which the node (cabinet) is associated.

In a multi-cabinet complex all the cabinets comprise a single system (complex) and are managed by a single console (the console on cabinet 0). Each cabinet, however, has its own console that can report diagnostic information, and there is still a console configuration entry for each cabinet. These consoles, however, are not normally accessed.

Node ID—The identification of the V-Class node. This number can be 0, 2, 4, or 6.

Jtag-hostname—Host name used by the SSP to communicate with the JTAG firmware on the associated node. JTAG IP addresses are 15.99.111.116 through 15.99.111.131.

Obp-hostname—Host name used for OBP communication (not normally referenced while administering the complex). OBP IP addresses are 15.99.111.166 through 15.99.111.181.

SSP-hostname—Local host name of the private/diagnostic LAN. The default host name for this interface is tsdart-d.

Console-port—Name of the physical connection to the node RS-232 console port. The port name is linked to a ttylink service entry via the file /spp/data/consoles.conf.

The `get_node_info` program extracts information from the `nodes.conf` file.

IMPORTANT

Most SSP configuration is carried out by `ts_config`. Sometimes problems are caused by manually editing the files maintained by `ts_config`, resulting in it not be able to properly parse the files on its own.

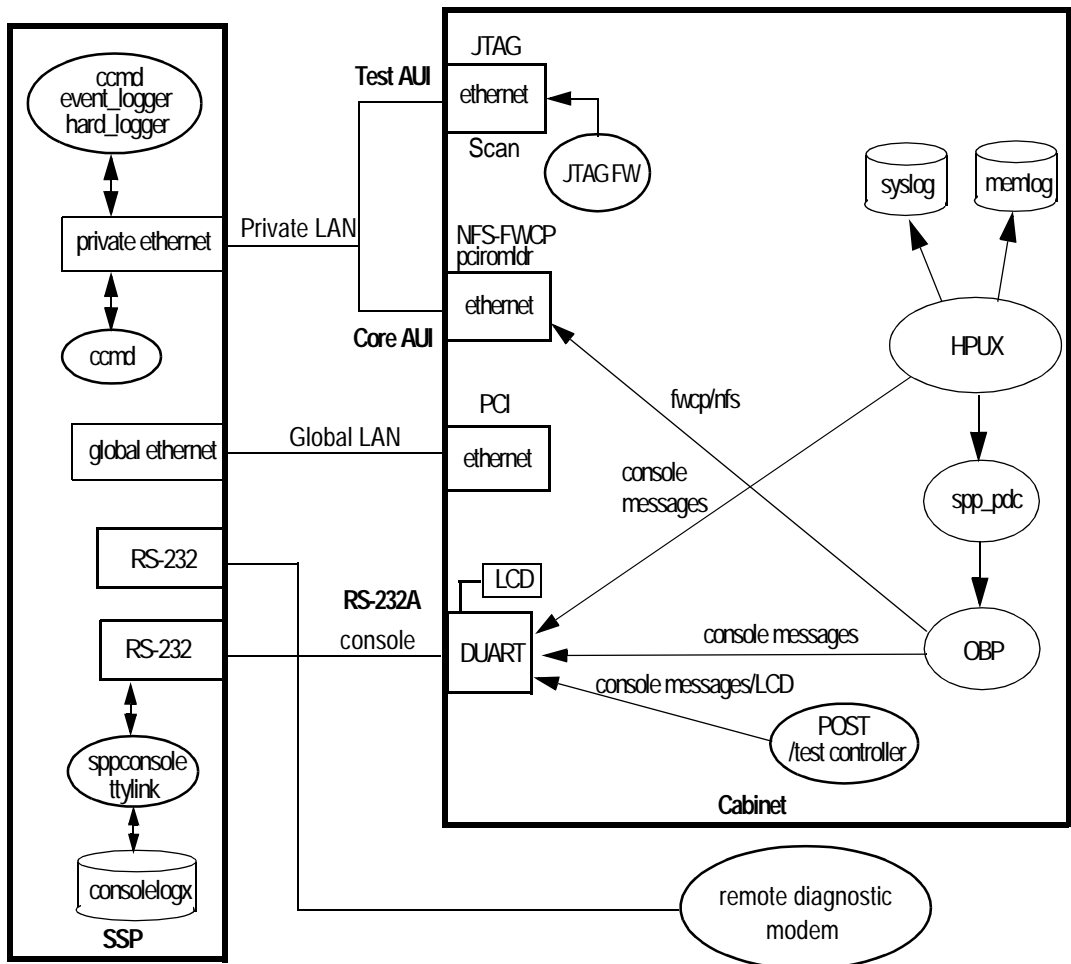
The `ts.install` script is designed not to run after initial installation to prevent inadvertent removal of changes made to the configuration files by `ts_config`. It can be forced to run, however, in order to put all the configuration files back to factory defaults. To rerun `ts.install`, execute the following command:

```
/spp/scripts/inst/ts.install
```

SSP-to-system communications

Figure 56 depicts the V-Class server to SSP communications using HP-UX.

Figure 56 SSP-to-system communications



A layer of firmware between HP-UX and OBP (Open Boot PROM) called `spp_pdc` allows the HP-UX kernel to communicate with OBP. `spp_pdc` is platform-dependent code and runs on top of OBP providing access to the devices and OBP configuration properties.

LAN communications

There are two ethernet ports located on the SCUB as shown in the diagram in the upper-left side of the node (dotted line) in Figure 56 on page 97. These comprise the “private” or diagnostic LAN. The JTAG port is used for scanning, and the NFS-FWCP port is used for downloading system firmware via `nfs` using the `fwcp` utility, via `tftp` using the `pdcf1` utility, downloading disk firmware using the `dfdutil` utility (`dfdutil` uses `tftp` for reading peripheral firmware), and loading Symbios FORTH code using the `pciromldr` utility. For more information on `dfdutil`, `tftp`, and `pciromldr`, see the appropriate man pages.

The configuration daemon, `ccmd`, which is located on the SSP obtains system configuration information over the private LAN from the JTAG port. It builds a configuration information database on the SSP. The board names and revisions, the device names and revisions, and the start-up information generated by POST are all read and stored in memory for use by other diagnostic tools.

IMPORTANT

Both the B180L and the 712 workstations must have two ethernet connections: one for the private LAN and one for the global LAN. These ports are different on each model of workstation. It is important that the installer connect the LAN cables to the correct connector on the SSP.

The SSP can be placed on the customers Local Area Network using the SSP's global ethernet.

SSP host name and IP addresses

The SSP software installation process assigns the correct IP address to the appropriate LAN device.

The SCUB IP address 15.99.111.116 is the first SCUB IP address available on the SSP. If more nodes are added, they are assigned 15.99.111.117, 15.99.111.118, and so on. `ts_config` keeps track of which addresses are assigned.

The SCUB IP host name is “*complexname-000n*,” where *n* is the node ID. The SSP supports multiple complexes (that is, multiple node 0s).

Serial communications

The DUART port on the SCUB provides an RS232 serial link to the SSP. Through this port HP-UX, OBP, POST (Power-On Self Test) and the Test Controller send console messages. The SSP processes these messages using the `sppconsole` and `ttylink` utilities and the `consolelogx` log file. POST and OBP also send system status to the LCD connected to the DUART. For more information on `sppconsole`, `ttylink`, and `consolelogx`, see the appropriate man pages.

NOTE

The second RS-232 port on the workstations are unused and not enabled at this time.

ccmd

`ccmd` (Complex Configuration Management Daemon) is a daemon that maintains a database of information about the V2500/V2600 hardware. `ccmd` also monitors the system and reports any significant changes in system status. It supports multiple nodes, multiple complexes and nodes that have the same node number.

There are two types of related information in the database: node information (node numbers, IP addresses and scan data) and configuration data which is initialized by POST. The node information is scanned from the hardware and processed with the aid of data files at `/spp/data`. The POST configuration data is required so that certain scan based utilities can emulate various hardware functions.

`ccmd` periodically sends out a broadcast to determine what nodes are available. If `ccmd` can not talk to a node that it previously reached, it sends a response to the console and the log. If it establishes or re-establishes contact, or if a node powers up, `ccmd` reads hardware information from the node and interrogates it through scan to determine the node configuration. From this data, a complete database is built on the SSP that will be used for all scan based diagnostics.

Once running, `ccmd` checks for power-up, power-down, reset, error, and environmental conditions on regular intervals. If at any time `ccmd` detects a change in the configuration, it changes the database, updates the `/spp/data/complex.cfg` file, sets up a directory in `/spp/data` for each complex and initializes a `node_#.pwr` file for each node in the complex specific subdirectory.

If `ccmd` detects an error condition, it invokes an error analysis tool (`hard_logger`) that logs and diagnoses error conditions. After an error is investigated, `ccmd` reboots the node or complex associated with the error. The reboot operation can be avoided with the use of the `autoreset [complex] on|off|chk` utility in `/spp/scripts`.

`ccmd` is listed in `/etc/inittab` as a process that should run continuously. It may be started manually, but since it kills any previous copies at start-up, diagnostic processes that may be running will be orphaned. Only one copy of `ccmd` may run on a SSP.

If started with no options, `ccmd` disassociates itself from the terminal or window where it was started. It instead reports to the console window and the file `/spp/data/ccmd_log`.

If `ccmd` is sent a `SIGHUP`, it regenerates the database.

All scan-based operations require `ccmd`. If `POST` is unable to run, then `ccmd` is not able to read configuration data and some system information is not accessible.

`ccmd` works in co-operation with most utilities to share a common ethernet port and Diagnostic (DART) bus. In general, the scan data is sent via UDP. The DART bus should be separate from any general purpose ethernet bus. If the DART bus is improperly set-up, `ccmd` cannot run properly.

Since `ccmd` can become corrupted by bad data, it may be necessary to kill the `ccmd` process to refresh the SSPs configuration image. Killing the `ccmd` process is not always enough. If the “heart beat” LED from the SCUB is not functioning then `ccmd` is unable to communicate with the system. A `ping` command to the SCUB will not be successful either. In this case, the system or node must be powered down to reset the SCUB and re-establish communication with the SSP.

xconfig

`xconfig` is the graphical tool that can also modify the parameters initialized by POST to reconfigure a node.

The graphical interface allows the user to see the configuration state. Also the names are consistent with the hardware names, since individual configuration parameters are hidden to the user. The drawback of `xconfig` is that it can not be used as a part of script-based tests, nor can it be used for remote debug.

`xconfig` is started from a shell. Information on node 0 is read and interpreted to form the starting X-windows display shown in Figure 57.

The `xconfig` window appears on the system indicated by the environmental variable `$DISPLAY`. This may be overridden, however, by using the following command:

```
% xconfig -display system_name:0.0
```

The `xconfig` window has two display views: one shows each component as a physical location in the server, the other shows them as logical names. Figure 57 and Figure 58 show the window in each view, respectively. To switch between views, click on the Help button in menu bar and then click the Change names option. See “Menu bar” on page 105.

Figure 57 xconfig window—physical location names

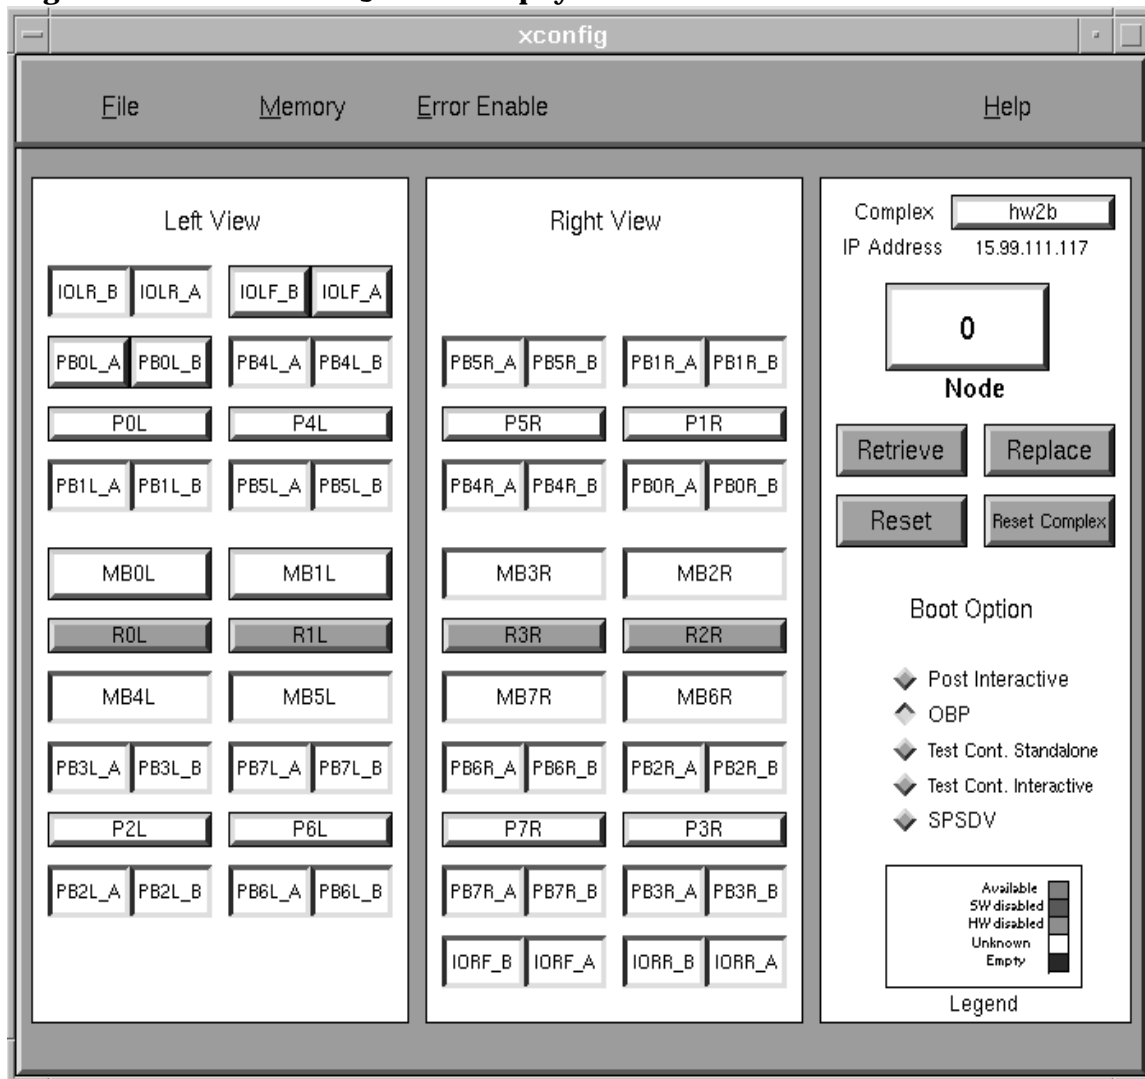
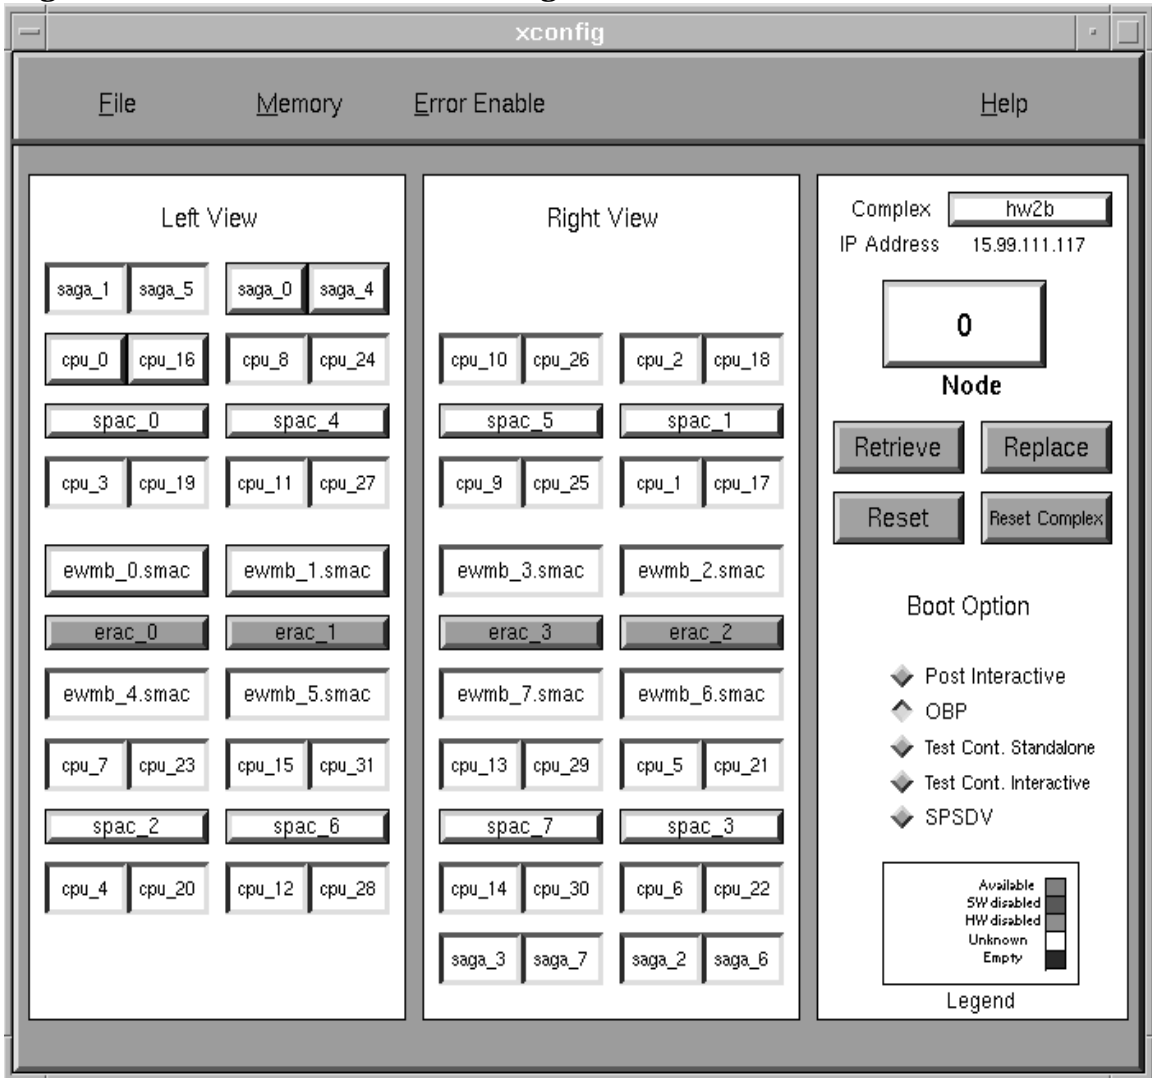


Figure 58 xconfig window—logical names



As buttons are clicked, the item selected changes state and color. There is a legend on the screen to explain the color and status. The change is recorded in the SSP's image of the node.

When the user is satisfied with the new configuration, it should be copied back into the node, and the node should be reset to enable the changes.

The main `xconfig` window has three sections:

- Menu bar—Provides additional capability and functions.
- Node configuration map—Provides the status of the node.
- Node control panel—Provides the capability to select a node and control the way data flows to it.


Menu bar

The menu bar appears at the top of the `xconfig` main window. It has four menus that provide additional features:

- File menu—Displays the file and exit options.
- Memory menu—Displays the main memory and CTI cache memory options.
- Error Enable menu—Displays the device menu options for error enabling and configuration.
- Help menu—Displays the help and about options.

The menu bar is shown in Figure 59.

Figure 59 `xconfig` window menu bar



The File menu provides the capability to save and restore node images and to exit `xconfig`.

The Memory menu provides the capability to enable or disable memory at the memory DIMM level by the total memory size and to change the network cache size on a multinode complex.


The Error Enable menu provides the capability to change a device's response to an error condition. This is normally only used for troubleshooting.

The Help menu provides a help box that acts as online documentation and also contains program revision information.

Node configuration map

The node configuration map is a representation of the left and right side views of a node as shown in Figure 60.

Figure 60 xconfig window node configuration map



The button boxes are positioned to represent the actual boards as viewed from the left and right sides. Each of the configurable components of the node is in the display. The buttons are used as follows:

- Green button—Indicates that the component is present and enabled.
- Red button—Indicates that the component is software disabled in the system.
- White button—Indicates that it is not possible to determine what the status of the component would be if POST were to be started.
- Blue box—Indicates that the component is either not present or fails the power-on self tests.
- Brown button—Indicates that POST had to hardware deconfigure this component in order to properly execute.
- Grey button—Indicates a hardware component that did not properly initialize.

The colors are shown in the legend box of the node control panel.

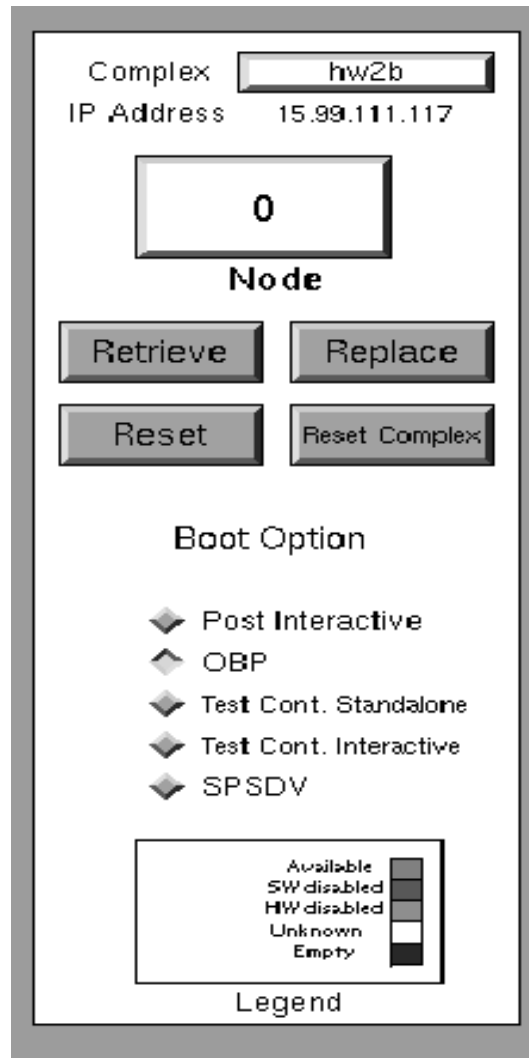
Components can change from enabled to disabled or disabled to unknown by clicking on the appropriate button with the left mouse button.

A multinode system requires an additional component on a memory board to enable the scalable coherent memory interface. This component can be viewed by right clicking the on the memory board button. The right mouse button toggles the memory board display between the memory board and the SCI device

Node control panel

The node control panel allows the user to select a node, select the stop clocks on an error function, select the boot parameters for a node and direct data flow between the node and the `xconfig` utility. It is shown in Figure 61.

Figure 61 xconfig window node control panel



The node number is shown in the node box. A new number can be selected by clicking on the node box and selecting the node from the pull-down menu. A new complex can be selected by clicking on the complex box and selecting it from the pull-down. A node IP address is displayed along with the node number and complex.

When a new node is selected and available, its data is automatically read and the node configuration map updated. The data image is kept on the SSP until it is rebuilt on the node using the Replace button. This is similar to the replace command on `sppdsh`.

Even though data can be rebuilt on a node, it does not become active until POST runs again and reconfigures the system. The Reset or Reset All buttons can be used to restart POST on one or all nodes of a system. A multinode system requires a reset all to properly function.

A Retrieve button is available on the node control panel to get a fresh copy of the parameters settings in the system. Clicking this button overwrites the setting local to the SSP and `xconfig`.

The Stop-on-hard button is typically used to assist in fault isolation. It stops all system clocks shortly after an error occurs. Only scan-based operations are available once system clocks have stopped.

The last group of buttons controls what happens after POST completes. The node can become idle or boot OBP, the test controller, or `spsdv`. The test controller and `spsdv` are additional diagnostic modes.

Configuration utilities

V2500/V2600 diagnostics provides utilities that assist the user with configuration management.

autoreset

`autoreset` allows the user to specify whether `ccmd` should automatically reset a complex after a hard error and after the hard logger error analysis software has run. `autoreset` occurs if a `ccmd_reset` file does not exist in the complex-specific directories

Arguments to `autoreset` arguments include `<complex_name> on` and `<complex_name> off` or `chk`.

The output of the `chk` option for a complex name of `hw2a` looks like:

```
Autoreset for hw2a is enabled.
```

or

```
Autoreset for hw2a is disabled.
```

NOTE

`autoreset` determines the behavior of `ccmd` when it encounters an error condition. `ccmd` makes its decision whether to reset a complex immediately after running `hard_logger`. Enabling `autoreset` after `hard_logger` has run does not reset the complex.

est_config

`est_config` is a utility that builds the node and complex descriptions used by `est`. `est_config` reads support files at `/spp/data/DB_RING_FILE`, reads the electronic board identifier (COP chip) and scans to completely describe the node or complex. It also uses the hardware database created by `ccmd`. The data retrieved is organized and sorted into an appropriate node configuration file in the `/spp/data/<complex name>` directory.

An optional configuration directory can be specified using the `-p` argument. `est_config` works across all nodes unless a specific node or complex is requested with the `-n` option.

NOTE

If there is a `node_#.pwr` file that is older than the `node_#.cfg` file, existing node configuration files do not need to be updated.

report_cfg

This utility generates a report summarizing the configuration of all nodes/complexes specified on the command line. The format of `report_cfg` is as follows:

```
report_cfg [<node id|complex> [<node id|complex> ...]]
```

`node id` may be a node number, IP name, or "all." If no node ID is specified, the utility defaults to all nodes in the current complex.

One or more of the options in Table 12 must be specified:

Table 12

report_cfg options

-d	Show all details
-s	Show summary only
-a	Show summary and details (same as -d and -s)
-A	Show ASIC detail
-i	Show I/O detail
-m	Show memory detail
-p	Show processor detail

If the `report_cfg` tool detects any nodes of complexes that contain SCA DIMMS and some memory boards that are not populated with STACS, it generates a report.

Example configuration report:

The system inventory has determined that you'll need to order 8 SCA Upgrade Kits in order to connect this cabinet with other SCA cabinets. These upgrade kits are available by additionally ordering opt. 010 of the required SCA HyperLink product (A5518A or A5519A). You may also have to order additional memory DIMMS, memory boards and or processor boards to meet the minimum requirements for a SCA configuration. Refer to the HP 9000 V-Class Ordering Guide for details.

Effects of hardware and software deconfiguration

`report_cfg` counts all processors, STACs, SMACs, SAGAs and ERACs if POST has not marked them as empty. This results in ASICs and processors being included in the summary count even though they may have failed or have been deconfigured by software. This is necessary because POST deconfigures STACs in a single node configuration. To allow the tool to count these ASICs, it must report all ASICs that are installed, not just those enabled by POST.

`report_cfg` includes all DIMMs that POST has not marked as empty. If the user deconfigures a SMAC with software or the SMAC fails the POST selftest, POST marks the DIMMs on that SMAC as empty. If POST has written valid size information into the BCM for a DIMM, `report_cfg` reports the physical size reported by POST.

For example, if a node has both 80- and 88-bit DIMMs, POST reconfigures the 88-bit DIMMs to behave as 80-bit DIMMs, and the system logically behaves as if it has all 80-bit DIMMs. `report_cfg`, however, distinguishes (using the physical attribute in the BCM) between the 80- and 88-bit DIMMs in its reports.

Another example would be a system that contains 16 GBytes of memory but half of the DIMMs are deconfigured by software. `report_cfg` still reports that the system contains 16 Gbytes of memory.

report_cfg summary report

To obtain a system summary report, use the `-s` option for the command. The following is a sample summary report by `report_cfg`:

`report_cfg -s`

Complex name	Complex serial number	Node ID
-----	-----	-----
hw4a	USR1234567	0
hw4a	USR1234567	2
Cabinets:	2	
Processors:	30	
Processor boards:	30 (30 singles, 0 duals)	
Memory boards:	16	
TAC chips:	16	
Enabled Memory (Mb):	8192	
88-bit 128Mbyte:	112	

report_cfg ASIC report

To obtain a report on the ASICs in a complex, use the `-A` option. The following is a sample ASIC report by `report_cfg`:

report_cfg -A

Complex	Node#	MIB COP		SCUB COP
hw2a	0	A5074-60002 00 a 3845	A5074-60003 00	b 3830
hw2a	2	A5074-60002 00 a 3840	A5074-60003 00	b 00XA

Complex	Node	Slot	+----- ASIC revisions -----+			
			PAC	MAC	TAC	RAC
hw2a	0	0	2	2	1	2
hw2a	0	1	2	2	1	2
hw2a	0	2	2	2	1	2
hw2a	0	3	2	2	1	2
hw2a	0	4	2	2	1	
hw2a	0	5	2	2	1	
hw2a	0	6	2	2	1	
hw2a	0	7	2	2	1	
hw2a	2	0	2	2	1	2
hw2a	2	1	2	2	1	2
hw2a	2	2	2	2	1	2
hw2a	2	3	2	2	1	2
hw2a	2	4	2	2	1	
hw2a	2	5	2	2	1	
hw2a	2	6	2	2	1	
hw2a	2	7	2	2	1	

report_cfg I/O report

To obtain an I/O report, use the `-i` option. The following is a sample I/O report by `report_cfg`:

report_cfg -i

Complex	Node#	MIB COP		SCUB COP
hw2a	0	A5074-60002 00 a 3845	A5074-60003 00	b 3830
hw2a	2	A5074-60002 00 a 3840	A5074-60003 00	b 00XA

Complex	Node#	I/O board		COP
hw2a	0	IORF_B	A5080-60001 00	a 3821
hw2a	0	IORF_A	A5080-60001 00	a 3821
hw2a	2	IOLF_B	A5080-60001 00	a 3821
hw2a	2	IOLF_A	A5080-60001 00	a 3821

report_cfg memory report

To obtain a report on the memory in a complex, use the `-m` option. The following is a sample memory report by `report_cfg`:

report_cfg -m

Complex	Node#	MIB COP		SCUB COP
hw2a	0	A5074-60002 00 a 3845	A5074-60003 00 b 3830	
hw2a	2	A5074-60002 00 a 3840	A5074-60003 00 b 00XA	

Complex	Node	Mem. Board	COP	80-bit			88-bit		
				3	1	2	3	1	2
hw4a	0	MB0L A5078-60003	01 a 00XB	2	8	6	2	8	6
hw4a	0	MB1L A5078-60003	01 a 00XB						
hw4a	0	MB2R A5078-60003	01 a 00X2						
hw4a	0	MB3R A5078-60003	01 a 00X2						
hw4a	0	MB4L A5078-60003	01 a 00XA						
hw4a	0	MB5L A5078-60003	01 a 00X2						
hw4a	0	MB6R A5078-60003	01 a 00XA						
hw4a	0	MB7R A5078-60003	01 a 00X2						
hw4a	2	MB0L A5078-60003	01 a 00XA						
hw4a	2	MB1L A5078-60003	01 a 3842						
hw4a	2	MB2R A5078-60003	01 a 00XA						
hw4a	2	MB3R A5078-60003	01 a 00XB						
hw4a	2	MB4L A5078-60003	01 a 3843						
hw4a	2	MB5L A5078-60003	01 a 00XD						
hw4a	2	MB6R A5078-60003	01 a 00XB						
hw4a	2	MB7R A5078-60003	01 a 00XD						

report_cfg processor report

To obtain a report on the processor in a complex, use the `-p` option. The following is a sample processor report by `report_cfg`:

report_cfg -p

Complex	Node#	MIB COP	SCUB COP
hw2a	0	A5074-60002 00 a 3845	A5074-60003 00 b 3830
hw2a	2	A5074-60002 00 a 3840	A5074-60003 00 b 00XA

Complex	Node#	Processor	COP	CPU rev
hw2a	0	PB0L_A	A5077-60005 00 a 00XA	2.0
hw2a	0	PB1R_A	A5492-60001 00 b 00XC	2.3
hw2a	0	PB1L_A	A5491-60001 00 a 00XA	2.0
hw2a	0	PB4L_A	A5492-60001 00 b 00XB	3.0
hw2a	0	PB5L_A	A5492-60001 00 a 00XB	2.0
hw2a	0	PB0L_B	A5077-60005 00 a 00XA	2.0
hw2a	0	PB1L_B	A5491-60001 00 a 00XA	2.0
hw2a	0	PB4L_B	A5492-60001 00 b 00XB	2.0
hw2a	0	PB5L_B	A5492-60001 00 a 00XB	2.0

Complex	Node#	Processor	COP	CPU rev
hw2a	2	PB2L_A	A5491-60001 00 a 00XA	2.0
hw2a	2	PB2R_A	A5492-60001 00 a 00XA	2.0
hw2a	2	PB3R_A	A5491-60001 00 a 00XA	2.0
hw2a	2	PB4L_A	A5492-60001 00 b 00XC	2.0
hw2a	2	PB5L_A	A5491-60001 00 a 00XA	2.0
hw2a	2	PB4L_B	A5492-60001 00 b 00XC	2.3

xsecure

`xsecure` is an application that helps make a V2500/V2600 class SSP secure from external sources. This tool disables modem and LAN activity to provide an extra layer of security for the V2500/V2600 system.

`xsecure` may be run as a command line tool or an windows-based application.

In secure mode, all network LANs other than the `tsdart` bus are disabled and the optional modem on the second serial port will be disabled. When in normal mode all networks and modems are re-enabled.

If the command line [-on | -off | -check] options are used, `xsecure` does not use the GUI interface. These options allow the user to turn the secure mode on, off or allow the user to check the secure mode status.

A simple button with a red or green secure mode indicator provides the user with secure mode status information. The red indicator shows that the secure mode process has begun. The label near the red button will inform the user when the SSP is secure. A green indicator and the appropriate label shows that the network is available and the SSP may be accessed through the ethernet port.

In order for `xsecure` to work properly the SSP, console cables, terminal mux and modems must be configured in specific ways. The SSP JTAG connections, OBP connections and an optional terminal mux must all be connected to the Diagnostic LAN and identified in the `/etc/hosts` file as `tsdart-d`. The `sppconsole` serial cable must be connected to serial port 1 and to node 0. An optional modem may be connected to serial port 2.

Different versions of the HP-UX operating system run on a V-Class server and its Service Support Processor. This section covers issues related to using HP-UX V11.0 and HP-UX V11.10 on V-Class servers. Multiple-cabinet server configurations and HP-UX SCA features require that HP-UX V11.10 be installed. Topics covered in this section include:

- HP-UX on the V2500/V2600
- Starting HP-UX
- Stopping HP-UX

HP-UX on the V2500/V2600

In general HP-UX administration tasks are performed on V-Class servers as they are on other HP servers. One difference is that V-Class servers run the HP-UX kernel only in 64-bit mode. This facilitates addressing the larger memory capacity available on the V2500/V2600. However, both 32-bit and 64-bit applications may be run simultaneously on the server.

Displaying System Information

This section gives brief information on the `model`, `top`, and `ioscan` HP-UX commands. More HP-UX information is available from the following Web site.

<http://docs.hp.com/>

The `model` command prints the hardware series and model for the machine it is issued on. On V2500/V2600 servers it displays information as shown in the command example below.

```
# model
9000/800/V2500
```

For details, see the *model(1)* man pages.

The `top` command displays information on the top processes (based on CPU use) on the system, and lists CPU utilization data for the system's processors. Because V-Class servers can have many processors, you may want to issue `top -h` when using this command. The `-h` option suppresses printing individual lines of processor information, instead printing only a one-line system average. This makes room on the screen for showing the active processes running on the system. For more information, see the *top(1)* man pages.

Listing the Server Hardware Configuration

The `ioscan` utility scans the system's hardware and lists all hardware found, including processors, memory, I/O devices, and interface cards. On V-Class servers, `ioscan` also reports PCI controllers and the core utility controller and its connections to the Service Support Processor. For more details, see the *ioscan (1M)* man pages.

On multiple-cabinet V2500/V2600 servers, the first component of the hardware path indicates which cabinet a hardware component resides upon.

Hardware on cabinet ID 0 is listed with a first hardware path field starting at 0, hardware on cabinet ID 2 is listed starting at 64, cabinet ID 4 starts at 128, and cabinet ID 6 starts at 192.

For example, a disk on cabinet ID 0 could have a hardware path of:

1/2/0.9.0

The above disk is on cabinet ID 0, is connected to PCI bus 1, slot 2, and has a target (SCSI ID) of 9. A disk at the same relative location but on cabinet ID 2 has a hardware path of:

65/2/0.9.0

The above disk is on cabinet ID 2, and is connected to PCI bus 65, slot 2, and has a target (SCSI ID) of 9. This corresponds to the top left PCI card cage of cabinet ID 2, as shown in Table 13.

The following table lists the hardware path numbering for key V2500/V2600 system components as they are numbered on the various cabinets in a multiple-cabinet V2500/V2600 server.

Table 13

Hardware Path Numbering for V2500/V2600 Cabinets

V2500/V2600 Cabinet ID	First Field of Hardware Path	Description of Hardware Component
Cabinet ID 0	0–7	PCI I/O bus bridges (card cages)
	8	Memory
	15	Core utilities board
	16–47	Processors (PA-RISC CPUs)
Cabinet ID 2	64–71	PCI I/O bus bridges (card cages)
	72	Memory
	79	Core utilities board
	80–111	Processors (PA-RISC CPUs)

V2500/V2600 Cabinet ID	First Field of Hardware Path	Description of Hardware Component
Cabinet ID 4	128–135	PCI I/O bus bridges (card cages)
	136	Memory
	143	Core utilities board
	144–175	Processors (PA-RISC CPUs)
Cabinet ID 6	192–199	PCI I/O bus bridges (card cages)
	200	Memory
	207	Core utilities board
	208–239	Processors (PA-RISC CPUs)

Configuring HP-UX for V-Class Servers

HP-UX V11.0 provides several tuned parameter sets that are useful on HP V-Class servers. Configuring HP-UX with these parameter sets can improve your V-Class server's performance when using it for different purposes, such as scientific, data processing, or mixed interactive use.

Using the SAM utility (`/usr/sbin/sam`), you can configure an HP-UX kernel for HP V-Class servers. To do so, select Kernel Configuration, then the Configurable Parameters subarea, and apply the tuned parameter set for your type of server use via the Actions menu.

HP-UX parameter sets

HP-UX kernel configurations are provided for the following types of V-Class server use:

- Scientific and technical use—Servers running applications that have very large data sets and may have long processing times. Examples include NASTRAN, Abaqus, mechanical and electrical design applications, and fluid dynamics applications.

The “V-Class Technical Server” tuned parameter set provides HP-UX kernel parameter settings for running such workloads on HP V-Class servers.

- **Dedicated commercial data processing use**—Servers whose use is restricted for online transaction processing (OLTP), running Oracle, and running other data processing workloads. These systems provide limited, if any, interactive user access.

The “OLTP/Database Server System” tuned parameter set provides a good HP-UX configuration for using HP V-Class servers for dedicated commercial data processing.

- **Mixed interactive and data processing use**—Servers used for interactive user log-ins, and for running OLTP/data processing workloads and miscellaneous other applications.

The “OLTP/Database Monolithic System” tuned parameter set is appropriate for mixed-use V-Class servers.

Multiple-cabinet kernel configurations

The following are notable initial kernel parameter settings for multiple-cabinet V2500/V2600 servers. These settings are provided when installing HP-UX 11.10 on multiple-cabinet servers.

- **maxuprc**—Maximum number of user processes.
Initial SCA value: 256
- **maxusers**—The MAXUSERS value, used in various kernel formulae.
Initial SCA value: 256
- **max_thread_proc**—Maximum number of threads allows in each process.
Initial SCA value: 500
- **maxswapchunks**—Maximum number of swap chunks. The size of each swap chunk is defined by the swchunk kernel parameter, which specifies the number of 1 KByte blocks.
Initial SCA value: 5000

NOTE

Note that the maxswapchunks parameter setting (and the swchunk setting) may need to be adjusted based on your system’s swap space and crash dump needs.

These tuned parameter sets are available from the SAM utility’s Configurable Parameters subarea, as described above, and are stored as files in the following directory.

`/usr/sam/lib/kc/tuned`

Refer to the SAM online help for examples and details on using kernel parameters.

Process and Thread “Gang Scheduling”

HP-UX V11.0 includes support for kernel threads and provides a “gang scheduling” feature for managing how threads belonging to the same process or application are executed.

The HP-UX gang scheduler permits a set of MPI processes, or multiple threads from a single process, to be scheduled concurrently as a group.

Gang scheduling is enabled and disabled by setting the `MP_GANG` environment variable to `ON` or `OFF`.

The gang scheduling feature can significantly improve parallel application performance in loaded timeshare environments that are oversubscribed. Oversubscription occurs when the total number of runnable parallel threads, runnable MPI processes, and other runnable processes exceeds the number of processors in the system.

Gang scheduling also permits low-latency interactions among threads in shared-memory parallel applications.

Only applications using the HP-UX V11.0 MPI or pthread libraries can be gang scheduled. Because HP compiler parallelism is primarily built on the pthread library, programs compiled with HP compilers can benefit from gang scheduling.

For more details, refer to the *gang_sched(7)* man page.

HP-UX 11.10 SCA Enhancements

Hewlett-Packard’s Scalable Computing Architecture (SCA) design allows for multiple resource localities (or “locality domains”) to be combined to form a single system running a single instance of the HP-UX operating system. One example of an SCA system is a multiple-cabinet HP V2500/V2600 server where each cabinet comprises a locality domain.

HP-UX V11.10 provides SCA enhancements. These include a utility (`mpsched`) as well as system calls and library routine extensions for controlling how processes and threads are launched, where they are placed for execution, whether they may migrate (move) during execution, and how they use memory. These SCA interfaces and programming

extensions also provide system inquiry features for retrieving information about the current hardware topology, as well as thread and process inquiry features. Both traditional system architectures as well as SCA systems are supported by the HP-UX 11.10 enhancements.

HP-UX SCA Features

HP-UX V11.10 SCA programming and launch features provide the following capabilities.

- **Inquiry Features**—Provide system hardware topology information, thread and process binding information, and thread and process launch policy information.

You can print or retrieve information about the system hardware configuration using the `mpsched` utility, the `mpctl` system call, and `pthread` library features. These interfaces also allow for retrieving information about thread and process bindings and launch policies.

- **Targeting and Binding**—Provide locality domain, processor, and memory targeting and binding.

You can specify the locality domain or processor on which a process runs by using the `mpsched` command, the `mpctl` system call, and the `pthread` interfaces `pthread_processor_bind_np` and `pthread_ldom_bind_np`.

You also can specify the locality from which a process' memory is allocated via the `mmap` and `shmget` system calls.

- **Launch and Scheduling**—Provide launch policies for threads and processes and support for “gang scheduling” of threads and processes.

You can set the launch policy for a process with the `mpsched` utility, the `mpctl` system call, and the `pthread` interface `pthread_launch_policy_np`. The launch policies supported in HP-UX 11.10 include the following.

- **None**—The system tries to launch threads and processes in a manner that provides the best performance, though applications should not rely on this providing any specific launch behavior. (This is the default policy for launching threads and processes.)
- **Round Robin**—Alternate among all localities until all localities have been selected once, then start over as needed.

- **Fill First**—Fill a locality first, then spill over to another locality, as needed. Once all localities are filled, start over as needed.
- **Packed**—Place all threads or processes in the same locality; do not spill over.
- **Least Loaded**—Place each thread or process in the locality that is least-loaded at the time of its creation.

Gang scheduling of threads and processes is supported through the `mpsched -g` option and the `MP_GANG` HP-UX environment variable. Details are provided in the *gang_sched(7)* man page and in the section “Process and Thread “Gang Scheduling”” on page 122.

More details about HP-UX programming, scheduling, and launch enhancements are available in the HP-UX V11.10 online man pages and in the *HP-UX SCA Programming and Process Management White Paper*.

Starting HP-UX

Bringing the 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 *Managing Systems and Workgroups*. Additional information is contained in the *V2500/V2600 SCA HP-UX System Guide*. This section describes:

A V-Class server consists of two systems:

- System Support Processor (SSP)
 - Boots the V-Class server
 - Monitors the V-Class server for hardware errors
 - Debugs a hung system
 - Runs HP-UX
- V-Class cabinet
 - Hosts OpenBoot PROM (OBP) software
 - Runs HP-UX

The main firmware interface, the OBP boot menu, provides a straightforward interface for managing a system before HP-UX boots. The OBP menu is available through the V-Class console interface.

In multiple-cabinet V2500/V2600 SCA systems, each V2500/V2600 cabinet runs its own copy of the firmware. Normally during booting you only need to interact with the OBP menu on cabinet ID 0, which serves as the “monarch” or master cabinet. The other cabinets (IDs 2, 4, and 6, if present) run the same set and version of firmware, but are considered to be “serf” cabinets under control of the monarch cabinet once OBP starts.

IMPORTANT

Only I/O devices on cabinet ID 0 may be used as the boot device for a V-Class server. Likewise, only logical volumes that reside entirely on cabinet ID 0 may be used as the boot device.

Start up, or boot, HP-UX after the operating system has been completely shut down or partially shut down to perform system administration tasks. The boot procedure differs according to the value of the Autoboot flag. See “Enabling Autoboot” on page 67 for information on how to set Autoboot. After you power-up your V-Class server, if Autoboot is set to:

- ON, OBP automatically starts HP-UX. (Press the **ESC** key within 10 seconds to interrupt the boot process and enter bootmenu commands.)
- OFF, the user must:
 - Start OBP at the SSP’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`

Power-On Sequence

Turn on the SSP power and allow it to boot before powering on the V-Class server’s cabinets. This allows the SSP to be used to monitor and control the V-Class server as it boots is used.

The following is the sequence for powering on an HP V-Class server and its SSP:

Step 1. Power on the SSP and allow it to boot HP-UX.

Step 2. Power on the V-Class server cabinets.

On multiple-cabinet V2500/V2600 SCA systems, the serf cabinets (cabinet IDs 2, 4, and 6, if present) should be powered on before the monarch cabinet (cabinet ID 0). This helps ensure that all cabinets synchronize following power on self test.

Step 3. Select a device on V-Class cabinet ID 0 from which to boot HP-UX, as needed.

The OBP menu’s SEARCH command searches for bootable devices connected on the cabinet. Only devices on cabinet ID 0 may be booted.

You can designate an I/O device to be the primary (PRI) or alternate (ALT) boot device by using the OBP menu’s PATH command.

Step 4. Issue the OBP menu's `BOOT` command to boot HP-UX on the V-Class server.

You can set the server to automatically boot HP-UX if you have also set a primary boot device (`PRI`). The OBP menu provides the `AUTO BOOT` option, which causes the server to automatically boot HP-UX from the primary boot device when `AUTO BOOT` is set to `ON`.

If a cabinet is already powered on before the SSP booted, the cabinet can be reset from the SSP after it boots, using the `do_reset` command.

Boot variables

Several variables affecting the boot process can be set using the HP boot (OBP) menu. For example, you can set the server to automatically proceed to boot HP-UX by setting the `AUTO BOOT` option to `ON` and setting the `PRI` boot path variable.

Boot variable settings are stored in non-volatile memory (NVRAM) residing on the utility board of each V-Class server cabinet. These variables are stored permanently until they are modified using the HP boot menu interface. Refer to Chapter 4, “Firmware (OBP and PDC)” for more information.

On multiple-cabinet V-Class servers the boot variable settings for the monarch cabinet (cabinet ID 0) determine the server's boot-time behavior.

You should ordinarily need to interact with only the monarch cabinet's OBP menu, although the `RC` (`RemoteCommand`) option is available for interacting with the OBP menus on other cabinets. See also “Assuming control of the console” on page 51.

Table 14 **Boot variables**

Variable	Description
AUTO BOOT [ON OFF]	If set to ON, the server automatically boots HP-UX from the primary (PRI) device during system startup or reset. When set to OFF the server boots to the OBP menu interface.
AUTO SEArch [ON OFF]	If set to ON, the server searches for and lists all bootable I/O devices.
AUTO Force [ON OFF]	If set to ON, then OBP allows HP-UX to boot even if one or more cabinets does not complete power on self test. When set to OFF, all cabinets must successfully pass power on self test for OBP to permit the server to boot HP-UX.
BootTimer [<i>seconds</i>]	Sets the number of seconds for the system to wait before booting. This is used to allow external mass storage devices to come online.
PAth [PRI ALT CON][<i>path</i>]	Displays and sets primary, alternate, and console hardware paths. Keyboard path is displayed but cannot be modified.
SECure [ON OFF]	If set to ON, enables secure boot mode. If set, the boot process cannot be interrupted. Only useful if autoboot is ON; the system will autosearch and autoboot.

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 the following message is displayed:

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

2. Improper file system shut down

The start-up process is interrupted:

```
/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 the user to run the `fsck` command manually. If `fsck` needs to be run 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 the problem will have to be solved before proceeding.

Stopping HP-UX

This section provides a brief overview of the process; for complete instructions, see *Managing Systems and Workgroups*. Additional information is contained in the *V2500/V2600 SCA HP-UX System Guide*.

Typically, the system is shut down to:

- Put it in single-user state so that the system can be updated or to check file 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 the file system. Use the `shutdown` or `reboot` commands.

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 a network service is used. The shutdown process logs the user out prematurely and returns control to the console. Run `shutdown` from the `sppconsole` window on the SSP.

See the shutdown man page for a complete description of the shutdown process and available options.

Rebooting the system

To shutdown HP-UX and reboot the V-Class server, perform the following steps:

- Step 1.** If the server is running HP-UX, log in to the server as root.
- Step 2.** Check activity on the server and warn users of the impending server reboot.

If HP-UX is hung, to reboot HP-UX you may need to reset the server. See “Resetting the V2500/V2600 server hardware” on page 134.

- Step 3.** Change to the root directory. Enter:

```
cd /
```

- Step 4.** Shut down the system using the `shutdown` or `reboot` command. Enter:

```
shutdown -r
```

or

```
reboot
```

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

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

See the `shutdown` and `reboot` man pages for complete descriptions of the commands and available options.

Shutting down the system

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

Step 1. Login to the server as root.

Step 2. Check activity on the server and warn users of the impending server shutdown.

Step 3. Change to the root directory. Enter:

```
cd /
```

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

```
shutdown
```

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

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

Step 5. Shut down and halt HP-UX using the `shutdown` or `reboot` command. Enter:

```
shutdown -h
```

or

```
reboot -h
```

Progress messages detailing system shutdown activities print to the terminal.

CAUTION

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

See the `shutdown` and `reboot` man pages for complete descriptions of the commands and available options.

Resetting the V2500/V2600 server hardware

The `/spp/bin/do_reset` command resets the V-Class hardware. The `do_reset` command is run from the Service Support Processor, causes OBP to reboot, and halts all activity on the V-Class server cabinets involved. For details see the `do_reset` man page on the Service Support Processor.

NOTE

Before running `do_reset` from the Service Support Processor you should shut down HP-UX running on the V-Class server, if possible, to avoid losing data.

Four levels of system reset are provided by `do_reset`, from level 1 (the default) to level 4 (Transfer of Control).

On multiple-cabinet V-Class server configurations, you can reset all cabinets or just selected cabinets with `do_reset`. Larger systems such as those with additional cabinets take more time to reset.

The `do_reset` command's syntax is as follows:

```
do_reset [node_id | all] [level] [boot_option]
```

If `do_reset` is specified with no arguments then the default level 1 reset of all cabinets is performed, rebooting the server to OBP.

Either a cabinet ID (*node_id*) or the keyword `all` can be specified to indicate which cabinets are to be reset. If a cabinet ID is specified, the reset level also must be specified.

Reset levels (*level*) of 1 through 4 are supported and are specified as numbers. Level 1, the default, provides a hard reset and reboots to OBP. A level 4 reset provides a Transfer of Control (TOC), equivalent to pressing the TOC button on a V-Class cabinets.

NOTE

Only one request for a level 4 reset (TOC) should be issued at a time. This ensures that the server properly completes the reset process.

Various *boot_options* are supported for level 1 resets. Normally only trained HP service personnel use these.

Examples

- `do_reset`
This performs a level 1 reset of all cabinets, resetting the server and rebooting to OBP.
- `do_reset all 4`

Performs a level 4 reset of all cabinets. This causes a Transfer of Control (TOC) that initiates a crash dump of the operating system, if crash dump is configured.

See the *savecrash(1M)* man page for crash dump details.

To reset the V-Class server hardware, perform the following steps:

Step 1. Shut down HP-UX on the V2500/V2600 server.

This involves logging in to the server and issuing the `shutdown -h` or `reboot -h` command. For details see the procedure “Shutting down the system” on page 133.

If the V-Class server is hung and you can not log in and shut down HP-UX, you can proceed with Step Two and may want to perform a level 4 reset at Step Three.

Step 2. Access a Service Support Processor login shell.

You can do this directly at the Service Support Processor workstation, or by remotely logging in with a `telnet` or `rlogin` session.

Step 3. Issue the `do_reset` command from a Service Support Processor login shell.

By default, `do_reset` performs a level 1 reset of all cabinets. See the *do_reset(1)* man page on the Service Support Processor for details.

A level 4 reset (`do_reset all 4`) performs a Transfer of Control (TOC) that resets the server and initiates the HP-UX crash dump process for it, if configured.

Pressing the TOC button on a V-Class server cabinet has the same effect as a level 4 `do_reset`.

Issue only one level 4 reset to ensure a proper TOC.

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 SSP automatically records information on environmental and system level events in several log files. See "SSP file system" on page 54 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 28 for more information.
- Step 3.** Record any relevant information contained in the log files in the `/spp/data/complex` directory on the SSP:
 - `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 SSP, 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 the Hewlett-Packard Customer Response Center.

HP-UX kernel configuration can affect performance. Refer to “Configuring HP-UX for V-Class Servers” on page 120. For more detailed information refer to *HP-UX 11.0 Configurable Kernel Parameters* and *HP V-Class Server HP-UX Configuration Notes* available at the following web site:

<http://docs.hp.com/hpux/os>

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 SSP, use one or both of the following utilities to communicate with the server:
 - ping
 - telnet
 - rlogin
- See the `ping`, `telnet`, and `rlogin` 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 SSP. See “Resetting the V2500/V2600 server hardware” on page 134.
 - Step 7.** Save the core dump file and contact the HP Response Center to have the core dump file analyzed. Refer to the service contract for the phone number of the Hewlett-Packard Response Center. See “Fast dump” on page 147 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 the system so that data can be recovered in the event of a disk head crash or similar situation. How frequently the backups are updated depends on how much data one can afford to be 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 /tmp/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 the problem to be classified 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.

WARNING

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 `ksh-shell` window. For more information about rebooting the system see “Rebooting the system” on page 146.

Step 6. If the problem reappears on the device there may be an interface card or 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 28 for more information.

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

- `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 ksh-shell window. For more information about rebooting the system see “Rebooting the system” on page 146.

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 the file systems, reboot the system and run the file system checker `fsck` to check and correct the problem. Follow all directions that `fsck` displays. Especially when the 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 146.

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 *Managing Systems and Workgroups*.

Recovery from other situations

When a problem appears with something other than that has previously been discussed or the problem can not be classified, proceed to “Rebooting the system” on page 146. Ensure that the exact text of the panic message is recorded for future troubleshooting purposes. See “Collecting information” on page 138 for further information.

Rebooting the system

Once a problem has been corrected, reset and reboot the system.

- Step 1.** Reset the V-Class server. See “Resetting the V2500/V2600 server hardware” on page 134.
- 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 *Managing Systems and Workgroups* for further instructions. If the problems were associated with the root file system, `fsck` will ask the operator to reboot the 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, this will not 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 it 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, it may be necessary 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. This may require the services of the Hewlett-Packard Response Center.

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, the operator 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 the operator 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. `libcrash(3)` 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 the operator next reboots 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.

IMPORTANT

Crashdump must be configured to dump on cabinet zero disks only.

It is important to have sufficient space to capture the part of memory that contains the instruction or data that caused the crash. More than one dump device can be defined 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 the 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).

New SCA-Extended Crash Dump Format

A new crash dump format is provided to support dumping selected memory on multiple-cabinet V2500/V2600 servers. This new SCA-extended crash dump format is used only on V2500/V2600 SCA systems and requires the crash dump utilities provided in the HP-UX 11.10 release.

Non-SCA systems, including single-cabinet V2500/V2600 servers and all other HP systems, use the non-SCA crash dump format. Unlike the SCA-extended crash dump format, the non-SCA crash dump format is backward compatible and does not require HP-UX 11.10 crash dump utilities.

For more information see the *Release Notes for HP-UX 11.10 SCA*.

Memory Dumped on V2500/V2600 SCA Servers

Crash dump on multiple-cabinet V2500/V2600 servers dumps memory from all cabinets, including the following:

All cabinet ID 0 memory, except memory dedicated as CTI cache.

All cabinet-private memory on serf cabinets (cabinet IDs 2, 4, and 6).

Crash dump does not dump memory used as CTI cache. The CTI cache is physical memory dedicated as a cache for internode communication, and is not made available to HP-UX or applications.

For V2500/V2600 SCA servers, the crash dump volume must reside entirely on cabinet ID 0.

To calculate an appropriate size for a V2500/V2600 SCA crash dump volume, estimate that you will need at most the following amount of space: the total amount of physical memory in the system, plus space to allow for dump headers and tables, minus the amount of memory dedicated as CTI cache, and minus the amount of memory that is kernel text replicated across cabinets.

The total amount of space required depends on whether a full dump or a selective dump is performed, and whether the crash dump is compressed or uncompressed.

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 the operator 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, copied 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, the operator can list which classes of memory *must always get dumped*, and which classes of memory *should not be dumped*. If both of these “lists” are left empty, HP-UX decides 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 a full dump has not been defined (in the kernel or at run time), the definitions can be overridden (within a ten second window) and a request for a full dump after a system crash can be performed. Likewise, if the cause of the crash is known, a request to not dump can be performed as well.

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 there is enough disk space and the fastest system recovery is required, 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, 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, only those pages that are endangered by paging activity can be saved.

Pages residing on the dedicated dump devices can remain there. It is possible to analyze memory dumps directly from the dedicated dump devices using a debugger that supports this feature. If, however, there is a need to send the memory dump to someone else for analysis, move the pages on the dedicated dump devices to the HP-UX file system area. Then 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 the part of memory that contains the instruction or piece of data that caused the crash is captured. The factors that must be considered 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 the system; depending on a number of factors, additional disk space greater than the amount of physical memory in the system may be needed

Dump definitions built into the kernel vs. defined at runtime

There are three places to 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 the dump devices for this system are defined with a total of 256-Mbytes of space in the kernel file and an additional 768-Mbytes of disk space in the `/etc/fstab` file, there would be 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 the 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 155.

Compressed save vs. noncompressed save

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

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, there is a 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), the operator can elect to have `savecrash` (the boot time utility that does the copy) compress the 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 to 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:

Recovering from failures

Abnormal system shutdowns

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:

$(6208 \times 4 \text{ Kbytes}) \times 1.25 = \text{approx. } 30 \text{ Mbytes}$

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, a new kernel must be built and the system rebooted 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

The `system` file can be edited and the `config` program used 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 the `system` file (the file that `config` uses to build the new kernel). This is usually the file `/stand/system` but it can be another file if that is preferred.

Dump to Hardware Device—For each hardware dump device to be configured 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 are required), 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, the above statement (**dump none**) must be used.

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 the `system` file has been edited, 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 there is not a concern about capturing a dump that occurs during the earliest stages of the boot process, 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

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 to be used as a dump device.

NOTE

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

The `crashconf` command

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, 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. Interrupt the dump at any time by pressing the **ESC** key. However, if a dump is interrupted, 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 67.

savecrash processing

During the boot process, a process called `savecrash` can be used 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. 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.

WARNING

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 either package and send it in for analysis or analyze it using a debugger.

NOTE

With HP-UX 11.0, it is possible to analyze a crash dump directly from dump devices. If, however, it needs to be saved to a tape or sent to someone, first copy the memory image to the HP-UX file system area.

Unless specifically disabled during reboot, the `savecrash` utility copies the memory image during the reboot process. The default HP-UX directory for the memory image in is `/var/adm/crash`. 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 devices are being used for both paging (swapping) and dumping, it is very important to not disable `savecrash` processing at boot time. If this is done, 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), copying the memory image to the HP-UX file system area can be delayed 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 copying the physical memory image from the dump devices to the HP-UX file system area has been delayed, run `savecrash` manually to do the copy when the system is running. Confirm that enough space to hold the copy in the HP-UX file system area has been configured before doing so.

If a partial save is being done, the only pages copied to the 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 the 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

Use `crashutil` to convert the file format when analyzing a crash dump on a computer running a different version of HP-UX than the V-Class server, or if the debugging tool does not recognize the specific format of the saved file.

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

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

<i>version</i>	Designates the version of the destination format.
<i>source</i>	Designates the pathname of the crash dump to be converted.

Recovering from failures

Abnormal system shutdowns

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. Contact the Hewlett-Packard representative for help in analyzing a crash dump.

A

LED codes

This appendix describes core utilities board (CUB) LED errors. The Attention LED on the core utilities board (CUB) turns on, and the Attention light bar on the front of the node flashes to indicate the presence of an error code listed in Table 15. Additionally, only the highest priority error is displayed. Once remedied, an error that is cleared may expose a lesser priority error.

Power on detected errors

This section describes core utilities board (CUB) LED errors from highest to lowest priority detected at power on. The Attention LED on the core utilities board (CUB) turns on, and the Attention light bar on the front of the node flashes to indicate the presence of an error code listed in Table 15. Additionally, only the highest priority error is displayed. Once remedied, an error that is cleared may expose a lesser priority error.

Errors are listed in sequence from the highest to lowest priority.

NOTE

Errors from LED hex code 00 through hex code 67 shut the system down, and errors from hex-code 68 through 73 leave the system up.

Table 15 CUB detects power on error

LED	Fault	Symptoms	Corrective action
00	3.3V error (highest priority)	<ol style="list-style-type: none">1. 5V is up 3.3V is not.2. SSP interface will not function.	Call the Response Center.
01	ASIC Install 0 (MIB)	<ol style="list-style-type: none">1. Incorrect rotation or part in one of the processor agent chip (PAC) sockets.2. Incorrect rotation or part in one of the routing (XBAR) attachment chip (RAC) sockets.	Call the Response Center.
02	ASIC Install 1 (EMB)	<ol style="list-style-type: none">1. Incorrect rotation or part in one of the memory access chip (MAC) sockets.2. Incorrect rotation or part in one of the toroidal access chip (TAC) on memory board (MB).	Call the Response Center.

LED	Fault	Symptoms	Corrective action
03	FPGA not OK	<ol style="list-style-type: none"> 1. Core Utilities Board (CUB) monitoring utilities chip (MUC) problem. 2. MUC cannot get correct program transfer from EEPROM on power up. 	<ul style="list-style-type: none"> • Cycle the node power using the Key switch. • Call the Response Center.
04	dc OK error (Upper Left)	<ol style="list-style-type: none"> 1. Power supply is reporting failure (dc OK) after keyswitch is turned on, but prior to CUB power on sequence. 2. This is the first of two or more supplies reporting failure. 	Call the Response Center.
05	dc OK error (Upper Right)	<ol style="list-style-type: none"> 1. Power supply is reporting failure (dc OK) after keyswitch is turned on, but prior to CUB power on sequence. 2. This is the first of two or more supplies reporting failure. 	Call the Response Center.
06	dc OK error (Lower Left)	<ol style="list-style-type: none"> 1. Power supply is reporting failure (dc OK) after keyswitch is turned on, but prior to CUB power on sequence. 2. This is the first of two or more supplies reporting failure. 	Call the Response Center.
07	dc OK error (Lower Right)	<ol style="list-style-type: none"> 1. Power supply is reporting failure (dc OK) after keyswitch is turned on, but prior to CUB power on sequence. 2. This is the first of two or more supplies reporting failure. 	Call the Response Center.

LED codes
Power on detected errors

LED	Fault	Symptoms	Corrective action
08-11	48V error NPSUL failure PWRUP=0-9	<ol style="list-style-type: none"> 1. Error occurs when 48 volt distribution falls below 42 volts during powerup state displayed. Powerup state indicates which loads are being turned on. 2. Excessive load on 48 volts due to an inadequate number of functioning 48 volt supplies or overload condition on 48V bus. 3. Possible node power supply (NPS) upper left failure. 	Call the Response Center.
12-1B	48V error NPSUR failure PWRUP=0-9	<ol style="list-style-type: none"> 1. Error occurs when 48 volt distribution falls below 42 volts during powerup state displayed. Powerup state indicates which loads are being turned on. 2. Excessive load on 48 volts due to an inadequate number of functioning 48 volt supplies or overload condition on 48V bus. 3. Possible node power supply (NPS) upper right failure. 	Call the Response Center.

LED	Fault	Symptoms	Corrective action
1C-25	48V error NPSLL failure PWRUP=0-9	<ol style="list-style-type: none"> 1. Error occurs when 48 volt distribution falls below 42 volts during powerup state displayed. Powerup state indicates which loads are being turned on. 2. Excessive load on 48 volts due to an inadequate number of functioning 48 volt supplies or overload condition on 48V bus. 3. Possible node power supply (NPS) lower left failure. 	Call the Response Center.
26-2F	48V error NPSLR failure PWRUP=0-9	<ol style="list-style-type: none"> 1. Error occurs when 48 volt distribution falls below 42 volts during powerup state displayed. Powerup state indicates which loads are being turned on. 2. Excessive load on 48 volts due to an inadequate number of functioning 48 volt supplies or overload condition on 48V bus. 3. Possible node power supply (NPS) lower right failure. 	Call the Response Center.

LED codes
Power on detected errors

LED	Fault	Symptoms	Corrective action
30-39	48V error (maintenance) no supply failure reported PWRUP=0-9	<ol style="list-style-type: none"> 1. Error occurs when 48 volt distribution falls below 42 volts during powerup state displayed. Powerup state indicates which loads are being turned on. 2. Excessive load on 48 volts due to an inadequate number of functioning 48 volt supplies or overload condition on 48V bus. 3. Possible node power supply (NPS) failure. 	Call the Response Center.
3A	48V Yo Yo error	<ol style="list-style-type: none"> 1. Core utilities board (CUB) lost and then regained 48V power without the machine being turned off or ac power failure. 2. Core utilities board (CUB) will display this error and not power on the system. 	<ul style="list-style-type: none"> • Cycle dc power to the node using the keyswitch to attempt to clear the Yo Yo bit. • Call the Response Center.
3B	MIB power fail (MIBPB)	<ol style="list-style-type: none"> 1. VDD (3.3V) error on MidPlane power board (MIBPB). 2. Midplane power fails and entire node will power down. 3. Core utilities board (CUB) still active. 	<ul style="list-style-type: none"> • Cycle dc power to the node using the keyswitch to attempt to clear the error. • Call the Response Center.
3C	Clock fail	<ul style="list-style-type: none"> • Core utilities board (CUB) monitors clock on MidPlane (MIB). 	Call the Response Center.

CUB detected memory power fail

This describes covers memory errors detected by the monitoring utilities chip (MUC) on the core utilities board after power-on.

Table 16 **CUB detects memory power fail**

LED	Fault	Symptoms	Corrective action
40	MB0L Power Fail	<ol style="list-style-type: none"> 1. 3.3V dropped below acceptable level. 2. Core utilities board (CUB) detected a power loss on reported memory board (MB). 3. Core utilities board (CUB) powers down the system 	Call the Response Center.
41	MB1L Power Fail		
42	MB2R Power Fail		
43	MB3R Power Fail		
44	MB4L Power Fail		
45	MB5L Power Fail		
46	MB6R Power Fail		
47	MB7R Power Fail		

CUB detected processor error

This section describes processor errors detected by the monitoring utilities chip (MUC) on the core utilities board after power-on.

Table 17 CUB detects processor power fail

LED	Fault	Symptoms	Corrective action
48	PB0L Power Fail	<ol style="list-style-type: none">1. 3.3V dropped below acceptable level.2. Core utilities board (CUB) detected a power loss on the reported processor board (PB).3. Core utilities board (CUB) powers down the system.	Call the Response Center.
49	PB1R Power Fail		
4A	PB2R Power Fail		
4B	PB3R Power Fail		
4C	PB4L Power Fail		
4D	PB5R Power Fail		
4E	PB6L Power Fail		
4F	PB7R Power Fail		
50	PB0R Power Fail		
51	PB1L Power Fail		
52	PB2R Power Fail		
53	PB3L Power Fail		
54	PB4R Power Fail		
55	PB5L Power Fail		
56	PB6R Power Fail		
57	PB7L Power Fail		

CUB detected I/O error

This section describes I/O errors detected by the monitoring utilities chip (MUC) on the core utilities board after power-on.

Table 18 **CUB detects I/O (IOB) power fail**

LED	Fault	Symptoms	Corrective action
58	Left Front I/O Board failure	<ol style="list-style-type: none"> 1. 3.3V or 5V dropped below acceptable level (+12V and -12V not monitored). 2. Core utilities board (CUB) detected a power loss on reported I/O board (IOB). 3. Core utilities board (CUB) powers down the system. 	Call the Response Center.
59	Left Rear I/O Board failure		
5A	Right Front I/O Board failure		
5B	Right Rear I/O Board failure		

CUB detected fan error

This section describes fan errors detected by the monitoring utilities chip (MUC) on the core utilities board after power-on.

NOTE Fan positions are referred to as viewed from the rear of the server.

Table 19 CUB detects fan power fail

LED	Fault	Symptoms	Corrective action
5C	Fan failure Upper Right	Sensor in the reported fan (as viewed from rear of system) determines fan failure.	Call the Response Center.
5D	Fan failure Upper Middle		
5E	Fan failure Upper Left		
5F	Fan failure Lower Right		
60	Fan failure Lower Middle		
61	Fan failure Lower Left		

CUB detected ambient air errors

This section describes air errors detected by the monitoring utilities chip (MUC) on the core utilities board after power-on.

Table 20 **CUB detects ambient air error**

LED	Fault	Symptoms	Corrective action
62	Ambient hot	<ol style="list-style-type: none"> 1. Ambient air too hot. 2. Core utilities board (CUB) powers down system. 3. Should have received “ambient air too warm” error 69 prior to this error. 	<ul style="list-style-type: none"> • Check site temperature. • Call the Response Center.
63	OVERTEM PMIB	<ol style="list-style-type: none"> 1. MidPlane (MIB) too hot. 2. Core utilities board (CUB) sensed overtemp on MidPlane power board (MIBPB) and powers down the system. 	<ul style="list-style-type: none"> • Check that airflow is not blocked. • Check fans. • Call the Response Center.
64	QUADRL 0	<ol style="list-style-type: none"> 1. Board overheated in Quadrant 0. 2. Core utilities board (CUB) sensed overtemp in Quadrant 0 and powers down the system. 	Call the Response Center.
65	QUADRU 1	<ol style="list-style-type: none"> 1. Board overheated in Quadrant 1. 2. Core utilities board (CUB) sensed overtemp in Quadrant 1 and powers down the system. 	Call the Response Center.
66	QUADLL 2	<ol style="list-style-type: none"> 1. Board overheated in Quadrant 2. 2. Core utilities board (CUB) sensed overtemp in Quadrant 2 and powers down the system. 	Call the Response Center.
67	QUADLU 3	<ol style="list-style-type: none"> 1. Board overheated in Quadrant 3. 2. Core utilities board (CUB) sensed overtemp in Quadrant 3 and powers down the system. 	Call the Response Center.

CUB detected hard error

This section describes hard errors detected by the monitoring utilities chip (MUC) on the core utilities board after power-on.

Table 21 **Hard error**

LED	Fault	Symptoms	Corrective action
68	Hard error (RAC) (PAC) (MAC) (TAC) (SAGA)	<ol style="list-style-type: none">1. Hard error lines to core utilities board (CUB) reported ASIC problem.2. Bit and hard error bus determine which ASIC to check	<ul style="list-style-type: none">• Read /spp/data/hard_list.• Call the Response Center.

CUB detected intake ambient air error

This section describes air intake errors detected by the monitoring utilities chip (MUC) on the core utilities board after power-on.

Table 22 **Ambient air (intake) error**

LED	Fault	Symptoms	Corrective action
69	Ambient air too warm is an environmental warning	Intake air through CUB too warm.	<ul style="list-style-type: none"> • Check site temperature and correct. • If the fault reoccurs when room temperature is within spec. call the Response Center

CUB detected dc error

This section describes dc errors detected by the monitoring utilities chip (MUC) on the core utilities board after power-on.

Table 23 **dc error**

LED	Fault	Symptoms	Corrective action
70	NPSUL failure (warning)	1. Node power supply (Viewed from Node front) failure reported. 2. Low-priority error for redundant power configurations.	Call the Response Center
71	NPSUR failure (warning)		
72	NPSLL failure (warning)		
73	NPSLR failure (warning)		

Index

Symbols

/spp directory, 4
/spp/bin, 55
/spp/data, 55
/spp/est, 56
/spp/etc, 54
/spp/firmware, 56
/spp/man, 56
/spp/scripts, 55
^E key sequence, 49

Numerics

10/100 Base T Ethernet, 12
1000 Base SX Gigabit
Ethernet, 12
32-bit applications, 118
64-bit applications, 118
712 workstation, 36, 98

A

Abaqus, 120
abnormal system shutdown, 147
accessing
I/O, 13
accounts, 37
acoustics, xviii
add terminal mux, 84, 85
applications
32-bit, 118
64-bit, 118
associated documents, xx
attention light bar, 31
auto command, OBP, 67
autoboot, 67, 126, 128
enable, 67
autoforce, 128
automatic core dump, 147
autosearch, 67, 128

B

B180L workstation, 36, 98

bcheckrc script
and file system, 128
and shutdown status, 128
BCIQ, xvii
binding
threads and processes, 123
blink, 33
boot menu
list of commands, 65
boot process output, 62
boot sequence, 60
boot variables, 127, 128
booting
I/O device on cabinet 0, 2

C

cabinet, 2
cabinet ID, 2
configurations, 18
numbering, 2, 119
cache, 11
CTI, 11, 12
cards
I/O supported, 12
physical access, 14
caution, defined, xv
ccmd, 71, 98, 100, 101
ccNUMA, 10
cmd, 54
command prompt. *see* boot menu
commands
see also utilities
autoboot, 67
blink, 33
crashconf, 147, 159
crashutil, 147
do_reset, 134
ioscan, 118
jf-ccmd_info, 53
model, 118
mpsched, 123
reboot, 132

savecrash, 147
set_complex, 46, 52, 53
shutdown, 132
sppconsole, 46
targeting nodes, 53
top, 118
complex, 98
name, 46
set_complex, 46, 52, 53
components of V-Class server, 2
configuration
node, 77–80
of HP-UX, 120, 121
utilities, 71–116
conserv, 54
console
see also SSP
assuming control, 51
changing connection, 52
commands, 49, 50
controlling remotely, 51
creating windows, 41
HP-UX, runs on, 125
starting, 45–49
commands, 45
starting from ts_config, 85
switch modes, 52
using, 45
watching, 50, 51
console port, 4
multiple-cabinet connections, 5
consolelog, 52, 55, 99
core dump, 147
Core Utility Board (CUB), 9
detected ambient air
errors, 175
detected dc error, 178
detected fan error, 174
detected hard error, 176
detected I/O error, 173
detected intake ambient air
error, 177

detected memory power
fail, 171
detected processor error, 172
core utility board (CUB), 36
CPU. *see* processor, 9
crash dump, 147
SCA, 149
crashconf, 159
crashutil, 147
creating new windows, 41
crossbar, 6
CTI (Coherent Toroidal
Interconnect), 15
cables, 15
cache memory, 11, 12
controllers, 6, 15
CTI cache, 105

D

DAT drive
eject button, 26
indicators, 25
data processing, 121
dc off, 23
dc on, 23
dcm, 55
DDS-3 DAT drive, 25
deconfigure node, 84
default passwords, 37
device files, 57
dfdutil, 98
diagnostic LAN, 4, 57, 98
multiple-cabinet connections, 5
digital apparatus statement, xvii
DIMM, 105
DIMMs, 11
direct memory access, 13
displays, 21
DMA, 13
do_reset, 55

Dual Universal Asynchronous
Receiver-Transmitter
(DUART), 99
dump, 147
configuring, 150–152
defining devices, 155–158
destination, contents, 148–149
order, 160
overview, 148
DVD-ROM drive, 24
busy indicator, 25
disk loading, 24
eject button, 25
illustrated, 24

E

eject button, 25
EMI, xvii
enable Autoboot, 67
environmental errors, 32
error codes. *see* LED errors
error indicator, 31
error information, 138
est, 55
event_log, 52, 56
event_logger, 55
Exemplar Routing Access
Controllers (ERACs), 6

F

failure information. *see also* logs
and LED errors
failures, recovery, 137
fast dump, 147
FCC, xvi
FDDI, 12
file system
and bcheckrc script, 128
firmware, 60, 97
Forbin Project, The, 47
force, 128
FORTH, 98

front panel
DAT drive, 25
fuse cautions, xix

G

gang scheduler, 122
GUI
xconfig window, 103, 104, 105
menu bar, 105
node configuration map, 106
node control panel, 108

H

hard_hist, 56
hard_logger, 55
hardware
listing, 118
topology inquiry support, 123
help, 67
Hewlett-Packard Response
Center, xxi
high leakage current, xviii
HP mode boot menu, 64
HP-UX, 36, 97, 99, 117
commands, 118
configuring, 120, 121
documentation Web site, 118
gang scheduler, 122
ioscan command, 118
model command, 118
rebooting after a system
panic, 146
shut down procedure, 132, 133
starting, 125
stopping, 130
system calls, 123
top command, 118
tuned parameter sets, 120, 121
HP-UX system
interruptions, 137
HP-UX, runs on test station, 125
HVD FWD SCSI, 12

HyperPlane Crossbar, 6

I

I/O

- controllers, 13
- listing, 118
- multiple-cabinet
 - numbering, 14
- numbering, 119
- physical access, 13
- supported cards, 12

indicator LEDs

- DAT, 25
- DVD-ROM, 24

indicators, 21

- DAT drive LEDs, 25
- dc on LED, 23
- light bar, 31

installation conditions, xix

interconnecting hardware, 6

interleaving of memory, 11

ioscan command, 118

IP address, 53, 98

IT power system, xviii

J

- jf-ccmd_info, 53
- JTAG, 53, 98

K

kernel

- configuration, 120
- threads, 122

key switch panel, 23

L

LAN

- 712/B180L, 57

launch policies, 123

LCD, 99

LCD (Liquid Crystal Display),
28

- codes, tables, 29–30
- message display line, 30
- node status line, 28
- processor status line, 28

LED errors

- core utility 01, 166
- core utility 02, 166
- core utility 03, 167
- core utility 04, 167
- core utility 05, 167
- core utility 06, 167
- core utility 07, 167
- core utility 08, 168
- core utility 09, 168
- core utility 10, 168
- core utility 11, 168
- core utility 12-1B, 168
- core utility 1C-25, 169
- core utility 26-2F, 169
- core utility 30-39, 170
- core utility 3A, 170
- core utility 3B, 170
- core utility 40, 171
- core utility 41, 171
- core utility 42, 171
- core utility 43, 171
- core utility 44, 171
- core utility 45, 171
- core utility 46, 171
- core utility 47, 171
- core utility 48, 172
- core utility 49, 172
- core utility 4A, 172
- core utility 4B, 172
- core utility 4C, 172
- core utility 4D, 172
- core utility 4E, 172
- core utility 4F, 172
- core utility 50, 172
- core utility 51, 172
- core utility 52, 172
- core utility 53, 172
- core utility 54, 172
- core utility 55, 172
- core utility 56, 172
- core utility 57, 172
- core utility 58, 173
- core utility 59, 173
- core utility 5A, 173
- core utility 5B, 173
- core utility 5C, 174
- core utility 5D, 174
- core utility 5E, 174
- core utility 5F, 174
- core utility 60, 174
- core utility 61, 174
- core utility 62, 175
- core utility 63, 175
- core utility 64, 175
- core utility 65, 175
- core utility 66, 175
- core utility 67, 175
- core utility 68, 176
- core utility 69, 177
- core utility 70, 178
- core utility 71, 178
- core utility 72, 178
- core utility 73, 178

LEDs

- attention light bar, 27
- CUB error, 32
- DC ON, 23
- libcrash, 147
- library routines
 - SCA extensions, 122
- light bar, 31
- locality domain, 122
 - binding, 123
 - launch policies, 123
- log files
 - event_log, 52
- logging in, 37
- logons, 37
- LVD Ultra2 SCSI, 12

LVM (Logical Volume Manager),
problems, 145

M

material handling
safety, xvi
media
DAT, 25
tape, 25
memory
80-bit DIMMs, 10
88-bit DIMMs, 10
board, 11
controllers, 6
CTI cache memory, 11, 12
interleaving, 11
latency, 10
numbering, 119
population, 10
supported DIMM sizes, 10
memory power fail, 171
MIB, 36
midplane, 36
migration of threads and
processes, 122
model command, 118
modem, 57
MPI
scheduling, 122
mpsched command, 123
mu-000X, 98
MUC
detected errors, 171
multiple-cabinet
cabinet IDs, 2
configurations, 18
console and diagnostic
connections, 4
CTI cable connections, 16
CTI controller, 6
I/O numbering, 14
memory configuration, 10

numbering, 119
Scalable Computing
Architecture, 1
server configurations, 117

N

NASTRAN, 120
network cache. *see* CTI cache
memory
node
see also complex
configure, 77–80
deconfigure, 84
reset, 82, 83
node routing board
MUC detected errors, 171
poweron detected errors, 166
node status line, 28
node. *see* cabinet
node_0.cfg, 55
notational conventions, xiv
note, defined, xv
NRB, 36
numbering
I/O devices, 14
of hardware components, 119
NVRAM, 127

O

OBP, 60, 71, 127, 128
commands
auto, 65, 67
autoboot, 128
autoforce, 128
autosearch, 128
boot, 65
boottimer, 65, 128
clearpim, 65
cpuconfig, 65
default, 65
display, 65
forthmode, 65

help, 69
io, 65
ls, 65
password, 65
path, 66, 128
pdt, 66
pim_info, 66
RemoteCommand, 66
reset, 66
restrict, 66
scsi, 66
search, 66
secure, 66, 128
shortcuts, 64
time, 66
version, 66
defined, 60
enabling Autoboot, 67
OLTP, 121
on/off switch, 23
Open Boot PROM (OBP), 97, 99,
109
operating system, 117
operator panel, 22
Oracle, 121
ordering
V-Class servers, 18
overview, 1

P

PA-RISC, 9
pce_util, 55
PCI
controller numbering, 13
controllers, 6
numbering, 14
physical location, 13
supported cards, 12
pcirom, 98
PDC, 60
performance problems, 139
POST, 97–102, 107, 109, 112

power, 23
powering down the system, 130
Power-On Self Test (POST), 28
private ethernet, 36
private LAN, 57
private LAN *see* diagnostic LAN
processor
 binding, 123
 numbering, 119
 PA-8200, 9
 PA-8500, 9
 PA-8600, 9
 processor agents, 6
Processor Dependent Code (PDC), 60
processor status line, 28
programming extensions
 SCA features, 122
prompt
 command, 64
pthread
 SCA features, 123
 scheduling, 122

R

radio frequency interference, xvi
Reader feedback, xxii
reboot, 132
rebooting, 146
recovering from failures, 137
remote console, 51
remove
 terminal mux, 85
report_cfg, 111
reset, 24, 134
reset node, 82, 83
Response Center, xxi
RISC processor architecture, 9
root, 37
RS-232, 99

S

safety and regulatory, xvi
 acoustics (Germany), xviii
 BCIQ (Taiwan), xvii
 digital apparatus
 statement, xvii
 EMI (European), xvii
 FCC notice, xvi
 fuse cautions, xix
 high leakage current, xviii
 installation conditions, xix
 IT power system, xviii
 material handling, xvi
 radio frequency
 interference, xvi
 VCCI, xvii
SAM utility, 120
savecore, 147
savecrash, 147
SCA, 1
 see also multiple-cabinet
 configuration
 split, 92–94
 HP-UX support, 122
 kernel configuration, 121
Scalable Coherent Interface, 15
Scalable Computing
 Architecture.*see* SCA
SCSI, 12
scub_ip address, 81, 82
 configure, 81, 82
server configurations, 18
Service Support Processor, 2, 3
 connections to V2500/V2600
 server, 4
 operations, 3
 ts_config utility, 11
 utilities board connection, 9
 xconfig utility, 11
set_complex, 46, 52, 53
shared-memory
 latency, 122
shortcuts, 64

shutdown, 132
shutdown command
 and remote systems, 130
 authorized users, 130
 considerations, 130
shutdown status and bcheckrc
 script, 128
SMP, 10
specify complex, 52
split SCA, 92–94
spp_pdc, 97
SPP_PDC (SPP Processor
 Dependent Code), 60
sppconsole, 55, 99
 commands, 49
 complex console, 39
 spy mode, 49
 starting, restarting, 45
 the command, 46
sppdsh, 71
sppuser, 37
 account, 4
 passwords, 37
 windows, 37
spy mode, 49
SSP, 41, 100
 cmd running on, 100
 configuration utilities, 71
 console window, 40
 see also console
 file system, 54
 logons, 37
 message output, 39
 message window, 40
 operation, 35–58
 root, 37
 sppuser, 37
 tcsh shell windows, 40
 test station console, 39
 windows, illustrated, 39
SSP-to-system, illustrated, 97
Stingray Core Utilities Board
 (SCUB), 98

-
- Stop-on-hard button, 109
 - supported I/O cards, 12
 - switches, 21
 - Symbios, 98
 - Symmetric Multi-Processing (SMP), 10
 - system
 - displays, 27
 - hangs, 140
 - logs, 52
 - panics, 141
 - file system problem, 144
 - interface card problem, 143
 - lan problem, 144
 - logical volume manager problem, 145
 - monitoring the system, 146
 - peripheral problem, 142
 - reboot procedure, 132, 146
 - reset, 24
 - reset procedure, 134
 - shutdown procedure, 133
 - shutdown, abnormal, 147
 - shutting down, 130
 - startup, 60
 - status, 99
 - T**
 - Tachyon Fibre Channel, 12
 - tape, 25
 - tcsh, 39
 - technical assistance
 - V-Class servers, xxi
 - terminal mux, 9
 - add/configure, 84, 85
 - remove, 85
 - test bus, 36
 - teststation *see* SSP
 - tftp, 98
 - threads
 - inquiry features, 123
 - launch policies, 123
 - scheduling, 122
 - TOC, 23, 24, 134
 - top command, 118
 - ts_config, 11, 71–96, 98
 - configuration procedures, 75–94
 - files, 95–96
 - operation, 73–75
 - starting, 72, 73
 - ttylink, 99
 - U**
 - UNIX. *see* HP-UX
 - upgrade JTAG firmware
 - JTAG, upgrade firmware, 75–77
 - USA radio frequency notice, xvi
 - using the console, 45
 - utilities
 - autoreset, 110
 - cmd, 71, 98, 100, 101
 - consolelog, 99
 - dfdutil, 98
 - est_config, 110
 - pcirom, 98
 - spp_pdc, 97
 - sppconsole, 99
 - sppdsh, 71
 - ts_config, 71–96, 98
 - xconfig, 71, 102–109
 - xsecure, 115
 - utilities board, 4, 9
 - multiple-cabinet connections, 5
 - numbering, 119
 - V**
 - V2500/V2600
 - differences, 9
 - V2500/V2600 server
 - overview, 1
 - variables, boot, 127, 128
 - VCCI, xvii
 - V-Class
 - components, 2
 - hardware configuration, 18
 - HP-UX configuration, 120
 - ordering, 18
 - Service Support Processor connection, 9
 - technical assistance, xxi
 - V2200 server, 9
 - V2250 server, 9
 - V2500/V2600 server, 1
 - volume group zero (vg00), 2
 - W**
 - warning, defined, xv
 - Web site, 18, 118
 - white paper
 - HP-UX SCA Programming and Process Management White Paper*, 124
 - windows, 39
 - workspace menu
 - V2500, 41–44
 - workstation
 - 712, 36, 98
 - B180L, 36, 98
 - differences, 98
 - X**
 - xconfig, 11, 71, 102–109
 - description, 102
 - menu bar, 105
 - node configuration, 106
 - node control panel, 108
 - window, 103, 104, 105
 - X-dimension CTI cables, 16
 - xsecure, 115
 - Y**
 - Y-dimension CTI cables, 16
-

Free Manuals Download Website

<http://myh66.com>

<http://usermanuals.us>

<http://www.somanuals.com>

<http://www.4manuals.cc>

<http://www.manual-lib.com>

<http://www.404manual.com>

<http://www.luxmanual.com>

<http://aubethermostatmanual.com>

Golf course search by state

<http://golfingnear.com>

Email search by domain

<http://emailbydomain.com>

Auto manuals search

<http://auto.somanuals.com>

TV manuals search

<http://tv.somanuals.com>